THE EFFECTS OF NOVEL AND FAMILIAR STIMULI ON POST-TRAINING SELF-REINFORCEMENT BEHAVIOUR

JULIAN DUNCAN GRAY

copy 1

THE EFFECTS OF NOVEL AND FAMILIAR STIMULI ON POST-TRAINING SELF-REINFORCEMENT BEHAVIOUR

© Julian Duncan Gray, B.A. (Hon.)

A Thesis submitted in partial fulfillment of

the requirements for the degree of

Master of Science,

Department of Psychology

Memorial University of Newfoundland

May 1976

St. John's Newfoundland

ABSTRACT

The present investigation explored the relationship between baseline self-reinforcement behaviour and subsequent post-training self-reinforcement rate. The relationship between baseline grouping and locus of control was also assessed.

Eighty-four subjects (forty-two males and forty-two females) were placed into "high" or "low" groups contingent upon their baseline scores, and were also identified as having either an internal or an external locus of control, as measured by Rotter's scale of Internal versus External Control (1966). Members of each baseline group then received training to one of three criteria (fifty per cent, one hundred per cent, or overtraining) on a four-choice verbal discrimination task using consonant-vowel-consonant trigrams.

In the post-training phase, self-reinforcement rate was assessed by presenting a random mixture of the familiar stimuli used in training and an equal number of novel stimuli.

Results indicated that the baseline groups tended to be affected differentially as as a function of training criteria, with low baseline subjects making the greatest increases in self-reinforcement under high training levels. Support was gained for the prediction that high levels of training would tend to reduce the difference in self-reinforcement styles of different baseline groups to familiar stimuli.

In the case of novel stimuli, no exposure was given prior to the post-training phase, and subjects responded to these stimuli, thus leading to a correlation between baseline scores and post-training

self-reinforcement scores for these items.

No relationship was found between locus of control and base-line self-reinforcement performance.

ACKNOWLEDGEMENTS

The author would like to offer sincere thanks to the Psychology Department of Vancouver City College who provided space and subjects for part of the research performed for this paper.

I would like to express my gratitude also to Dr. Albert Kozma of the Psychology Department at the Memorial University of Newfoundland. Dr. Kozma's patient encouragement greatly facilitated the production of this thesis, and for his supervision I am most grateful.

Lastly, my thanks go to Barbara Ennis, the long-suffering typist who produced the final draft of this paper under the considerable pressures of a heavy workload and a fast-approaching deadline.

TABLE OF CONTENTS

TABLE OF CONTENTS

LIST OF TABLES

INTRODUCTION

A basic premise of the operant conditioning paradigm is that past reinforcement history influences the probability that a given response will be emitted. Thus, a behaviour which has led to reward in the past will be repeated with greater probability than one which has formerly led to non-reward or aversive consequences.

The natural environment is replete with examples of reinforcement principles in action as, for instance, in the case of the grading system employed by educational institutions and the use of common courtesy. High grades and praise or thanks can function to make their antecedent behaviours more likely in the future.

There remains, however, a class of behaviour which appears to be lacking any external source of reinforcement. An example of this class is the learning demonstrated in programmed-teaching or computer-assisted education. Feedback is provided, but there appears to be no other immediate source of reinforcement. In such cases, Skinner (1953) has proposed that self-reinforcement takes place.

The basis of the self-reinforcement construct is that reinforcement is available to the subject but is administered contingent upon performance criteria. Thus, in a programmed-learning setting, the subject would self-reinforce only when feedback from the programme indicates that such reinforcement is appropriate. Such use of self-reinforcement would presumably lead to a strengthening of the emitted response and would facilitate the learning process.

Masters and Mokros (1974) have suggested that "contingent self-evaluation and application of verbal labels to one's behaviours" is the common process for self-reinforcement in the natural setting. Generally, the process would take the form of a covert activity with the verbal labels functioning as secondary reinforcers. For laboratory scrutiny, however, it has been necessary to devise tasks in which the essentially covert process is translated into overt responses which can be recorded by an experimenter. To this end, three paradigms have been employed: temptation, vicarious learning and directed learning.

In the temptation paradigm, subjects are placed in a situation where self-gratification can be exercised apparently unobserved by the experimenter. The nature of this procedure has restricted its use to the observation of inappropriate or non-contingent self-gratification. Using this method, Kanfer (1966) proposed that children rated by their teacher as being in the upper half of their class would show greater adherence to the experimenter's criteria for self-reward and would give fewer inappropriate self-reinforcements than children in the lower half. It was further hypothesized that this adherence would be greater in higher grades than in lower grades. Five hundred forty-three children in grades 2 to 8 were given fifteen opportunities to guess a number between one and one hundred. The number was apparently picked from a box containing one copy of each number. If they guessed correctly, the subjects were to indicate this in a booklet provided for such a purpose.

The children were divided into three groups. The Point Reinforcement group received no apparent extrinsic reinforcement, but members of the Candy Reinforcement group were told that at the end of the experiment they would receive a candy for each correct guess.

The Model Reinforcement group saw a model guess correctly and receive a candy for three out of five demonstration trials, and was reinforced in the same manner as the Candy Reinforcement group.

It was assumed that since the probability of a correct response was one in one hundred for each trial, and much less for all trials, any self-reinforcement would be inappropriate or "cheating". The first trial was counted as a practise trial and the remaining fourteen trials were scored.

The results indicated that grades differed significantly under all treatment conditions. Positive responses decreased consistently with age. There was no significant increase in self-reinforcement by the Candy Reinforcement group over the Point Reinforcement group, indicating that candy alone was not sufficient to increase "cheating" responses. However, subjects who saw the model self-reinforce at a sixty per cent rate tended strongly toward the sixty per cent rate in their own self-reinforcement behaviour.

The comparison of upper and lower halves of the classes (as rated by the teachers) yielded the finding that children in the upper halves gave fewer self-reinforcement under Candy and Model conditions than their lower-grouped class mates. Under the Point Reinforcement condition, however, the upper and lower halves of the classes responded alike. Further analysis of the data showed that the motivational group differences resulted from increased self-rewarding by the lower groups under Candy and Model re-inforcement conditions. This finding suggested that children in the upper half were less likely to be influenced by situational variables.

While the dependent variable in this experiment was non-contingent self-reinforcement, thereby placing the study in the "Temptation" category, the Model Reinforcement condition exemplified the vicarious learning paradigm employed by a number of researchers (for example, Bandura and Kupers 1964, and Marston 1965). The method involves the use of models who are seen by subjects to self-reinforce at different rates and for varying criteria.

The Bandura and Kupers (1964) study examined the effects of a model's criterion for administering self-reinforcement upon the subsequent self-reinforcing behaviour of a child who observed the model engaged in a motor task. Models, who varied in age and sex, were observed by the subjects to take candy from a bowl on half of the task trials. Since the experimenter controlled the apparent accuracy of the models, it was possible to make it appear that some models reinforced themselves for a low and others for a high performance criterion. The results indicated that the children imitated the model's perceived accuracy for the administration of self-reinforcement.

The Marston (1965) study examined the effect of imitation upon the frequency of self-reinforcement by one hundred thirty-five male undergraduate students who performed a verbal task and self-administered a vocal, social reward. Models and subjects gave word associations to different lists of stimuli and reinforced themselves by saying "Good" whenever they considered themselves to have made the "most popular association".

Baseline measures of self-reinforcement were taken with subjects giving word associations and overt self-reinforcement. The subjects were then divided into groups for comparison of the level of involvement with a model. There were three categories of involvement: observation without concurrent word association by the subject, observation with concurrent word association by the subject but no overt self-reinforcement, and overt self-reinforcement. Where concurrent word associations were given by the subject, these were alternated with those of the model in four blocks of ten trials. For each Level of Involvement group, there were three subgroups, the first of which heard the model increase his rate of self-reinforcement in a monotonic function over the four blocks, the increase being from two to eight self-reinforcements out of ten trials. The second subgroup heard the model self-reinforce at a steady rate of two out of ten, while a control group heard no model.

Subsequent to the observation phase, all groups received forty trials without a model, giving word associations and overt self-reinforcement as in the baseline phase. In a further twenty trials, the subjects heard another person give word associations and covert self-reinforcement and were asked to say "Good" whenever the responses met their own criteria for the "best popular response".

For the group which gave word associations concurrently with the model, it was possible to assess the immediate effect of the model. The control group, having no model exhibited a significantly higher rate of self-reinforcement than two groups which had seen a model. The two model groups did not differ significantly from each other. It appeared, therefore, that both models had a depressing effect upon self-reinforcement. Analysis by blocks revealed that while Block one relationships were as indicated above, by Block four the low-modelling (two self-reinforcements in ten trials) group had decreased sufficiently to fall significantly below the high-modelling (two to ten self-reinforcements in ten trials) group. In Block one, both groups were depressed below baseline levels of self-reinforcement. The low-modelling group continued to decrease in rate of self-reinforcement to a level slightly above that of its model, while the high-modelling group increased its rate but never matched its model.

When the groups were compared on the rate of self-reinforcement during the four blocks of postobservation trials, it was found that there was a significant effect for the modelled self-reinforcement rate, but that the level of involvement of a subject with a model did not appear to alter the model's degree of influence.

There were no significant changes over blocks, indicating a continuing effect of the model over forty trials.

The data for reinforcement of another person resembled that of the postobservation self-reinforcement trials, with significant model effects, but no significant differences for the levels of involvement.

The Marston (1965) study supported the results of Bandura and Kupers (1964) by demonstrating the marked effect of a model on the self-reinforcement behaviour of an observer. It further demonstrated that this effect is not determined by the subject's overt behaviour during observation, because silent observation was as effective as overt word association and self-reinforcement in alternation with the model.

It can be seen that while the temptation paradigm allows observation of non-contingent or inappropriate self-gratification, the vicarious learning paradigm allows a certain degree of control over the subject's perception of criteria for contingent self-reinforcement. More precise control may be exerted by the use of the directed learning paradigm.

Within the framework of directed learning, a subject is trained to a specified criterion which permits a degree of predictability about the emission of the learned response. If this predictability is related to a subject's certainty or confidence in the accuracy of the response, his self-reinforcement behaviour might be expected to vary as a function of his degree of training. This was investigated by Kanfer, Bradley and Marston (1962) who taught eighty male undergraduates a visual discrimination task to two levels of training. Green or red lights were used to indicate to the subject whether he had made a correct or an incorrect response, respectively on each trial. The low-learning group received twenty-five trials, while the high learning group received fifty trials. In Phase II, both groups were subdivided, half of each receiving twenty self-reinforcement test trials and half an unrelated filler task.

During the test trials, the subject guessed the right-left location of a particular stimulus type before they saw the stimulus card. Upon presentation of the card, they could activate the green light if they considered their guess correct.

The authors, in summarizing their results, state:

> The number of correct self-reinforcing responses was significantly higher with longer acquisition, and both types of incorrect self-reinforcement (failure to reinforce and reinforcing of incorrect responses) were lower with shorter training. The accuracy of (self-reinforcement) responses was related to amount of learning.

In a later study, Bartol and Duerfeldt (1970) made the further discovery that post-training self-reinforcement reflected the subject's self-reinforcement rate before training. Kozma and Easterbrook (1974) brought these findings together, suggesting that self-reinforcement rate is a joint function of prior external feedback (training) and baserate.

In order to test this idea, they established baseline self-reinforcement rate on a four-choice verbal discrimination task and then brought the subjects to either forty per cent, sixty per cent, or eighty per cent training criteria on the same task. A post-training phase then assessed the effects of training on self-reinforcement tendencies. They hypothesized that training would increase self-reinforcement rate only if a subjects training criterion exceeded his baseline score. Their prediction was borne out by the data for low baseline subjects, whose self-reinforcement rate increased significantly more than either medium or high baseline subjects in the forty per cent and sixty per cent conditions. Under the eighty per cent criterion, both low and medium baseline subjects showed significantly greater increases than high baseline subjects. It was noted that high baseline subjects had a much greater tendency to administer self-reinforcement after incorrect responses than did medium and low baseline subjects, who made progressively more errors of omission, failing to self-reinforce after correct responses.

The high rate of incorrect self-reinforcement in high baseline subjects was a phenomenon common to all training conditions. The main change for these subjects was the greater number of correct choices, for which self-reinforcement was appropriate, in the higher training groups. The tendency of low and medium baseline subjects was to omit self-reinforcement when it would have been appropriate and this tendency diminished as training levels increased. The authors suggested that this reduction could be related to an increase in confidence on the part of low and medium baseline subjects.

Kozma and Kerwin (1975) used a directed-learning paradigm with forty per cent and eighty per cent training criteria to investigate the effects of confidence on subjects' self-reinforcement behaviour. They found that subjects who exhibited high levels of self-reinforcement during baseline rated themselves as being more confident than those who had low baserates. Further, higher confidence scores followed training to the eighty per cent criterion than to the forty per cent criterion. Their results, while correlational, and therefore not capable of confirming causality, do appear to lend support for the idea that confidence and self-reinforcement are directly related.

If confidence is the key variable, and if it is amenable to alteration by training, it should be possible to increase the self-reinforcement behaviour of low baseline subjects to a level similar to that of their high baseline counterparts by employing high training criteria, as suggested by the Kozma and Kerwin (1975) findings. To this end, the present study, using the directed learning program, utilized fifty per cent, one hundred per cent and overtraining criteria.

It was suggested that high training criteria would result in more pronounced changes in the self-reinforcement rate of low baseline subjects than they would in the case of high baseline subjects. Thus, the greatest effect would occur for a group whose baserate was low and whose training level was high.

Also of concern in the present study was the generality of the training effect. As noted, it was anticipated that the differential between the self-reinforcement rates of high and low baseline subjects would be reduced or eliminated by training subjects to a high criterion. However, this effect might apply only to the stimuli employed in the training phase. As has already been shown, self-reinforcement behaviour is amenable to standard operant conditioning techniques (Skinner, 1953 and other studies cited) and it therefore seemed reasonable to assume that the training stimuli function as discrete discriminative stimuli which lead to the emission of self-reinforcement behaviour. Consequently, if a new set of stimuli were to appear in the post-training phase, the subject would be unprepared to make intelligent discriminations on these items. He would therefore be in the same uninformed state that he was in during the baseline phase and might reasonably be expected to exhibit a similar self-reinforcement pattern to that which he followed during the baseline procedure.

Such reasoning finds some support in the pronounced baseline differences found in both the Bartol and Duerfeldt (1970) and Kozma and Easterbrook (1974) studies. Such differences suggest a past reinforcement history which would outweigh the effects of training in the present procedure.

Consequently, a mixture of the training stimuli and an equal number of novel stimuli was presented to the subjects in the post-training phase. It was hypothesized that training effects would be apparent for the familiar stimuli used in training, but that a return to baserate would be shown for novel stimuli.

Skinner (1953) has suggested that when an individual's reinforcement history is long enough, certain types of response become the *modus operandi* of his behaviour. In humans, such patterns are frequently labelled as personality traits. If self-reinforcement rate in the baseline phase is the product of a long reinforcement history, it might also be considered to reflect a personality trait. If not actually a discrete characteristic, it might still be related to some other, perhaps more fundamental trait.

Following a similar rationale, Bellack (1972) investigated the relationship between self-reinforcement and a subject's locus of control as measured by the Rotter (1966) Intraversion-Extraversion Scale. People scoring at the low end of the Rotter Scale are identified as "internals" and believe that reinforcement results from their own activity and that they determine their own fate. Those at the high end are "externals" and believe that reinforcement is essentially independent of their own behaviour and that their destiny is controlled by fate or powerful outside forces (Bellack, 1972, Rotter 1966). The use of self-reinforcement is also dependent upon perceived behavioural contingencies and might be expected to relate to locus of control. Consequently, "internal" subjects could be expected to emit self-reinforcement responses at a high rate while the reverse would be true for "external" subjects.

Various investigators, cited by Bellack (1972) have demonstrated the relationship between locus of control and such phenomena as: the amount of knowledge tuberculosis patients on the same ward had about their course of treatment; the degree of familiarity reform-school inmates had with parole and institutional rules; and the amount of involvement in civil rights movements of Negro college students. In all cases, it was "internal" subjects who were more knowledgeable or more active. These data suggest that "internal" subjects take a more active role in their affairs and perhaps use self-reinforcement to a greater extent than "external" subjects.

Bellack (1972) used the directed learning procedure and taught sixty-eight subjects a four-choice verbal discrimination task to a sixty per cent criterion. He then presented half of the subjects with ten trial blocks using the ten training stimuli while the other half of the subjects received a set of different stimuli. During this phase, self-reinforcement was observed. Finally, the Rotter scale was administered. It was hypothesized that "(1) there would be no differences between internals and externals during initial training; (2) internals would use more self-reinforcement than externals during the self-reinforcement phases of the study, and (3) there would be an interaction between use of self-reinforcement and test stimuli, internals having a larger self-reinforcement increment over externals with (unfamiliar) generalization stimuli, for which there would be less environmental feedback on which to judge performance".

Of the three hypotheses, only the first, relating to initial training period, was supported. There was no significant difference in acquisition rate between the internal and external groups. The lack of support for the remaining hypothses is of perhaps greater interest, however. Examination of the data suggests that behaviour during the self-reinforcement phase was closely tied to training experience. This is not surprising in the light of the more recent Kanfer, Bradley and Marston (1962) and Kozma and Easterbrook (1974) studies. It would appear that training effects were sufficient to obscure the differences between internal and external subjects.

The present study therefore proposed to relate locus of control to baseline self-reinfrocement. It was hypothesized that since strong baseline effects do exist (Kozma and Easterbrook, 1974) and that since baseline is measured before training and is therefore uncontaminated by training effects, the optimal point at which to assess the relationship between self-reinforcement behavior and locus of control would be in the baseline phase.

In summary, the hypotheses of the present study were:

1. High training criteria would reduce or eliminate the difference in self-reinforcement tendencies between high and low baseline subjects for familiar stimuli.

2. Baseline differences would be reflected in the response to novel stimuli used in the post-training phase.

3. A negative relationship would exist between baseline self-reinforcement scores and scores on the Rotter scale; low baseline self-reinforcement subjects would have high (external) scores on the Rotter scale and high baseline self-reinforcement subjects would have low (internal) scores.

METHOD

Subjects

Participants in the present study were forty-two male and forty-two female undergraduate students of Vancouver (British Columbia) City College and Memorial University of Newfoundland.

Apparatus and Materials

In front of the subject was a small (10cm. x 10cm.) panel with a button at the centre of each quadrant and a lamp (G.E. # 1892) at the top-centre. In front of the experimenter was a similar panel with lamps (G.E. # 1892) replacing the buttons on the subject's panel and a lamp (G.E. # 1892) at the top-centre. Each button on the subject's panel was wired to a lamp on the experimenter's panel. The lamps at the top-centre of each panel were operated by a hand-held button (the "correct" button on the end of a two-metre cord.

Each of the twenty-five stimulus cards (7.5 cm. x 12.5 cm.) displayed a medium-association (forty-eight per cent to fifty-two per cent, Archer, 1960) consonant-vowel-consonant trigram at the centre of each quadrant. Thus, one hundred trigrams were assigned randomly to their respective quadrants.

Five stimulus cards were used as a pre-baseline practise set. Ten cards were used throughout the baseline, training and post-training phases, and comprised the "familiar" set. On each "familiar" card, one trigram was randomly designated as being correct. During the post-training phase, the "familiar" cards were shuffled together with the remaining ten "novel" cards. The trigrams on each card are listed in

Appendix C.

All experimental instructions were presented on typed sheets which are reproduced in Appendix B.

Design

The design consisted of two baseline groups. Kozma and Easterbrook (1974) found that baseline scores were evenly distributed into three groups by using baseline scores of less than seven, seven to thirteen, and thirteen to twenty self-reinforcements in twenty trials. In order to highlight the baseline effect, only subjects with baseline scores of seven or less ("low") and thirteen or more ("high") were included in the present study. Subjects with scores between seven and thirteen were taken through the experimental procedure but were not included in the analysis.

The study employed three learning criteria in the training phase: fifty per cent, one hundred per cent and overtraining. In the overtraining condition, subjects were given one additional block of trials after reaching the one hundred per cent level of competence.

During the post-training phase, two types of stimulus were employed; "familiar" and "novel".

Thus, the design employed a 2 (baseline) x 3 (training) x 2 (stimulus type) factorial analysis of variance with repeated measures on the last factor (Winer, 1962).

Baseline and locus of control data were analyzed seperately.

Procedure

The subject and the experimenter were seated at adjacent sides of a table and a 40 cm. x 60 cm. partition obscured each other's view of the other side.

The experimenter first asked the subject to complete the Rotter Scale (Appendix A) by selecting one statement from each pair of statements. The subject's choice was to be based on his opinion of which statement was more descriptive of himself. There was no further verbal exchange between the subject and the experimenter until the scale was completed.

The subject was then asked to read the instructions for the baseline phase. These instructions requested each subject to select the correct trigram from four printed on each card and to indicate his choice by pushing a button corresponding to the selected trigram. Then, if he felt that he had made a correct choice, the subject was to indicate it by pushing the "correct" button.

Many subjects requested more information about the basis for their decisions. They were told that this choice was part of their task and were then encouraged to continue.

The practise set was then presented, one card at a time. For each card, the subject was given as much time as he desired to make a decision. Following the presentation of the last card, the experimenter said "OK?" to which most subjects replied "Yes" or, "I guess so". Any subject wo asked for further information was given the same advise as before.

The baseline phase followed. Cards of the "familiar" set were shuffled, presented one at a time, shuffled, and presented again. Since there were ten cards and each card was presented twice, the baseline consisted of twenty trials.

Immediately following the baseline phase of the study, the subject was given the instructions for the training phase. In this section of the experiment, cards of the main set were presented as before. They were shuffled after each block of ten trials in order to randomize the sequence of presentation. During the training phase, external feedback on the accuracy of a subject's choice was provided. Thus, the experimenter operated the "correct" button when the subject chose a trigram which had been previously designated as correct by the experimenter.

Three training groups were used, as described, and subjects were assigned to these on a random basis. For the fifty per cent and one hundred per cent criterion groups, training was terminated at the end of the block in which criterion learning was demonstrated. Subjects in the overtraining group were given one additional block after reaching the one hundred per cent level.

At the conclusion of training, each subject was handed the post-training instructions. Post-training was similar to baseline with the exception that cards of the "novel" set were now shuffled in with the "familiar" cards. After all twenty cards had been presented, they were shuffled and presented again. Post-training scores were, therefore, out of twenty for both the "familiar" and the "novel" sets.

Throughout the procedure, the experimenter recorded the code number of each card presented and all the subject's responses. No time limits were placed upon the subject, and the average time per subject was one hour.

RESULTS

Locus of Control and Baseline Self-Reinforcement

The first phase of the analysis involved the calculation of a Pearson r for the scores on baseline and the Rotter scale. The result ($r = .054$) failed to show a significant relationship between the two sets of scores and further analysis of these data was abandoned.

Analysis of Post-Training Self-Reinforcement Scores

Post-training self-reinforcement scores for all treatment conditions are given in Table 1. A 2(Baseline) x 3(Training) x 2(Stimulus Type) analysis of variance with repeated measures on the last factor yielded significant Baseline ($F(1,36) = 12.05$; $p < .05$) and Stimulus Type ($F(1,36) = 64.25$; $p < .05$) main effects. Significant effects were also found for the Baseline x Training and the Baseline x Stimulus Type interactions, ($F(2,36) = 5.53$; $p < .05$ and $F(1,36) = 6.00$; $p < .05$, respectively. Table 2 summarizes these findings.

The significant Baseline effect reflected the overall tendency of high baseline subjects to self-reinforcement at a higher rate than low baseline subjects during the post-training phase.

An inspection of the Baseline x Training interaction scores (Table 3) suggested that the main source of Baseline differences was in the fifty per cent training group. To test this notion, a Newman-Keuls procedure (Kirk, 1968) was employed (Table 4). The data revealed that significant differences ($p < .05$) existed for comparisons between subjects in the low baseline, fifty per cent training group and all other groups. For high baseline subjects in the fifty per cent training group, significant ($p < .05$) differences existed when they were compared

TABLE 1

Summary of Baseline Means with Post-Training Means of Self-Reinforcement on Novel and Familiar Stimuli for all Levels of Baseline and Training

Baseline	Training	Baseline	Familiar	Novel
High	100^+	14.4	14.1	8.0
	100	16.3	19.9	9.6
	50	17.7	18.0	14.1
Low	100^+	0.9	19.1	2.6
	100	3.0	18.0	5.9
	50	1.0	10.7	1.3

TABLE 2

Summary of Analysis of Variance using Novel and Familiar Stimuli as Repeated Measures

Source	SS	df	MS	F
Between	2056.20	41	50.15	1.52
Baseline (B)	398.68	1	398.68	12.05*
Training (T)	100.67	2	50.34	1.52
BT	366.00	2	183.00	5.53*
Error (between)	1190.86	36	33.08	
Within	3458.50	42	82.35	2.66
Stimulus Type (S)	1991.44	1	1991.44	64.26*
BS	186.01	1	186.01	6.00*
TS	100.67	2	50.34	1.62
BTS	64.67	2	32.34	1.04
Error (within)	1115.71	36	30.99	

*$P < .05$

TABLE 3

Post-Training Self-Reinforcement

Cell Totals for Baseline and Training Levels

		TRAINING		
		100+	100	50
BASELINE	HIGH	155	206	225
	LOW	152	167	84

with all but the high baseline, one hundred per cent group. Further, there were no significant differences between baseline groups in either the one hundred per cent or the overtraining conditions.

The significant Stimulus Type effect showed that subjects issued self-reinforcement at a greater rate for "familiar" stimuli than they did for "novel" stimuli. Analysis of the Baseline x Stimulus Type scores (Table 5) indicated that subjects in both baseline groups responded to "familiar" stimuli in a similar manner. There was, however, a significant difference ($p < .05$) in their response patterns to "novel" stimuli. The subjects with high baseline scores emitted self-reinforcement at a higher rate than did their low baseline counterparts. Both baseline groups responded differentially to "familiar" and "novel" stimuli.

TABLE 4

Comparison of Post-Training Reinforcement Means for Baseline and Training Groups. A Newman-Keuls Procedure has been used to rank and compare cells shown in Table 3

Group	Description Baseline-Training	1	2	3	4	5	6	Critical Values
1	Low - 50%	-	68*	71*	83*	122*	141*	64.8
2	Low -100%		-	3	15	54	73*	61.8
3	High -100%			-	12	51	70*	58.8
4	Low -100%				-	39	58*	53.4
5	High -100%					-	19	43.7
6	High - 50%						-	

TABLE 5

Summary of F Scores for Cell Comparisons in the Baseline x Stimulus Type Interaction

Comparison between High and Low Baseline	
Familiar	.97
Novel	27.33*
Comparison between Novel and Familiar Stimuli	
High Baseline	23.24*
Low Baseline	82.16*

*P < .05

DISCUSSION

Locus of Control

The data suggested that locus of control is unrelated to baseline self-reinforcement behaviour. This finding, along with Bellack's (1972) discovery that locus of control is unrelated to post-training self-reinforcement behaviour is difficult to explain. The studies cited by Bellack (1972) support the notion that "internal-subjects, while controlling their own destiny, do so by exploring and becoming involved with their environment. It seems possible, in this light, that "internal" subjects might be at a loss to deal with a situation in which there is virtually total ambiguity. They could therefore be expected to adopt a variety of strategies, and appear to have done so. In the case of "external" subjects, it would appear that they fail to seek out relevant stimulus characteristics, preferring to rely on feedback from external sources.

Because of this reliance, their behaviour might be unpredictable in an ambiguous task situation.

In summary, while "internal" subjects may be seeking, but unable to find, stimulus cues in an ambiguous situation, "external" subjects may not be seeking for cues at all. In both cases, the lack of input appears to have led to an outcome which was unrelated to locus of control and perhaps had more to do with past reinforcement history. Marston (1964), as noted earlier, was also unsuccessful in finding a personality correlate of self-reinforcement behaviour.

Familiar Stimuli

The data showed that training criteria had a strong influence on the subjects' post-training self-reinforcement behaviour. In the case of high baseline subjects, by whom self-reinforcement had been administered at a high rate prior to training, the effects of even high training criteria were not dramatic because there was not very much opportunity for an increase in self-reinforcement rates.

During the baseline phase, there was no "correctness" or "incorrectness" inherent in the task, and the decision of each subject to reinforce or not to reinforce himself was perhaps related to an internal state of the subject. Kozma and Kerwin (1975) have chosen the notion of confidence to describe this internal state. High baserate subjects might be conceptualized as being a residually confident group, and training would therefore do little to increase this state.

Low baseline subjects would, under the confidence notion, be described as less confident than their high baseline counterparts because their rate of self-reinforcement was low. Such cautious behaviour would be altered when training increased a subject's certainty about the correctness of his responses. Thus, as the data revealed, low baseline subjects would increase their rate of self-reinforcement to a level which closely approximated their certainty, or degree of training. Under the fifty per cent criterion, the subjects increased their self-reinforcement rates much less than the low baserate subjects who were trained to the higher criteria. Such findings are in agreement with those of Kozma and Easterbrook (1974) who found that post-training rates approximated training criteria when the criteria were greater than the subjects' baserate.

In the present study, one score (for the high baserate, overtraining group) did not conform to the overall patterns of the hypothesis and the results, being somewhat lower than expected. Inspection of the data reveals that the baserate for this group was also atypically low. While the mean scores of the fifty and one hundred per cent training groups were 16.3 and 17.7, respectively, the overtraining group mean was only 14.4. The cutoff point for the high baseline scores was 13, so the 14.4 mean for the overtraining group represents a cluster of scores at the lower extremity of the high baseline continuum.

The fact that subjects with this abnormally low range of scores were placed in one training group may be attributed to the random assignment of subjects to training groups. From the rest of the data in this study, and from the findings of Kozma and Easterbrook (1974) it seems improbable that such a result is typical. A replication of the present procedure should be performed in order to clarify this point.

Novel Stimuli

The confidence model described above can be applied to the data for post-training self-reinforcement responses to novel stimuli. Low baseline subjects would be viewed as lacking in confidence in the accuracy of their choices. Regardless of their training criterion for the familiar stimuli, the subjects' confidence remained virtually unaffected with regard to the novel stimuli. Accordingly, the analysis showed that low baserate subjects tended to revert to their baseline self-reinforcement patterns when confronted with unfamiliar stimuli in the post-training phase.

The slight increase in post-training self-reinforcement scores on novel stimuli obtained in the current investigation might bear further investigation. It is possible that, while confidence appeared to be an item-specific phenomenon, there was some general effect related to the task itself. Thus, low baseline subjects may have been more test anxious than high baseline subjects and this fear may have abated as the task continued. The result of fear reduction would be the increased administration of self-reinforcement noted in the data.

High baseline subjects confronted with novel post-training stimuli exhibited a general decrease in self-reinforcement rate. This effect was minimal in the low training group but was quite apparent for subjects in the one hundred per cent and the overtraining groups. As noted, high baseline subjects can be conceived of as being residually confident. The reduction in the self-reinforcement rate of this group to novel stimuli during the post-training phase was, therefore, an interesting phenomenon.

At this point, there is no reason to believe that subjects in the two baseline groups differ in their adaptability to environmental contingencies. It is reasonable, however, to imply that the groups differ in the way that they deal with environmental situations. The low baseline (low confidence) group appears to advance with caution until adequate information is supplied. The high baseline subjects seem to advance into new situations with the attitude that they will emit responses of a certain type unless the environment clearly demands responses of a different type.

In this study, the provision of two completely different sets of stimuli in the post-training phase probably highlighted the demand characteristics of each set. After the training phase, all subjects were at least fifty per cent familiar with the training stimuli. Those subjects who had received high degrees of training were even more familiar with the training stimuli. Nevertheless, all subjects could be expected to differentiate between the two stimuli sets. This high degree of discriminability, particularly for subjects in the one hundred per cent and the overtraining group, appears to have been sufficient to cue a change in the response strategy of high, baseline subjects, resulting in the post-training phase. This is a new finding in the study of self-reinforcement.

The fact that high baseline subjects did not exhibit a greater reduction in self-reinforcement rate may be the result of the failure of the discriminability and demand characteristics of the two sets of stimuli to outweigh the effects of "residual confidence". The balance between these components is a relationship which invites further investigation.

The confidence model, together with a maxmal discrimination factor, appears to account for most of the data in the present study. It is however, of interest to note that the confidence, as implied by baseline performance was not related to the subjects' perceived locus of control, as measured by the Rotter (1966) scale. It is apparent that further research is necessary to explain the paradoxical observations.

Summary

1. The hypothesis that high training criteria would reduce or

eliminate the differences in post-training self-reinforcement rates to familiar stimuli of high and low baseline subjects was supported.

2. The hypothesis that subjects would tend to reflect their baserates in their responses to novel post-training stimuli was supported.

3. No relationship was found between a subject's perceived locus of control and his baseline self-reinforcement rate. An explanation of this finding was suggested.

REFERENCES

Archer, C.J. A revaluation of the meaningfulness of all possible consonant vowel consonant trigrams. Psychological Monographs, General and Applied, 1960, 74, #497.

Bandura, A ., and Kupers, C. Transmission of patterns of self-reinforcement through modeling. Journal of Abnormal and Social Psychology, 1964, 69, 1-9.

Bartol, G. H., and Duerfeldt, P.H. Self reinforcing behaviour: The effects of base rate and dependency. Journal of General Psychology, 1970, 83, 151-161.

Bellack, A. Internal vs External Locus of Control and the use of Self-reinforcement. Psychological Reports, 1972, 31, 723-733.

Kanfer, F.H., Bradley, M.M., and Marston, A.R. Self- reinforcement as a function of degree of learning. Psychological Reports, 1962, 10, 885-886.

Kanfer, F.H. The influence of age and incentive conditions on children's self-rewards. Psychological Reports, 1966, 19 263-274.

Kirk, R.E. Experimental Design, Procedures for the Behavioural Sciences. Belmont, California: Brooks/Cole Publishing Company, 1968, 91-93.

Kozma, A. and Easterbrook, P. Effects of baseline self-reinforcement behaviour and training level on post-training self-reinforcement behaviour. Journal of Experimental Psychology, 1974, 102 (2) 256-259.

Kozma, A., and Kerwin, S. The effects of a monetary incentive on the accuracy of post-training self-reinforcement behaviour. Canadian Journal of Behavioural Science, 1975, 7(4), 420-428.

Marston, A.R. Personality variables related to self-reinforcement. Journal of Psychology, 1964, 58, 169-175.

Marston, A.R., Imitation, self-reinforcement and reinforcement of another person. Journal of Personality and Social Psychology, 1965, 2, 2, 255-261.

Masters, J.C., and Mokros, J.R. Self-reinforcement processes in children. In W.H. Reese (Ed.) Advances in Child Development and Behaviour. New York: Academic Press, 1974, 151-187.

Rotter, Julian B. Generalized expectancies for internal versus external control of reinforcement. Psychological Monographs, 1966, 80, 1, #609.

Skinner, B.F. Science and Human Behaviour. New York: MacMillan, 1953.

Winer, B.J. Statistical Principles in Experimental Design. New York: McGraw -Hill Book Company, Inc., 1962, 337-353.

APPENDIX A – Intraversion-Extraversion Scale

1. a. Children get into trouble because their parents punish them too much.
   b. The trouble with most children nowadays is that their parents are too easy with them.
2. a. Many of the unhappy things in people's lives are partly due to bad luck.
   b. People's misfortunes result from the mistakes they make.
3. a. One of the major reasons why we have wars is because people do not take enough interest in politics.
   b. There will always be wars, no matter how hard people try to prevent them.
4. a. In the long run people will get the respect they deserve in this world.
   b. Unfortunately, an individual's worth often passes unrecognized no matter how hard he tries.
5. a. The idea that teachers are unfair to students is nonsense.
   b. Most students don't realize the extent to which their grades are influenced by accidental happenings.
6. a. Without the right breaks one cannot be an effective leader.
   b. Capable people who fail to become leaders have not taken advantage of their opportunities.
7. a. No matter how hard you try some people just don't like you.
   b. People who can't get others to like them don't understand how to get along with others.

8. a. Heredity plays a major role in determining one's personality.
   b. It is one's experience in life which determine what they like.

9. a. I have often found that what is going to happen will happen.
   b. Trusting to fate has never turned out as well for me as making a decision to take a definite course of action.

10. a. In the case of the well prepared student there is rarely, if ever, such a thing as an unfair test.
    b. Many times exam questions tend to be so unrelated to course work that studying is really useless.

11. a. Becoming a success is a matter of hard work, luck has little or nothing to do with it.
    b. Getting a good job depends mainly on being in the right place at the right time.

12. a. The average citizen can have an influence in government decisions.
    b. This world is run by the few people in power, and there is not much the little guy can do about it.

13. a. When I make plans, I am almost certain that I can make them work.
    b. It is not always wise to plan too far ahead because many things turn out to be a matter of good or bad fortune anyhow.

14. a. There are certain people who are just no good.
    b. There is some good in everybody.

15. a. In my case getting what I want has little or nothing to do with luck.
    b. Many times we might just as well decide what to do by flipping a coin.

16. a. Who gets to be boss often depends on who was lucky enough to be in the right place first.

    b. Getting people to do the right things depends upon ability, luck has little or nothing to do with it.

17. a. As far as world affairs are concerned, most of us are the victims of forces we can neither understand, nor control.

    b. By taking an active part in political and social affairs the people can control world events.

18. a. Most people don't realize the extent to which their lives are controlled by accidental happenings.

    b. There is really no such thing as "Luck".

19. a. One should always be willing to admit mistakes.

    b. It is usually best to cover up one's mistakes.

20. a. It is hard to know whether or not a person really likes you.

    b. How many friends you have depends upon how nice a person you are.

21. a. In the long run the bad things that happen to us are balanced by the good ones.

    b. Most misfortunes are the result of lack of ability, ignorance, laziness or all three.

22. a. With enough effort we can wipe out political corruption.

    b. It is difficult for people to have control over the things politicians do in the office.

23. a. Sometimes I cannot understand how teachers arrive at the grades they give.

23. b. There is a direct connection between how hard I study and the grades I get.

24. a. A good leader makes it clear to everybody what their jobs are.

b. A good leader expects people to decide for themselves what they should do.

25. a. Many times I feel that I have little influence over the things that happen to me.

b. It is impossible for me to believe that chance or luck plays an important role in my life.

26. a. People are lonely because they do not try to be friendly.

b. There is not much use in trying too hard to please people, if they like you, they like you.

27. a. There is too much emphasis on athletics in high school.

b. Team sports are an excellent way to build character.

28. a. What happens to me is my own doing.

b. Sometimes I feel I do not have enough control over the direction my life is taking.

29. a. Most of the time I cannot understand why politicians behave the way they do.

b. In the long run the people are responsible for bad government on a national as well as on a local level.

APPENDIX B - Instructions

Pre-Baseline:

Thank you for helping us! We are most anxious to determine the average way in which people respond to the task that you are about to be given. Your score will contribute to that average.

When you finish reading this, tell the experimenter and he will then show you a card which looks like this:

Word #1	Word #2
Word #3	Word #4

Your job is to choose one of the words (actually nonsense syllables) and indicate your choice by pushing one of the four buttons in front of you. You will notice that the buttons are arranged in the same way as the words on the card. Just push the button which is in the same position as your chosen syllable.

Next, if you think your choice was correct, push the button on the end of the cord. It will light the bulb in front of you.

Don't worry if you do not immediately understand the problem. You will soon catch on.

Remember; choose a word by pushing a button, then put the light on in front of you with the other button if you think you were right.

Let the experimenter know when you have finished this.

Training:

We are going to change things a little now. First of all, please give the "correct" button to the experimenter.

Thank you. We are now going to continue as before, but the experimenter will push the "correct" button to let you know when you have chosen the right word. Your job is to learn the correct syllables for each card as well as you can.

If you are ready, tell the experimenter, please.

Post-Training:

Remember the first part of this experiment? We shall now repeat it. You make choices, as you have done all along, but the experimenter will give the "correct" button to you. Once again, you are to turn on the "correct" light whenever you think your choice is the right one. When you are ready, let me know.

APPENDIX C

Consonant-Vowel-Consonant Trigrams used as Stimuli. Listed in clockwise order from upper left corner on each card.

Practise

1.	FID	SEB	JOQ	PAJ
2.	POH	CEP	TAQ	JEY
3.	MUZ	VID	QAL	NEF
4.	LEH	MEZ	KYX	VEP
5.	YUC	NUW	PYR	MIF

Familiar

1.	SAH	KES	LEK	GEY
2.	CYS	KYL	QIK	VOD
3.	BYR	MYK	SYX	DOH
4.	CIF	FOW	VIZ	BYG
5.	DYR	POZ	KAC	ROH
6.	SOQ	KIZ	BAR	QAD
7.	NOH	HAX	FAW	FOT
8.	LAH	NYK.	LOH	GUK
9.	DAK	MUW	YOM	WOH
10.	NOY	BYC	HEJ	REW

NOVEL

1.	NAZ	CUK	HEF	DUQ
2.	MYC	NUS	WEL	NUZ
3.	FAH	POY	LUF	HIV
4.	VUL	LEB	SUF	QIP
5.	JOP	JIV	JYN	BEM
6.	GYL	ROQ	WYF	PEZ
7.	DIB	BIV	KYN	LEZ
8.	DUH	TAY	VOS	ZEK
9.	SYG	JIS	MOX	KUL
10.	SUX	SAZ	HEB	CYR