Temporal Effects of Reinforcement: A Negative "Frustration" Effect¹

J. E. R. STADDON

Duke University

Reinforcement omission in situations that show positive goal-gradients elevates subsequent responding (positive "frustration" effect). In this experiment, in a situation showing a negative goal-gradient, reinforcement omission depressed subsequent responding (negative "frustration" effect). A simple hypothesis in terms of discriminative after-effects of reinforcement accounts for both effects. The same interpretation is adequate to describe frustration effects in runways and avoids problems faced by frustration theory. Some further tests of the hypothesis are proposed.

If food is occasionally omitted in the mid-goalbox of a double runway apparatus, hungry rats will run faster down the second alley on those trials than on trials when they are fed in the mid-goalbox (the frustration effect, Amsel & Roussel, 1952). This effect has been attributed to an active, motivational effect of nonreward (*frustrative nonreward*, Amsel, 1958). Similarly, if a brief blackout of the same duration as food presentation, and therefore identical to it except for the actual operation of the food magazine and light, is occasionally presented in lieu of reinforcement on fixed-interval schedules, response rate is much higher during the following interval (omission effect). This effect has been interpreted in terms of the inhibitory after-effects (temporal inhibition) of the stimulus initiating the fixed interval (Innis & Staddon, 1969; Staddon & Innis, 1966, 1969). Despite these differing explanations and procedures the similarity between these two effects suggests that essentially the same mechanism may be involved. This similarity (reviewed in Staddon & Innis, 1969) includes the fact that starting time measures are most sensitive to both effects, that the effects persist over many days, that the onset of the effects following reinforcement omission is immediate, that both

¹ Supported by NIMH Grant MH 14194, NSF Grant GB 8504, and grants from Duke University. I thank Nancy K. Innis and John Kello for assistance and critical comments. Requests for reprints should be sent to the author, Department of Psychology, Duke University, Durham, North Carolina 27706.

effects are inversely related to detention time following nonreward, and that the effects do not depend on initial training with invariable reward. The present paper outlines a unified approach and tests one implication of it.

"Frustration" effects have been demonstrated in a number of situations: runways, fixed-interval schedules, and fixed-ratio and spaced-responding (DRL) schedules (Davenport, Flaherty, & Dyrud, 1966; Davenport & Thompson, 1965). In all these situations the dependent variable is positively accelerated for a time following the delivery of reward. For brevity this phase of positively accelerated responding will be termed a "positive goal-gradient," since it shares many features of that venerable notion. Thus, in runways the profile of running speed is typically an inverted U or V shape, depending on the length of the runway, with an accelerative phase at the beginning followed by a decelerative phase at the end (Hull, 1934; Rashotte & Amsel, 1968). In the double runway the accelerative phase of responding in the second runway immediately follows either food or nonreward in the mid-goalbox. Similarly, animals on fixed-interval schedules typically pause and then respond faster and faster after the beginning of each interval ("scallop"); on fixed-ratio schedules they show a similar but less pronounced acceleration ("pause and run"). On spaced responding (DRL) schedules, conditional probability of responding ("IRTs/op.," Anger, 1956) rises following either a response or a reinforcement (Staddon, 1965). A similar analysis applies to the "runway DRL" procedures of Logan (e.g., Logan, 1968). In experiments where these procedures have been used to demonstrate the frustration effect this accelerative phase of responding is always immediately preceded by either reward or nonreward. Moreover, this contiguous relationship seems to be essential: a long timeout or detention time separating nonreward from the period when responding occurs reduces or abolishes the effect (MacKinnon & Amsel, 1964; Davenport, Flaherty, & Dyrud, 1966; Staddon & Innis, 1969). Further, in situations where there is no goal-gradient, in the present sense, there is no frustration effect. Examples are those variable-interval schedules that produce a truly constant response rate, where neither presentation of a brief blackout in lieu of reinforcement nor reduction in magnitude of reinforcement has an effect on subsequent responding (Revusky, personal communication, and unpublished observations in our laboratory), and fixed-interval schedules following scopolamine injection (which eliminates "scallop") (Innis & Staddon, 1969).

These considerations suggest that there may be a relationship between the existence of a positive goal-gradient, initiated by reward, and the finding of a positive frustration effect. An explanation for this relationship

in terms of generalization of stimulus after-effects will now be described, using the fixed-interval case as an example. The applicability of the approach to runway situations will then be discussed, followed by an account of the prediction tested by this experiment.

If reinforcement is occasionally completely omitted once animals have learned to respond on a fixed-interval schedule, they continue to respond after the time when reinforcement would ordinarily have occurred and maintain a rate of responding during the following interval which is as high as or higher than the usual "running" rate (Staddon & Innis, 1969). This pattern occurs the first time reinforcement is omitted and is maintained essentially intact, even when occasional reinforcements continue to be omitted (cf. Staddon, 1967). Under these conditions the procedure is effectively a mixed fixed-interval x , fixed-interval $2x$, in the terminology of Ferster and Skinner (1957). The high rate of responding that occurs over a fixed interval when reinforcement is completely omitted is most simply considered as an expression of "what is learned" under the fixed-interval procedure. Since there is no stimulus change at the moment when reinforcement would normally have occurred, it makes little sense to talk of this behavior as reflecting an effect of nonreward, either as an external stimulus event or even as a relationship between such an event and a memory or expectation generated by the animal's past history. To an approximation, the animal has learned the rule: "pause after reinforcement (or whatever stimulus normally initiates the fixed interval), then respond at an approximately constant rate until the next reinforcement" (cf. Staddon, 1967; Schneider, 1969). If reinforcement is omitted, responding occurs for a longer time than usual, and since there is no pause during the second "interval," rate over that interval is higher than usual (the omission effect).

This pattern of behavior following complete reinforcement omission constitutes the baseline or null case with respect to which the effects of nonreward, defined as a stimulus presented in lieu of reinforcement, must be assessed. The comparison is clear-cut: if a brief blackout (similar to the conditions prevailing during reinforcement) is presented in lieu of reinforcement, response rate over the following interval, though elevated with respect to rate following reinforcement, is *depressed* by comparison with the rate following complete reinforcement omission. Thus, blackout must be considered to have inhibitory after-effects under these conditions, although it is evidently not as inhibitory as reinforcement. In these terms, the omission effect must be attributed not to an activating effect of nonreward (blackout), but rather to a difference between the inhibitory after-effects of reinforcement and the stimulus presented in lieu of reinforcement.

This reasoning is similar in many respects to earlier interpretations of the frustration effect in terms of inhibitory after-effects (e.g., Seward, Pereboom, Butler, & Jones, 1957). It differs in that no position is taken on the mechanism whereby reinforcement acquires inhibitory after-effects. It may involve satiation to some degree, reflect the tendency of consummatory behavior to terminate appetitive sequences (Craig, 1918; Sherrington, 1906), be analogous to Pavlov's "inhibition of delay" (Hinrichs, 1968; Staddon, 1969), or be a purely associative adaptation to absence of reinforcement in the period just after reinforcement, analogous to the delay of reward gradient (Dews, 1962). It is important to emphasize that since these after-effects are not found under every reinforcement schedule (variable-interval schedules often provide counter-examples), they cannot represent an entirely unconditioned effect of food and must depend on schedule parameters. Therefore they will be considered as instances of discriminative control, in the usual sense. This should not be taken to imply that reinforcement acts just like a "neutral" stimulus; it is for one thing much more salient. And the capacity of reinforcing stimuli to elicit unconditioned (consummatory) reactions may also play a role in acquisition of discriminative control.

The essential features of the present approach, which may be applied to either inhibitory or excitatory after-effects, are contained in the following strong hypothesis concerning the mechanism of "frustration" effects due to reinforcement omission:

The effects of reinforcement omission in a given situation depend entirely upon the after-effects of reinforcement in that situation. The effects on subsequent responding of a stimulus presented in lieu of reinforcement (i.e., nonreward) will be of the same kind as the effects of reinforcement but generally of smaller magnitude.

This hypothesis is nothing more than the principle of stimulus generalization, applied to stimulus after-effects rather than to responding in the presence of the stimulus. It is simply assumed that there will be many situations in which reinforcement has acquired some measure of control over behavior that follows it; if reinforcement is omitted the behavior to be expected will depend both on the kind of control usually exerted by reinforcement and on how similar the stimulus presented in lieu of reinforcement is to reinforcement. "Similarity" is a complex notion and cannot always be estimated *a priori* (cf. Shepard, 1965); and the problem is even more complicated for stimuli (like food) that are not "neutral." Nevertheless, as a general rule it seems safe to expect that the effect of the stimulus presented in place of reinforcement will be weaker than the effect due to reinforcement, although by how much may be difficult to foresee.

Thus, the present hypothesis attributes the close association between positive goal-gradients and positive frustration effects to the common involvement of inhibitory after-effects.

This approach can be applied to free-operant situations other than fixed-interval without great difficulty. Thus, reinforcement omission on fixed-ratio schedules should produce a smaller elevation in responding than on fixed-interval, since the postreinforcement pause is generally a smaller fraction of the interreinforcement interval; this appears to be the case (Davenport & Thompson, 1965). On DRL schedules the probability of a response is low following either a response or a reinforcement. However response latency following reinforcement is usually longer than following a response, which accounts for the frequent finding that reinforcements tend to occur in runs on DRL schedules (cf. Farmer & Schoenfeld, 1964; Ferraro, Schoenfeld, & Snapper, 1965; Kelleher, Fry, & Cook, 1959). The finding of a modest omission/frustration effect in a DRL situation (Davenport, Flaherty, & Dyrud, 1966) is simply another way of expressing this difference in inhibitory after-effects between reinforcement and a response.

The main problem in applying this approach to runway data is profound ignorance concerning the after-effects due to reward; in the runway there is no simple operation analogous to complete reward omission which does not at the same time present the animal with stimuli that can be considered as sources of frustrative nonreward. This is because in a runway the animal constantly alters his own stimulating environment merely by making the measured response (i.e., running). Since the response one is measuring is at the same time affecting the stimuli whose controlling effects we are trying to measure in terms of that same response, isolation of the effects of any stimulus in the runway is difficult. In the absence of any clear test one is free to describe the frustration effect either in terms of an excitatory effect of nonreward, an inhibitory effect of reward, or both. On grounds of parsimony either of the first two is preferable to the last, and the same consideration favors the inhibitory view, since it can deal both with free-operant and with runway data. However, since an argument by parsimony alone is far from convincing, especially in view of the weight of empirical evidence that speaks in favor of the frustration approach, some space will be devoted to a brief discussion of the relevant runway literature.

Runway Results and Frustration Theory

McHose and Ludvigson (1965) investigated the effect of pretraining two groups of rats with 2- and 10-pellet rewards, respectively, in the mid-goalbox (G1) of a double runway apparatus. After pretraining, both groups received some trials with 2-pellet rewards in G1. Both groups

showed the same second alley (A2) running speeds on these trials, contrary to frustration theory, which would suggest that the group pretrained on 10 pellets should run faster than the group pretrained on 2 pellets. The authors conclude: "The drive reduction hypothesis (Seward, Pereboom, Butler, & Jones, 1957), that larger G1 reward magnitudes produce greater hunger reduction and consequently slower speeds in A2 most readily fits the present incomplete reward reduction data." (McHose & Ludvigson, 1965, p. 495). Their data are, in fact, compatible with a variety of descriptions in terms of inhibitory after-effects, including the one proposed here. In further support of the analogy between free-operant and runway data a similar inverse relationship between response rate and antecedent reward magnitude has been demonstrated in fixed-interval schedules (Staddon, 1969), and here also the relationship appears to be independent of experience with the various magnitudes.

In a more extensive experiment, similar to the McHose and Ludvigson study, Daly (1968) also found little evidence for an effect of shift in reward magnitude on either A1 or A2 speeds (starting, runway, or goal measure). She writes: "Groups showing significant changes in performance following the shift in reward [in G1] were compared with their control groups at the shifted-to level. None of these differences was significant across all 36 shift trials ($p > .05$)." (p. 434). She did, however, find that groups shifted from 15 pellets in G1 to either 0 or 1 pellet ran faster in A2. She rejects a version of the Seward *et al.* (1957) demotivation hypothesis on the grounds that it would predict that the effects should occur throughout training (they developed with training), that they would appear in all segments of the alley (the higher rate following smaller reward was most apparent in the start measure in A2) and that running speeds should always be slower in the second alley than in the first. None of these objections applies to the after-effects view proposed here, which assumes that the effects develop with training and are inversely related to distance from reward (or nonreward), hence largest in the start measure. The present view says nothing about the relative speeds in the two alleys, since these presumably depend on a number of unknown factors such as the salience of the stimuli (food vs. being placed in the start box) at the beginning of each alley, the unconditioned (consummatory) behaviors elicited by these stimuli, the effect of running speed in the first alley on speed in the second (if any), and so on.

The alternative explanation offered by Daly for her results implies some kind of comparison between reward magnitudes in G1 and G2. Thus the difference in A2 performance of the groups that received 15 and 1 pellets, respectively, in G1 (all groups received 6 pellets in G2) is interpreted as follows: "The 6 pellets following 15 pellets depressed perfor-

mance in the second alley (inhibitory effect); the small 1-pellet reward prior to the 6-pellet reward elevated performance in the second alley (excitatory effect); and these effects increased over trials and occurred closest to the frustrating reward." (p. 436). Apart from the complexity of this interpretation, its ad hoc character, and the mixture of contemporaneous and historical causality (the 6-pellet reward in G2 depresses performance in A2—an effect of past history, since G2 is reached after traversing A2—but the 1-pellet reward in G1 immediately energizes subsequent behavior in A2—a contemporaneous effect), it is also hard to test: the obvious control group is one where the rats receive 15 pellets in both G1 and G2, but under these conditions running speed would certainly be higher in A2, because of the larger reward in G2, and there is no obvious way of separating this energizing effect from the effect due to a comparison between G1 and G2 rewards postulated by Daly. She recognizes these difficulties, but instead of abandoning the approach that led to them concludes pessimistically that: ". . . the present data and associated analysis strongly imply that appropriate control groups cannot be run in the double runway for shifts to nonzero reward magnitudes." (p. 436).

An experiment by Wagner (1959) is often cited as crucial for the energizing view of nonreward. The important groups in Wagner's experiment are group N, that was never rewarded in G1, and group P (the frustration group) that was rewarded on 50% of trials in G1. Wagner found that on nonreward trials group P ran faster in A2 than group N. This outcome is usually interpreted as indicating that nonreward is energizing for group P rats, that are sometimes rewarded in G1, but not for group N rats, that never are. However there are a number of alternative interpretations that may weaken this inference. For example, group N rats run more slowly down the first alley than do rats that are rewarded in G1 (Daly, 1968). If the inhibitory effect of the stimuli in G1 on nonreward trials is relative to the entering speed, group P animals will be less inhibited than group N, yielding a higher speed for that group in A2. A second possibility is to emphasize that reward is a very salient stimulus. It may be that inhibitory after-effects acting to depress running speed in the early part of A2 will be due to the most salient stimulus experienced by the animal in G1, in line with current work on overshadowing of control by one stimulus by control due to another more salient stimulus (Pavlov, 1927; Kamin, 1969). Thus reward may overshadow goalbox stimuli for group P but not, of course, for group N, since these animals never received reward in G1. Thus on nonreward trials group P will be free of inhibitory control while group N will not, leading to the usual difference in running speeds favoring group P.

Still another tack is to pursue the analogy with fixed-interval schedules.

In these terms each alley of the double runway is considered as equivalent to a fixed interval. The point to be demonstrated is that a group trained like Wagner's group N, with interpolated timeouts at the end of the first (shorter) fixed interval, will make more responses during the second interval than animals trained throughout with 100% timeouts at the end of the shorter interval (analogous to Wagner's group P). Given certain reasonable assumptions about the nature of FI responding this can be shown graphically, but the exposition is lengthy. The crux of the matter is that animals exposed to what is effectively a mixed FI, FI schedule tend to develop a pattern of responding (pause and "running" rate) appropriate to the shorter of the two intervals, rather than to the longer or to an average (Ferster & Skinner, 1957; Staddon, 1967). Thus animals on a procedure analogous to group N will develop a postreinforcement pause appropriate to the *sum* of the short and long intervals, since this will be the shortest interreinforcement interval experienced by this group (it is assumed that the timeout has a negligible effect), while the animals trained like Wagner's group P will develop a pattern of responding appropriate to the shorter interval. As a consequence of the different discriminations developed by these two groups, there will usually be a difference favoring group P in the total number of responses made during the long interval following reinforcement omission at the end of the short. For the same reason the group N animals will generally respond more slowly during the short interval (cf. the runway data of Daly, 1968). Thus within the free-operant context it is possible to derive results resembling Wagner's from quite a different set of assumptions. Since there is little temptation to invoke a frustration explanation here, it may be equally inappropriate for the runway analog.

The point, of course, is not that there are sufficient data available to rule out a frustration interpretation for Wagner's results, but rather that there are enough uncertainties about this interpretation to maintain our interest in the possibility of an alternative. It seems clear that Wagner's group N may not be an appropriate control group for deciding between excitation due to nonreward vs. inhibition due to reward as explanations of the frustration effect.

Amsel (1958) proposed that the introduction of delayed reward in a situation where previously it had been immediate may be a frustrative event, in his sense. Nevertheless a series of careful attempts to demonstrate an elevation in running speed in A2 following delay of reward in G1 has been almost completely unsuccessful (McHose, 1966; Sgro, Glotfelty, & Podlesni, 1969; Sgro, 1969). Another experiment examining the effect both within and between subjects, however, has demonstrated a slight depression of A2 start speed following delayed reward in G1 (Sgro, Glotfelty, & Moore, 1970). These data pose no problems for the inhibi-

tory after-effects view since the temporal relationship between reward and running in A2 (and thus the opportunity for exercise of inhibitory control by reward on A2 start speed) is unimpaired by the delay operation. The reason for the slight depression in A2 starting speed following delay is not obvious but may be related to the possibility mentioned earlier that the inhibitory effects of stimuli in G1 are relative to the entering speed; following delay and detention, either in G1 or before entering G1, the effects of this variable may be attenuated, leading to a greater inhibitory effect on subsequent A2 starting speed.

Two experiments by Peckham and Amsel (1967) afford a clear test of the frustration hypothesis. They used two double runway apparatuses, with identical, gray second alleys and goalboxes but either a white or a black first alley-goalbox combination. For one color-control group a white alley-goalbox combination was associated with an 8-pellet reward in G1, a black combination with a 2-pellet reward (and 8B-2W for the other color-control group). Both groups received 2 pellets in G2. The rats were trained for four trials per day for 64 days with equal numbers of trials with the white and with the black alleys, presented in a counterbalanced sequence. They were then tested for 12 days with only 50% reward in G1. In data averaged across all 12 test days, the first experiment showed, as predicted, a substantially larger frustration effect on trials when the 8-pellet reward was omitted (white alley and goalbox) than when the 2-pellet reward was omitted (black alley and goalbox). This was true for both color-control groups. However, this large effect was due almost entirely to lower start and running speeds following the larger reward, rather than to higher speeds following its omission. A small difference between mean nonreward speeds (N_2 vs. N_8) was not statistically significant. These data are in perfect agreement with the inhibitory after-effects view, especially as the difference in start and running speeds following the large and small rewards had been present during training. However, the second experiment, in all essential respects apparently identical to the first, did show a significant difference between 8- and 2-pellet nonreward start and running speeds. This result cannot be explained by an inhibitory after-effects view without ad hoc modification. As the authors correctly state: "The unequivocal test for the frustration interpretation is the comparison of N_8 and N_2 in the test phase." (Peckham & Amsel, 1967, p. 189). Unfortunately there are a number of features of this second experiment which render it less than conclusive. For example, considering just the start measure, where the frustration effect is usually largest, the difference between N_2 and N_8 in terms of this measure was due to the *absence* of a frustration effect on 2-pellet trials for the 8W-2B color-control group (Peckham & Amsel, 1967, Fig. 1). The N_2 vs. N_8 difference was much

smaller for the other color-control group, since there was a frustration effect on both 2- and 8-pellet nonreward trials. The N_8 vs. N_2 difference is probably not significant if this group alone is considered. Yet the inclusion of the first group means that absence of a frustration effect, looked at in one way, forms part of the justification for frustration theory, looked at in another.

In addition, an earlier presentation of the first experiment (Peckham & Amsel, 1964) shows data for the test trials in three 4-day blocks. Again just considering the starting measure, these data show that the modest N_8 vs. N_2 difference decreased across trials not via a decrease in N_8 , as would be predicted by frustration theory, but by an increase in N_2 . Thus, although these experiments provide some (perhaps the only) evidence for a truly excitatory effect of reward omission, in a within-subjects situation, there seem to be enough uncertainties of interpretation to make one hesitant to suspend the whole structure of frustration theory by this particular thread.

Lott (1967) has drawn attention to the "paradoxical" status of frustrating events within Neo-Hullian behavior theory. A frustrating event is one that by definition has been paired with reinforcement and thus should be mildly positive in its effects; yet frustration theory assimilates frustrating events and aversive events. The after-effects view does not encounter this problem since it views a frustrating event as a stimulus having (discriminative) after-effects by virtue of its similarity to another discriminative stimulus (reward). There is nothing paradoxical about a stimulus having both discriminative and reinforcing properties. Testing for these discriminative properties is an extinction operation, however (like any generalization test), and if prolonged will naturally result in the stimulus in question becoming aversive, in exactly the same way that any stimulus in the presence of which responding has been extinguished becomes aversive (Terrace, 1966). The effects observed on the first few reinforcement omissions reflect the discriminative rather than the aversive properties, however, since the latter develop only gradually; thus the frustration effect appears at once (Amsel & Roussel, 1952), but behavioral contrast, one indicator of aversiveness due to extinction, takes some time to develop (Reynolds, 1961).

The frustration effect and the effects of partial reinforcement in acquisition and extinction are all assumed by frustration theory to be due to the same mechanism. Yet injections of amobarbital which abolish the latter two effects leave the frustration effect unchanged (Gray, 1969; Ison, Daly, & Glass, 1967; Ludvigson, 1967). This result poses no difficulties for the inhibitory after-effects notion since amobarbital, unlike scopolamine, has no effect on the form of the fixed-interval scallop

(although overall rate is reduced, Dews, 1964), indicating that it does not abolish these effects. This result does suggest, however, that the frustration effect may actually be unrelated to the other phenomena discussed by frustration theory. It raises questions about the capacity of the theory to provide an accurate picture of the relationships among these different effects. Since the integrative capacity of frustration theory is usually considered to be one of its main virtues, this result is perhaps even more damaging than the existence of alternative explanations for a number of reinforcement omission experiments.

In summary, frustration theory has difficulties even with some runway data and seems largely irrelevant to an understanding of free-operant reinforcement omission experiments. The interpretation in terms of inhibitory after-effects avoids these difficulties, and there are alternative interpretations for experiments such as Wagner's (1959) which have been used to justify an "active" view of nonreward. It should also, perhaps, be emphasized that the frustration effect in the double runway is almost always small (less than 10% of baseline), and, in fact, many experimental groups fail to show any effect at all (e.g., Daly, 1968; Peckham & Amsel, 1967); consequently the amount of inhibitory after-effect due to reward that is necessary to yield a frustration effect in the runway is modest indeed, far short of the substantial postreinforcement inhibition in the fixed-interval situation. It must also be noted that there is a dearth of empirical evidence on the stimulus factors that control running within the various segments of double and single runways. Informal observations, such as the finding that if the retrace door at the entrance to G2 in the double alley is accidentally left closed well-trained rats will often run into it, strongly suggest that at least at asymptote rats are little guided by runway cues, but rather run off a learned "program" as they progress through the alley—perhaps in much the same way that a would-be pianist at first is guided by each individual note to be played, but later runs off whole sequences almost independently of sensory feedback (Lashley, 1951). More evidence is needed, but this kind of mechanism is obviously quite compatible with the notion that at asymptote the animal's behavior in the alley is controlled largely by salient events (e.g., reward) at the beginning of the alley, with runway cues exerting relatively little effect once running has begun.

A Prediction

The arguments already presented suggest that the effects of reinforcement omission on fixed-interval schedules can most simply be understood in terms of generalization decrement in the discriminative after-effects of reinforcement. In this case, and in the case of most reinforce-

ment schedules where special steps are not taken to achieve the contrary result, reinforcement has inhibitory after-effects, i.e., there is a positive goal-gradient in the sense defined earlier. The ubiquity of positive goal-gradients follows from two facts: that animals tend to respond most rapidly at the time or place when reinforcement is most probable, and that under most operant procedures, where the performance of a definite response is necessary for reinforcement to occur, the probability that such a response will be reinforced almost inevitably increases with time. Neither of these two statements is universally true, and animals do not invariably show positive goal-gradients. Nevertheless it is noteworthy that even variable-interval schedules, which are usually thought of as making reinforcement completely unpredictable, are rarely programmed in a completely random fashion; and indeed if the usual kind of finite film strip is used, cannot be totally random since there will always be a given longest interreinforcement interval. Thus the animal's response rate will usually tend to rise at long postreinforcement times, as demonstrated in an extensive series of experiments by Catania and Reynolds (1968). The opposite strategy, for response probability to decrease with time, is also maladaptive, since an extra-long interreinforcement interval may then lead to extinction (cf. Staddon, 1967). The ubiquity of positive goal-gradients, both in free-operant situations and elsewhere, makes it hard to obtain an adequate test of the hypothesis stated earlier, and perhaps explains to some extent the plausibility of emotional or motivational explanations for the frustration effect: if an effect is found under a wide variety of circumstances it is reasonable to suspect that it is not due to purely discriminative factors. The problem is that the positive goal-gradient reflects a common, and indeed hard to avoid, property of instrumental conditioning situations. The only common exception to this rule is the class of variable-interval schedules that does generate a constant rate of responding as a function of postreinforcement time (i.e., zero goal-gradient). As already mentioned, reinforcement omission under these conditions fails to produce an increase in subsequent responding. However, the argument would be greatly strengthened if a situation could be found to generate a *negative* goal-gradient (excitatory after-effects), i.e., a tendency for the measured response to *decrease* as a function of postreinforcement time. This would then provide a completely symmetrical alternative to the fixed-interval case and should produce exactly the opposite effects of reinforcement omission, namely, a lowered response rate (negative frustration effect). Such a procedure, together with variable-interval and fixed-interval schedules, would cover all possible monotonic goal-gradients and their associated omission effects and provide a strong test of the present view, within the free-operant context. The present experiment

was a preliminary exploration of the properties of such a schedule and confirmed predictions from the after-effects hypothesis.

METHOD

Subjects

Five White Carneaux pigeons were used, two without previous experimental experience (85, 93) and three with experience in fixed-interval reinforcement-omission experiments. All birds were maintained at 80% of their free-feeding weights throughout the experiment.

Apparatus

A standard, single-key, operant conditioning apparatus for pigeons was used. Effective key pecks produced an audible "feedback" click from the transparent response key. An in-line projector (Industrial Electronics Engineers) was mounted close behind the response key. White noise masked most extraneous sounds, and scheduling and recording apparatus was located in an adjoining room. Data were recorded on digital and print-out counters and a cumulative recorder.

Procedure

After the naive birds had been trained to eat from the food magazine and shaped to peck a translucent response key, all birds were given five sessions, each comprising 60 reinforcements (reinforcement was 3-sec access to mixed grain), on a variable-interval 47-sec schedule.² The response key was transilluminated with pink light. The birds were then given two sessions of a modified procedure during which the key was pink for 1 min following each reinforcement and changed to white thereafter. Whether the key was pink or white, reinforcements were scheduled by the variable-interval timer. However, the key color signified the behavior to be reinforced: in the presence of the pink key (postreinforcement time < 1 min), *key-pecking* was reinforced; in the presence of the white key (postreinforcement time > 1 min), *not-pecking*, for at least 10 sec, was required to obtain reinforcements set up by the variable-interval programmer.³ Thus, if the animal ceased responding 1 min after each reinforcement, he could collect the next reinforcement as soon as it became available. After 2 days with this cued procedure the birds were given a

² The VI tape comprised the following sequence of interreinforcement intervals: 6, 20, 12, 68, 15, 106, 12, 49, 68, and 114 sec.

³ In the terminology of Ferster and Skinner (1957) this is a multiple VI-VI DRO schedule. The uncued procedure is a mixed VI-VI DRO.

further 19 days of the same reinforcement contingencies but with the response key green throughout (uncued procedure).

Starting on the 20th day of the uncued procedure four test sessions were given (on or about days 20, 25, 28, and 31). Between tests the birds were run on the uncued procedure. During test sessions the variable-interval reinforcement sequence had interreinforcement intervals of variable duration always separated by intervals of fixed duration (about 30 sec), thus: V, 30, V, 30, etc.⁴ Reinforcements were omitted at the end of 50% of the variable-duration intervals (determined by an irregular sequence) during the first and third tests (omission sessions). A brief blackout, of the same duration as reinforcement, was the stimulus presented in lieu of reinforcement. Test sessions began with a response-contingent reinforcement and both the number of responses and duration of the actual (as distinct from the scheduled) interreinforcement interval were recorded on printout counters for every interval. The dependent variable was response rate over the fixed interreinforcement interval that began with either reinforcement or blackout. Response rate over the same intervals during the second and fourth test sessions, when reinforcements were not omitted (control sessions), provided a control for differences due to sequential effects.

RESULTS

Figure 1 shows segments of cumulative response records taken from day 19 of the uncued procedure for the five pigeons. All the records clearly show the monotonically decreasing rate of responding (as a function of postreinforcement time) generated by this schedule. Moreover, the similarity of form between these records and the familiar fixed-interval "scallop" is quite striking; turning the Figure upside down shows this resemblance.

Table 1 summarizes the mean results for all four test sessions and shows the individual results for the first two sessions. The Table shows response rate over the fixed intervals following reinforcement (R) and blackout (N for nonreinforcement) at the end of the preceding variable intervals. The data are divided into two classes depending upon whether N or R was delivered for pecking (if the variable interval was less than 1 min) or not-pecking (if the variable interval was greater than 1 min). Data

⁴ The following sequence of variable-duration interreinforcement intervals was used [(N) indicates an interval ending in nonreinforcement during omission tests]: 4t(N), 6t, 9t(N), 4t, 12t(N), 18t(N), 2t, 12t, 24t, t, 6t(N), 2t(N), 9t, 18t, 24t(N), t(N), 4t, 2t(N), 24t, 12t(N), t, 9t(N), 4t(N), 2t, t(N), 18t, 18t(N), 9t, 24t(N), 6t(N), 6t and 12t sec. Variable intervals were separated by fixed intervals of 6t sec; t was 6.3 sec on the first two tests, 5.0 sec on the second two. Two cycles of this sequence made up a session.

FIG. 1. Cumulative response records taken from day 19 of the uncued procedure for the five birds. Diagonal blips indicate reinforcements.

for the control sessions show the response rates during fixed intervals corresponding to the N and R intervals of the reinforcement omission sessions.

Response rate during fixed intervals following N or R delivered for not-pecking (i.e., at postreinforcement times greater than 1 min, right two

TABLE 1
Responses per Minute During Fixed Intervals Following Reinforcement (R) or Blackout (N) Received Either for Pecking or Not-Pecking During Omission Test Sessions^a

Bird No.	Condition	Pecking		Not-pecking	
		R	N	R	N
95	Omission	48.6	24.2	76.4	46.4
	Control	39.0	32.0	57.3	56.3
96	Omission	70.6	59.9	91.6	38.6
	Control	46.5	59.2	30.0	35.2
85	Omission	63.2	67.1	85.9	47.8
	Control	53.9	59.7	76.8	74.4
56	Omission	40.4	38.9	55.7	23.2
	Control	60.8	38.3	80.6	81.8
93	Omission	27.9	28.4	43.9	21.4
	Control	28.7	25.7	34.5	47.4
Mean (1)	Omission	50.1	43.7	70.7	35.5
	Control	45.8	43.0	55.8	59.0
Mean (2)	Omission	48.3	37.7	69.4	12.7
	Control	51.5	52.5	66.1	64.0

^a Control session data are for comparable intervals when no reinforcements were omitted. Individual data are shown for the first omission and control sessions, mean data for all four test sessions.

columns) was always *lower* following N than following R. However, responding following N was higher than if no stimulus at all was presented, since all the birds responded slowly or not at all at postreinforcement times greater than 1 min (cf. Fig. 1). Thus, cumulative records for omission test sessions showed brief bursts of responding after most blackouts. Differences between responding following N and R were much smaller when these events occurred following a key-peck (two left columns). Control session data showed no reliable differences between response rate during the fixed intervals corresponding to N and R intervals, ruling out sequential effects as explanations of differences. The second omission test was similar to the first but showed even less responding following N for not-pecking (bottom two rows).

DISCUSSION

These results are not consistent with a simple frustrative nonreward interpretation, since responding following nonreinforcement (blackout) was not enhanced; it was either depressed (if blackout followed not-pecking) or showed little change (if blackout followed pecking) in comparison with rate following reinforcement. Moreover, a more sophisticated interpretation in terms of an enhanced "response of not responding" will not do because when blackout was received for not-pecking, subsequent responding though depressed with respect to rate following reinforcement was elevated with respect to responding when neither reinforcement nor blackout occurred. It was also elevated with respect to the rate just preceding blackout (since blackout was received for not-pecking). A similar argument can also be made in the case of fixed-interval schedules: response rate following a brief blackout presented in lieu of reinforcement is depressed with respect to rate over the same interval when no stimulus is presented in lieu of reinforcement (cf. Staddon & Innis, 1969). These relationships are summarized in the schematic cumulative response records of Fig. 2. In neither case, therefore, is it correct to speak of nonreinforcement as enhancing the prevailing (in the sense of just-preceding) mode of responding.

Both fixed-interval schedule results and the results of this experiment are consistent with the hypothesis offered earlier: a stimulus presented in lieu of reinforcement (nonreinforcement) produces after-effects similar to reinforcement but of smaller magnitude (i.e., the principle of stimulus generalization applied to stimulus after-effects). In most reinforcement omission experiments, nonreinforcement involves presentation of a stimulus quite similar to reinforcement, e.g., a brief blackout in lieu of food presentation in the dark, as in this experiment, or an empty goalbox, as in runway experiments. Thus, if reinforcement depresses responding

FIG. 2. Schematic cumulative response records illustrating the results of either no stimulus (O), nonreinforcement (N, a stimulus similar to reinforcement), or reinforcement (R), presented at points A on procedures producing either a positive goal-gradient (+G), e.g., a fixed-interval 1-min schedule, or a negative goal-gradient (-G), e.g., the uncued procedure of this experiment. The three curves starting at points A show the pattern of responding typically following each event.

that follows it, as in fixed-interval schedules (inhibitory after-effects), nonreinforcement should be inhibitory also, but less strongly, yielding an elevation in response rate apparently due to nonreinforcement (a positive "frustration" effect). Similarly, if reinforcement elevates responding that follows it (excitatory after-effects or temporal excitation), as in this experiment, nonreinforcement should be less excitatory yielding a depression in responding (a negative "frustration" effect).

The difference between the effects of nonreinforcement at short and long postreinforcement times (i.e., for pecking and not-pecking) is also comprehensible in these terms. Response rate here was high at short postreinforcement times and low at long times. Presentation of reinforcement at short postreinforcement times therefore produced little change in rate, since responding was already occurring at a high rate. *A fortiori* nonreinforcement at short postreinforcement times also produced little change in rate. Thus, little difference between the effects of reinforcement and nonreinforcement was observed when these events occurred at short postreinforcement times (i.e., for pecking). On the other hand, presentation of reinforcement at long postreinforcement times, when behavior was occurring at a low rate, produced a large elevation in subsequent responding, with nonreinforcement again producing a proportionately smaller temporal excitatory effect. The difference in the increments in responding

produced by these two events was therefore much more obvious when they occurred at long postreinforcement times (i.e., for not-pecking) than at short.

Finally, the differences between the first and second omission test sessions are also consistent with this analysis. The similarity between the after-effects of reinforcement and blackout in this situation presumably owes something to the similarities between these two events. One similarity is their invariable pairing prior to the first omission test. This invariable association was weakened on the first test, however, which afforded an opportunity for discrimination and may thus have increased the dissimilarity between the two events. Since reinforcement has excitatory after-effects in this situation, increasing dissimilarity between reinforcement and blackout implies weakened excitatory after-effects due to the blackout in the second test by comparison with the first, as was in fact observed.

There is no simple procedure in the runway which is equivalent to the uncued procedure of this experiment, and thus there is no opportunity for a strong test of the after-effects hypothesis within the runway methodology. Less powerful tests are available, however, although some pilot exploration may be necessary to decide on appropriate parameter values. The assumption to be tested is that stimuli in G1 have inhibitory after-effects, so that (after training) a reduction in the number of stimuli in G1 should either produce no change in A2 start speed, or an increase, but never a decrease: there will be no change if the stimulus removed is not salient and thus exerts little discriminative control, and there will be an increase if the stimulus does exert some (inhibitory) control. Thus the stimuli to be manipulated should approach food reward in salience if an interpretable result is to be obtained. Two possibilities are the food cup, which is strongly associated with food, and mild electric shock through the floor of G1. In the first case, following the usual training with reward in both G1 and G2, elimination of both food and food cup should yield a larger elevation in A2 start speed than elimination of food alone (there may be no difference if the food cup exerts no control, in which case no conclusion can be drawn). This result is counter to frustration theory, which would predict a smaller frustration effect when both food and food cup are eliminated because the situation is now less like a "situation where [reward] had been present previously" (Amsel, 1958, p. 102). In the second case rats could be trained to run to G1 where they would receive both food and mild foot-shock (presumably a salient stimulus). They could then be tested with food omitted, shock omitted, and food and shock omitted. The after-effects view would again predict that if anything the rats should run faster in A2 when both food and shock are omitted than when only food is omitted; but frustration theory would predict the

opposite, since both shock and nonreward are aversive events, and thus the two in combination should produce the greatest elevation in running speed.

In the absence of direct experimental tests of this sort one can only argue from parsimony and by analogy. Nevertheless, both the arguments here presented and the results of this experiment cast considerable doubts on the correctness of emotional or motivational interpretations of the frustration effect, in runways as well as in free-operant experiments.

REFERENCES

- AMSEL, A. The role of frustrative nonreward in noncontinuous reward situations. *Psychological Bulletin*, 1958, **55**, 102-119.
- AMSEL, A., & ROUSSEL, J. Motivational properties of frustration: I. Effect on a running response of the addition of frustration to the motivational complex. *Journal of Experimental Psychology*, 1952, **43**, 363-368.
- ANGER, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, 1956, **52**, 145-161.
- CATANIA, A. C., & 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 (Whole Part 2).
- CRAIG, W. Appetites and aversions as constituents of instincts. *Biological Bulletin of the Marine Biological Laboratory*, 1918, **34**, 91-107.
- DALY, H. B. Excitatory and inhibitory effects of complete and incomplete reward reduction in the double runway. *Journal of Experimental Psychology*, 1968, **76**, 430-438.
- DAVENPORT, J. W., FLAHERTY, C. F., & DYRUD, J. P. Temporal persistence of frustration effects in monkeys and rats. *Psychonomic Science*, 1966, **6**, 411-412.
- DAVENPORT, J. W., & THOMPSON, C. I. The Amsel frustration effect in monkeys. *Psychonomic Science*, 1965, **3**, 481-482.
- DEWS, P. B. The effect of multiple S^A periods on responding on a fixed-interval schedule. *Journal of the Experimental Analysis of Behavior*, 1962, **5**, 369-374.
- DEWS, P. B. A behavioral effect of amobarbital. *Naunyn-Schmiedeberg's Archiv für Experimentelle Pathologie und Pharmakologie*, 1964, **248**, 296-307.
- FARMER, J., & SCHOENFELD, W. N. Effects of a DRL contingency added to a fixed-interval reinforcement schedule. *Journal of the Experimental Analysis of Behavior*, 1964, **7**, 391-399.
- FERRARO, D. P., SCHOENFELD, W. N., & SNAPPER, A. G. Sequential response effects in the white rat during conditioning and extinction on a DRL schedule. *Journal of the Experimental Analysis of Behavior*, 1965, **8**, 255-260.
- FERSTER, C. B., & SKINNER, B. F. *Schedules of reinforcement*. New York: Appleton-Century-Crofts, 1957.
- GRAY, J. A. Sodium amobarbital and effects of frustrative nonreward. *Journal of Comparative and Physiological Psychology*, 1969, **69**, 55-64.
- HINRICHS, J. V. Disinhibition of delay in fixed-interval instrumental conditioning. *Psychonomic Science*, 1968, **12**, 313-314.
- HULL, C. L. The rat's speed-of-locomotion gradient in the approach to food. *Journal of Comparative Psychology*, 1934, **17**, 393-422.
- INNIS, N. K., & STADDON, J. E. R. Scopolamine and reinforcement omission on fixed-interval schedules. *Psychonomic Science*, 1969, **14**, 43-45.

- ISON, J. R., DALY, H. B., & GLASS, D. H. Amobarbital sodium and the effects of reward and nonreward in the Amsel double runway. *Psychological Reports*, 1967, **20**, 491-496.
- KAMIN, L. J. Predictability, surprise, attention, and conditioning. In B. A. Campbell and R. M. Church, (Eds.), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969.
- KELLEHER, R. T., FRY, W., & COOK, L. Interresponse time distribution as a function of differential reinforcement of temporally spaced responses. *Journal of the Experimental Analysis of Behavior*, 1959, **2**, 91-106.
- LASHLEY, K. S. The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior*. New York: Wiley, 1951.
- LOGAN, F. A. Frustration effect following correlated non-reinforcement. *Journal of the Experimental Psychology*, 1968, **78**, 396-450.
- LOTT, D. F. Secondary reinforcement and frustration: A conceptual paradox. *Psychological Bulletin*, 1967, **67**, 197-198.
- LUDVIGSON, H. W. A preliminary investigation of the effects of sodium amytal, prior reward in G1, and activity level on the FE. *Psychonomic Science*, 1967, **8**, 115-116.
- MACKINNON, J. R., & AMSEL, A. Magnitude of the frustration effect as a function of confinement and detention in the frustrating situation. *Journal of Experimental Psychology*, 1964, **67**, 468-474.
- MCBOSE, J. H. Delay increase and subsequent responding. *Psychonomic Science*, 1966, **5**, 213-214.
- MCBOSE, J. H., & LUDVIGSON, H. W. Role of reward magnitude and incomplete reduction of reward magnitude in the frustration effect. *Journal of Experimental Psychology*, 1965, **70**, 490-495.
- PAVLOV, I. P. *Conditioned reflexes*. (Translated by G. V. Anrep) London: Oxford University Press, 1927.
- PECKHAM, R. H., & AMSEL, A. Magnitude of reward and the frustration effect in a within-subjects design. *Psychonomic Science*, 1964, **1**, 285-286.
- PECKHAM, R. H., & AMSEL, A. Within-subject demonstration of a relationship between frustration and magnitude of reward in a differential magnitude of reward discrimination. *Journal of Experimental Psychology*, 1967, **73**, 187-195.
- RASHOTTE, M. E., & AMSEL, A. Transfer of slow-response rituals to extinction of a continuously rewarded response. *Journal of Comparative and Physiological Psychology*, 1968, **66**, 432-443.
- REYNOLDS, G. S. Behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1961, **4**, 57-71.
- SCHNEIDER, B. A. A two-state analysis of fixed-interval responding in the pigeon. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 677-687.
- SEWARD, J. P., PEREBOOM, A. C., BUTLER, B., & JONES, R. B. The role of prefeeding in an apparent frustration effect. *Journal of Experimental Psychology*, 1957, **54**, 445-450.
- SGRO, J. A. Complete removal and delay of sucrose reward in the double alleyway. *Journal of Comparative and Physiological Psychology*, 1969, **69**, 442-447.
- SGRO, J. A., GLOTFELTY, R. A., & MOORE, B. D. Delay of reward in the double alleyway. *Journal of Experimental Psychology*, 1970, **84**, 82-87.
- SGRO, J. A., GLOTFELTY, R. A., & PODLESNI, J. A. Contrast effects and delay of reward in the double alleyway. *Psychonomic Science*, 1969, **16**, 29-31.
- SHEPARD, R. N. Approximation to uniform gradients of generalization by monotone transformations of scale. In D. J. Mostofsky (Ed.), *Stimulus generalization*. Stanford: Stanford University Press, 1965. Pp. 94-110.

- SHERRINGTON, C. S. *The integrative action of the nervous system*. New Haven: Yale University Press, 1906.
- STADDON, J. E. R. Some properties of spaced responding in pigeons. *Journal of the Experimental Analysis of Behavior*, 1965, **8**, 19-27.
- STADDON, J. E. R. Attention and temporal discrimination: factors controlling responding under a cyclic-interval schedule. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 349-359.
- STADDON, J. E. R. Inhibition and the operant. A review of "Sensory Inhibition" by G. V. Békésy, and "Mach Bands: Quantitative Studies on Neural Networks in the Retina" by Floyd Ratliff. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 481-487.
- STADDON, J. E. R. Effect of reinforcement duration on fixed-interval responding. *Journal of the Experimental Analysis of Behavior*, 1970, **13**, 9-11.
- STADDON, J. E. R., & INNIS, N. K. An effect analogous to "frustration" on interval reinforcement schedules. *Psychonomic Science*, 1966, **4**, 287-288.
- STADDON, J. E. R., & INNIS, N. K. Reinforcement omission on fixed-interval schedules. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 689-700.
- TERRACE, H. S. Stimulus control. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application*. New York: Appleton-Century-Crofts, 1966. Pp. 271-344.
- WAGNER, A. R. The role of reinforcement and nonreinforcement in an "apparent frustration effect." *Journal of Experimental Psychology*, 1959, **57**, 130-136.

(Received July 17, 1969)