*SUCCESSIVE DISCRIMINATION TRAINING WITH
EQUATED REINFORCEMENT FREQUENCIES:
FAILURE TO OBTAIN BEHAVIORAL CONTRAST¹*

R. A. BOAKES, M. S. HALLIDAY, AND J. S. MOLE

UNIVERSITY OF SUSSEX

In two experiments, pigeons were trained on two-component multiple schedules in which responding in one component (S_1) was always maintained by a variable-interval schedule. In Experiment I, low response rates were reinforced in the second (S_2) component for six master subjects. This schedule was adjusted to equate reinforcement frequencies in the two components. These subjects were compared to yoked partners, for which reinforcement in the S_2 component was made available on a variable-interval schedule whose value was determined by the master subjects. A similar procedure was used in Experiment II, where the S_2 schedule for master subjects made reinforcers contingent on the absence of responding. No evidence was found in either experiment for a behavioral contrast effect in the S_1 component attributable to response reduction in the S_2 component. A reliable contrast effect was obtained from a group of pigeons given extinction conditions in the S_2 component, which was compared to a group maintained throughout on a multiple variable-interval schedule. The results suggest that previous indications of behavioral contrast in similar situations were probably caused by uneven reinforcement distributions or reflect uncontrolled fluctuations in response rates.

Key words: behavioral contrast, adjusted DRL schedule, adjusted DRO schedule, yoked controls, discrimination learning, multiple schedules, key pecking, pigeons

Various authors have claimed that a reduction of responding in one component of a multiple schedule is an important factor in producing positive behavioral contrast in the other component (Freeman, 1971; Mackintosh, 1975, p. 379; Terrace, 1972). Some of the strongest evidence for this view comes from studies of discrimination training under conditions in which reinforcement frequencies in the two components of a multiple schedule are equated. For example, Terrace (Experiment III, 1968) trained six pigeons to peck on a variable-interval schedule in the presence of one stimulus (S_1) and then introduced a second stimulus (S_2) that was associated with a schedule that differentially reinforced low rates of responding (DRL schedule). By continually adjusting the DRL schedule, he was able to produce a situation in which approxi-

mately equal rates of reinforcement were associated with S_1 and S_2 . The DRL schedule decreased responding in three subjects; at the same time, S_1 rates increased in these subjects. A similar procedure was used by Weisman (Experiment II, 1969), all four of whose subjects showed an increase in S_1 rates as S_2 rates declined under multiple variable-interval DRL conditions (*mult VI 1-min DRL*).

A schedule in which the delivery of reinforcement is contingent on the absence of responding for some specified period of time (DRO schedule) results in greater reductions in responding than a DRL schedule. Weisman (1970) used a DRO condition in one component of a multiple schedule, adjusting the time parameter so that the average reinforcement frequency obtained on the DRO schedule was equal to that obtained on the variable-interval schedule of the S_1 component. The results were very similar to those of his earlier DRL study.

Thus, all three studies in which a successive discrimination had been established in the absence of any difference in reinforcer frequency obtained the same result, namely that response reductions in the S_2 component are accompanied by increased responding in the S_1 com-

¹This research was supported by a grant from the Medical Research Council to the first two authors and the participation of J. S. Mole was made possible by his tenure of an MRC Research Studentship. We are grateful for the assistance of Ingrid Rose in carrying out the second experiment. Reprints may be obtained from R. A. Boakes, Laboratory of Experimental Psychology, University of Sussex, Brighton, Sussex, BN1 9QG, England.

ponent. Despite this consistency, a number of questions should be raised before concluding that response reduction is a critical factor in producing behavioral contrast.

The conclusion that Terrace and Weisman draw from their data appears to be contradicted by the results of experiments by Bloomfield (1967a), Nevin (1968), and Reynolds (1961). In these studies, DRO and DRL schedules in S_2 , while significantly suppressing rates of responding in S_2 , failed to increase response rates in S_1 . The argument advanced by Terrace against accepting these results at their face value is that in none of these experiments was any attempt made to match the frequencies of reinforcement in S_1 and S_2 : typically, the frequency of reinforcement in S_2 was considerably greater than in S_1 . It is not clear why, if reduction of responding is the critical determinant of contrast, this should matter. More important, however, one may wonder how successful Terrace and Weisman were in actually equating *local* rates of reinforcement in S_1 and S_2 . This point is best illustrated by considering the possible effect of Terrace's (1968) adjustment procedure, where the values of the DRL schedule were changed from session to session. Since response rates in the DRL component were likely to have declined within a session, it is probable that few reinforcements were delivered in the S_2 component during the early part of each session and that a large number were clumped together in the final period of a session. Consequently, contrast effects may have occurred in some subjects because the conditions in the S_2 component were effectively like extinction for a large part of each session. Even though Weisman (1969, 1970) adjusted DRL and DRO parameters before each stimulus period, it is also possible in his experiments that, though average frequencies of reinforcement were equated, the distributions were very different in the two components.

A second possible reason for the apparent discrepancy between the results of the two sets of experiments arises from differences in procedure. In Bloomfield (1967a), Nevin (1968), and Reynolds (1961) subjects were given repeated cycles of baseline and experimental conditions. Terrace (1968) and Weisman (1969, 1970), on the other hand, gave only an initial series of baseline sessions, followed by experimental sessions. They assessed the oc-

currence of contrast by comparing experimental performance in the S_1 component with that obtaining during the initial baseline sessions. Although this is a common procedure, it may not be appropriate for assessing relatively small amounts of contrast (Mackintosh, 1975, p. 369); it is possible that many of the S_1 increases reported in the latter studies were no greater than increases that might have occurred if the initial baseline conditions had been maintained.

The present experiments were undertaken to determine whether behavioral contrast occurred in such situations when the effects of possible differences in reinforcement distribution were controlled for by the use of yoked subjects. A control group was trained on a *mult VI VI* schedule throughout each experiment to assess the general increases in responding that occur with prolonged training.

EXPERIMENT I EFFECTS OF AN ADJUSTED DRL SCHEDULE

METHOD

Subjects

Twenty-four male homing pigeons, approximately 18 months old, were maintained at 80% of their free-feeding weights throughout the experiment. They had no experimental history.

Apparatus

Four standard pigeon chambers, measuring 30 by 30 by 30 cm, were housed in separate sound-attenuating cubicles. In each, a translucent response key 2 cm in diameter was mounted centrally 19 cm above floor and could be transilluminated with either red or green light. An aperture below the key gave access to a Gerbrands hopper containing polished mixed grain. A 24-V dc 1.5-W white bulb mounted centrally on the opposite wall at a height of 27 cm served as the houselight. All walls were painted matt black.

Masking noise was present throughout all sessions and experimental conditions were controlled by conventional electromechanical programming equipment.

Experimental Design

Following pretraining, there were three phases to the experiment: an initial baseline

phase in which all subjects received identical training on a multiple variable-interval variable-interval schedule, an experimental phase, and a final baseline phase, where conditions were the same as in the initial phase. At the end of the first baseline phase, subjects were assigned equally to one of four conditions, in such a way as to match baseline performance.

In the DRL condition, the schedule was *mult* VI 1-min DRL t , in which the value of t was adjusted so as to match reinforcement densities in the two components. In the YOKE condition, the schedule was *mult* VI 1-min VI x , where the availability of reinforcement in the second component was yoked to that of a DRL subject. In the EXT condition, the schedule was *mult* VI 1-min EXT and in the CONTROL condition it was *mult* VI 1-min VI 1-min; that is, conditions stayed constant in all three phases. These conditions are summarized in Table 1.

Subjects were run in six batches of four birds and during the experimental phase each batch contained a subject in each condition. Within each condition, the red light was associated with the S_1 component for half the subjects and with the S_2 component for the remaining subjects. The assignment of DRL and YOKE subjects to pigeon chambers was exactly balanced.

Procedure

Magazine training was carried out and pecking at a white key was trained in a separate chamber. Subjects were then transferred

to the experimental chamber, where two pretraining sessions were given. In the first of these sessions, the first 40 responses were each reinforced and subsequently, reinforcement was available on a *mult* VI 30-sec VI 30-sec schedule. The same schedule was maintained throughout the second of these sessions.

During pretraining and the rest of the experiment, stimulus conditions were as follows. Each session contained 31 stimulus periods of 1-min duration separated by 5-sec intervals of blackout when the response key was inoperative. S_1 and S_2 components occurred in irregular order with the constraint that no more than three S_1 or three S_2 components could occur in succession and that the S_1 component was always present during the first stimulus period of each session. No data were recorded in the first period. Reinforcement consisted of presenting the magazine for 4 sec. The house-light was on during a stimulus period, except when the magazine was operated.

In the 14 sessions of initial baseline training, pecking was reinforced on a VI 1-min schedule in both components. The variable-interval schedule was that described by Catania and Reynolds (1968) as the "Harvard Golden Tape". Median rates of responding and relative rates in the S_1 component were calculated for the final five sessions of this phase and subjects were matched on these two variables when assigned to the four conditions of the experimental phase. In all conditions, reinforcement continued to be available on the VI 1-min schedule in the S_1 component.

Table 1

Conditions and number of sessions in the three phases of both experiments. The methods used for determining the values of the t and x parameters of the DRL, DRO, and Yoked VI schedules are described in the text. Multiple schedules were used throughout and the S_1 component was always VI 1-min, so that only the S_2 schedule is shown in the table.

Phase	Experiment I		Experiment II		Both Experiments	
	DRL ($N = 6$)	YOKE ($N = 6$)	DRO ($N = 6$)	YOKE ($N = 6$)	EXT ($N = 12$)	CONTROLS ($N = 12$)
1 Initial baseline (14 sessions)	<i>mult</i> VI 1-min VI 1-min for all subjects in this first phase					
2 Experimental (12 sessions)	DRL t	VI x	DRO t	VI x	EXT	VI 1-min
3 Final (10 sessions)	VI 1-min	VI 1-min	EXT	EXT	VI 1-min	VI 1-min

For the DRL subjects, a response in the S_2 component was reinforced if at least t sec had elapsed since the previous response or since the start of the stimulus period, where t was adjusted as follows. At the end of each stimulus period, the reinforcement rates in the two components for that session were compared: if the S_2 reinforcement rate was higher than the S_1 rate, then t was increased by 2 sec, and if it was lower than the S_1 rate, t was decreased by 1 sec. The starting value of t in the first session was 4 sec and in following sessions was the final value of the previous session, less 4 sec. Limits of a maximum of 20 sec and a minimum of 3 sec were set on the value of t .

Whenever in the S_2 component t sec had passed since the last response of the DRL subject, reinforcement also became available to its YOKE partner. Thus, in the S_2 component, responding by YOKE subjects was reinforced at variable intervals determined by the matched DRL subject. For EXT subjects, no reinforcement was available in the S_2 component during the experimental phase; for CONTROL subjects, conditions remained identical to those of the previous phase.

In the final phase, the conditions for CONTROL subjects still remained unchanged, *i.e.*, mult VI 1-min VI 1-min, and all other subjects were returned to these conditions.

RESULTS

The adjustment procedure for matching reinforcement densities in the S_1 and S_2 components for the DRL subjects was reasonably effective. Mean intervals between successive reinforcers were computed separately for both components over the first six (E1 to E6) and the final six sessions (E7 to E12) of the experimental phase. These are shown for each DRL subject in Figure 1. It can be seen that, except for the final six sessions of P135 and the initial six sessions of P142, the overall densities of reinforcement in S_1 and S_2 were very close. The mean interreinforcer intervals for the YOKE subjects were in general close to those of their DRL partners. In the S_2 component, the median difference between DRL and YOKE partners in obtained interreinforcer intervals was 1 sec.

The distributions of reinforcers received by the DRL subjects are shown by the histograms in Figure 1. The variable-interval schedule produced distributions in which approxi-

mately 50 to 60% of S_1 periods contained at least one reinforcer ("Rft"), 20 to 25% contained no reinforcer, but were preceded by an S_1 period with reinforcement ("0"), 8 to 12% contained no reinforcer and were preceded by a single period without a reinforcer ("1"), and so on. It can be seen that only rarely did a run of four or more S_1 periods without reinforcement ("3" and "4+") occur. Reinforcers in the S_2 component were distributed less regularly: there was a lower percentage of periods containing at least one reinforcer ("Rft") and runs of five or more S_2 periods without a reinforcer ("4+") were relatively frequent, especially during the first six sessions.

The performance of each DRL subject, together with that of its YOKE partner, is shown in Figure 2. Discrimination between S_1 and S_2 developed clearly in four DRL subjects, was slight for P142, and nonexistent for P132. It can be seen that baseline response rates were well matched in the pairs of DRL and YOKE subjects and that the DRL schedule was effective in reducing response rates in the S_2 component, relative to both the prior baseline level and to the S_2 rates of the YOKE partner, in all pairs except, again, P132 and P139. Despite the reduction of S_2 responding in the majority of DRL subjects, their S_1 rates were

Fig. 1. Experiment I: mean interreinforcement intervals and distribution of reinforcement averages over the first (E1 to E6) and second (E7 to E12) halves of the experimental phase for each DRL subject. Columns labelled "Rft" indicate the percentage of stimulus periods containing at least one reinforcer, those labelled "0" the percentage of periods containing no reinforcer and preceded by a period containing at least one reinforcer, those labelled "1" periods containing no reinforcer and preceded by one period containing no reinforcer, and so on.

Fig. 2. Experiment I: rates of responding during the last five sessions of initial baseline training and all subsequent sessions for each pair of DRL and YOKE subjects.

no higher than those of their YOKE partners and, particularly in the initial few sessions, tended to be considerably lower.

The median response rates in the final five sessions of the experimental phase were compared to those of the final five sessions of the preceding baseline condition for each subject. The change in rates of responding by DRL subjects in the S_2 component, a median decrease of 27 responses per minute (a median percentage decrease of 43% of the baseline level), was reliably greater (*t*-test, one-tail, $p < 0.025$) than the change, a median increase of 10 responses per minute (24% increase over baseline level), shown by the YOKE subjects. Rates of responding by the DRL subjects in the S_1 component increased by 10 responses

per minute (14% increase) and this was not reliably different (*t*-test, $p > 0.10$) from the median increase of 28 responses per minute (46% increase) shown by the YOKE subjects in this component.

As can be seen in Figure 2, in the final phase of the experiment the re-instatement of baseline conditions increased S_2 response rates of the DRL subjects. It did not produce any consistent changes in the S_1 response rates of the DRL subjects, but five of the six YOKE subjects decreased their S_1 rates. The median decrease in S_1 rate between the last five sessions of the experimental phase and the last five sessions of the final phase was seven responses per minute (11%) for the YOKE subjects and zero for the DRL subjects. This

difference between the groups failed to reach significance (t -test, two-tail, $p > 0.05$). Of the four YOKE subjects whose S_1 rates increased markedly (by more than 20 responses per minute) during the experimental phase only three, P121, P134, and P140, showed a subsequent decline in the final baseline phase.

The results from the EXT and CONTROL subjects are considered here, together with those from subjects given identical conditions in the second experiment, which is described below. The median response rates of the 12 EXT subjects in the two experiments are shown for each session in Figure 3. A typical behavioral contrast effect is seen, in that response rates increased rapidly from an initial baseline level of fewer than 50 responses per minute to almost 80 responses per minute by the end of the second phase (*mult* VI 1-min EXT) and declined during the final phase when returned to the *mult* VI 1-min VI 1-min schedule. An increase in S_1 rates during the second phase occurred for 11 subjects; only eight showed a decline from the end of the second phase to the start of the final phase.

The evidence for behavioral contrast from the results shown in Figure 3 for the EXT subjects looks less impressive when compared to those from the 12 CONTROL subjects main-

tained on the *mult* VI 1-min VI 1-min schedule throughout. The equivalent data for these subjects are also shown in Figure 3. The groups were compared by determining for each subject the difference between the median S_1 rate over the final five sessions of the initial baseline phase and the median over the final five sessions of the second phase. Though the median increase of 31 responses per minute showed by the EXT subjects was reliably greater (Mann-Whitney, one-tail, $p < 0.05$) than the increase of six responses per minute shown by the control subjects, there was considerable individual variability, as seen in Figure 4. Thus, in the EXT group, four subjects showed only moderate increases, of fewer than 20 responses per minute, and for one subject, the S_1 rate *decreased* by more than 20 responses per minute. Similarly, though the S_1 rates of the majority of CONTROL subjects changed little during the second phase, increases of more than 20 responses per minute occurred for two subjects and a further subject showed a greater increase in rate than any of the EXT subjects. The results from other studies, which are also shown in Figure 4, are discussed below.

Changes in response rate occurring during the transition from the second to the final

Fig. 3. Median rates of responding in S_1 periods for EXT subjects ($N = 12$) and CONTROL subjects ($N = 12$) in both experiments over all experimental sessions.

Fig. 4. Distributions of increases in S_1 response rates from the final five sessions of the initial baseline phase to the final five sessions of the experimental phase for the EXT and CONTROL subjects of both experiments.

phase were similarly assessed by comparing for each subject the median rate over the final five sessions of the second phase with the median for the following five sessions. The median decrease of 10 responses per minute in the EXT group was not reliably different from the median of zero responses per minute in the CONTROL group (Mann-Whitney, $p > 0.10$). Even when this comparison was restricted to the six EXT subjects whose S_1 rates had increased by more than 30 responses per minute in the second phase, their subsequent decreases in response rate (median 16 responses per minute) was not reliably different from that shown by the six subjects (median four responses per minute) with the largest increase in S_1 rates among the CONTROL subjects during the second phase.

The main reason for including these two groups in the present study was to provide norms for assessing the performance of the DRL subjects in Experiment 1 and the DRO subjects in Experiment 2. Because of the considerable variability of the DRL subjects, overall the changes in their S_1 rates did not differ reliably from either those of the CONTROL group or of the EXT group.

EXPERIMENT II EFFECTS OF AN ADJUSTED DRO SCHEDULE

Experiment I showed that an adjusted DRL schedule produced reinforcement distributions different from those obtained on a variable-interval schedule. It seemed likely that this would also be the case when responding was reduced by an adjusted DRO schedule. If

so, then possible behavioral contrast effects in multiple schedules involving such a component (Weisman, 1970) should also be assessed in comparison with yoked control subjects. Such a comparison was made in this second experiment.

Furthermore, since DRO schedules lead to a greater and more rapid reduction of responding than DRL schedules, any effect of response reduction *per se* should be more apparent here than in Experiment I.

METHOD

Subjects

Twenty-four male homing pigeons, approximately 18 months old, were maintained at 80% of their free-feeding weights throughout the experiment. They had no experimental history.

Apparatus

As in Experiment I.

Experimental Design

The form of the first two phases was similar to that of Experiment I, in that after all subjects had been given initial baseline training, they were assigned to four equal groups in such a way as to match performance at the end of the initial baseline phase. Two of these groups, EXT and CONTROL, were given identical conditions to those of the equivalent groups in Experiment I. In the DRO condition, the schedule in the second phase was *mult VI 1-min DRO t*, where, as with the previous DRL condition, the value of t was adjusted so as to maintain equal densities of reinforcement in the two components. In the YOKE condition, the schedule was *mult VI 1-min VI x*, where the availability of reinforcement in the S_2 component was determined by the behavior of a subject's yoked DRO partner.

The only change in design from Experiment I was in the third phase, where the schedule for all DRO and YOKE subjects was *mult VI 1-min EXT*. EXT and CONTROL subjects were returned to baseline conditions in this phase, as in Experiment I. The sequence of conditions for each of the four groups is summarized in Table 1.

Subjects were run in six batches of four birds and during the second phase, each batch contained a subject in each condition. Within

each group, the red light was associated with the S_1 component for half the subjects and with the S_2 component for the other half. The assignment of DRO and YOKE subjects to pigeon chambers was exactly balanced.

Procedure

For all subjects, the procedures for pre-training, initial baseline training, and assignment to conditions in subsequent phases were exactly as in Experiment I. Similarly, the temporal parameters and stimulus conditions of the experimental sessions remained as before throughout the experiment.

In the second phase, under DRO conditions a reinforcer was delivered in the S_2 component as soon as t sec had elapsed since the previous response or since the start of the stimulus period, where t was adjusted as follows. Reinforcement rates in the two components were compared at the end of each stimulus period: if the S_2 reinforcement rate was higher than the S_1 rate, then t was increased by 3 sec; if it was lower than S_1 rate, t was decreased by 2 sec. The starting value of t in the first session was 4 sec and in following sessions was the final value of the previous session, less 4 sec. Limits of a maximum value of 32 sec and a minimum value of 3 sec were set on t . This is a more rapid adjustment procedure than that used before. This change was made because it seemed likely that the magnitude and speed of changes in response rates would be greater with the DRO condition than with the previous DRL condition.

Whenever in the S_2 component a reinforcer was delivered to a DRO subject, the next response by its YOKE partner was reinforced, providing that this response occurred within the same stimulus period.

RESULTS

The results of the adjustment procedure for matching overall reinforcer densities in the S_1 and S_2 components for the DRO subjects can be seen in Figure 5, which displays data on mean interreinforcer intervals and the distribution of reinforcers in the same way as Figure 1. The procedure was not quite so effective as that of Experiment I, in that with some subjects, e.g., P193 and P178, it allowed the frequency of reinforcers in the S_2 component to exceed that in the S_1 component. However, the greatest discrepancy between S_1 and S_2

reinforcer densities occurred with P218, which began to respond slowly and erratically in both components; consequently, it missed many reinforcers from the S_1 variable-interval schedule, while reinforcers were regularly delivered on the S_2 DRO schedule.

The mean interreinforcer intervals of the YOKE subjects were close to those of their partners, as in the previous experiment. Since in the S_2 component the delivery of a reinforcer to a DRO subject occurred immediately t sec had elapsed, whereas for its yoked partner reinforcement was contingent on the next response, the interreinforcer intervals were consistently longer for the latter subjects. However, the median difference was only 3.5 sec and the largest difference was 6.5 sec.

As can be seen from Figure 5, the distributions of reinforcers determined by the adjusted DRO schedule in the S_2 component were in all but two cases very different from the distributions generated by the S_1 variable-interval schedule. Furthermore, the form of the distributions changed considerably as discrimination training continued. For all subjects except P178, reinforcers were distributed very unevenly during the first few sessions, in that long runs of stimulus periods without a reinforcer ("4+"), and consequently stimulus periods with many reinforcers, were relatively

Fig. 5. Experiment II: mean interreinforcement intervals and distribution of reinforcement averaged over the first (E1 to E6) and second (E7 to E12) halves of the experimental phase for each DRO subject. The columns are labelled as in Figure 1.

frequent. In later sessions, when little responding occurred during S_2 periods, the distribution of reinforcers became more even, with a reinforcer occurring in 70 to 80% of the S_2 stimulus periods.

Rates of responding in both components by each pair of DRO and YOKE subjects are shown in Figure 6. Five DRO subjects showed a clear discrimination in that, while responding was maintained in the S_1 component, it had almost ceased in the S_2 component by the end of the second phase. The exception, as noted above, was P218, which in many sessions failed to respond in either component. Even though the reduction in S_2 response rates of DRO subjects was much more pronounced than for the DRL subjects of Experiment I, for only one subject, P193, was there a clear increase of S_1 rates above the prior baseline level and for only one pair, P207 and P220,

was the rate in the S_1 component consistently higher for the DRO bird than for its YOKE partner.

Quantitative comparisons were made in the same way as in Experiment I. In the S_2 component, rates of responding by the DRO subjects showed a median decrease of 41 responses per minute (95% decrease) from their baseline level by the final five sessions of the second phase, which was reliably different (t-test, one-tail, $p < 0.025$) from the median increase of one response per minute (2% increase) in the response rates of the YOKE subjects over the same period. In the S_1 component, the comparable figures were a median decrease of three responses per minute (7% decrease) for the DRO subjects, which did not differ reliably (t-test, one-tail, $p > 0.10$) from the increase of two responses per minute (6%) increase of the YOKE subjects.

Fig. 6. Experiment II: rates of responding during the last five sessions of initial baseline training and all subsequent sessions for each pair of DRO and YOKE subjects.

In the final phase, in which both DRO and YOKE subjects were in extinction in S_2 , S_1 response rates increased above the level of the previous phase for four DRO subjects and for all YOKE subjects except P220. As seen in Figure 6, this change in S_1 rates tended to be smaller for DRO subjects, whose median increase was nine responses per minute (19%), than for YOKE subjects, whose median increase was 19 responses per minute (36%), though this difference was not significant (t -test, $p > 0.10$).

The performance of the YOKE subjects in the second phase differed from that of the comparable subjects in Experiment I, as can be seen from a comparison of Figures 2 and 6: the subjects yoked to DRL partners tended to respond faster in the S_1 component than in the S_2 component, whereas the opposite occurred with the YOKE subjects. In both experiments, the YOKE subjects' rates of responding in the S_1 component, relative to their rates in the S_2 component ($P_1/P_1 + P_2$), were calculated for the last five sessions of the second phase of the experiment. The median value of this measure was 0.53 for the YOKE subjects in Experiment I, which is reliably greater (Mann-Whitney, two-tail, $p < 0.05$) than the same measure for the present YOKE subjects (median of 0.44).

Fig. 7. Relative rates of responding in S_1 periods ($P_1/P_1 + P_2$) as a function of relative rates of reinforcement ($R_1/R_1 + R_2$) for the YOKE subjects in both experiments.

This difference could have arisen because of different distributions of reinforcement in the S_2 component in the two experiments or because of a difference in the overall rates of reinforcement. The latter seemed likely to be an important variable, since over the same period the relative rates of reinforcement in the S_1 component ($R_1/R_1 + R_2$) for the YOKE subjects of Experiment I were higher (Mann-Whitney, two-tail, $p < 0.05$) than those of Experiment II. Over all 12 YOKE subjects, a high correlation ($r = 0.73$, $p < 0.025$) was found between relative response rate and relative reinforcement rate. These data are shown in Figure 7, together with the regression equation.

As with the DRL subjects of Experiment I, the data from the DRO subjects were compared with those from the EXT and CONTROL subjects. The median decrease in the S_1 response rates of the DRO subjects was reliably different from the increase in S_1 rates shown by the EXT subjects (Mann-Whitney, one-tail, $p < 0.025$), but did not differ significantly from the CONTROL group.

DISCUSSION

At the beginning of the paper, we noted an apparent discrepancy arising from research on the effects on responding in the unchanged component of a multiple schedule (S_1 component) of reducing response rates in the second component (S_2 component) by introducing either a DRL or DRO schedule. One set of studies (Bloomfield, 1967a; Nevin, 1968; Reynolds, 1967), implies that any changes in S_1 rates are attributable to changes of reinforcement density in the S_2 component, independent of any changes of response rate in this component. In a second set of experiments, designed to evaluate the effects of S_2 response reduction when overall reinforcement densities were equated (Terrace, 1968; and Weisman, 1969 and 1970), increases in S_1 response rates were obtained that were attributed to the reduction of S_2 responding.

We suggested two possible reasons for this apparent discrepancy. One was that the two sets used different procedures, and therefore effectively different criteria for assessing possible contrast effects. The other was that the procedures used by Terrace and Weisman for equating overall reinforcement densities might

well have produced very different distributions of reinforcement in the two components, so that any apparent contrast effect could be produced by this difference rather than by the reduction of responding in the S_2 component. The present research tends to confirm both suggestions. The data from the EXT and CONTROL subjects imply that the criterion used by Terrace and Weisman for classifying increases in S_1 rates as instances of behavioral contrast, namely comparison with a single prior baseline level of responding, is inappropriately low compared with the criterion implicit in the procedures of Bloomfield, Nevin, and Reynolds, where repeated cycles of baseline and experimental conditions were given. This point is discussed in more detail below. Second, it was found that adjustment of both DRL and DRO schedules, so as to equate average reinforcement density in the two components, produced marked differences in the distributions of reinforcement. When such differences were controlled for by comparison with the yoked control subjects, it was found that, though response rates in the S_2 components were reduced to a significantly greater extent in DRL and DRO subjects than in their YOKE partners, there was not even a tendency for S_1 rates to be higher in the DRL and DRO subjects.

Nevertheless, it might be claimed that the present experiments failed to replicate Terrace (1968) and Weisman (1969, 1970) and therefore that the present analysis bears only indirectly on their results. At first sight, this has some plausibility, in that the evidence for the occurrence of behavioral contrast in the DRL subjects of Experiment I was rather slight and in the DRO subjects of Experiment II almost nonexistent. In contrast, both Terrace and Weisman claim to have obtained contrast under comparable conditions. However, closer examination shows that the present results are in fact very comparable to the earlier ones. Similar analyses to those employed above were carried out on the data from these earlier studies. For the six DRL subjects of Terrace (1968), the median increase in S_1 responding was approximately eight responses per minute (12% increase) above the prior baseline level and for the four DRL subjects in Weisman (1969) 20 responses per minute (34% increase). In Experiment I, the comparable figures for the six DRL subjects were

10 responses per minute (14% increase). Other comparisons also show a great deal of similarity in the results obtained from the three experiments and; since all indicate considerable individual variability among the small number of subjects employed, they suggest strongly that little weight should be attached to apparent differences.

A similar comparison between the performance of the six DRO subjects of Experiment II and that of the four subjects in Weisman (1970) suggests a similar conclusion. However, though there is no significant difference between the two groups, it is noticeable that, whereas three of four of Weisman's subjects showed clear increases in S_1 rates, only one of the present six subjects did so. It is probably pertinent that in Experiment II, reinforcement rates in the S_2 component tended to be somewhat lower than in the S_1 component, whereas the opposite was true in Weisman's study. This raises a further question of comparison between the studies: how well did the present adjustment procedures succeed in equating reinforcement densities as compared to the less formal approach employed in the earlier experiments?

The use of an explicit schedule for adjusting DRL and DRO parameters, as in the present experiments, has certain advantages over more informal procedures. For example, it allows complete automatic control and exact replicability of such experiments. At the same time, it is potentially less flexible and consequently may be less effective in matching the average reinforcement density of the variable-interval component. Terrace (1968) did not present detailed results relevant to this question, but reported that in general, reinforcement densities were somewhat higher during the DRL component. With the one exception noted above, the schedule used to adjust the DRL parameter in Experiment I was almost as successful in matching reinforcement densities as was Weisman (1969), who did provide details of interreinforcement intervals.

A similar comparison between the effects of the schedule used to adjust the DRO parameter in Experiment II and by Weisman (1970) showed that this schedule was less effective. It seemed subsequently that minor changes, such as raising the maximum limit on the time interval, might well have improved the schedule. Thus, it appears quite feasible to use com-

pletely automated control of such experiments and obtain as close a matching of reinforcement densities as that obtained using subjective judgement.

The extent to which the present procedures were successful in equating reinforcement densities is only indirectly related to the principal concern of the two experiments, which was the comparison between the DRL or DRO subjects and their yoked partners. Since within each pair both average density and distribution of reinforcement were equated, any possible effect of S_2 response reduction in giving rise to contrast should be detectable in terms of a higher S_1 response rate in the master subject. As we have seen, no such effect was detected. In Experiment II, where the amount of S_2 response reduction was greater than in Experiment I, the tendency was for S_1 rates to be lower for the DRO subjects than for the YOKE subjects.

There is still the question of whether the unusual distribution of reinforcement occurring in the S_2 component had any effect on S_1 responding. In this context, it is interesting that the results for the YOKE subjects, as shown in Figure 7, suggest that reinforcement distribution did have important effects on the interactions between the two components. Matching of relative response rates and relative reinforcement rates is not normally observed in multiple variable-interval schedules, unless extremely short component durations are used (Shimp and Wheatley, 1971; Todorov, 1972). Although the data on such interactions with component lengths of 1 min in Shimp and Wheatley (1971) are very limited, the overall results of their parametric study and that of Todorov (1972) indicate that the function relating relative response rates to relative reinforcement rates for such schedules normally has a slope of considerably less than 45° . The data from Reynolds (1963) for 3-min components, as presented by Herrnstein (1970), support this. In all the above studies, the variable-interval schedules in the two main components were constructed with similar distributions of reinforcements. Consequently, the slope of the function in Figure 7 suggests that the performance of the YOKE subjects was not only sensitive to average reinforcement rates but was also affected by the distribution in the S_2 component. For example, the S_1 rates of P121, P134, and P141 were

considerably higher than would be expected if reinforcement had been available on a standard variable-interval schedule in the S_2 component.

Distributions of reinforcement can be roughly ordered from highly regular to highly irregular. A highly irregular distribution is one in which both very long and very short intervals occur relatively frequently. In terms of the representation of reinforcement distributions used in Figures 1 and 5, such distributions would be indicated by a U-shaped pattern. As is seen from these figures, many of the obtained distributions were of this form, particularly in the initial half of the second phase (E1 to E6) in Experiment I. At the other extreme is a schedule containing neither very long nor very short intervals, which would produce a high frequency of stimulus periods containing one reinforcer, as shown by the "Rft" column in Figures 1 and 5. Many of the distributions for the second half of the second phase (E7 to E12) in Experiment II were of this form. The distribution produced by a standard variable-interval schedule is between these extremes.

A number of factors made it difficult to isolate any possible effects of type of distribution from the present data. These factors included the confounding of distribution differences and differences in relative rates of reinforcement, as shown in Figure 7, and the fact that for many subjects the distributions changed markedly during the experiment phase. Consequently, the failure of various *post hoc* analyses to reveal significant correlations between S_1 responding and various measures of the obtained distributions in the S_2 component was not surprising. Nevertheless, certain aspects of the data suggested what the outcome of an experiment explicitly designed to investigate such effects might be. For example, considering just the first half of the experimental phase (E1 to E6) the two pairs in Experiment I showing the largest apparent contrast effects (DRL:P135/YOKE:P121 and DRL:P131/YOKE:P122 in Figure 2) had the highest frequency of long runs of S_2 periods without reinforcement, while overall reinforcement densities in the two components matched well (see Figure 1). Observations of this kind suggested that, with reinforcement densities held constant, the direct manipulation of the S_2 distribution would lead to increases in S_1 re-

sponse rates with a very irregular distribution and decreases with a very regular distribution in the S_2 component.

Returning to the second theme of this paper, the question of what criteria to apply in classifying increases in response rate as instances of behavioral contrast, we turn now to the data from the EXT and CONTROL subjects. Though a significantly greater increase in S_1 responding was shown by the EXT subjects in comparison to the CONTROL subjects, there appeared to be an unusual amount of variability in both groups. However, a detailed comparison with other published data suggested that such variability is in fact typical. To make such a comparison, we searched for all studies in this Journal, based on research in other laboratories, that fulfilled the following criteria: (a) individual data were presented; (b) *mult* VI 1-min EXT conditions were introduced after at least 14 sessions of training on a *mult* VI 1-min VI 1-min schedule; (c) there was no experimental history before initial baseline training; and (d) no correction procedure was employed. This search yielded 10 subjects, three from Bloomfield (1967*b*) and seven from Richards (1972), whose data are plotted in Figure 4 and are seen to be as variable as those of the present EXT subjects. Similar variability is seen in studies that failed some of the above criteria, as may be seen, for example, from examination of the results of Dukhayyil and Lyons (1973), Rilling, Askew, Ahlskog, and Kramer (1969), and Weisman and Palmer (1969).

Our results allow a comparison of three possible criteria for assessing contrast effects. The most widely used criterion, that of an increase above a prior baseline level, was shown to be inadequate by the results for the CONTROL subjects, as can be seen in Figure 4. Because there were only 14 sessions in the initial baseline phase of the experiments, as in Terrace (1968), it is possible that this criterion would have been more satisfactory had 20 sessions of such training been given, as in Weisman's experiments (1969, 1970). Examination of the data from the CONTROL subjects did not confirm this possibility: when the median response rate over the sixteenth to twentieth *mult* VI 1-min VI 1-min sessions was compared to that for the following five sessions for each subject, increases were shown by six subjects, of which one exceeded 18 responses per

minute and one 50 responses per minute. A similar comparison suggested that 26 sessions of baseline training would have been no more satisfactory: in this case, increases were shown by seven subjects, of which three exceeded 15 responses per minute.

A great number of experiments on behavioral contrast, including earlier studies from this laboratory (*e.g.*, Halliday and Boakes, 1971), have used a small number of subjects and departure from an initial baseline level as the criterion for the effect. The implicit assumptions behind such an experimental approach are that individual variability is relatively small under *mult* VI EXT conditions and that in the absence of any change of conditions only minor fluctuations in response rate occur on a *mult* VI VI schedule once a dozen or so sessions have been given. The present results show that these assumptions are ill founded. It is possible that the degree of variability shown by the CONTROL subjects was unusual. However, the comparable data of Gonzalez and Champlin (1974)—authors who make a similar point about the proper control procedures for studying behavioral contrast—suggest that this was not the case.

The second within-subject criterion, that of a decrease in response rate on return to baseline conditions, also appears to be inadequate for the present kind of situation. A similar conclusion was reached by Bloomfield (1967*b*), although his results suggest that within-subject criteria become more satisfactory as cycles of *mult* VI EXT and *mult* VI VI conditions are repeated. It thus seems clear that the inclusion of "no change" groups is necessary for experiments that do not involve repeating conditions. Until it becomes standard practice to include such controls, a number of issues, such as the debate over the disappearance of contrast with extended training (*e.g.*, Hearst, 1971; Selekman, 1973), are likely to remain unsolved.

In conclusion, the present experiments provide no evidence for supposing that response reduction by such means in multiple schedules, where average reinforcement frequencies in the two components are equated, produces an increase in response rate in the other component. Where such increases have been reported, it seems likely that they were caused by an uneven distribution of reinforcements or

that they reflect chance fluctuations in response rates.

REFERENCES

- Bloomfield, T. M. Behavioral contrast and relative reinforcement in two multiple schedules. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 151-158. (a)
- Bloomfield, T. M. Some temporal properties of behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 159-164. (b)
- Catania, A. C. and Reynolds, G. S. A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 327-383.
- Dukhayil, A. and Lyons, J. E. The effect of overtraining on behavioral contrast and the peak-shift. *Journal of the Experimental Analysis of Behavior*, 1973, 20, 253-263.
- Freeman, B. J. Behavioral contrast: reinforcement frequency or response suppression? *Psychological Bulletin*, 1971, 75, 347-356.
- Gonzalez, R. C. and Champlin, G. Positive behavioral contrast, negative simultaneous contrast and their relation to frustration in pigeons. *Journal of Comparative and Physiological Psychology*, 1974, 87, 173-187.
- Halliday, M. S. and Boakes, R. A. Behavioral contrast and response independent reinforcement. *Journal of the Experimental Analysis of Behavior*, 1971, 16, 429-434.
- Hearst, E. Contrast and stimulus generalization following prolonged discrimination training. *Journal of the Experimental Analysis of Behavior*, 1971, 15, 355-363.
- Herrnstein, R. J. On the law of effect. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 243-266.
- Mackintosh, N. J. *The psychology of animal learning*. London: Academic Press, 1975.
- Nevin, J. A. Differential reinforcement and stimulus control of not responding. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 715-726.
- Reynolds, G. S. Behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1961, 4, 57-71.
- Reynolds, G. S. Some limitations on behavioral contrast and induction during successive discrimination. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 131-139.
- Richards, R. W. Reinforcement delay: some effects on behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1972, 17, 381-394.
- Rilling, M., Askew, H. R., Ahlskog, J. E., and Kramer, T. J. Aversive properties of the negative stimulus in a successive discrimination. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 917-932.
- Selekman, W. Behavioral contrast and inhibitory stimulus control as related to extended training. *Journal of the Experimental Analysis of Behavior*, 1973, 20, 245-252.
- Shimp, C. P. and Wheatley, K. L. Matching to relative reinforcement frequency in multiple schedules with a short component duration. *Journal of the Experimental Analysis of Behavior*, 1971, 15, 205-210.
- Terrace, H. S. Discrimination learning, the peak-shift, and behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 727-741.
- Terrace, H. S. Conditioned inhibition in successive discrimination learning. In R. A. Boakes and M. S. Halliday (Eds.), *Inhibition and learning*. London: Academic Press, 1972, Pp. 99-119.
- Todorov, J. C. Component duration and relative response rates in multiple schedules. *Journal of the Experimental Analysis of Behavior*, 1972, 17, 45-49.
- Weisman, R. G. Some determinants of inhibitory stimulus control. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 443-450.
- Weisman, R. G. Factors influencing inhibitory stimulus control: differential reinforcement of other behavior during discrimination training. *Journal of the Experimental Analysis of Behavior*, 1970, 14, 87-91.
- Weisman, R. G. and Palmer, J. A. Factors influencing inhibitory stimulus control: discrimination training and prior non-differential reinforcement. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 229-237.

Received 21 October 1974.

(Final Acceptance 2 March 1976.)