7922405

DAUGHTRY, TIMOTHY CARR
EFFECTS OF RATE, PATTERNING, AND CONTINGENCY OF REINFORCEMENT UPON PERCEIVED CONTROL AND LEARNED HELPLESSNESS.

THE UNIVERSITY OF NORTH CAROLINA AT GREENSBORO, PH.D., 1979

University Microfilms International 300 N. ZEEB ROAD, ANN ARBOR, MI 48106

EFFECTS OF RATE, PATTERNING, AND CONTINGENCY
OF REINFORCEMENT UPON PERCEIVED CONTROL
AND LEARNED HELPLESSNESS

by

Timothy Daughtry

A Dissertation Submitted to
the Faculty of the Graduate School at
The University of North Carolina at Greensboro
in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Greensboro
1979

Approved by

Dissertation Adviser

APPROVAL PAGE

This dissertation has been approved by the following committee of the Faculty of the Graduate School at the University of North Carolina at Greensboro.

Dissertation Adviser ____________________

Committee Members ____________________

____

____

____

Dec 11, 1978
Date of Acceptance by Committee

Sept 29, 1978
Date of Final Oral Examination

DAUGHTRY, TIMOTHY. Effects of Rate, Patterning, and Contingency of Reinforcement upon Perceived Control and Learned Helplessness. (1978)
Directed by: Dr. Richard L. Shull. Pp. 78.

The effects of rate, patterning, and contingency of reinforcement upon ratings of perceived control and subsequent performance on anagrams were explored. Sixty college students learned a pretreatment task and received either contingent or noncontingent reinforcement. Under each of these conditions, either a high or low rate of reinforcement was given. The noncontingent conditions were further subdivided into yoked (increasing) and random (unchanging) patterns of reinforcement. It was found that higher rates of reinforcement were perceived as more controllable than lower rates regardless of actual contingency. It was concluded that human judgements of control are not based on the controllability of outcomes by responding. None of the pretreatments affected subsequent anagram performance. Several possible reasons for the failure to find an effect on anagram performance were discussed.

ACKNOWLEDGEMENTS

The author gratefully acknowledges the assistance and advice of Dr. Richard L. Shull, who served as committee chairman, and of Drs. Kendon Smith, P. Scott Lawrence, John Seta, and Charles Church who served as committee members.

Of course, J. Vernon Odom's assistance as a second experimenter was invaluable in conducting the experiment. Gratitude is also expressed to Nancy Elliott for assistance in computer programming and to Dr. David Soderquist for consultation on data analysis.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Page

INTRODUCTION

The perception of causation of environmental events often has a notably immediate and convincing quality about it: a moving object strikes a stationary one, and the latter is quite readily seen as being "caused" to move by the former. In situations in which the causative agent of such mechanical actions is our own behavior, our perception of personal causation is again so immediate as to arouse little interest under normal circumstances: one's hand pushes an object and the object moves, with the amount of movement being commensurate with the amount of effort exerted. So compelling, in fact, is the perception of causation in such cases that Gestalt theorists often considered it to be innate. But what of situations in which the action is not a simple mechanical one, or when the correspondence between behavior and outcome is complex, rather than a strictly one-to-one correspondence? Consider for example the perception of causation in the act of summoning an elevator by pressing a button. The elevator may arrive after a considerable delay, or it may arrive quickly. Complicating matters is the fact that the elevator may arrive even when the button has not been pressed. Yet, with repeated experience in such situations, most people develop the belief that they have some degree of control over the arrival of the elevator. Obviously, social learning variables (e.g. instructions) have been excluded from the foregoing example, but it perhaps illustrates the fact that the perception of causality is indeed a complex affair. Furthermore, review of a number of current psychological journals reveals that the study of beliefs about personal control, and the effects of these beliefs on behavior, is receiving a considerable amount of attention from researchers and theorists.

As it happens, however, it is possible to find quite divergent theoretical and empirical work concerning the nature and accuracy of human judgements of complex relationships. On the one hand, Seligman's theory of "learned helplessness" (Seligman, 1975) has at its core the tenet that animals and humans develop relatively veridical internal representations of the degree of relatedness obtaining between environmental stimuli, or between their own behavior and their outcomes. This theory also contends that, within limits, these perceptions transfer to situations beyond the ones in which they were formed. On the other hand, Jenkins and Ward (1965) and Smedslund (1963) have found that human judgements of contingencies are far from accurate, with the inaccuracies following a consistent pattern. The present research examines the effect of rate, patterning, and contingency of reinforcement upon judgements of reinforcemental control in human subjects, and the effects of these judgements upon performance in a subsequent task.

Seligman's theory of learned helplessness has received considerable attention of late, and it will be presented here in some detail. One possible reason for the interest focused on this theory is that the arguments surrounding it are reminiscent of the Tolman-Hull disagreements of some years ago. First, Seligman's theory is explicitly a cognitive one at a time when cognitive explanations of psychological phenomena are rising in prominence. In contrast, other theorists, perhaps best exemplified in a thorough review by Levis (1976), have mounted S-R interpretations of the helplessness phenomena. The issue of "what is learned?" (a cognitive representation of environmental contingencies versus specific behaviors) is very much alive in discussions of Seligman's work. Another issue, reminiscent of the "molecular-molar" question, is whether the effective control of behavior is best analyzed at a trial-by-trial level (contiguity) or whether the animal in some way averages events over time (contingency). The present experiment is not an attempt to resolve such issues as whether an S-R or a cognitive interpretation has more utility. Theorists of either persuasion could probably construct convincing accounts of the results. The present experiment, rather, focuses on critical variables which influence what is learned.

I shall begin by reviewing Seligman's theory and the experimental work with animals which has been offered as supportive of the theory, following that by a review of attempts to replicate the learned helplessness effects using human subjects. Then I will present the work of Jenkins and Ward (1965) and others who have explored judgements of contingency made by human subjects. The latter studies antedate the original work in which Seligman begins to formulate the notion of helplessness, but their results suggest variables which could indeed be used in exploring the boundaries of the helplessness effect. After reviewing these studies, the rationale for the present experiment will be presented.

Learned Helplessness: The Essential Theory and Findings

A central concept in Seligman's theory is that of the instrumental conditioning space (Figure 1), which relates the probability of reinforcement to the presence or absence of responding. Along the abscissa, the probability of reinforcement, given that a response has occurred, varies from zero (extinction) to 1.0 (continuous reinforcement). Of course, situations also exist in which the presentation of reinforcers depends on the absence of responding, and these situations are represented along the ordinate. Points within the conditioning space reflect various combinations of reinforcement probabilities for responding and not responding. On a schedule described by point "b," for example, reinforcement probability is .75 for responding and .25 in the absence of responding. Seligman's primary interest, however, is in situations described by points along the diagonal (e.g., point _a_). It is readily apparent that, at any point along the

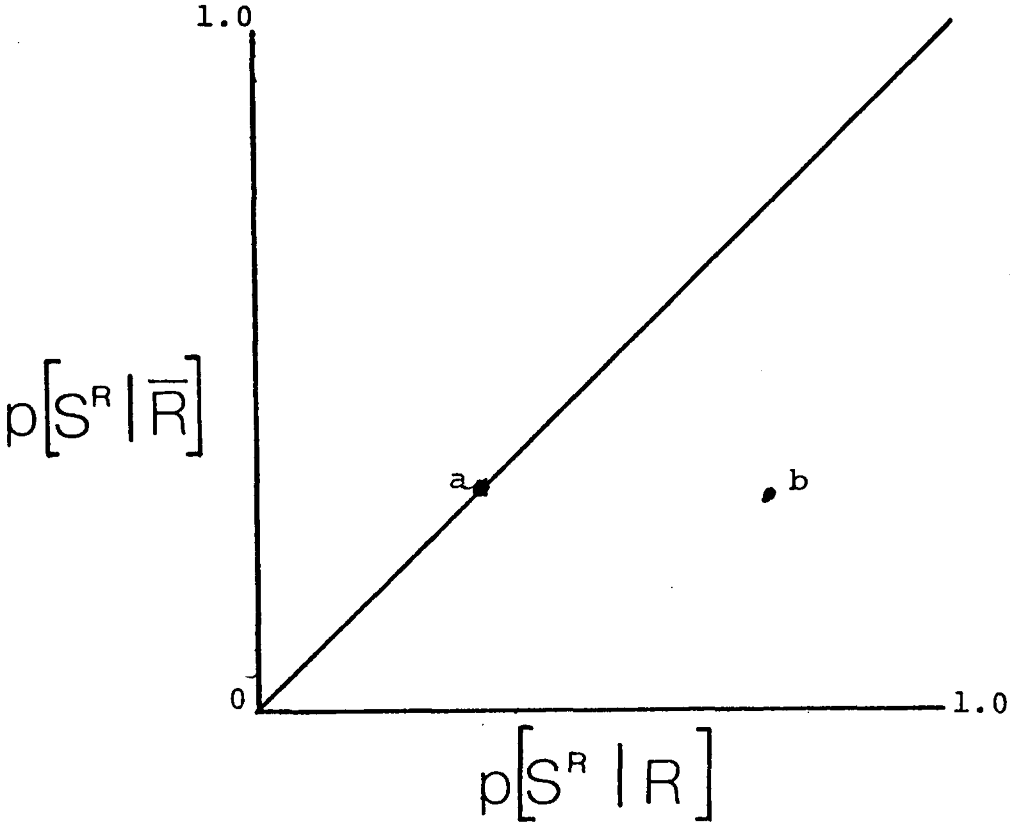

Figure 1. Instrumental Training Space (Adapted from Seligman, 1975)

diagonal, the probabilities of reinforcement in the presence and in the absence of responding are identical. Reinforcement is uncontrollable along this diagonal, in that neither responding nor refraining from responding alters reinforcement probability. According to Seligman (1975), the organism forms internal representations (expectancies) reflecting the response-reinforcer relationships to which it is exposed. When exposed to an uncontrollable reinforcer (any point along the diagonal), the organism learns that its responses do not affect reinforcement probability. This expectation removes the incentive for active responding, and the animal becomes "helpless."

Seligman (1975) has summarized the experimental design and essential findings which have contributed to the development of this theoretical approach. An essential element in learned helplessness experimentation is the use of the triadic design. One group of subjects receives some outcome contingently upon responding, i.e., the outcome is controllable. A second group is yoked to the first, receiving the same number and pattern of outcomes, but independently of responding. A third group receives no pretreatment. All three groups are then given some test in which the outcomes are controllable. The yoked procedure is necessary to demonstrate that it is uncontrollability of the outcome, rather than simple exposure to it *per se*, which accounts for any differences on test trials.

The actual helplessness pretreatment used by Seligman and his colleagues involved strapping mongrel dogs in a harness and exposing them to 64 unsignalled, inescapable 6.0-m.A. shocks. Control animals received either no pretreatment or controllable shock. A day later, all dogs were tested on 10 trials of escape-avoidance learning in a standard 2-way shuttlebox. Helplessness was defined as the failure of the group receiving inescapable shock to learn the escape-avoidance task; their latencies for barrier-jumping were significantly longer than those of the other two groups, and they failed to escape at all on a greater number of trials. Further, it was reported that dogs in the inescapable shock condition often sat passively and accepted shock on the test trials (Overmier & Seligman, 1967; Seligman & Maier, 1967). These deficits in response initiation, according to Seligman, demonstrated a motivational deficit. He also suggested that helplessness involved a cognitive deficit, in that helpless dogs who did occasionally jump the barrier and escape shock on one trial reverted to passivity on subsequent trials. With naive dogs, on the other hand, the occurrence of an escape response was a reliable predictor that further escapes would occur. On Seligman's analysis, prior learning that responding and shock termination are independent proactively interfered with learning the new contingency between shock offset and jumping the barrier. Maier and Seligman (1976) also present an emotional deficit, defined as the physiological effect (ulcers, etc.) of exposure to uncontrollable aversive events.

A variety of experimental variations on the above procedure have been performed. Seligman and Maier (1967) "immunized" one group of dogs by giving them escapable shock in the shuttle box before exposure to the inescapable shock in the harness, and found that this procedure mitigated the effect of the inescapable shocks: immunized dogs learned the test task as rapidly as did naive dogs. Overmier (1968) demonstrated that inescapable preshock interfered with avoidance learning alone if escape was blocked. Overmier and Seligman (1967) attempted to refute the argument that helplessness might be due to the adventitious reinforcement of responses during pretreatment which were incompatible with escape-avoidance. They curarized dogs during pretreatment in an effort to prevent the development of overt patterns of responding, and they still found the interference with escape-avoidance learning which typifies learned helplessness. Also, Seligman, Maier, and Geer (1968) found that forcibly exposing helpless dogs to the prevailing contingency by dragging them across the barrier eventually alleviated helplessness. It was reasoned that the dogs' expectation that responding and shock termination were independent was gradually altered in this manner. Finally, Seligman, Maier, and Solomon (1971) and Seligman (1975) review studies which replicated the helplessness results with a variety of species, including cats, rats, fish, and man.

At this point, it will be well to review several important points regarding the above conceptualization of the learning process and the portions of this conceptualization which are explicitly included in the statements defining helplessness _per se_. At several points, it is stated or implied that organisims simply learn about contingencies. That is, organisims are sensitive to points throughout the conditioning space (Figure 1). The clearest statements to date, however, of the helplessness phenomena are more restricted. It is stated that ". . . the first step in the theory is that the organism acquires an expectation of response-outcome independence, when outcomes are uncontrollable" (Maier and Seligman, 1976, p. 18). The interference with learning and performance, then, is produced because there is no expectation that responding will produce relief. Though the more general statement that organisms are responsive to points throughout the conditioning space is not included in these statements regarding helplessness, they are strongly implied, and are, I believe, appropriate for inclusion in discussions of Seligman's work.

A second point needing emphasis is that the two steps in the theoretical statement, that organisms actively learn about uncontrollability and that this learning interferes with subsequent learning, _can_ be seen as independent. It is possible, for example, for an organism to develop an expectation about contingency but, for various reasons, not to transfer that expectation to other situations. As Maier and Seligman (1976) note, the factors influencing the development of the expectation and the factors influencing transfer need to be explored as boundary conditions of helplessness theory.

Most learned helplessness studies with animals have employed uncontrollable aversive events. It is worth mentioning that a few studies have examined the effects of uncontrollable appetitive events, though they have not proven notably supportive of the learned helplessness theory. Engberg Hansen, Welker, and Thomas (1973) found slower acquisition of an autoshaped key-peck in pigeons exposed to prior noncontingent food compared to pecking in control subjects. In discussing the Engberg et al. findings, Gamzu, Williams, and Schwartz (1973), however, noted that the removal during the test phase of the treadle used by birds in the contingent food group also removed the stimulus for competing behaviors and, thus, rendered questionable the adequacy of this control group. Such competing behaviors could have slowed acquisition of key-pecking and thus eliminated the difference between the contingent and noncontingent groups. Hulse (1974) found that prior noncontingent food retarded the acquisition of bar pressing in rats. He also found, however, that if the probability of noncontingent food gradually decreased during the pretreatment, the learned helplessness interference effect was not found. Much of the learned helplessness research has been done with human subjects, however, and it is to the human analogue studies that we now turn. These studies will illustrate the variety and complexity of research in the helplessness area. Particular emphasis will be given to studies in which rate and patterning of reinforcement may be illustrated as important variables.

Learned Helplessness: Human Analogue Studies

Maier and Seligman (1976) restrict their formal statements to situations involving uncontrollable *aversive* events. Some human studies involve uncontrollable aversive noise or shock, and others employ insoluble problems as an analogue of *appetitive* (nonaversive) events. Certainly, it can be maintained that failure to receive an expected positive reinforcer is aversive, and that the above dichotomy is really one of degree. Seligman's frequent discussion of helplessness effects with nonaversive stimuli, and his own research in the area, however, should suffice as reasons for including such studies on an equal footing with those using aversive events.

With human subjects, insoluble problems are often used to induce perceptions of helplessness. An insoluble problem can be described by a point at the origin of the instrumental conditioning space, as the probability of reinforcement is zero regardless of the nature of the subject's response. In an early study, Fosco and Geer (1971) gave four groups of subjects varying numbers of insoluble problems. Subjects attempted to guess the correct combination of button presses in order to avoid being shocked. On trials designated as insoluble, the subject was shocked regardless of the combination of his button presses. On soluble test problems, it was found that errors increased significantly with increased prior experience with insoluble problems, a result which was interpreted as consistent with the learned helplessness model. A noteworthy interpretative problem, also cited in a review by Wortman and Brehm (1975), is that subjects receiving more insoluble problems also received more shocks; we are unable to tell whether insolubility alone is sufficient to account for the results.

Another early experiment in this area was reported by Thornton and Jacobs (1971). Some subjects could avoid shock on a choice reaction time task, while others received unavoidable shock as they worked on the task. All subjects were then tested on another task in which shock could be avoided by pressing the correct combination of buttons. The group which had received avoidable shock on the earlier task responded faster on the test task than those who had received uncontrollable shock. The authors noted that the group receiving uncontrollable shock performed as well as a group which received no shock during the first phase. The obtained differences, then, appear to represent a facilitation effect after receiving avoidable shock rather than interference after unavoidable shock. Furthermore, an interpretative problem arises in that, during instructions for the pretreatment, subjects in each condition were informed as to whether their responding would affect the occurrence of shock. Thus, it is difficult to determine the extent to which the instructional set alone, independently of shock contingency, could account for the results.

Thornton and Jacobs (1972) used a similar version of the above pretreatment task, and then assessed the effect of these experiences on a subsequent test of intellectual performance. Instead of the expected interference in the group receiving uncontrollable shock, however, they found facilitation of intellectual performance. They were able to replicate this surprising result in a second experiment. No clear reason could be cited for this pattern of results, but it was suggested that the dissimilarity between the pretreatment and test tasks might have served to disrupt the generalization of helplessness. The realization that the second task was controllable, they suggested, might have increased the motivation to do well for subjects in the uncontrollable shock condition. It was also suggested that a Hullian drive concept might be invoked as an explanatory device, in that increased stress from uncontrollable shock might have increased drive level and thus facilitated performance. At any rate, the results suggest some possible limitations on generalization of helplessness.

An experiment which did find substantial transfer from a pretreatment using shock to a dissimilar test task was performed by Glass and Singer (1972). Subjects in the No Perceived Avoidance condition were led to believe that they received shocks because of their failure to solve puzzles (actually insoluble). Subjects in the Perceived Avoidance condition believed that their successes on puzzles (actually soluble) allowed them to avoid shocks. Though subjects in both conditions received a similar number of shocks, those in the No

Perceived Avoidance condition displayed interference on test tasks (proofreading and the Stroop Color Word Test). Post-experimental ratings by subjects suggested that perceptions of helplessness mediated the transfer effects.

Roth and Bootzin (1974) specifically addressed the issue of generalization of helplessness to a task which was explicitly dissimilar to the pretreatment task. They used random (noncontingent) reinforcement on a concept learning task as the pretreatment, rather than the familiar electric shock. As was the case for Thornton and Jacobs (1972), a reversal of the predicted helplessness effect was found: subjects receiving prior noncontingent feedback showed more active attempts at control, on a subsequent task, than subjects who had experienced contingent reinforcement or no pretreatment in the first phase. The authors suggest that a curvilinear relationship might exist between amount of exposure to noncontingent reinforcement and degree of active attempts at control: organisms may struggle harder to gain control when first exposed to noncontingent reinforcement, and with continued exposure, may eventually give up. Whether facilitation or helplessness is found would then depend on the amount of exposure to noncontingent reinforcement. Roth and Kubal (1975) varied amount of exposure to noncontingent reinforcement as well as importance of the pretreatment task. For subjects working on an important task during the pretreatment, the curvilinear relationship was found: more exposure to noncontingency increased the probability of helplessness and decreased the probability of facilitation. Only facilitation effects were found for a task presented as unimportant to subjects. Relevant to the present experiment is the fact that random reinforcement, rather than yoked, was used in the noncontingent reinforcement condition. The potential importance of a random (unchanging) and yoked (increasing) pattern of noncontingent reinforcement will be discussed shortly.

Thornton and Powell (1974) report a series of experiments employing variations on the design used by Thornton and Jacobs (1971; 1972). The importance of instructions regarding shock contingency was revealed in an experiment on alleviation of helplessness. Subjects who were informed that shock was controllable in the test task did not show helplessness following noncontingent shock, though helplessness did emerge for subjects receiving noncontingent shock without such "alleviation" treatment. These results suggest a limitation on the transfer of helplessness effects from pretreatment to test task; i.e., factors which increase the perceived controllability of the test task may reduce the probability of helplessness.

Hiroto (1974) applied the triadic design used in the earlier animal studies to noise-escape in humans. On a number of test measures, subjects who had received prior inescapable noise were inferior to those who had received escapable noise or no noise. Extending this procedure, Hiroto and Seligman (1975) performed a series of experiments designed to assess the generalization of helplessness. As these experiments are particularly illustrative of the question of rate and patterning of noncontingent reinforcement, they will be discussed in the following paragraphs at some length.

Two kinds of tasks, labelled as "Instrumental" (the noise-escape task of Hiroto, 1974) or "Cognitive" (a discrimination task using "correct-incorrect" feedback), were employed in the pretreatment phase. For the Instrumental task, escape from aversive noise was either contingent or noncontingent (yoked to contingent). Additionally, a feedback light informed subjects in the contingent condition whether noise offset occurred because of their response or because the noise had "timed out." Subjects in the noncontingent condition always received the light which indicated that noise offset was caused by the timer, not by their response. In the Cognitive pretreatment, feedback was either contingent (veridical) or noncontingent (a prearranged, random schedule of "correct-incorrect" feedback). Three discrimination problems were used. In addition to the trial-by-trial feedback, subjects also tried to report the solution of the problem at the completion of each one. Most of the subjects receiving contingent veridical feedback successfully reported the solution. Subjects in the noncontingent condition were always told that their solution was incorrect.

The test tasks on which helplessness was later assessed were also labelled as "Instrumental" (a different noise-escape task) or "Cognitive" (a series of five-letter anagrams). Four experiments were performed in all, pairing each of the two pretreatments with each of the two test tasks (Instrumental-Instrumental, Instrumental Cognitive, Cognitive-Instrumental, Cognitive-Cognitive). For all combinations except Cognitive-Cognitive, the authors report helplessness effects (more errors and longer response latencies) following the noncontingent pretreatments. A subsequent experiment, adding a fourth problem to the Cognitive pretreatment, found the predicted interference on the anagram test task. The authors interpreted these findings as being indicative of a wide degree of generalization of the helplessness effect, and suggested that such an expectancy of response-reinforcer independence could lie at the core of human depression. It is pointed out in the report, however, and deserves reiteration here, that subjects did perceive all of the tasks as part of the same experiment. Such a perception of commonality could mediate generalization, even between tasks which are quite different in content. A more impressive demonstration of generalization would involve some unobtrusive measure of transfer, taken in a situation which is perceived as distinct and separate from the one in which the pretreatment was administered.

With respect to generalization across tasks of different modalities, which was the essential focus of the Hiroto and Seligman (1975) study, it should be pointed out that the Instrumental-Cognitive distinction is a questionable one. All tasks involved the shaping of a response, whether verbal or manual, by the consequence of the response. It is not clear, in other words, why the tasks involving verbal responses were labelled as cognitive. The more relevant point for the experiment to be reported here, however, is the manner in which noncontingent reinforcement was administered. Consider the question of rate of reinforcement, for example, in the Instrumental tasks. Recall that a feedback light was used to differentiate noise offset which occurred because of the "timing out" of the apparatus and that which occurred because the subject had responded correctly. Also recall that the noncontingent condition involved presenting the "incorrect" light (noise offset because of timer) after each response. Though the noncontingent condition was yoked to the contingent condition for total exposure to noise, the noncontingent condition received only the "incorrect" light. Granted, the procedure is consistent with an essential point in defining uncontrollability--that variations in responding for subjects in the noncontingent condition did not alter the probability of noise offset or the occurrence of the "incorrect" light. Such a procedure is actually defined by the origin or zero point of the instrumental training space (Figure 1), in that the probability of reinforcement is zero, regardless of the subject's responses. It is unclear, however, whether the perception of control (assessed in postexperimental ratings of helplessness) and the interference on the second task results from uncontrollability or from the zero rate of reinforcement. In other words, we are unable to predict whether differences would emerge between contingent and noncontingent reinforcement conditions if the rates of reinforcement were equal.

The above arguments relate to the question of how organisms learn. If contingency is the critical variable, and people actually learn about contingency, then people receiving noncontingent reinforcement at a nonzero rate would detect the noncontingency. That is, differences in perceived control would occur between groups receiving contingent and noncontingent reinforcement at an equal rate. The implication is that all points along the uncontrollability diagonal (Figure 1) would be equivalent in producing helplessness effects. If, on the other hand, people learn the frequency of reinforcement, then contingent and noncontingent reinforcement would be seen as equally controllable, provided the rates of reinforcement were equal. Evidence will be presented shortly supporting the idea that it is frequency, not contingency, that subjects learn. The equivalency of points along the uncontrollability diagonal, then, deserves empirical investigation.

A second feature of the Hiroto and Seligman (1975) study was the practice of giving random feedback on three discrimination (Cognitive) problems, thus rendering the problems insoluble. Though the trial-by-trial random feedback included 50% "correct," the subjects in the noncontingent condition, when they attempted to state the solution to the problem, were always told that they were incorrect. If the effective reinforcement is the feedback on each attempt to state the solution, then rate of reinforcement and noncontingency are again confounded. Hence, the foregoing argument about the Instrumental task applies as well to the Cognitive task. A further feature of random reinforcement is that it is unchanging; the probability of reinforcement early in training is not appreciably different from the probability late in training. With contingent reinforcement, on the other hand, reinforcement tends to increase in density as the task is learned. Thus, with yoked noncontingent reinforcement, the density or rate of reinforcement increases throughout the task, though the lack of contingency is constant. This point has potential importance for two reasons. One is that, though random and yoked procedures may be equated for overall reinforcement rate, it is entirely possible that subjects do not weigh recent and remote rate equally. Secondly, most trial-and-error tasks outside of experimental settings involve an increasing pattern of success (for example, learning to ride a bicycle). If generalization from such tasks to the experimental setting occurs, it could differentially influence the perception of increasing (yoked) and unchanging (random) reinforcement schemes; the yoked pattern could appear more controllable.

Thus, the work of Hiroto (1974) and Hiroto and Seligman (1975) supports the prediction that uncontrollability leads to helplessness in a limited sense. It has not, however, been convincingly demonstrated that uncontrollability, apart from rate of reinforcement, is the crucial factor.

To date, only one study has explicitly compared yoked and random reinforcement schemes. Eisenberger, Park, and

Frank (1976) gave social approval to children either contingently or noncontingently as the children worked a task. Noncontingent reinforcement was either random or yoked. On a subsequent task, it was found that the contingent approval group learned faster than the noncontingent groups, which did not differ from each other or from two no-pretreatment groups. The failure to find differences in performance between the noncontingent and no-pretreatment groups was interpreted as possibly resulting from some inadequacy in the procedure used to induce helplessness.

One experiment has been performed which explicitly varied rate of noncontingent reinforcement. Benson and Kennelly (1976) were interested in the fact that the Hiroto and Seligman (1976) method of using insoluble discrimination problems necessarily involves some aversive stimulation in the form of failure. Their argument, essentially, was that the amount of aversive stimulation from insoluble problems could be more important than the uncontrollability of reinforcement. Their design was similar to that of Hiroto and Seligman, with the addition of a group receiving 100% correct feedback both on individual trials and on attempts to state the solution to the problem. In this latter condition, they argued, reinforcement is uncontrollable, but the aversive stimulation arising from "incorrect" feedback is removed. Interference on anagram performance was found, as expected, in the standard insoluble condition. The 100%-correct group, however, did not show interference. It was argued that the failure of noncontingent positive feedback to produce helplessness undermines Seligman's theory. It is especially interesting that a postexperimental questionnaire on attributions suggested that the 100%-correct subjects and the standard insoluble-condition subjects had perceived reinforcement as uncontrollable.

The experiment described above is particularly interesting in that it is the only one to date to vary the rate of noncontingent reinforcement. Though the authors focused on the aversiveness of the insoluble problem, rather than on the question of how people store information about contingency, their results clearly suggest that the effects of uncontrollability may be quite different depending on whether the rate of reinforcement is zero or 100%. Of course, their 100%-correct treatment involved an unchanging pattern of reinforcement. Thus, we are unable to tell whether it was the noncontingency or the unchanging pattern which is reflected in the similar ratings of uncontrollability in the standard insoluble and 100%-correct groups.

Taken as a whole, then, the experiments just reviewed with human subjects do not really provide conclusive support for helplessness theory. There is some evidence that perceived helplessness transfers to subsequent tasks. What has not been explored is whether the perception of helplessness develops from exposure to noncontingency or from receiving consistent failure. The equivalency of points along the uncontrollability diagonal (Figure 1) remains to be examined.

We now turn to the second general area of research mentioned at the outset of this paper, those studies exploring the perception of contingency. As these studies form a foundation for examining the importance of rate of reinforcement in perceived control, they will be presented in some detail.

Perceived Control and Contingency

In contrast to the studies of learned helplessness, the following studies have not been performed with a common theoretical groundwork to lend them a sense of continuity. Granting the variability in experimental procedures, however, it is still possible to discern general trends which are relevant to the learned helplessness literature, and which also have a bearing on the criticisms urged earlier in this paper.

Inhelder and Piaget (1958) suggested that a concept of correlation in children becomes developed at about 14-15 years of age. They showed children an array of cards on which faces appeared, with each face having one of the four possible combinations of brown or blonde hair and brown or blue eyes. It was argued that a concept of contingency or correlation existed when the child based his statements about the relation between hair and eye color upon the difference between the total number of "confirming" cases (in this instance, Blonde hair/blue eyes and brown hair/brown eyes) and the total number of "nonconfirming" cases (blonde hair/brown eyes and brown hair/blue eyes). Jenkins and Ward (1965) argue, however, that these results are probably not particularly representative in that a small number of stimuli was used, they were all presented to the subject at the same time, and he could arrange them into groups as he liked. Another objection was to the definition of correlation used by Inhelder and Piaget. Consider the 2x2 matrix in Figure 2 for example. There are more confirming (a and d) than nonconfirming (b and c) cases, yet, in Figure 2, there is no relationship between hair and eye color. The probability of blonde hair given blue eyes (8/10) is not different from the probability of blonde hair given brown eyes (4/5). Hence, Jenkins and Ward argue that the definition of correlation used by Inhelder and Piaget applies only in the special case in which one of the two variables has both states occurring with equal frequency.

Smedslund (1963) examined the understanding of correlation in adults. In his first experiment, nursing students were shown one of five packs of cards, each of which contained letters representing symptoms and diagnoses. The subjects' task was to determine whether a relationship existed between two particular letters designating one symptom and one diagnosis. It was found that, over a range of objective relationships, the subjects tended to report correlations depending on the absolute number of mutual occurrences of

	blonde hair	brown hair
blue eyes	a = 8	b = 2
brown eyes	c = 4	d = 1

Figure 2. 2x2 Table from Jenkins and Ward (1965)

the symptom and diagnosis, and tended to overlook the number of times that one occurred without the other. Hence, subjects often perceived a relationship where there was none. A second experiment, modified to simplify the task, supported these findings. In concluding, Smedslund states

> The apparent main finding of these experiments, then, is that normal adults with no training in statistics do not have a cognitive structure isomorphic with the concept of correlation. Their strategies and inferences typically reveal a particularistic, nonstatistical approach, or an excessive dependence on the frequency of ++ instances. (Smedslund, 1963, p. 170)

The above studies involved the judgement of a relationship between events which were external to the subject. Jenkins and Ward (1965) examined ability of subjects to judge the relationship between their own behavior and their feedback on a task. If the 2x2 matrix in Figure 2 is transformed to represent two responses (R_1 and R_2) and two outcomes (0_1 and 0_2), the index of contingency (ΔP) may be written

$$|p(0_1/R_1) - p(0_1/R_2)| = \Delta P$$

That is, the extent to which 0_1 is controlled by responding depends on the difference in the conditional probabilities. No difference ($\Delta P = 0$) represents no control (the "helplessness" diagonal). When $\Delta P = 1$, the outcome is perfectly controlled by responding. In an impressive series of experiments, subjects worked five problems presented according to a 5x6 Latin Square design. On two of the problems, the outcome was contingent on the responses [$P(0_1/R_1) - P(0_1/R_2) = .8 - .5 = \Delta P$ of .3 for problem X; $.8 - .2 = \Delta P$ of .6 for problem Y].

Three of the problems were noncontingent (random feedback with varying proportions of reinforcement). The problems were worked on an apparatus containing two response keys (R_1 and R_2) and an outcome light. In one set of instructions, the subject was told to "score" as often as possible by pressing one of the keys on each trial, and the outcome light yielded either "score" or "no score" as feedback. In a second set of instructions ("control"), the subject tried to learn to control the appearance of two neutral geometric stimuli by his responses on the two keys. The latter condition was included to control for the possibility that the "score" condition, with an obviously preferred outcome, might produce biased judgements of control. If the score light appeared frequently, even if noncontingently, the subject might restrict his responses to the key which happened to precede the score and thus develop a spurious sense of control. The "control" instructions, with its neutral outcomes, were expected to lead to more equal response sampling and more accurate judgements of control. The actual judgements of control were made after each problem on scales of 0-100 (no control - perfect control). It was pointed out in the instructions that zero control could be the correct answer for any of the problems. In addition to the above conditions, some subjects were active and others were spectators.

It was found that neither of the instructional conditions (score versus control) resulted in judgements of control which were consistent with the actual ΔP index. The same effect occurred in the active and spectator conditions. In fact, for all conditions, the primary feature which predicted the judgement of control on a given problem was the number of times the score light had come on. Notably, the overall correlation between number of scores and ratings of control was .70. Analysis based on the response-outcome frequencies actually obtained by the subject revealed no real departure from the above results.

In two further experiments, Jenkins and Ward varied the instructions, training, and nature of the questions used in assessing control. Across a variety of manipulations, ratings of control tended to follow number of successes ("scores"), and was unrelated to the objective degree of correlation. Only after extensive pretraining was it possible to disrupt the dependency of the judgements of control on number of successes, though the accuracy of the judgements remained poor.

In short, the results of Jenkins and Ward (1965) suggest, as did those of Smedslund (1963), that adults do not have very well developed judgement of contingency. Subjects may perceive a high degree of correlation or contingency, even in the absence of a relationship, if the absolute number of successes or "confirming cases" is high. Such results seem to question a basic assumption of learned helplessness, that subjects perceive noncontingency accurately (Maier and Seligman, 1976).

Ward and Jenkins (1965) examined the influence of the mode of presentation of the information in the judgement task. Subjects estimated the degree of control over rainfall exerted by cloud seeding. The data were presented to the subject in a set of problems, each varying the number of days seeding was present and absent, and the number of days rain was present and absent. Some subjects received the information serially, some in table form, and some received the serial presentation followed by a tabular summary. It was found that the group receiving only the table, without the serial display, tended to follow judgement rules that were more logical than those used by the groups receiving serial presentation alone or both serial and summary. These results, they conclude, are consistent with the idea that adults do not make correlational judgements very accurately when the information is presented serially, as it usually is in most learning situations. Judgements only become more rationally sound when organization, such as a summary table. is imposed on the data.

Thus far, then, the general pattern of results suggests that humans are rather unskilled at judging the degree of correlation between series of events external to themselves, or between their behavior and their outcomes. Some evidence exists, however, which suggests that, at least in some situations, humans can in fact make relatively accurate judgements of correlation. Beach and Scopp (1966) gave subjects 10 decks of cards, with each card containing two numbers from 1-10. The decks represented correlations which ranged from -.85 to +.85. The subjects' task was to rate the degree of relationship between the numbers after looking through the deck. It was found that the proportion of subjects making optimal inferences (according to a Bayesian approach) increased with the size of the correlation. Furthermore, the subjects' confidence ratings increased in a like manner. Citing the earlier evidence that humans have difficulty in judging contingencies, Erlick (1966) argues that humans typically learn correlations between events having more than two states of occurrence. Subjects watched two dials, each of which could assume each of five values. The values on the two dials varied across several series, such that the correlations ranged from -1.0 to 1.0. Though there were discrepancies, the overall pattern was for the mean correlational estimates to follow the objective correlation. Peterson and Beach (1967), in a brief review of these studies, suggest that the reliance on only one cell of the 2x2 matrix, as found in Smedslund (1963) and Jenkins and Ward (1965), may be limited to such simple cases. Humans become better "intuitive statisticians," they argue, in more complex cases.

More recently, Estes (1976) has presented data on probability learning which serve to illuminate some factors which are relevant to the present discussion. Subjects were shown data from "opinion polls" about political candidates, and were asked to make predictions about outcomes. The candidates were represented by letters (A_1 vs. A_2, A_3 vs. A_4, etc.), and on each observation trial a tally appeared after the "winner." A variation involved a simulated poll about product preferences, but the procedure was essentially the same. Even with relatively small probability differences between alternatives (e.g., .46 vs. .54), the typical finding was that subjects predicted correctly from 77% to 87% of the time. In most experiments of this type, however, both the frequency and the probability of winning for a given alternative are both useful for predicting the winner. Particularly interesting results are obtained, however, when an alternative's probability of winning is not consistent with its frequency of winning outcomes. If A and B are presented together, with A winning 75 times and B winning 25 times, A has a .75 probability as of being a winner. If C and D are presented 200 times, with each winning 100 times, each has a .50 probability of winning. When A and C are pitted against each other on a test trial, however, A has a higher probability of winning, though C has a higher overall frequency of wins. In such a case, it was found that subjects select the alternative having the higher winning frequency, even if its probability of winning is lower. It appears then, that subjects store information about the relative frequency of winning for an alternative, and rely less on probability in making their predictions.

Estes explored the encoding processes in more detail, assessing the effects of specifically instructing subjects to attend to outcome classes. Some subjects pronounced the name of the winning alternative, some pronounced the name of the loser, and still others pronounced both. These attention instructions clearly affected choice behavior, in that the variance accounted for by frequency of wins and losses depended on whether wins or losses, respectively, were pronounced during observation trials. Estes concludes that

> unless constrained by special instructions to attend to losing outcomes, subjects tend to ignore losses, store information in memory almost exclusively in terms of relative frequency of winning outcomes, and make predictions on the basis of this stored information. (Estes, (1976), p. 48-49)

Generalization across such varying procedures is necessarily quite tenous, but a brief summary may be helpful. It appears that, at least with a 2x2 contingency table, subjects do not judge correlations very accurately (Jenkins & Ward, 1965; Smedslund, 1963). This holds for both judgements about correlations between external events and for correlations between one's behavior and its outcome. Performance is improved by summary, rather than serial, presentation of the information (Ward & Jenkins, 1965). There is some evidence that subjects rely heavily on "positive" data (e.g., joint occurrences of events to be judged, or successful outcomes), and are less influenced by negative data (absence of one or both outcomes, losing outcomes). Furthermore, this reliance on positive outcomes can be disrupted by special training or instructions (Estes, 1976; Jenkins & Ward, 1965). Finally, judgements of correlation may be more accurate in situations which are more complex than those represented by a 2x2 table (Beach & Scopp, 1966; Erlick, 1966). Even assuming an appropriate degree of caution, however, it is possible to discern implications which the contingency judgement studies may have for the study of learned helplessness.

Overview: Rate, Patterning, and Contingency of Reinforcement

Having just reviewed evidence that the frequency of success, independently of success probability, influenced the judgement of control and predictions of human subjects, I will now discuss the implications of these data for learned helplessness theory. Clearly, there are discrepancies between Seligman's position that people learn about contingencies and the evidence just presented that judgements of control depend on frequency of reinforcement ("success"). How do we account for these discrepancies?

Referring to the instrumental training space of Figure 1, we recall that Seligman has argued that organisms are sensitive to points throughout the space, not just the abscissa and ordinate which have tended to attract the attention of researchers. There is little discussion of the possible mechanism by which organisms store and "compute" the relationships between behavior and outcome, but it is clear that contingency, rather than contiguity, is seen as the controlling variable. Seligman's studies have demonstrated that people exposed to noncontingent reinforcement showed interference on subsequent tasks, presumably because they generalized their expectation of response-outcome independence from pretreatment to test task. The data presented by Jenkins and Ward (1965), however, suggest that perceived control would increase as the frequency of noncontingent reinforcement increased, i.e., moving along the uncontrollability diagonal from lower left to upper right, we should expect perceived control to increase.

I have already stated that Seligman's procedure for exposing subjects to noncontingency is essentially a complete failure procedure, and that failure is represented by a point at the origin of the training space. On this analysis, the results presented by Seligman are consistent with those of Jenkins and Ward (1965), in that a zero rate of reinforcement resulted in little perceived control.. The interpretative problem is that Seligman predicts these results based on the absence of contingency, whereas Jenkins and Ward would predict the results based on the absence of reinforcement. What is needed, then, is an empirical study in which predictions based on contingency and those based on frequency would differ. A helplessness study using rates of noncontingent reinforcement at intermediate points along the uncontrollability diagonal would meet such a need. If helplessness results from noncontingency, then a high rate of noncontingent reinforcement would lead to perceived noncontingency and interference on the anagram task relative to a group receiving an equally high rate of contingent reinforcement. If helplessness results from a low frequency of reinforcement, the group receiving a high rate of noncontingent reinforcement should not differ from the contingent reinforcement group.

I have already outlined the potential importance of reinforcement patterning over time. There is also a need, then, for evaluations of yoked (increasing) and random (unchanging) patterns of reinforcement.

In the present experiment, human subjects learned a pretreatment task under either a high or a low rate of reinforcement. Within each reinforcement level, some subjects received reinforcement contingently upon responding, while others received noncontingent reinforcement. Figure 3 represents the reinforcement patterns used in the experiment. Subjects learning the task with a high rate of contingent reinforcement generated a learning curve and pattern of reinforcement similar to that depicted by line A. Contingent reinforcement at the low rate generated a pattern similar to line C. Subjects receiving yoked reinforcement also received patterns A and C

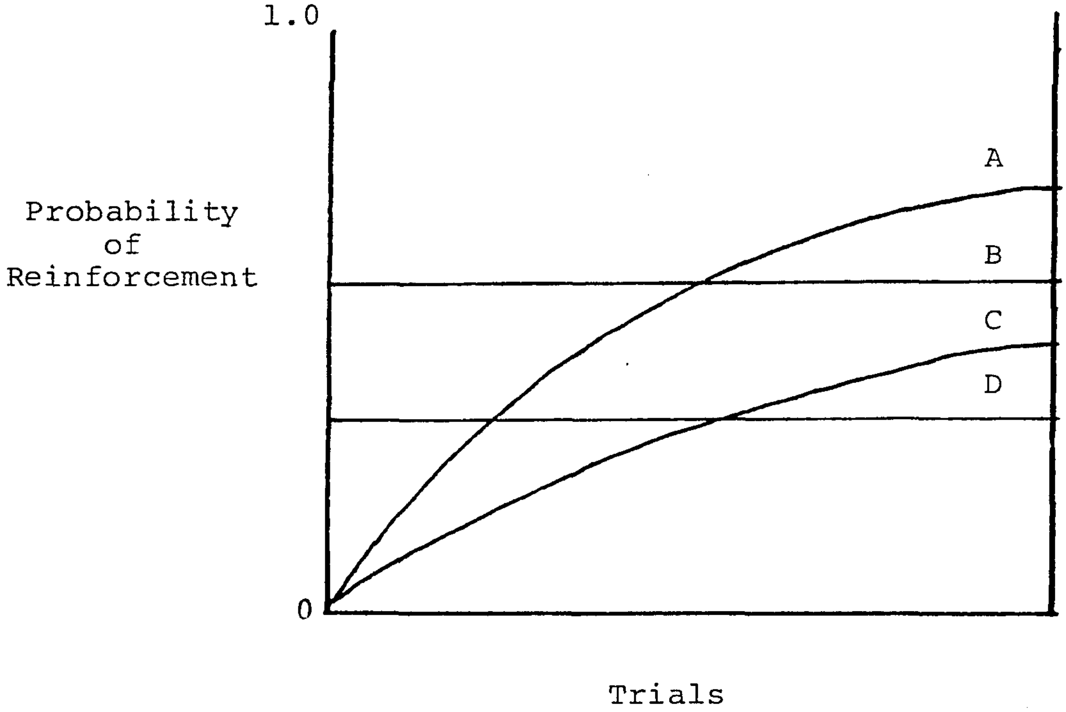

Figure 3. Schematized Reinforcement Patterns Used with Pretreatment Task

(high and low rate), but these reinforcement patterns were independent of responding. If subjects are, in fact, sensitive to the contingency between responding and outcome, the rate of noncontingent reinforcement should not influence perception of control. That is, if points along the uncontrollability diagonal in Figure 1 produce equivalent perception of control, then both yoked conditions (A and C) should yield equivalent perceptions of control. If, on the other hand, rate of reinforcement affects perception of control independently of actual contingency, then the yoked condition receiving pattern A should perceive more control than the yoked condition receiving pattern C. Similar reasoning holds for the random noncontingent reinforcement patterns depicted by B (high rate) and D (low rate). If subjects judge control based on actual noncontingency, B and D should yield equivalent perception of control. Again, more perceived control in B than in D would suggest that rate of reinforcement influences perception of control independently of actual control. Finally, inclusion of both yoked and random noncontingent conditions allows evaluation of the effects of patterning of reinforcement. In the yoked condition, reinforcement increases throughout the task, while random reinforcement is unchanging. A tendency for subjects to weigh recent rate more heavily than remote rate could produce more perceived control in the yoked treatment than in the random one.

Maier and Seligman (1976) have argued cogently for research which defines the boundary conditions of the helplessness effect. The present study can be seen as exploring rate of reinforcement and patterning as boundary conditions.

METHOD

Overview

Essentially, the experiment involved two levels of reinforcement (High and Low), and three levels of controllability/patterning. The latter dimension may be conceptualized as "similarity to controllability." Thus, maximal similarity exists, of course, in the groups receiving contingent reinforcement. The yoked groups received a patterning of reinforcement which was comparable to the contingent condition, but which differed in actual controllability. Finally, the groups receiving random reinforcement differed from the contingent groups in both patterning of reinforcement and controllability. Hence, the pretreatment formed a 2x3 factorial design. Addition of a no-pretreatment control in the test phase would have resulted in an awkward 2x3+1 design. To facilitate analysis, each reinforcement level had its own no-pretreatment control group. Hence, a 2x4 factorial design was used for the test phase. This method of artificially crossing the experimental and control conditions, when the control would not otherwise fit the design, was suggested by Himmelfarb (1975). Figure 4 summarizes the design for the pretreatment and test phases. Two male experimenters were used, with each running half of the subjects in each condition.

Pretreatment

	Contingent	Yoked	Random
High Rate			
Low Rate			

Test Task

	Contingent	Yoked	Random	Control
High Rate				
Low Rate				

Figure 4. Summary of Experimental Design for Pretreatment and Test Phases (Experimenters and Trial Block omitted for clarity of presentation)

Subjects

A total of 84 students, participating for partial credit in introductory psychology, served as subjects. Subjects were randomly assigned to conditions, with the restriction that each group have 10 subjects. Apparatus failure or experimenter error resulted in the dropping of 4 subjects from the pretreatment phase, leaving an N of 60 for pretreatment. In the test phase, 20 subjects served as no-pretreatment controls, resulting in an N of 80. A total of 19 males and 61 females participated. The groups are designated as follows: High Contingent, Low Contingent, High Yoked, Low Yoked, High Random, Low Random, and No-Pretreatment.

Apparatus

Standard electromechanical equipment was used. In pretreatment, the subject's panel was mounted on a plywood stand and stood on a table in front of the subject. The control rack and an event recorder were situated behind a screen approximately 2 meters from the subject. The control rack was encased in sound-attenuating material.

The subjects' panel consisted of three lights spaced 7.7 cm. apart horizontally and a response button 9 cm. below the third light. The first light was white and had the word "Ready" printed under it. The middle light (green) was labelled "Start"; and the third (red), "Score."

Materials

Ratings of perceived control during the pretreatment phase were made on individual sheets of paper contained in a binder. On each sheet was a horizontal line 20.3 cm. long, with 100 equal units indicated by vertical slashes. Under the horizontal line, the numbers 0-100 were printed in units of 10. Above the line, the question "How much control do you have over the score light?" was printed. The words "no control" were printed at the "0" end of the scale, and "complete control" appeared under the "100."

The anagram task in the test phase consisted of the 20 five-letter anagrams used in Seligman's work. The anagrams were each printed on a card, and all cards were contained in a binder. Each anagram had the same solution pattern, 5-3-1-2-4 (e.g., E R L K C).

Procedure

All subjects participated individually. Each subject was greeted at the door and escorted to the table by the experimenter. For all except No-Pretreatment subjects, it was explained that the experiment consisted of two tasks, learning to press the button in such a way as to operate the score light, and solving a series of anagrams or "word puzzles." For the pretreatment task, it was explained that the Ready light meant that a trial was about to begin. During the Start light (one trial) the subject was to press the button once. It was explained that, if the Score light flashed at the end of the trial, the subject had "hit" the target interval. The subject was to try to locate the target interval and then hit it on as many trials as possible. The subject was cautioned that the target interval was a small one.

It was then explained that the Experimenter would occasionally stop the apparatus and ask for a rating. The binder with rating sheets was handed to the subject and the rating procedure was explained:

> As you can see, you are asked to rate the degree of control that you feel you have developed over the score light at that time. A rating is made simply by drawing a slash through the line at whatever point you feel is appropriate. For example, a rating of 0 would mean that, at that point, you feel you have found no way to influence whether or not the score light will flash on any given trial. A rating of 100 means that you can make the score light flash on any trial by responding in a certain way, and that you also can respond in such a way as to ensure that the score light will not flash if you should choose to keep it from doing so. Ratings in between represent varying degrees of ability to control the score light. That is, you feel to some degree that you could make the score light flash or not flash, but that you don't feel that you could completely control it.

The subject was then allowed to watch the apparatus for 3 trials without responding, to ensure that he understood the operation. The subject then made an initial rating, followed by a rating at the end of each 25 trials. The actual operation of the apparatus was as follows: The Ready light operated for 3 sec., followed by the Start (trial) light, which stayed on for 5 sec. The target interval was .63 sec. long, and was located 3.13 seconds into the trial. No stimulus changes corresponded to the target interval. Button presses during the target interval were recorded as "correct" responses for Contingent subjects only.

For High Contingent subjects, each correct response was reinforced. For Low Contingent, 60% of the correct responses were reinforced. For Noncontingent subjects, the experimenter secretly used a switch to operate the score light. The yoking procedure was accomplished by using the event record from each Contingent subject and giving that pattern of reinforcement to a yoked subject. The number of reinforcements for each Contingent subject was recorded, and a randomized schedule based on this number was given to a Random subject. For all noncontingent subjects, a scheduled reinforcement was delivered only if the subject responded on that trial. If no response occurred, the scheduled reinforcement was delivered on the next trial on which a response occurred. The pretreatment consisted of 75 trials with a mean intertrial interval of 9.6 sec. and a range of 6-12 sec.

Upon completion of the pretreatment task, the subject's panel was removed and he was given the binder with the anagrams. The instructions, adapted from Hiroto and Seligman (1975), were as follows:

> This task involves finding the solution to some anagrams. As you may know, anagrams are words with the letters scrambled. The problem is for you to unscramble the letters so that they form a word. As soon as you've found the word, tell me what it is, and then wait until I tell you to go on to the next one. Though these anagrams are not terribly easy, they do form words. One final point is that there may be a pattern or principle by which to solve the anagrams, but that's all I can say.

The Experimenter then moved behind a screen and proceeded with the task. The subjects' response and latency were recorded. If the subject had failed to solve the anagram after 60 sec., a latency of 60 sec. was recorded and the subject was asked to go on to the next anagram.

Subjects in the No-Pretreatment condition received only the instructions for the anagram task. They were not informed that there had been a previous task for other subjects.

Dependent variables were the latency of response, number of failures to solve the anagrams, and trials to criterion defined as the number of trials required to reach 3 consecutive responses in less than 15 seconds each. These are the dependent variables used by Hiroto and Seligman (1975).

Upon completion of the anagram task, all subjects were then asked to complete two brief questionnaires, one for each task. The questions were designed to assist in determining whether any suspiciousness had emerged, and to see if any consistent factors could help in data interpretation. No special instructions were given for these questionnaires.

All subjects were then given a thorough debriefing, and comments about the experiment were solicited. The debriefing emphasized the potential importance of perception of control in understanding certain clinical problems. It was stressed that, on the pretreatment task, we were interested in how certain patterns of events influenced perception of control. This explanation was given to focus attention on our testing of patterns in the environment, rather than on assessing the individual's ability to perceive control accurately. It was stressed that neither task reflected in any way on the subjects' ability or intelligence. Subjects were then shown the operation of the control apparatus. A written version of the debriefing was mailed to all subjects at the end of the experiment.

RESULTS

Independent Variables

It will be recalled that the Contingent, Yoked, and Random conditions, summing across High and Low Rate of reinforcement, were conceptualized as forming a dimension of "similarity to controllability." In discussing the analysis to follow, this similarity factor will be referred to simply as "Groups," and the High-Low Rate factor will be referred to as Rate. Also, 3 blocks of 25 trials and 4 ratings (1 prior to the task and 3 during the task) were used, and this factor will be referred to as "Blocks." A fourth factor in the analysis was "Experimenter." Accordingly, a 2x3x2x3 analysis of variance (ANOVA) for repeated measures was performed on the obtained reinforcements variable. Remember that, as Yoked and Random conditions were matched for reinforcement to the Contingent condition, no Groups difference in reinforcement was expected. The critical differences, then, were for Rate and Blocks (to insure the increasing reinforcement pattern for Contingent and Yoked). A summary of the analysis is presented in Table 1.

As planned, reinforcement did not differ among Groups. Obtained reinforcement increased for both Contingent and Yoked groups from block 1 to block 2 and from block 2 to block 3. For these groups, block 3 also had more reinforcement than block 1. Reinforcement did not change across blocks in the Random condition. Contingent, Yoked and Random conditions did not

TABLE 1

Summary of ANOVA for Rate of Obtained Reinforcement over Blocks

Source	DF	MEAN SQUARE	F
Rate (R)	1	1855.98	45.67 *
Groups (G)	2	0.34	0.01
Experimenters (E)	1	47.02	1.16
R x G	2	0.51	0.01
R x E	1	80.00	1.97
G x E	2	0.04	0.00
R x G x E	2	0.32	0.01
Error	48	40.64	
Blocks (B)	2	256.54	28.81 *
R x B	2	62.87	7.06 *
G x B	4	67.70	7.60 *
B x E	2	9.77	1.10
B x R x G	4	16.06	1.80
B x R x E	2	45.35	5.09 *
B x G x E	4	0.34	0.04
B x R x G x E	4	6.77	0.76
Error	96	8.90	

* $p < .05$

differ from each other at each block. These relationships were reflected in a Group x Block interaction ($F = 7.60$, $df = 4, 91$, $p < .05$) and a Newman-Keuls post-hoc analysis. The means for this interaction are presented in Table 2.

The Rate manipulation was successful, in that the High Rate condition received more reinforcement than Low Rate at each block. These results were consistent for both experimenters. A Newman-Keuls test revealed that, for High Rate (Experimenter 1), obtained reinforcement increased from block 1 to block 2, though blocks 2 and 3 did not differ. No such increase occurred for the Low Rate condition with Experimenter 1. For Experimenter 2, High Rate, there was an increase in reinforcement from blocks 1 and 2 to block 3, though blocks 1 and 2 did not differ. As with Experimenter 1, the Low Rate conditions did not show increased reinforcement across blocks. These differences were reflected as a Rate x Block x Experimenter interaction ($F = 5.09$; $df = 2, 96$; $p < .05$). The means for Rate and Block are presented for each experimenter in Table 3.

Analysis of Perceived Control Ratings

The ratings taken prior to initiation of the pretreatment task did not differ for the High and Low Rate conditions. The effect of the differing rates of reinforcement once the task began, however, was that subjects receiving the High Rate of reinforcement perceived more control over the reinforcement than those receiving the Low Rate, (Figure 5). This difference was

TABLE 2

Means (M) and Standard Deviations (SD) for Group X Block Interaction on Rate of Reinforcement

			Block		
			1	2	3
Group	Contingent	M	6.00	9.40	12.15
		SD	3.91	6.09	6.64
	Yoked	M	6.00	9.00	12.35
		SD	3.99	5.94	6.71
	Random	M	9.25	9.40	9.15
		SD	4.79	5.02	5.34

TABLE 3

Means (M) and Standard Deviations (SD) for Rate X Block X Experimenter Interaction on Rate of Reinforcement

Experimenter	Rate		Block		
			1	2	3
1	High	M	9.47	13.67	13.60
		SD	3.68	4.32	4.37
	Low	M	6.07	6.47	8.93
		SD	3.63	5.19	5.70
2	High	M	8.80	12.07	16.79
		SD	5.61	4.79	4.90
	Low	M	4.00	4.97	5.53
		SD	3.70	2.39	3.70

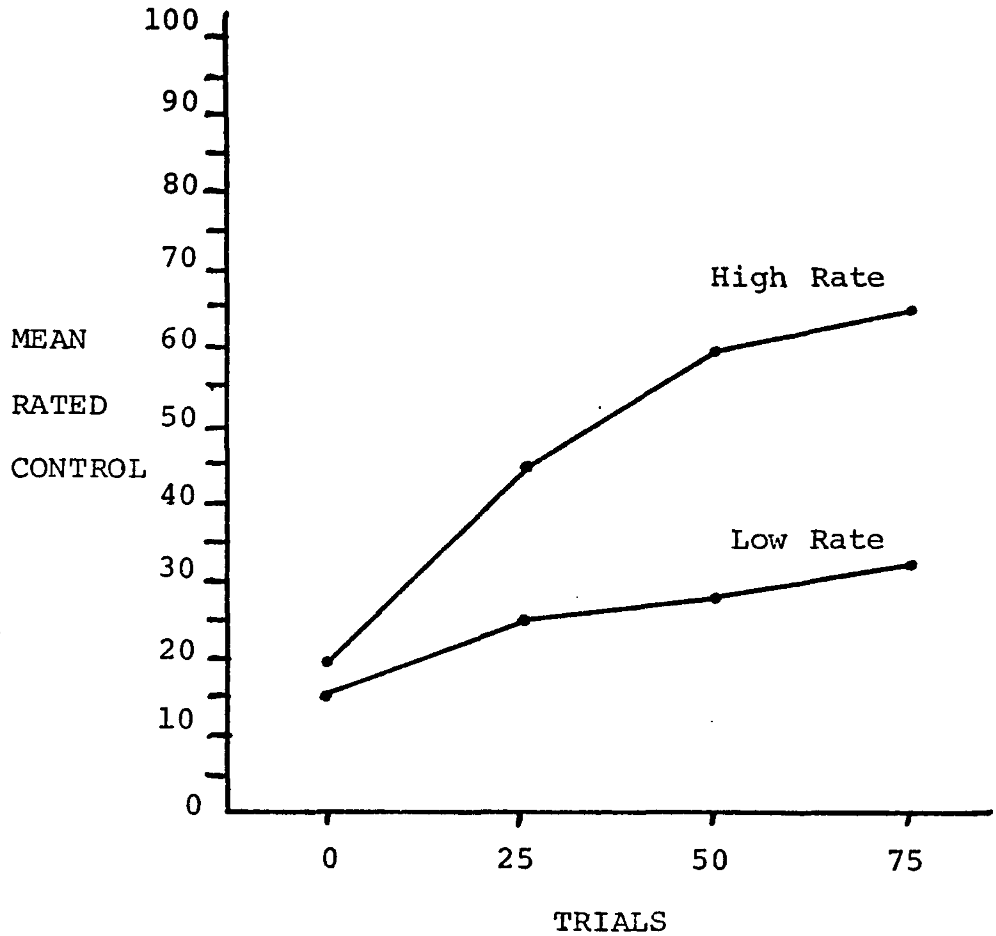

Figure 5. Rated Control as a Function of Trials

substantiated by a significant Newman-Keuls post hoc test on the ratings at blocks 2 and 3. Though the High and Low Rate mean ratings did not differ on the second rating (after the first block of trials), the High Rate mean was 18.23 points higher than the Low Rate mean, and this difference was barely less than the Newman-Keuls critical value for that contrast (18.26). On these ratings of perceived control, it did not matter whether reinforcement was Contingent, Yoked, or Random (the main effect for Group was nonsignificant). The dependency of perceived control upon rate of reinforcement is seen as a significant Rate x Block interaction on the ANOVA ($F = 5.46$; $df = 3, 151$; $p < .05$). Furthermore, at High Rate, all ratings taken during the pretreatment task were higher than the one taken prior to the task. For Low Rate, however, the ratings did not change from the first rating through the last. Table 4 presents the summary of the ANOVA on the ratings of perceived control.

Anagram Task

The different pretreatment conditions had no effect on anagram performance as reflected in latency of response and number of errors. On these measures, none of the pretreatment conditions differed from each other or from the No Pretreatment control group. A correlational analysis of anagram performance as a function of rated control likewise revealed no effects. A summary of the multivariate ANOVA performed on these measures is presented in Table 5. The means are presented in Tables 6 and 7.

TABLE 4

Summary of ANOVA on Rated Degree of Control at each Block

Source	DF	MEAN SQUARE	F
Rate (R)	1	27820.75	22.47 *
Group (G)	2	466.01	0.38
Experimenter (E)	1	3511.30	2.84
R x G	2	170.42	.14
R x E	1	1382.39	1.12
G x E	2	1185.43	.96
R x G x E	2	856.53	.69
Error	48	1238.32	
Blocks (B)	3	10217.64	22.40 *
R x B	3	2488.58	5.46 *
B x G	6	637.49	1.40
B x E	3	376.40	0.83
R x B x G	6	261.59	0.57
R x B x E	3	438.46	0.96
B x G x E	6	529.34	1.16
R x B x G x E	6	266.34	0.58
Error	144	456.06	

* $p < .05$

TABLE 5

Summary of MANOVA* for Latency and Number of Errors Measures of Anagram Performance

Source	DF	F
Rate (R)	2, 63	.20
Group (G)	2,124	1.03
Experimenter (E)	2, 63	.09
R x G	6,124	.09
R x E	2, 63	1.19
G x E	6,124	1.51
R x G x E	6,124	.81

* Tested using Hotelling-Lawley Trace

TABLE 6

Means (M) and Standard Deviations (SD) for Latency (in Sec.) Measure of Anagram Performance (Summed across Experimenters for Clarity of Presentation)

Rate		Group			
		Contingent	Yoked	Random	Control
High	M	27.69	27.93	24.03	29.58
	SD	15.00	13.79	13.66	13.27
Low	M	25.70	23.73	23.21	31.03
	SD	9.77	13.83	15.71	7.97

TABLE 7

Means (M) and Standard Deviation (SD) for Errors Measure of Anagram Performance (Summed across Experimenters for Clarity of Presentation)

Rate		Group			
		Contingent	Yoked	Random	Control
High	M	6.70	8.20	5.80	8.20
	SD	4.90	4.98	4.16	4.78
Low	M	6.40	7.40	5.90	8.60
	SD	3.17	5.87	5.30	4.35

On the trials to criterion measure, however, subjects in the Contingent condition under Experimenter 2 took more trials to reach criterion than those under Experimenter 1. On this same measure, however, the means for each experimenter did not differ in the Yoked, Random, or No-Pretreatment groups. These differences emerged as a Group x Experimenter interaction on the ANOVA ($F = 3.09$; $df = 1, 38$; $p < .05$). The summary of this ANOVA is presented in Table 8. The means for each Experimenter at each Group are seen in Table 9.

Postexperimental Questionnaire

On the postexperimental questionnaire, employed to probe factors which could be helpful in interpretation, only two items were significant. One item asked subjects to weight the relative influence on their performance of ability, effort, task characteristics, and luck (Weiner, Frieze, Kukla, Reed, Rest, and Rosenbaum, 1971). These ratings were done with the restriction that their total be 100%. Accordingly, an arcsin transformation was performed on the scores. On the pretreatment task, subjects in the High Rate condition tended to rate ability higher than those in the Low Rate condition. Also, on a 31-point scale, High Rate subjects indicated more interest than Low Rate subjects in participating in another experiment using the pretreatment task ($F = 7.50$; $df = 1, 48$; $p < .01$).

TABLE 8

Summary of ANOVA on Trials to Criterion Measure of Anagram Performance

Source	DF	MEAN SQUARE	F
Rate (R)	1	5.91	.35
Group (G)	3	42.98	2.55
Experimenter (E)	1	18.46	1.09
R x G	3	10.30	.61
R x E	1	9.29	.55
G x E	3	52.14	3.09 *
R x G x E	3	23.21	1.38
Error	38	16.87	

* $p < .05$

TABLE 9

Means (M) and Standard Deviations (SD) for Group X Experimenter Interaction on Trials to Criterion ANOVA

Group	Experimenter	N		
Contingent	1	8	M SD	8.13 2.53
Contingent	2	6	M SD	14.83 6.49
Yoked	1	4	M SD	11.50 2.89
Yoked	2	8	M SD	10.13 3.64
Random	1	8	M SD	7.63 4.03
Random	2	6	M SD	9.33 3.56
Control	1	7	M SD	14.43 4.79
Control	2	7	M SD	12.14 3.48

Response Patterning in Pretreatment and Perception of Control

During each trial in the Pretreatment, an event recorder reflected the temporal location of each response. The recorder was programmed to divide each trial into .63 sec. intervals, and to reflect which interval contained the response.

An an index of response patterning, the Standard Deviation (SD) of the response intervals was computed for each subject over each block of 25 trials. Hence, for subjects whose responding centered around a narrow band of time, this index would approach zero. Subjects who showed no consistent patterning in responding would produce relatively greater indices.

The above measure of response patterning during each block was then correlated with the rating of control given at the end of that block. For the combined Contingent groups, the correlation between response patterning and rating of control was -.64 ($p < .01$). The same correlation for all Noncontingent groups was -.30 ($p < .01$). Furthermore, there was a tendency for response variability to decrease as reinforcement increased. This tendency was reflected by a correlation of -.74 ($p < .01$) for the Contingent conditions and -.33 ($p < .01$) for the Noncontingent conditions.

DISCUSSION

Perception of Control Ratings

If actual controllability of reinforcement on the pretreatment task were crucial in determining ratings of control, the ratings for the Contingent condition would be higher than those for Yoked and Random. Furthermore, it is likely that, using our instructions about the meaning of control, the High Contingent group would give higher ratings than the Low Contingent group. That is, though both groups could produce or avoid the score light by varying their responses, the High Contingent group could produce it with more certainty than the Low Contingent group. The ratings for the Contingent, Yoked, and Random groups, however, did not differ. That is, actual controllability and patterning of reinforcement did not determine the ratings. In fact, only the Rate variable showed any consistent effect on ratings of control. As in the Jenkins and Ward (1965) study, ratings of control increased with more reinforcement, regardless of the actual contingency of reinforcement. That is, frequency of success, independently of contingency, influenced the ratings of control. This effect cannot be the result of chance bias in the ratings, as the two Rate conditions did not differ on the first rating (taken prior to initiation of the task).

A theoretical point of view emphasizing contiguity, of course, would note that noncontingency is defined from the experimenter's point of view, i.e., the experimenter administers reinforcement without regard to the nature of the response emitted by the subject. From the subject's point of view, however, the case may be entirely different. A subject, especially one motivated to obtain reinforcement, may be expected to increase the frequency of a behavior which happens to precede reinforcement. Some of the subjects in the noncontingent condition did, in fact, show evidence of such "superstitious" learning. That is, the standard deviation of response intervals was quite small.

Furthermore, there was a tendency for higher ratings of control to occur in subjects with less response variability, as reflected in the -.30 correlation between response variability and perceived control in the noncontingent conditions. I am hesitant, however, to attribute to superstitious learning (in the traditional sense) the failure for noncontingent reinforcement to be perceived as noncontingent. One reason for such hesitancy is the fact that findings of superstitious learning have not fared well under recent theoretical and empirical scrutiny (e.g. Staddon and Simmelhag, 1971). Perhaps a more compelling reason is that there was insufficient response stereotypy to suggest that superstitious learning had occurred. On the other hand, it is possible that superstitious learning, not in the sense of responding with a consistent latency, but in the sense of response distribution, did occur. For example, if response distribution tended to follow the reinforcement distribution of early trials, this effect could be seen as a form of superstitious learning. Though response stereotypy would not be as evident in the latter case, such an effect would be evidence of control of responding by noncontingent reinforcement. In fact, it could be argued that the ratings of control in the noncontingent conditions reflect in themselves a kind of superstition.

The question of response sampling and perceived control is an important one and deserves elaboration. In the present experiment, reduced response variability was correlated with more perceived control and with higher reinforcement rates for both contingent and noncontingent conditions. In the noncontingent conditions, it would be possible for a contingency between responding and outcome to develop. Consider a hypothetical case, for example, in which a subject made most of his responses in a narrow band of time, with very few responses occurring outside of this predominant band. A contingency could exist if the probability of reinforcement in the narrow, predominant band happened to differ from that for responses outside of this band. Such a difference could develop if response variability was very small, as in the case just described, owing to chance fluctuations in the prearranged schedule. That is, with a very small number of responses outside of the predominant band, there is less likelihood that the response-outcome probabilities for those few responses will approach the objective response-outcome probabilities. The subject may not, in other words, have sampled enough to get a reliable estimate of the objective response-outcome contingency. As response variability increases, there is more opportunity for the subject to find out that the reinforcement probability for all response bands is equal.

Two points may be made in addressing the above problem. One is that, though the negative correlation between rate of reinforcement and response variability was statistically significant in the noncontingent conditions, this relationship leaves much variability unaccounted for. That is, the relationship is not strong enough to suggest that response sampling was greatly reduced, even in the high rate condition. As the objective reinforcement probabilities for all response bands were equal in the noncontingent conditions, then, it is unlikely that subjects were basing their judgements of control on the ability of variations in responding to produce variations in reinforcement probability. Secondly, Jenkins and Ward (1965) found that subjects receiving noncontingent reinforcement still judged moderate degrees of control even when the experiment was designed to produce equal sampling of all response-outcome probabilities. It appears, then, that the development of unexpected response-outcome contingency is not a likely explanation for the failure of noncontingent reinforcement to be perceived as such. Future experiments using noncontingent reinforcement, however, may do well to use procedures in which the obtained degree of contingency may be readily evaluated.

Of course, a possibility worth considering is that the method used to assess perception of contingency is not a particularly sensitive one. A different phrasing of the control question, different instructions, or some other procedural variation may have reflected more sensitivity on the part of subjects to the fact that the score light occurred independently of their responding. Open-ended questioning at the conclusion of the experiment did indicate that some subjects were able to state that the light came on after widely different responses, yet they gave non-zero ratings of control. It is notable, however, that such comments also occurred for subjects who had received contingent reinforcement.

It must also be remembered that the instructions for the ratings included the idea that control meant being able to produce <u>or</u> avoid the score light by variations in responding. If subjects had been, as it were, correlating the variation in responding with variations in reinforcement, they would have reached the conclusion that their behavior had no consistent effect on the probability of the score light. Hence, these ratings would have been lower than those for the contingent conditions. With respect to the ratings of control in the pretreatment, then, we may conclude, as did Jenkins and

Ward (1965), that the subjects' concept of control was not based on statistical relatedness. This conclusion is also consistent with that of Smedslund (1963), who found that subjects' judgements of relatedness between events (symptoms and diseases) was based only on the frequency of conjunction of the events, and that the frequency of disjunction was not considered. The present experiment confirmed these findings, and, further, suggested that degree of similarity to actual control did not affect the ratings.

How, then, do these results relate to the empirical findings and theoretical statements of helplessness theory? With respect to the empirical results, it seems apparent that the practice of failing subjects on all discrimination problems essentially amounts to using zero reinforcement. Also, it will be remembered that the noise-escape studies used a feedback light which indicated whether the noise offset was a result of responding or whether the noise terminated on its own. Subjects in the noncontingent groups received the "failure" light on each trial. Again, this procedure places the subject in a situation represented by the origin or zero point of the instrumental training space. Postexperimental ratings in such studies have found that the noncontingent groups typically rate themselves as more helpless than the contingent groups. The results of the present experiment suggest that the differences in perceived control could reflect differences in reinforcement rate more than differences in controllability *per se*.

The first portion of helplessness theory, as stated earlier, asserts that exposure to uncontrollability may result in perception of uncontrollability. Maier and Seligman (1976) also suggest that research is needed to explore the factors which determine when uncontrollability is perceived as such. The present results indicate quite clearly that rate of reinforcement is one of the variables which will influence the perception of control. Particularly interesting is the fact that rate of reinforcement is equally important in determining the perception of contingent and noncontingent reinforcement. An interesting possibility is raised by the results of Benson and Kennelly (1976), who added an always-correct condition to the triadic design used by Seligman. Benson and Kennelly found no interference on anagram performance following 100% noncontingent reinforcement, though they did find interference after pretreatment comparable to that used by Seligman. Of particular relevance to the pretreatment in the present study, however, was the finding that the failure group and the always-correct group both perceived reinforcement as uncontrollable. Of course, caution must be used in comparing results, owing to differences in tasks and in methods used to assess perception of control. The interesting possibility is suggested, though, that perception of control may be an inverted-U function of amount of reinforcement, with zero and 100% noncontingent reinforcement being perceived as uncontrollable, and intermediate amounts being erroneously perceived as somewhat controllable.

The second portion of helplessness theory asserted that, once uncontrollability is perceived, it will interfere with subsequent performance. Again, it is mentioned that the limits of this transfer need to be explored.

Clearly, the anagram phase of the present experiment was not consistent with the above assertion. The latency and number-of-errors measures did not reflect the effects of pretreatment. That is, though there were differences in rated control during the pretreatment, these differences were not reflected among the experimental groups or between the experimental and No-Pretreatment conditions. The one difference which did emerge on the trials-to-criterion measure was between experimenters within the contingent group. This difference was unexpected and is difficult to interpret. Certainly, the finding cannot be predicted from any of the theoretical positions discussed in this paper.

We now turn to possible reasons for the failure to find a transfer of the perception of control in pretreatment to performance on the anagram task. Certainly, the first possibility which must be considered is that the hypothesis of transfer between pretreatment perceptions of controllability and test task performance is false, i.e., that the differences between perceived control reflected in the pretreatment ratings of the High and Low Rate groups did not affect the anagram performance. If there were a strict relationship between initial perception of control and subsequent performance, we should expect the lower perception of control in the Low Amount group to lead to interference in anagram performance relative to the High Amount group.

Recall that the most recent statement of helplessness theory by Maier and Seligman (1976) stated that, once uncontrollability is perceived, it interferes with subsequent performance. Whether the present results can fairly be applied to the stated helplessness theory, then, depends on what is meant by perceived uncontrollability. The final mean rating of control for subjects in the High Rate group was 62.90, which can be seen as reflecting a moderate degree of control. The Low Rate ratings were considerably lower (30.67). As the probing procedures used for measuring perceived control were very different than those used in other helplessness experiments, we cannot be sure whether this latter perception, admittedly at the lower end of the scale, is low enough to be expected to interfere with subsequent performance.

Of course, caution must be used in making interpretations of such negative findings. A clear possibility is that something inherent in the task or procedures could have mitigated the effects of pretraining perceptions and caused the failure of these perceptions to transfer. For example, Klein, Fencil-Morse, and Seligman (1976) found that a group receiving instructions to attribute failure to the difficulty of the task did not show the interference effect after uncontrollable failure. This attributional effect appeared only for depressed subjects, however, and did not occur for nondepressed subjects. The present study did not explicitly vary depression, so extrapolation from the Klein et al., results is difficult. It appears, however, that the subject's attributions for his performance in pretreatment could be an important variable in determining whether transfer will occur.

Another procedural difference deserving attention is that ratings of control were taken during the pretreatment. If the act of making ratings during the pretreatment sensitized the subjects to the dimension of controllability, they may have been more alert to the solvability of the anagrams and thus removed any differences among pretreatment groups or between pretreatment and no-pretreatment control groups.

Finally, social interaction variables between the experimenters and subjects cannot be ruled out. Such variables are difficult to quantify and include in describing procedures.

The fact remains, however, that the present experimenters interacted with subjects in a relatively neutral manner, i.e., periodically asking the subject to make a rating. The interactions in Seligman's experiments with insoluble discrimination problems involved repeatedly telling the subjects that they were not correct in their estimate of the correct value. Such differences in experimenter behavior could conceivably affect the stressfulness of the task and, hence, performance on the task.

In fairness, then, the above criticisms suggest caution in interpreting the failure to find interference here as contrary to helplessness theory. The reasons for caution essentially involve the fact that the present experiment did not demonstrate the ability to produce any kind of transfer from pretreatment to test task. If, for example, a group had been used which received complete failure on the pretreatment (with very low control ratings) and this group had shown transfer, we could be more certain that the absence of transfer in the other groups reflected the effects of the reinforcement levels used. Failure to find transfer in such a group would strengthen interpretations based on sensitization, task characteristics, and so forth.

With these cautions in mind, then, what can be said about the relevance of these results for helplessness theory? Maier and Seligman (1976) have stated that perception of noncontingency can result from exposure to noncontingent reinforcement. They further suggest that clarification is needed with respect to variables which determine whether or not noncontingency will be perceived. The finding that perceived control increases with increased reinforcement, regardless of actual contingency, suggests an important boundary condition for the first part of helplessness theory. It is suggested that subjects are more likely to develop the belief that reinforcers are uncontrollable when those reinforcers are infrequent. Apparently, uncontrollability _per se_ is not sufficient for the development of a perception of uncontrollability. The finding of Benson and Kennelly (1976) that 100% noncontingent reinforcement did not produce interference of subsequent performance, even though comparisons with the failure group revealed that both perceived reinforcement as uncontrollable, is apparently at odds with the present results. We cannot say whether or not a group receiving the score light on every trial would give ratings lower than those receiving an intermediate amount, but an empirical test of such a proposition should be enlightening. With respect to the failure to find transfer, a tentative conclusion is that intermediate reinforcement amounts do not produce sufficient perception of uncontrollability to cause interference on the test task.

In developing the position in the present paper that rate of reinforcement must be equal for an unbiased evaluation of the effects of contingent and noncontingent reinforcement, I have chosen to work within the limitations of the yoked control procedure. Church (1964) has argued cogently, however, that individual subject differences in sensitivity to reinforcement do not effect experimental and yoked groups equally, and that the bias is in favor of superior responding in the experimental group. Such differences, he suggests, can complicate interpretation of results in which a yoked control has been employed. Further work in helplessness theory might well explore alternatives to the yoked control procedure, such as single subject designs. Another possibility, especially in exploring the reinforcement patterning dimensions, would be to employ experimenter-determined noncontingent reinforcement schedules rather than relying on patterns generated by subjects in the contingent conditions. Even in the absence of Church's arguments, the former procedure has the advantage of reducing the variability generated by subjects in the contingent groups, and, consequently, rendering more detectable the differences among experimental conditions.

The present work has perhaps suggested more questions than answers. Though extensions of helplessness theory to clinical areas such as anxiety and depression are interesting, it is clear that more understanding of the laboratory phenomenon of helplessness is needed before such extensions can be fully evaluated. Parametric work along the lines of the experiment reported here is needed. For example, the extremes of the uncontrollability diagonal (noncontingent failure and noncontingent success) need to be explored in conjunction with the intermediate range. If transfer is found under some conditions, systematic variation of such variables as experimenter "warmth," attributions for failure, and similarity of pretreatment and test situations need to be explored to assess the robustness of the transfer effect. Another possible potent variable would be expectation for performance level during the pretreatment phase. Manipulation of this expectation should alter the effectiveness of obtained reinforcement. That is, the effectiveness of a particular rate of reinforcement would depend on the expected rate of reinforcement for the task.

BIBLIOGRAPHY

Beach, L. R., & Scopp, T. S. Inferences about correlations. Psychonomic Science, 1966, 6, 253-254.

Benson, J. S., & Kennelly, K. J. Learned helplessness: The result of uncontrollable reinforcements or uncontrollable aversive stimuli? Journal of Personality and Social Psychology, 1976, 34, 138-145.

Church, R. M. Systematic effect of random error in the yoked control design. Psychological Bulletin, 1964, 62, 122-131.

Eisenberger, R., Park, D. C., & Frank, M. Learned industriousness and social reinforcement. Journal of Personality and Social Psychology, 1976, 33, 227-232.

Engberg, L. A., Hansen, G., Welker, R. L., & Thomas, D. R. Acquisition of key-pecking via autoshaping as a function of prior experience: "Learned laziness?" Science, 1973, 178, 1002-1004.

Erlick, D. E. Human estimates of statistical relatedness. Psychonomic Science, 1966, 5, 365-366.

Estes, W. K. The cognitive side of probability learning. Psychological Review, 1976, 83, 37-64.

Fosco, E., & Geer, J. H. Effects of gaining control over aversive stimuli after differing amounts of no control. Psychological Report, 1971, 29, 1153-1154.

Gamzu, E., Williams, D. A., & Schwartz, B. Pitfalls of organismic concepts: "Learned laziness?" Science, 1973, 101, 367-368.

Glass, D. C., & Singer, J. E. Urban stress: Experiments on noise and social stressors. New York: Academic Press, 1972.

Himmelfarb, S. What do you do when the control group doesn't fit into the experimental design? Psychological Bulletin, 1975, 82, 363-368.

Hiroto, D. S. Locus of control and learned helplessness. Journal of Experimental Psychology, 1974, 102, 187-193.

Hiroto, D. S., & Seligman, M. E. P. Generality of learned helplessness in man. *Journal of Personality and Social Psychology*, 1975, *31*, 311-327.

Hulse, S. H. Patterned reinforcement. In G. H. Bower & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 7). New York: Academic Press, 1974.

Inhelder, B., & Piaget, J. *The growth of logical thinking from childhood to adolescence*. New York: Basic Books, 1958.

Jenkins, H., & Ward, W. Judgement of contingency between responses and outcomes. *Psychological Monographs*, 1965, *79* (1, Whole No. 594).

Klein, D. C., Fencil-Morse, E., & Seligman, M. E. P. Learned helplessness, failure, and the attribution of failure. *Journal of Personality and Social Psycholocy*, 1976, *33*, 508-516.

Levis, D. J. Learned helplessness: A reply and alternative S-R interpretation. *Journal of Experimental Psychology: General*, 1976, *105*, 47-65.

Maier, S. F., & Seligman, M. E. P. Learned helplessness in theory and evidence. *Journal of Experimental Psychology: General*, 1976, *105*, 3-46.

Overmier, J. B. Interference with avoidance behavior: Failure to avoid traumatic shock. *Journal of Experimental Psychology*, 1968, *78*, 340.343.

Overmier, J. B., & Seligman, M. E. P. Effects of inescapable shock upon subsequent escape and avoidance learning. *Journal of Comparative and Physiological Psychology*, 1967, *63*, 23-33.

Peterson, C. R., & Beach, L. R. Man as an intuitive statistician. *Psychological Bulletin*, 1967, *68*, 29-46.

Roth, S., & Bootzin, R. R. Effects of experimentally induced expectancies of external control: An investigation of learned helplessness. *Journal of Personality and Social Psychology*, 1974, *29*, 253-264.

Roth, S., & Kubal, L. Effects of noncontingent reinforcement on tasks of differing importance: Facilitation and learned helplessness. *Journal of Personality and Social Psychology*, 1975, *32*, 680-691.

Seligman, M. E. P. Helplessness: On depression, development, and death. San Francisco: W. H. Freeman & Co., 1975.

Seligman, M. E. P., & Maier, S. F. Failure to escape traumatic shock. Journal of Experimental Psychology, 1967, 74, 1-9.

Seligman, M. E. P., Maier, S. F., & Geer, J. The alleviation of learned helplessness in the dog. Journal of Abnormal and Social Psychology, 1968, 73, 256-262.

Seligman, M. E. P., Maier, S. F., & Solomon, R. L. Unpredictable and uncontrollable aversive events. In F. R. Brush (Ed.), Aversive conditioning and learning. New York: Academic Press, 1971.

Smedslund, J. The concept of correlation in adults. Scandinavian Journal of Psychology, 1963, 4, 165-173.

Staddon, J. E. R., & Simmelhag, V. L. The "superstition" experiment: A reexamination of its implications for the principles of adaptive behavior. Psychological Review, 1971, 78, 3-43.

Thornton, J. W., & Jacobs, P. D. Learned helplessness in human subjects. Journal of Experimental Psychology, 1971, 87, 367-372.

Thornton, J. W., & Jacobs, P. D. The facilitating effects of prior inescapable, unavoidable stress on intellectual performance. Psychonomic Science, 1972, 26, 185-187.

Thornton, J. W., & Powell, G. D. Immunization to and alleviation of learned helplessness in man. American Journal of Psychology, 1974, 87, 351-367.

Ward, W. C., & Jenkins, H. M. The display of information and the judgment of contingency. Canadian Journal of Psychology/Review of Canadian Psychology, 1965, 19, 231-241.

Weiner, B., Frieze, I., Kukla, A., Reed, L., Rest, S., & Rosenbaum, R. M. Perceiving the courses of success and failure. Morristown, N.Y.: General Learning Press, 1971.

Wortman, C. B., & Brehm, J. W. Responses to uncontrollable outcomes: An integration of reactance theory and the learned helplessness model. In L. E. Berkowitz (Ed.), Advances in experimental social psychology. New York: Academic Press, 1975.