ELSEVIER

Reexamining the frustration effect in rats: Aftereffects of surprising reinforcement and nonreinforcement[☆]

Steven C. Stout, Robert L. Boughner, and Mauricio R. Papini[*]

Department of Psychology, Texas Christian University, P.O. Box 298920, Fort Worth, TX 76129, USA

Abstract

The reinforcement-omission effect (ROE; also called frustration effect), or greater response strength immediately after nonreinforcement (N) than reinforcement (R), has been traditionally interpreted in terms of one of two factors: transient facilitation after N induced by primary frustration or transient suppression after R induced by postconsummatory processes. Three instrumental lever-pressing experiments with rats demonstrated that the ROE can be caused by either factor in isolation, or by both acting simultaneously. The distribution of trials and the interval between N or R and the target response determine which factor would cause the ROE. Both aftereffects decay in time, but the after-N process decays at a slower rate than the after-R factor.
© 2003 Elsevier Inc. All rights reserved.

In 1952 Amsel and Roussel reported a runway experiment that demonstrated (or so they thought) that the surprising omission of an appetitive reinforcer was followed by behavioral invigoration. In that experiment, rats received training in a double

[*] Corresponding author. Fax: 1-817-257-7681.
E-mail address: M.Papini@tcu.edu (M.R. Papini).

runway apparatus in which two runways were connected in series. Following a phase of continuous reinforcement, the first goal box was shifted to a partial reinforcement schedule in which the occurrence of reinforcement (R) and nonreinforcement (N) was unpredictable (i.e., surprising). The second goal box was reinforced throughout the experiment. The basic observation was that rats ran faster in the second alley after N than after R in the first alley. Amsel and Roussel (1952) interpreted this apparent response facilitation after N as reflecting emotional invigoration of behavior induced by surprising nonreward (Papini & Dudley, 1997). They claimed that a negative affective state of primary frustration triggered by surprising N and with drive-inducing properties was responsible for the invigoration of the second instrumental response. Their article generated a substantial amount of research that, in our opinion and that of others (cf. Amsel, 1992; Couvillon, Nagrampa, & Bitterman, 1994; Scull, 1973), did not provide conclusive evidence in support of Amsel and Roussel's claim. Following Amsel and Roussel's interpretation, this phenomenon was labeled the frustration effect, but the inconclusive results of the research that followed seemed to suggest that the frustration effect does not necessarily reflect the effect of frustration. Here we adopt the more neutral label of reinforcement-omission effect (ROE; Kello, 1972) to refer to a learning phenomenon defined as greater response strength (e.g., higher rate, shorter latency, etc.) immediately following N trials than immediately following R trials in a partial reinforcement situation.

A basic question about the ROE is the extent to which it depends on the surprisingness of the reward omission. Amsel and Ward (1965) addressed this issue by adding discriminative stimuli to the first runway of a double-runway apparatus that signaled R and N outcomes (i.e., an A+/B- discrimination), while keeping the second runway under continuous reinforcement. In this manner, the surprisingness of N trials was reduced over the course of training as the rats learned to discriminate A from B. Responding in the second runway changed in a manner consistent with a surprising nonreward account of the ROE (Amsel, 1962). First, early in acquisition running was slow and no ROE was observed. Second, the ROE emerged with the development of rapid approach to both A and B. Third, as running speed declined on B-trials and the behavioral discrimination progressed, the ROE also declined and eventually disappeared. This sequence was observed in individual animals, as well as in the group. Hug (1970) replicated the same sequence using a single alternation schedule of R and N trials, instead of external discriminative stimuli. Similarly, Terrace (1972) reported the dissipation of wing flapping and other affective reactions in pigeons following extended discrimination learning, which he interpreted in much the same way: Once nonreward ceases to be surprising, it no longer elicits affective aftereffects.

Additional evidence for the influence of surprising nonreward on the ROE was provided by experiments in which the effective degree of reward loss was manipulated by variables affecting the strength of the instrumental behavior. Some such manipulations included reinforcer magnitude (Daly, 1968), amount of training prior to the shift to partial reinforcement (Amsel & Ward, 1965), similarity of cues in the first and second goal boxes (Amsel & Hancock, 1957), level of deprivation (McHose & Ludvigson, 1964), and incomplete reductions in reinforcer magnitude (Bower,

1962). A review by Scull (1973) concluded that many (but not all) of these manipulations produced effects that were consistent with Amsel's (1958, 1992) interpretation of the ROE. Also consistent with such interpretation is the finding that lesions of the amygdala, a structure known to be involved in fear conditioning and arousal (Kapp, Whalen, Supple, & Pascoe, 1992), eliminate the ROE (Henke, 1977).

Despite seemingly extensive supportive evidence, Seward, Pereboom, Butler, and Jones (1957) noted a simple problem with defining the ROE as a stronger response after N than after R: it could just as easily reflect response suppression after R, instead of response facilitation after N. They suggested that food consumption induces a transient decrease in hunger motivation that does not depend upon reward expectancy. This idea, briefly discussed also by Amsel and Roussel (1952), has figured prominently in theories of instrumental behavior (e.g., Catania, 1973), and is a potential confound that deserves careful consideration.

Wagner (1959) introduced a commonly cited control for postconsummatory response suppression. Initially, one group of rats was continuously reinforced in the first goal box, while a second group was never reinforced in the first goal box. In the following phase, the continuously reinforced group was shifted to partial reinforcement in the first goal box, while the other group continued to receive no reward. Both groups were consistently reinforced in the second goal box, as usual. Second-alley running speeds of the partially reinforced animals were higher after N than after R trials, providing conventional evidence for the ROE. Moreover, the partially reinforced rats also ran faster in the second alley following N trials than did the rats in the never-rewarded control group, which, in turn, ran at about the same speed as the partial animals after R. The main problem with this "Wagner control" is that it has not always yielded evidence of response facilitation (see Scull, 1973). Despite criticism, Wagner's (1959) experiment reinforced the interpretation that the ROE stems from response invigoration following surprising N (see Mackintosh, 1974). Most subsequent researchers dispensed entirely with control conditions and simply used the after-N vs. after-R, within-subject comparison to assess the ROE, citing Wagner's results as a justification. The lack of control conditions implies that the source of the ROE in most of the published literature is unresolved. The fact remains that the Amsel–Roussel within-subject design can never logically exclude a postconsummatory aftereffect on its own.

There seem to be at least two techniques to control for response suppression after R as a source of the ROE. One technique involves varying reward expectancy between groups while not confounding other important parameters of training, such as reinforcement frequency. In this case, behavior following expected nonreward is compared with behavior following surprising nonreward. The results obtained with this technique have been ambiguous. Dudley and Papini (1995, 1997) found results consistent with response invigoration after N, whereas Scull, Davies, and Amsel (1970) found results consistent with response suppression after R. Stout (2001) noted that these experiments were different in terms of the distribution of trials. Whereas Dudley and Papini used intertrial intervals (ITIs) that averaged 90 s, Scull et al. used 5-s long ITIs. In one experiment, Stout (2001, Experiment 3) used a procedure similar to that used by Scull et al. (1970), but with an ITI of 30 s. This lengthening of the

ITI was also suggested by the notion that the marginal utility of a reward increases as reward frequency decreases (e.g., Fantino & Preston, 1989). Thus missing a reward under conditions of lower reward density might be more emotionally arousing than under the relatively massed conditions of previous experiments. Rats received 60 trials per session, each starting with the protraction of a lever that provided food reinforcement on 50% of the trials according to a fixed-interval 16-s schedule. For one group, N and R were unpredictable (partial reinforcement), whereas for the other group, a discriminative stimulus (light) provided information on the forthcoming outcome (discrimination training). Accordingly, rats in the partial reinforcement condition failed to respond differentially as a function of forthcoming outcome (0.77 vs. 0.78 responses/min before N and R, respectively), but rats in the discrimination training condition responded differentially (0.36 vs. 0.51 responses/min before N and R, respectively). These data provided operational evidence to define the conditions of reinforcement as surprising in the partial reinforcement case, but as expected in the discrimination training case. The main comparison involved segregating trials in each group according to whether the preceding trial was N or R. An analysis of performance during the initial 5 s of each trial (i.e., when performance is most likely to be affected by the preceding outcome), indicated that the ROE occurred in both groups. However, the difference in performance after-N vs. after-R was greater in the partial reinforcement condition (0.70 vs. 0.21 responses/min) than in the discrimination training condition (0.42 vs. 0.16 responses/min), and there was significantly higher responding after N in the partial reinforcement condition than in the discrimination training condition.

Another technique useful to control for response suppression after R as a source of the ROE consists of manipulating the interval between the outcome (e.g., N or R) and the opportunity to respond in a target trial (the so-called midtrial interval, MTI). Dickinson and Scull (1975) suggested that response invigoration after N can be distinguished from suppression after R by determining whether the dissipation of the ROE produced by increasing the MTI occurs because of a decrease in behavior after N, or because of an increase after R. Several experiments explored the effects of MTI length on the ROE with mixed results. Using rats, MacKinnon and Amsel (1964) and Dudley and Papini (1995) found results consistent with response invigoration following surprising nonreward. However, others have found results consistent with postconsummatory response suppression in monkeys, rats (Davenport, Flaherty, & Dryud, 1966; Dickinson & Scull, 1975), and pigeons (Papini & Hollingsworth, 1998).

The ROE, as empirically defined, is a robust phenomenon. To the authors' knowledge, no published study in which the strength of behavior after N was compared with that after R has failed to find evidence of the ROE. Unresolved, in the majority of cases, is whether this empirical ROE reflects facilitated responding after N, the only case that would justify categorizing the ROE as an aftereffect of surprising nonreward. The best evidence supporting this possibility comes from studies that have used adequate controls to discard response suppression after R. The expectancy-variation and MTI-variation techniques considered previously are the most demanding tests for a surprising nonreward interpretation of the ROE, but they have provided

ambiguous evidence. One possibility, suggested by Stout's (2001) experiment, is that the ROE is caused both by suppression after R *and* by facilitation after N, and that the extent to which each process contributes to the ROE is dependent upon the parameters of training. Accordingly, the following experiments were based on the assumption that the ROE results from these two processes, one involving suppression of behavior after R, probably produced by response competition or a transient reduction in food motivation, and the other involving facilitation of behavior induced by surprising nonreward. The present experiments used variations in the MTI to dissociate after-N from after-R processes. The evidence suggests that the aftereffects of nonreinforcement and reinforcement decay at different rates, with the after-N process decaying more slowly than the after-R process.

Experiment 1

The present experiment used the MTI technique to assess the size of the ROE under different MTI values. By all accounts, the ROE should decrease in size as the MTI increases due to either the decay of the after-R process or of the after-N process (see Dudley & Papini, 1995). However, by assessing the relative contribution of each aftereffect one could potentially determine the source of the ROE. In pigeons, for example, the ROE is eliminated when the interval increases up to 20 s, under conditions of training similar to those used in the present experiment. However, the elimination of the ROE is entirely attributable to an increase in after-R performance (a postconsummatory factor), as pigeons exhibit no detectable change in performance after N (Papini & Hollingsworth, 1998; Stout, Muzio, Boughner, & Papini, 2002).

In the present experiment, each trial included two components differentiated by a correlated cue and the enforcement of different reinforcement schedules. For each trial, the first component was a fixed-interval 16-s cycle signaled by either a steady or flashing light (plus the extension of the lever) and ending in reinforcement on a random 50% of the cycles. (An additional group receiving discrimination training on this fixed-interval component was included; however, because these animals failed to discriminate, their data were excluded from all analyses.) The second component started after a 2-s MTI and involved the extension of the lever (no other stimuli were used in this component); reinforcement was administered according to a variable-interval 20-s schedule. This component always lasted for at least 5 s, thus providing a period over which to measure lever-pressing responses immediately after the surprising outcomes of the first component. The variable-interval component provided a common target that approximates the classic double runway experiments in which the second alley was always rewarded (e.g., Amsel & Roussel, 1952). Rats received extensive training under these conditions and then were tested with two MTI values: 2 and 20 s. The question of interest was whether the ROE was reduced or eliminated as the MTI was lengthened by an increase of performance after R (thus implicating a postconsummatory mechanism), by a decrease in performance after N (thus implicating primary frustration), or both.

Method

Subjects

Subjects were 16 male Wistar rats about 110-days old at the start of the experiment. The subjects had previously served as the low magnitude control group in a consummatory contrast experiment. Their previous training involved drinking 2% sucrose solution during 5-min long sessions for 15 days. While this prior training occurred in the same conditioning boxes used for the present experiment, the animals were not given lever-pressing training or reinforced with food pellets. About three weeks intervened between the end of the consummatory response experiment and the present experiment. Rats were maintained at 85% of their ad libitum body weight from the start of the previous experiment through measured feeding at least 20 min following each session.

Apparatus

Six standard operant conditioning boxes (MED Associates) were used, each enclosed in a sound-attenuating chamber. Four of these six boxes measured 20.1-cm wide, 28-cm long, and 20.5-cm high (interior dimensions). The floor of the boxes was a grid floor of stainless steel rods 0.4 cm in diameter and spaced 1.6 cm apart. A food cup was located on the front wall of the chamber, 2 cm above the floor. A retractable lever was located 1 cm to the left of the feeder, 7 cm above the floor. Two of the six boxes had the following dimensions: 23.5-cm wide, 29-cm long, and 19-cm high. The grid floor of these two boxes was composed of aluminum rods 0.2 cm in diameter and 1 cm apart. The food cup was located on the front wall of the chamber, 2 cm above the floor. A retractable lever was located 1 cm to the left of the feeder, 6 cm above the floor. Each box was equipped with pellet dispensers that delivered 45-mg Noyes pellets (A/I rodent diet). A light on the roof of the sound-attenuating chamber provided diffuse illumination (GE 1820). Speakers outside the operant box but inside the sound-attenuating chamber provided white noise that, along with a fan, created a SPL of 7 dB. A computer located in an adjacent room administered experimental events and recorded the number and time of lever presses.

Procedure

The rats were randomly assigned to the four boxes (i.e., irrespective of their box assignment in the previous experiment). Rats received two 20-min sessions of *context habituation*. Animals were exposed to the conditioning boxes with the houselight on, but without any delivery of food. The lever remained retracted during these sessions.

Pretraining started the following day. All subjects were pretrained to press the lever with a mixed Pavlovian-instrumental procedure. A trial started with the extension of the lever, and a pellet was delivered 20 s later if the animal did not press. If the subject pressed, the lever was retracted and one pellet was immediately delivered. The ITI averaged 70 s (range: 10–130 s), and there were 10 trials per session over a total of 5 sessions. The subjects rapidly achieved successful performance on this task. By session 5 all subjects were receiving all 10 pellets instrumentally. Session 6 involved 60 progressive-interval trials. Progressive interval trials were grouped in

pairs. The MTI (within the pair) was 2 s and the ITI (between pairs) averaged 30 s (range: 15–45 s); these values were maintained in the rest of this experiment. The interval between lever extension and the first press that resulted in reinforcement was lengthened by a 2.5% increment on each trial, starting with a value of 1 s and until the fixed-interval value reached 5 s on trial 20, where it remained until the end of the session.

Baseline training started the next day and lasted for 10 sessions. During baseline sessions, each trial included a pair of components (fixed-interval 16 s and variable-interval 20 s schedules), each ending in reinforcement. Thus the sequence of trials in these baseline sessions was: Fixed-interval 16 s, 2-s MTI, variable-interval 20 s, 30-s ITI, fixed-interval 16 s, and so forth. Trials began with the insertion of the lever. In the fixed-interval component, the first response after 16 s had elapsed following lever insertion resulted in the retraction of the lever and the immediate delivery of one pellet. The lever was inserted again 2 s later and reinforced on a modified variable-interval schedule. The computer sampled from a string of values derived from a Flescher and Hoffman (1962) series modified so that no interval was shorter than 5 s. The average interval of the series was 20 s. Truncating the series guaranteed that each trial would have a 5-s period free from reinforcement during which to measure responding. The variable-interval and the ITI were generated by sampling without replacement from a list of discrete values. This resulted in experimental sessions of similar duration for all the animals. Each session consisted of 30 trial pairs (i.e., 30 fixed-interval components and 30 variable-interval components).

Training started the following day and continued for an additional 30 sessions. The four rats assigned to each box were ranked on the basis of their average response rate during the variable-interval component on all but the first pretraining session (a recording error resulted in data loss), and randomly assigned in matched pairs to two groups ($n = 8$). As mentioned previously, the rats assigned to a discrimination group failed to discriminate and thus were excluded from all analyses; this report then concentrates exclusively on the performance of rats that received partial reinforcement on the fixed-interval component of the pair. The experimental phase was identical to continuous reinforcement pretraining except for two changes. First, there were 32 trial pairs per session (rather than 30). Of these trials, 16 fixed-interval components ended in nonreward and 16 in reward. Second, fixed-interval lever extension overlapped with the presentation of either a flashing (0.5 s on, 0.5 s off) or a steady light stimulus located above the lever. The flashing and steady light stimuli were not correlated with the forthcoming outcome and were presented to match exposure to these cues with the discrimination group. For each session, each rat received the following fixed-interval 16 s trials: 8 flashing N trials, 8 steady N trials, 8 flashing R trials, and 8 steady R trials. Nonrewarded fixed-interval components ended with the lever automatically retracted 16.75 s following insertion. This was roughly the time it took the animals to earn the pellet on reinforced trials.

MTI testing provided the main results of this experiment during the final 4 sessions of the experiment. The procedure used in the first and fourth sessions was the same used in the previous sessions, that is, with an MTI of 2 s. In the second and third sessions, the MTI was increased to 20 s, while all other aspects

of the training procedure remained constant. Thus the sequence of MTI values were 2, 20, 20, and 2 s.

Throughout this experiment, the dependent measures were the mean response rate over the entire fixed-interval cycle and the mean response rate over the first 5 s of the variable-interval component. Response rates refer to responses per minute. Analyses of variance were computed on these values using the SPSS software package. The α value was set at the 0.05 level in all the statistical tests reported in this paper.

Results and discussion

Fig. 1 presents the results of the fixed-interval (A) and variable-interval (B) components during the baseline and training phases. A one-way analysis of fixed-interval performance during baseline sessions yielded no effect, $F < 1$, whereas that of variable-interval performance indicated a significant decrease in performance across sessions, $F(8, 56) = 4.03$. The decline in response rate in the variable-interval component probably reflects the development of a temporal discrimination of the initial 5 s, which were always nonrewarded.

Training performance during the fixed-interval component (Fig. 1A) provided no evidence that rats anticipated the forthcoming outcome, as expected for a partial reinforcement situation, $F(1, 7) = 1.75$. There was a significant decline in responding over sessions, $F(29, 203) = 1.96$, but the forthcoming outcome by session interaction was not significant, $F(29, 203) = 1.31$.

Immediately following the shift from baseline (continuous reinforcement) to training (partial reinforcement) in the fixed-interval component, a strong ROE was observed in the variable-interval component (Fig. 1B). The ROE persisted over the entire training phase. Response rate following nonrewarded fixed-interval components was substantially higher than response rate following rewarded fixed-interval components: 42.9 after N and 25.3 after R (mean responses per minute for the entire training phase). A Prior Outcome (After N, After R) × Session analysis of variable-interval performance indicated a significant ROE, $F(1, 7) = 15.82$. There was also a significant change across sessions, $F(29, 203) = 2.02$, and a significant interaction of session and prior reinforcement condition, $F(29, 203) = 1.61$. The ROE emerged on the first postshift session, $F(1, 7) = 30.28$.

There was no clear difference in the fixed-interval component as a function of forthcoming outcome during the 4 sessions of MTI testing. The average response rates were 40.8 after R and 41.0 after N, and this difference was not significant, $F < 1$. The response rate during the variable-interval component at each MTI value is depicted in Fig. 2. Increasing the MTI from 2 to 20 s eliminated the ROE and, then, restoring the MTI to 2 s reinstated the effect. The elimination of the ROE at 20 s in sessions 2 and 3 occurred because of a simultaneous decrease in responding after N and an increase in responding after R. A Prior Outcome (After N, After R) × MTI (2 s, 20 s) analysis confirmed these conclusions. For this analysis, data were averaged over sessions involving equal MTI values. The effect of prior outcome was significant, $F(1, 7) = 8.75$, but was qualified by a significant prior outcome by MTI interaction, $F(1, 7) = 8.38$. Decomposition of the interaction confirmed that

A

Fixed Interval Component
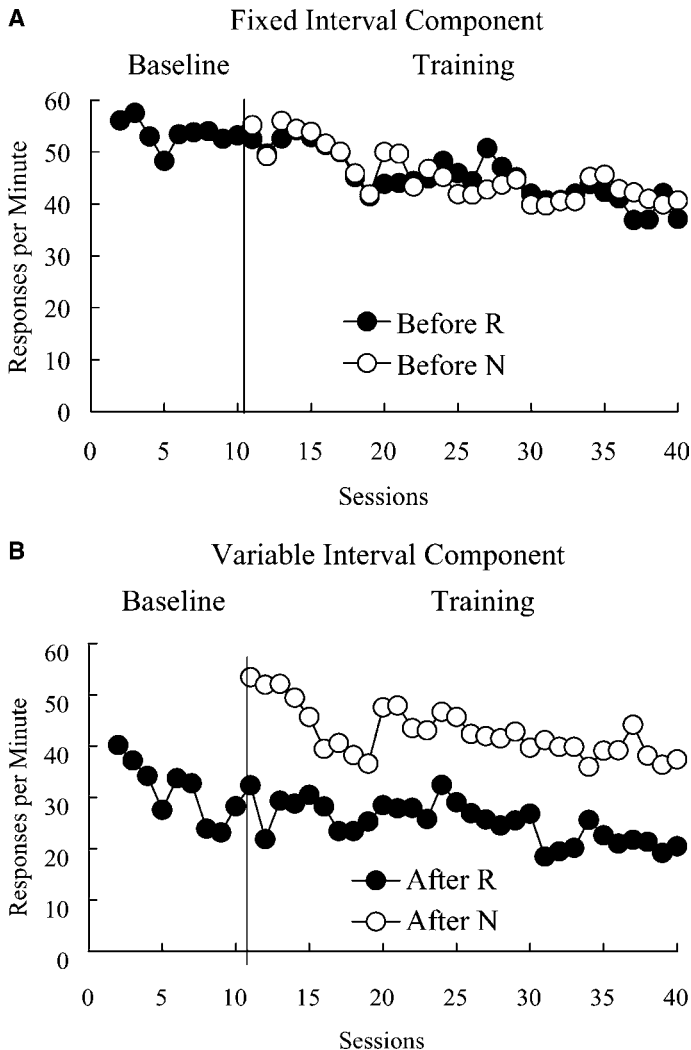

B

Variable Interval Component

Fig. 1. Response rate performance as a function of forthcoming outcome (A) and preceding outcome (B). R: reinforcement, N: nonreinforcement. Data for the fixed-interval component (A) were averaged for the entire cycle across sessions. All trials ended in R during baseline sessions. A 50% schedule of partial reinforcement was applied to the fixed-interval component during training. Data for the variable-interval component (B) were averaged for the initial 5-s period and plotted across sessions. All the variable-interval cycles ended in R (i.e., a common target), but differed in terms of their preceding outcome. The B shows that the ROE emerged the first session in which partial reinforcement was introduced. Data from Experiment 1 with an average intertrial interval of 30 s and a 2-s midtrial interval.

it was caused by the presence of an ROE at 2 s, $F(1, 7) = 8.71$, but not at 20 s, $F < 1$. Furthermore, analysis of behavior at each level of outcome shows that the ROE with the 2-s MTI was the result of predominantly one effect. After-N responding
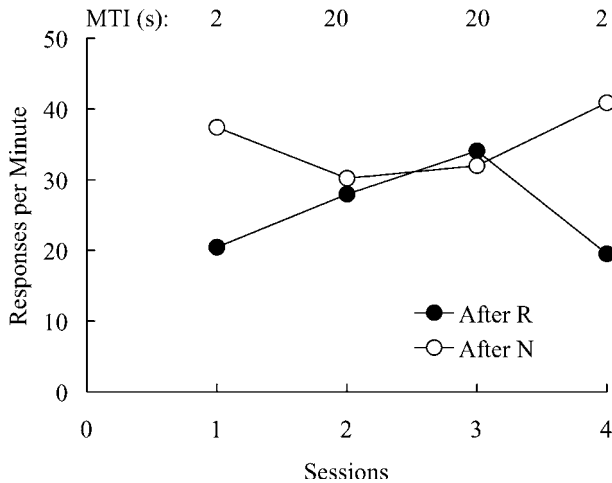
Fig. 2. Response rate during the initial 5-s period of the variable-interval 20-s trial as a function of pre-ceding outcome (R: reinforcement, N: nonreinforcement). In sessions 1 and 4 the midtrial interval (MTI) was 2-s long, whereas in sessions 2 and 3 the MTI was 20-s long. The results suggest that both after-R and after-N processes contribute to the ROE. Data from Experiment 1 with average intertrial intervals of 30 s.

significantly decreased from 39.14 to 31.08 responses per min as the MTI was length-ened from 2 to 20 s, $F(1,7) = 5.76$. After-R responding increased from 19.98 to 31.00 responses per min as the MTI was increased from 2 to 20 s, but this difference failed to achieve significance, $F(1,7) = 2.29$.

Experiment 2

Experiment 1 provided evidence consistent with the assumption that the ROE re-sults from the contribution of an after-N process that decays in about 20 s, at least under the present training conditions. The results also suggest the potential presence of a second component, one occurring after R and also decaying completely in about 20 s. Using the same MTI technique, the present experiment aimed at providing information on the relative time course of these aftereffects.

Method

The subjects and apparatus were the same as in Experiment 1. The training pro-cedure and dependent variables were also those of the previous experiment, with the following exceptions. The MTI was varied at four levels. The time values used in the previous experiment, 2 and 20 s, were kept as end points. The first value was doubled to produce the second interval (4 s) and doubled again to produce the third interval (8 s). The MTI values were chosen under the assumption that most psychological processes decay exponentially (e.g., short-term memory; Roberts & Grant, 1976);

it was then predicted that most of the change would occur during the earlier portions of the MTI. The sequence of MTI values was balanced across animals over the four sessions using an incomplete intragroup counterbalancing procedure. There were four possible sequences that matched the position of each value in the sequence: 2-4-8-20, 8-2-20-4, 4-20-2-8, and 20-8-4-2 (each MTI value appears once in each position and is preceded and followed by all the other values an equal number of times). Two rats were randomly assigned to each of these sequences. Moreover, each conditioning box was assigned the same number of subjects run in each of the four possible sequences of MTI values. Each rat was trained in a given MTI value during one session.

Results and discussion

A Forthcoming Outcome (Before N, Before R) × MTI (2, 4, 8, 20 s) analysis was computed on fixed-interval data to determine the effects of partial reinforcement training. As expected, there was no evidence that trial outcomes were discriminated, $F(3, 21) = 1.57$, or that the variation in MTI values had any impact on fixed-interval performance, $F < 1$; the interaction effect was also nonsignificant, $F(3, 21) = 1.29$.

The main results of this experiment are presented in Fig. 3. As in the previous experiment, the ROE observed at 2 s dissipated as the MTI was lengthened up to 20 s both because responding following nonreward declined and because responding following reward increased. Statistical analyses confirmed these observations. A Prior Outcome (After N, After R) × MTI (2, 4, 8, 20 s) analysis indicated a significant effect of outcome on behavior, $F(1, 7) = 6.16$. This effect interacted with MTI,
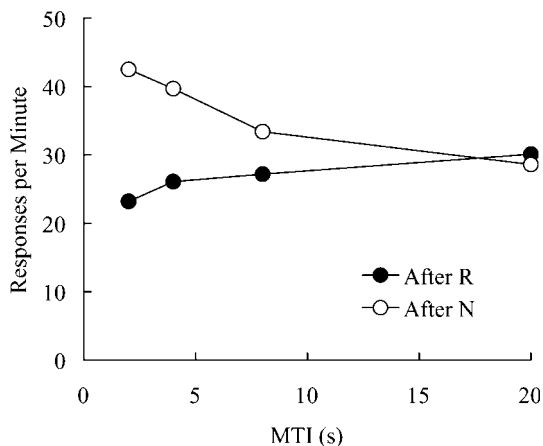

Fig. 3. Response rate during the initial 5-s period of the variable-interval 20-s trial as a function of the length of the midtrial interval (MTI; values: 2, 4, 8 or 20 s) and preceding outcome (R: reinforcement; N: nonreinforcement). All animals received 50% partial reinforcement training during the preceding fixed-interval cycle. This figure shows that performance changes faster after R than after N cycles. Data from Experiment 2 with an average intertrial interval of 30 s.

$F(3, 21) = 7.70$, which is consistent with the reduction of the ROE as the MTI was lengthened. Repeated-measure analyses at each MTI value revealed that the ROE was significant at 2 s, $F(1, 7) = 6.78$, and at 4 s, $F(1, 7) = 8.72$, marginal at 8 s, $F(1, 7) = 3.75$, $p < 0.10$, and nonsignificant at 20 s, $F < 1$. That MTI had no overall effect on behavior, $F < 1$, is consistent with the fact that, as the MTI was lengthened, after-N behavior decreased across the four MTI values while after-R behavior increased, thus reducing this main effect.

A repeated-measure analysis restricted to behavior after N confirmed a significant decrease in response rate across MTI values, $F(3, 21) = 6.39$. Pairwise comparisons computed with the LSD technique indicated the following results. After N, performance was significantly higher in the 2-s MTI than in either the 8- or 20-s MTI values, and higher also in the 4-s MTI than in the 20-s MTI value. All the other comparisons yielded nonsignificant differences. A similar analysis restricted to behavior after R indicated that the increase was not significant $F < 1$. This pattern implicates a process that enhances responding following N trials and that is largely dissipated in 8 s. There was a tendency of performance to increase as a function of MTI length following R trials, but this postconsummatory effect was not substantiated statistically.

Experiment 3

Experiment 2 provided support for the notion that the ROE is caused by a process triggered by nonreward that facilitates responding and decays in time. In Experiment 3, the MTI technique was applied to the question of whether the nature of the ROE varies as a function of the distribution of trials. Stout (2001) found a stronger ROE in a group receiving partial reinforcement training than in one receiving discrimination training with relatively long ITIs, but ROEs of similar size with relatively short ITIs. In such a contrast, nonreinforcement can be considered surprising in the case of partial reinforcement training (provided rats show no evidence of anticipating forthcoming outcomes; see Fig. 1), but expected in the case of discrimination training (provided rats show evidence of discrimination). These results suggest that otherwise similar ROEs found in the partial reinforcement conditions may nonetheless be based on different underlying processes as a function of the distribution of trials. If this is correct, rats receiving partial reinforcement training under either spaced or massed conditions should exhibit different dynamic changes after N and after R when tested with the MTI technique. Specifically, whereas groups should not differ in their after-R profiles, only the group trained under relatively spaced conditions should exhibit response decrement across MTI values after N.

Method

Subjects and apparatus

The subjects were 40 male Wistar rats, about 150-days old at the start of the experiment, maintained as described in previous experiments. These rats had

received previous consummatory training involving exposure to sucrose solutions under conditions similar to those described in Experiment 1. These rats were experimentally naïve in terms of their exposure to the stimuli used in the present experiment (i.e., lever, light, and pellets). The conditioning boxes were the same four boxes described in Experiment 1 with the following changes. A second lever was installed, positioned at the other side of the feeder cup and at the same distance from the cup and floor as the first lever.

Procedure

Rats were randomly assigned to the boxes, irrespective of their prior assignments. *Context habituation* involved two 20-min sessions of exposure to the conditioning environment. The house light remained on during these sessions, but no food was presented and the levers remained retracted.

Pretraining started the next day. In the following 7 sessions, lever pressing was developed using the mixed Pavlovian-instrumental training described in Experiment 1. There were 10 trials per session in sessions 1–5, and 60 trials per session in sessions 6 and 7. In each session, half the trials involved the right lever and the other half the left lever; right-lever and left-lever trials were randomly intermixed during the session. A trial started with the insertion of one lever; a pellet was delivered contingent upon a lever press or at the end of a 20-s period, whichever occurred first. Trials were separated by an average ITI of 70 s (range: 10–130 s). In the following 4 sessions, rats received training in a progressive interval schedule similar to that described in Experiment 1. By the fourth session, all rats were responding on a fixed-interval 16-s schedule on both levers. The fixed-interval 16-s schedule was maintained during 5 additional sessions. A fixed ITI of 2 s was kept during these 9 sessions. This completed pretraining.

Training lasted for 40 additional sessions. The procedure was analogous to that used in Experiment 1. Each trial had two components separated by an MTI of 2 s. The first component was a fixed-interval 16-s schedule that ended in N or each R, in a random half of the trials. The second component was a variable-interval 20-s schedule that always ended in reinforcement. Rats were randomly assigned to two groups ($n = 10$) depending on the length of the ITI. In Group 2, trials (each trial involved a fixed-interval and a variable-interval component) were separated by a mean ITI of 2 s (range: 1.5–2.5 s). In Group 45, trials were separated by a mean ITI of 45 s (range: 30–60 s). This experiment also involved discrimination training for two additional groups ($n = 10$) and, as in Experiment 1, these rats failed to discriminate. Thus, their performance was not included in any of the analyses presented below.

MTI testing was administered during the last 8 sessions. The same MTI values used in the previous experiment were tested again in the present experiment: 2, 4, 8, and 20 s. A given value remained constant for any given session. Thus, there were 4 possible sequences that matched the position of each value in the sequence: 2-4-8-20, 8-2-20-4, 4-20-2-8, and 20-8-4-2 (each MTI value appears once in each position). Rats were randomly assigned to one of these sequences and were run through the entire sequence twice. The main data were obtained by averaging the response rates (responses per min) for the entire fixed-interval component and for the initial 5 s of

the variable-interval component across the two sessions allocated to each of the four MTI values. Other procedural parameters were as described in Experiment 1.

Results and discussion

The performance of the groups trained with the short (2 s) and long (45 s) ITIs during acquisition was similar. Regarding the fixed-interval component, the overall mean response rates for the 40 sessions of training were the following. In Group 2, 43.5 and 42.9 responses per min were recorded in trials that ended in N and R, respectively. In Group 45, 39.5 and 39.5 responses per min were recorded also for trials ending in N and R, respectively. There was no statistical evidence in either group that rats could discriminate the forthcoming outcome of fixed-interval components and thus such outcomes were surprising.

The variable-interval component also demonstrated similar performance across groups, including the size of the ROE. The overall acquisition performance yielded the following response rates: 44.4 and 23.3 responses per min after N and after R, respectively, in Group 2, and 39.4 and 22.1 responses per min after N and after R, respectively, in Group 45. An ITI × Prior Outcome (After N, After R) × Session analysis indicated only significant effects for the prior outcome, $F(1, 702) = 119.14$, and session, $F(39, 702) = 2.65$. All other main effects and interactions failed to achieve significance, $Fs \leqslant 1.19$. As was the case in Experiment 1 (see Fig. 1B), ROEs were evident with both ITI conditions from the very first session in which rats were exposed to partial reinforcement. The average response rates for after-N and after-R trials in this first session were, respectively, 32.0 and 17.5 for Group 2, and 40.5 and 22.7 for Group 45. An ITI × Prior Outcome analysis indicated a significant ROE, $F(1, 18) = 54.08$, but nonsignificant ITI and interaction effects, $Fs(1, 18) < 2.26$.

The main results come from the MTI testing phase in which rats were exposed to four MTI values. Performance during the fixed-interval component was not affected by the forthcoming outcome across all MTI values and groups. An ITI (2, 45) × Forthcoming Outcome (Before N, Before R) × MTI (2, 4, 8, 20 s) analysis indicated a nonsignificant main effect for outcome, $F < 1$. However, there was a significant triple interaction, $F(3, 54) = 5.02$, and a significant ITI by MTI interaction, $F(3, 54) = 2.94$, both capturing a crossing over of the groups: Group 2 performed generally above Group 45, except at the 8-s MTI value where its performance was below that of Group 45. All the other effects were nonsignificant. Separate analyses for each group at each MTI value indicated that none of the effects of forthcoming outcome reached a significant level, $Fs(1, 9) < 5.06$. As expected, then, outcomes were also surprising during the MTI testing phase.

Performance during the variable-interval component as a function of group, preceding outcome, and MTI value is plotted in Fig. 4. The ROE dissipated as the MTI increased up to 20 s in both groups, although the difference across trial types continued to be present in Group 2. Both after-R functions exhibited an increase across MTI values that was similar across groups, suggesting that the post-consummatory factor was not modulated by the ITI manipulation. The after-N function was essentially flat for Group 2, but it showed an increase followed by
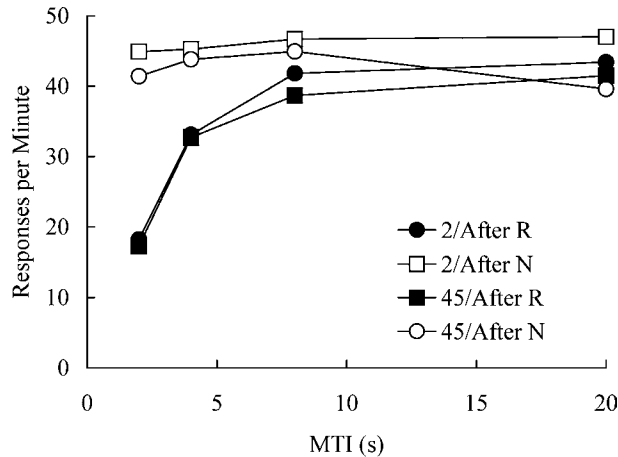
Fig. 4. Response rate during the initial 5-s period of the variable-interval 20-s trial as a function of the length of the midtrial interval (MTI; values: 2, 4, 8 or 20 s), intertrial interval (2 or 45 s), and preceding outcome (R: reinforcement, N: nonreinforcement). Both Groups 2 and 45 received 50% partial reinforcement training during the preceding fixed-interval cycle. This figure shows that whereas the after-R process occurs under both short and long intertrial intervals, the contribution of the after-N process was observed only in the long intertrial interval group. Data from Experiment 3.

a decrease in rate for Group 45. Statistical analyses of these data confirmed these descriptions.

An overall analysis involving ITI (2, 45 s) × Prior Outcome (After N, After R) × MTI (2, 4, 8, 20 s) indicated a significant interaction between MTI and prior outcome, $F(3, 54) = 25.92$, which describes the gradual disappearance of the ROE as the MTI increased. Also significant were the main effects of prior outcome, $F(1, 54) = 38.15$, and MTI, $F(3, 54) = 11.45$. All the other effects were nonsignificant, $Fs < 1$. A series of statistical tests were also conducted to determine the source of these differences. First, the performance of each group at each MTI value was analyzed separately. In Group 2, performance after N was significantly different from after R at MTI values of 2, 4, and also at 20 s, $Fs(1, 9) > 6.76$; at the 8-s MTI, the difference was only marginally significant, $F(1, 9)4.94, p = 0.053$. Thus, although the size of the difference between after-N and after-R performance decreased in Group 2, the ROE continued to be present even at the MTI value of 20 s. In Group 45, the ROE was present at 2-, 4-, and 8-s MTIs, $Fs(1, 9) > 5.12$, but not at the 20-s MTI, $F(1, 9) = 1.94$.

Second, each function was analyzed separately across the four MTI values to determine pairwise differences. In Group 2, after N, the MTI effect was nonsignificant, $F < 1$. Therefore, the after-N function in Group 2 was essentially flat. In Group 45, after N, the MTI effect was larger, but still failed to reach significance, $F(1, 9) = 2.50$. Finally, the after-R functions exhibited a significant MTI effect in both groups, $Fs(1, 9) > 9.61$. In Group 2, LSD tests indicated that the 2-s MTI was significantly different from the other values and that the 4-s MTI was significantly different from the 20-s MTI; all other comparisons were nonsignificant.

In Group 45, performance under the 2-s MTI was significantly more suppressed than under the other MTI values, which, in turn, did not differ from each other.

The results of the present Group 45 differed in two respects from those of the previous experiment (this comparison seems appropriate given the similarity in ITI values: 45 and 30 s; compare Figs. 3 and 4). First, the after-N function in Group 45 changed relatively less than the corresponding function in the previous experiment. Second, the after-R function in Group 45 changed more than that in the previous experiment. The source of these differences remains undetermined. On the other hand, there was agreement across experiments in two features. First, the after-N function decayed more slowly than the after-R function. Such differential decay observed in both experiments supports the hypothesis that two distinct processes contribute to the ROE. Second, the 4- and 8-s MTI values produced a ROE that was entirely attributable to the invigoration of lever-pressing performance after N. Thus there are conditions under which the ROE is primarily determined by the aftereffects of reward omission. The present experiment also showed that the MTI technique is more sensitive than a direct assessment of performance during training trials. In this case, the MTI technique uncovered a group difference in ROE as a function of trial distribution that was not obvious during the training phase.

General discussion

The present results support the view that the ROE observed in rats can result from two independent processes, namely, facilitation of behavior following nonrewarded trials and suppression of behavior following rewarded trials. The degree to which either process, or both together, control responding in a given situation depends upon the parameters of training. These findings make two important contributions.

First, the results reported in this paper clarify a long-standing confusion about the influence of surprising nonreward on ongoing behavior. The confusion, first clearly raised by Seward et al. (1957), lies in distinguishing between the facilitation and suppression of postoutcome instrumental behavior in a within-group comparison. The best available results supporting the presence of a facilitatory effect of surprising nonreward on ongoing performance come from between-group experiments in which the amount and frequency of reinforcement is equated across groups (Dudley & Papini, 1995, 1997; Stout, 2001). The techniques used in those experiments, however, do not allow for a concurrent assessment of both after-N and after-R effects, thus leaving unanswered the question of the relative contribution of these aftereffects to the ROE. In contrast, the MTI technique used in the present experiments provided information on the role of these aftereffects, as well as of their rates of decay. Amsel and Roussel (1952) had originally assumed that the invigorating effects of primary frustration decay in time and provided evidence that lengthening the time during which the rat was restrained in the first goal box reduced the ROE, as manifested in running speeds in the second runway. Other experimenters also studied this decay assumption under several conditions (e.g., Davenport et al., 1966; Dudley & Papini,

1995). However, the MTI manipulation can provide additional information if the direction of the decaying trace is taken into account (Dickinson & Scull, 1975). This is so because lengthening the MTI may eliminate the ROE either by a decrease in after-N performance or by an increase in after-R performance. Determining which of these two outcomes occurs allows for the identification of the process responsible for the ROE. The present experiments show that both processes can contribute to the ROE and that they can be dissociated by their differential rate of decay. The MTI technique provided evidence that the ROE may depend only on response facilitation after N (all experiments), or suppression after R (Experiment 3, 2-s ITIs).

The MTI technique also allowed for a clarification of the relationship between trial distribution and the ROE. As was pointed out in the introduction, comparisons between conditions involving surprising vs. expected N have produced evidence of after-N facilitation in experiments with long, but not short ITIs (Dudley & Papini, 1995, 1997; Scull et al., 1970; Stout, 2001). As Experiment 3 demonstrated, massed training conditions produce a ROE that is essentially dependent exclusively on the aftereffects of reinforcement. Amsel's (1992) frustration theory can be easily accommodated to account for the prevalence of response invigoration after N when relatively spaced conditions are used. According to frustration theory, the ROE reflects a transient invigoration of prepotent behavior by an affective state of primary frustration possessing drive-inducing properties. The magnitude of primary frustration is held to be a direct function of the discrepancy between the magnitude of the reward expected and the amount actually obtained. As it is, this assumption applies equally well to spaced and massed training conditions. However, the dependency of response facilitation upon spaced training conditions can be accommodated if it is assumed that the effects of surprising nonreward may be attenuated by the presence of a still active representation of the reward, as it might be expected under massed conditions of training. Imagine a pair of successive trials, an R trial followed by an N trial. If the presentation of food in the first trial produces a decaying trace, then the trace may still be active at the time of nonreinforcement in the second trial if the ITI is short, but not if it is long. In the short-ITI case, the impact of surprising nonreward would be attenuated relative to that of the long-ITI condition.

The present results also leave us with a few questions for future research. One of these questions concerns the interpretation of published results. As stated previously, most published studies reporting the ROE do not control for the possibility that the source of the effect lies in the aftereffects of reinforcement, rather than nonreinforcement. There is a tendency to consider R trials as a baseline against which to compare the performance in after-N trials, an assumption demonstrated to be wrong by the present results. A second unresolved problem concerns the issue of the aftereffects of expected nonreward, as in a situation involving A+/B− discrimination training. Future work should replicate the common-target, MTI-extension design of Experiment 3 using stimuli that can be discriminated by rats. Evidence for surprising-nonreward elevation of performance would be found if a ROE attributable to an after-N process (as confirmed by MTI extension) were eliminated by discriminated performance in the preceding fixed-interval cycle. Third, also unresolved is the issue of the appropriate characterization of the after-R process. Whereas the prevalent view

suggests that rats undergo a transient decrease in hunger motivation as a result of having eaten a small piece of food (i.e., hunger demotivation; Seward et al., 1957), other possibilities cannot be easily discarded. For example, rats may engage in a variety of postconsummatory behaviors that might directly interfere with the target response, including chewing, swallowing, and grooming. Such competing responses may be directly activated by recent feeding, by motivational changes induced by recent feeding, or by the animal's assessment of a low reward probability characteristic of the period that follows food delivery (e.g., Staddon, 1977; Timberlake, 1994). Finally, there are discrepancies between previously published results and the present experiments that cannot be resolved with the available evidence. For example, Dickinson and Scull (1975, Experiment 1) also reported that the ROE was dissipated by an increase in the MTI from 3 to 60 s, using training conditions analogous to those of the present experiments and rats as subjects. Yet they found only evidence of after-R suppression. A potentially important procedural difference between Dickinson and Scull's experiment and the present ones relates to the assessment of the aftereffects. In the Dickinson and Scull experiment, the aftereffects of N and R were measured in terms of performance during the entire second component of the trial, which was a fixed-interval 60-s schedule. In contrast, the present experiments assessed the aftereffects of N and R in terms of the response rate during the initial 5 s of a variable-interval 20-s schedule. The choice of a fixed-interval schedule may not be the ideal given that stimuli associated with the onset of the cycle may acquire inhibitory control over the response (Mackintosh, 1974). Such an effect would tend to obscure any facilitatory aftereffect of nonreinforcement. Furthermore, averaging performance over the entire second component may "contaminate" the assessment of aftereffects with anticipatory responses under the control of forthcoming outcomes (Williams, 1983). Such anticipatory responses may be expected to elevate response rate thus countering any decay in after-N facilitation.

Another contribution of the present experiments stems from the potential of the MTI technique for a comparative analysis of surprising nonreward phenomena. Research with a variety of vertebrates has shown a species discontinuity in the emergence of phenomena that depend on surprising nonreward (Bitterman, 2000; Papini, 2002). For example, a shift from a more preferred or a larger reward to a less preferred or smaller reward is typically followed by an abrupt deterioration of instrumental and consummatory behavior in a variety of mammals (see Papini, Mustaca, & Bitterman, 1988, for an example involving marsupials), a phenomenon called successive negative contrast. Similar shifts in reward quality or magnitude lead to a behavioral disruption that is either gradual or not present at all in nonmammalian vertebrates (see Papini, 1997, for an example involving pigeons). One possible explanation of this apparent species divergence in learning mechanisms suggests that the brain processes underlying the processing of surprising nonreward may have evolved in the ancestors of modern mammals and thus may not be present in other vertebrate lineages (Papini, 2002, 2003). This hypothesis predicts that the absence of successive contrast effects in nonmammalian species would be accompanied by ROEs not dependent on after-N processes, if at all present. In a parallel series of experiments involving pigeons and using the same techniques described in the present studies with

rats, Stout et al. (2002) reported evidence that the ROE is entirely attributable to suppression of instrumental behavior immediately following reinforcement. For example, just as in rats, lengthening the MTI eliminates the ROE; however, this is accomplished exclusively by an increase in responding after R (see also Papini & Hollingsworth, 1998). Analogous experiments in other vertebrates, combined with brain manipulations, will be needed to test the hypothesis that the affective correlates of surprising nonreward occur only in mammals.

References

Amsel, A. (1958). The role of frustrative nonreward in noncontinuous reward situations. *Psychological Bulletin, 55*, 102–119.

Amsel, A. (1962). Frustrative nonreward in partial reinforcement and discrimination learning: Some recent history and a theoretical extension. *Psychological Review, 69*, 306–328.

Amsel, A. (1992). *Frustration theory: An analysis of dispositional learning and memory*. Cambridge, UK: Cambridge University Press.

Amsel, A., & Hancock, W. (1957). Motivational properties of frustration: III. Relation of the frustration effect to antedating goal processes. *Journal of Experimental Psychology, 53*, 126–131.

Amsel, A., & Roussel, J. (1952). Motivational properties of frustration: I. Effect on a running response of the addition of frustration to the motivational complex. *Journal of Experimental Psychology, 43*, 363–368.

Amsel, A., & Ward, J. S. (1965). Frustration and persistence: Resistance to discrimination following prior experience with the discriminanda. *Psychological Monographs, 79*.

Bitterman, M. E. (2000). Cognitive evolution: A psychological perspective. In C. Heyes, & L. Huber (Eds.), *The evolution of cognition* (pp. 61–79). Cambridge, MA: MIT Press.

Bower, G. H. (1962). The influence of graded reductions in reward and prior frustrating events upon the magnitude of the frustration effect. *Journal of Comparative and Physiological Psychology, 55*, 582–587.

Catania, A. C. (1973). Self-inhibiting effects of reinforcement. *Journal of the Experimental Behavior, 19*, 517–526.

Couvillon, P. A., Nagrampa, J. A., & Bitterman, M. E. (1994). Learning in honeybees (*Apis mellifera*) as a function of sucrose concentration: Analysis of the retrospective effect. *Journal of Comparative Psychology, 108*, 274–281.

Daly, H. B. (1968). Excitatory and inhibitor effects of complete and incomplete reward reduction in the double runway. *Journal of Experimental Psychology, 76*, 430–438.

Davenport, J. W., Flaherty, C. F., & Dryud, J. P. (1966). Temporal persistence of frustration effects in monkeys and rats. *Psychonomic Science, 6*, 411–412.

Dickinson, A., & Scull, J. (1975). Transient effects of reward presentation and omission on subsequent operant responding. *Journal of Comparative and Physiological Psychology, 88*, 447–458.

Dudley, R. T., & Papini, M. R. (1995). Pavlovian performance of rats following unexpected reward omissions. *Learning and Motivation, 26*, 63–82.

Dudley, R. T., & Papini, M. R. (1997). Amsel's frustration effect: A Pavlovian replication with control for frequency and distribution of rewards. *Physiology and Behavior, 61*, 627–629.

Fantino, E., & Preston, R. (1989). Choice and foraging: The effects of accessibility on acceptability. *Journal of the Experimental Analysis of Behavior, 50*, 395–403.

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

Henke, P. G. (1977). Dissociation of the frustration effect and the partial reinforcement extinction effect after limbic lesions in rats. *Journal of Comparative and Physiological Psychology, 91*, 1032–1038.

Hug, J. J. (1970). Frustration effect after development of patterned responding to single alternation reinforcement. *Psychonomic Science, 21*, 61–62.

Kapp, B. S., Whalen, P. J., Supple, W. F., & Pascoe, J. P. (1992). Amygdaloid contributions to conditioned arousal and sensory information processing. In J. P. Aggleton (Ed.), *The amygdala: Neurobiological aspects of emotion, memory, and mental dysfunction* (pp. 229–254). New York: Wiley-Liss.

Kello, J. E. (1972). The reinforcement-omission effect on fixed-interval schedules: Frustration or inhibition? *Learning and Motivation, 3*, 138–147.

McHose, J. H., & Ludvigson, H. W. (1964). Frustration effect as a function of drive. *Psychological Reports, 14*, 371–374.

MacKinnon, J. R., & Amsel, A. (1964). Magnitude of the frustration effect as a function of confinement and detention in the frustrating situation. *Journal of Experimental Psychology, 67*, 468–474.

Mackintosh, N. J. (1974). *The psychology of animal learning.* London: Academic Press.

Papini, M. R. (1997). Role of reinforcement in spaced-trial operant learning in pigeons (*Columba livia*). *Journal of Comparative Psychology, 111*, 275–285.

Papini, M. R. (2002). Pattern and process in the evolution of learning. *Psychological Review, 109*, 186–201.

Papini, M. R. (2003). Comparative psychology of surprising nonreward. *Brain, Behavior and Evolution, 62*, 83–95.

Papini, M. R., & Dudley, R. T. (1997). Consequences of surprising reward omissions. *Review of General Psychology, 1*, 175–197.

Papini, M. R., & Hollingsworth, P. R. (1998). Role of nonreinforcement in the fixed-interval performance of pigeons. *Psychonomic Bulletin and Review, 5*, 84–90.

Papini, M. R., Mustaca, A. E., & Bitterman, M. E. (1988). Successive negative contrast in the consummatory responding of didelphid marsupials. *Animal Learning and Behavior, 16*, 53–57.

Roberts, W. A., & Grant, D. S. (1976). Studies of short-term memory in the pigeon using the delayed matching to sample procedure. In D. L. Medin, W. A. Roberts, & R. T. Davis (Eds.), *Processes of animal memory* (pp. 79–112). Hillsdale, NJ: Erlbaum.

Scull, J. W. (1973). The Amsel frustration effect: Interpretations and research. *Psychological Bulletin, 79*, 352–361.

Scull, J. W., Davies, K., & Amsel, A. (1970). Behavioral contrast and frustration effect in multiple and mixed fixed-interval schedules in the rat. *Journal of Comparative and Physiological Psychology, 71*, 478–483.

Seward, J. P., Pereboom, A. C., Butler, B., & Jones, R. B. (1957). The role of prefeeding in an apparent frustration effect. *Journal of Experimental Psychology, 54*, 445–450.

Staddon, J. E. R. (1977). Schedule-induced behavior. In W. K. Honig, & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 125–152). Englewood Cliffs, NJ: Prentice-Hall.

Stout, S.C., (2001). Aftereffects of surprising nonreward on operant performance of rats and pigeons. Unpublished doctoral dissertation, Texas Christian University, Fort Worth.

Stout, S. C., Muzio, R. N., Boughner, R. L., & Papini, M. R. (2002). Aftereffects of the surprising presentation and omission of appetitive reinforcers on key-pecking performance in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 28*, 242–256.

Terrace, H. S. (1972). By-products of discrimination learning. *Psychology of Learning and Motivation, 5*, 195–265.

Timberlake, W. (1994). Behavior systems, associationism, and Pavlovian conditioning. *Psychonomic Bulletin and Review, 1*, 405–420.

Wagner, A. R. (1959). The role of reinforcement and nonreinforcement in an apparent "frustration effect". *Journal of Experimental Psychology, 57*, 130–136.

Williams, B. (1983). Another look at contrast in multiple schedules. *Journal of the Experimental Analysis of Behavior, 39*, 345–384.