ON THE NATURE OF NON-RESPONDING IN
DISCRIMINATION LEARNING WITH AND
WITHOUT ERRORS¹

H. S. TERRACE

COLUMBIA UNIVERSITY

In human subjects, discrimination learning with errors results in active responding incompatible with the reinforced response. The direction of such incompatible behavior is opposite to that of the reinforced response. Responding occurs only during the stimulus correlated with extinction. The frequency of active non-responding is maximal shortly after the start of discrimination training (the time at which the frequency of errors decreases most rapidly) and approaches zero as discrimination training continues. The magnitude of behavioral contrast is not related systematically to the number of errors. Instead it is related directly to the frequency of active non-responding. Active non-responding appears to be motivated by the aversiveness of self-produced frustration, in the sense that active non-responding allows the subject to avoid the aversiveness of non-reinforced responding.

Pavlov's classical experiments on extinction and discrimination learning were the first of innumerable demonstrations that withholding reinforcement decreases the frequency of a conditioned response (Pavlov, 1927). This empirical generalization is, of course, consistent with reinforcement theory. Just as reinforcement should increase response strength, the removal of reinforcement should decrease response strength. What replaces behavior weakened by extinction has yet to be shown. Since there is no reason to believe that conditioning is an exception to the rule that nature abhors a vacuum, an account of the behavior that replaces the extinguished response is of fundamental importance for conditioning theory.

Some of the characteristics of the behavior that replaces a conditioned response are suggested by experiments that have compared discrimination performance following learning with and without errors (Terrace, 1963, 1972a). Two groups of subjects are trained to discriminate two stimuli, one correlated with reinforcement (S+) and one correlated with extinction (S-). One group is trained by the traditional method that allows many unreinforced responses (errors) to occur in the presence of S-.

Another group is trained by a fading procedure that starts discrimination training with a problem simple enough not to occasion any errors. Discrimination training progresses gradually to the same terminal problem learned by the first group. Most subjects trained by the fading procedure make few or no errors.

Among the many differences in performance that have been observed following discrimination learning with and without errors (Terrace, 1972a), the three most relevant to a characterization of non-responding are (i) "emotional" responses in the presence of S-, e.g., wing-flapping in pigeons (Terrace, 1966), (ii) an increase in the rate of responding to S+, referred to as "behavioral contrast" (Reynolds, 1961; Terrace, 1966), and (iii) the learning of a new response that allows the subject to escape from S- (Terrace, 1971). These phenomena, which are observed only after discrimination learning with errors, but which are not permanent characteristics of discrimination performance (Terrace, 1972a), suggest that the occurrence of non-reinforced responding temporarily energizes behavior. A similar conclusion, based upon a series of experiments on frustrative non-reward was reached earlier by Amsel (1958, 1962, 1972). Azrin, Hutchinson, and Hake (1966) noted that an alternation of periods of reinforcement and extinction results in aggression during extinction. Their results are also consistent with the generalization that non-reinforced responding increases activity level.

¹This research was in part supported by grants from the National Science Foundation (GB-30781) and from the National Institutes of Health (HD-00930). Reprints may be obtained from H. S. Terrace, Columbia University, 360 Schermerhorn Hall, New York, N.Y. 10027.

The standard method of training a successive discrimination allows the experimenter to observe only the topography of the response reinforced in the presence of S+. A simple change in procedure, which requires that the subject maintain continuous contact with the response manipulandum, allows the experimenter to observe the topography of other responses. This method has been used in at least one study of the extinction of leg flexion in the dog following avoidance conditioning (Brogden, Lipman, and Culler, 1938). In this experiment, the right foreleg rested on a tambour, which could record both flexion and extension. During the first phase of the experiment, the flexion response was conditioned by an avoidance paradigm. Even though extensions of the leg were not required by the avoidance contingency, they nevertheless occurred, albeit at relatively low frequency. During extinction, the frequency of extension was greater than it was during conditioning. Brogden *et al.* concluded that "despite the virtual cessation of gross leg-flexion . . . extinction . . . proves to be far indeed from complete inaction" (p. 114).

The literature on discrimination learning based upon positive reinforcement does not appear to contain any observations of behaviors that occur in the presence of S-. An obvious problem is getting an animal to maintain constant contact with the manipulandum without resorting to aversive control or to some sort of restraining device. The present experiment circumvented this problem by using human subjects, who had no difficulty complying with instructions to hold onto the response manipulandum throughout their experimental training. The manipulandum, a joystick, allowed the experimenter to record behavior both during S+ and S-. Forward motion of the joystick operated a microswitch. This defined a "positive" response, one that could be followed by reinforcement in the presence of S+, both during baseline and during discrimination training. Movement of the joystick in the opposite direction activated a force transducer. The output of the force transducer provided a measure of behavior in a direction opposite to that required for reinforcement. For this reason, displacements of the joystick that activated the transducer will be referred to as "antagonistic". Even though the main reason for using human subjects was to enable the experimenter to monitor behavior in S- as

well as in S+, it is also of interest to observe human performance on a successive discrimination and to see how such performance compares with that of animal subjects.

METHOD

Subjects

Eight undergraduates (five males, three females) were recruited through the student newspaper.

Procedure

Each subject was instructed to attend to a screen located 1.8 m from the front edge of his chair. Subjects were also told that they could earn dimes by operating a joystick located to the right of the chair whenever a stimulus appeared on the screen and that it was important to keep their hand on the joystick at all times, even when not operating it. The dimensions of the room housing the subjects' chair, the joystick, and the stimulus screen were 3.2 (w) by 5.4 (l) by 2.5 (h) m. This room was visually and acoustically isolated from the rest of the laboratory.

The instructions read to each subject were as follows:

In this experiment you can earn money by pulling the lever to the right of the chair towards you. You may sit in the chair in whatever position makes you most comfortable. Please look at the screen and the panel beneath it throughout the session. You should only pull the lever towards you when the square is projected on the screen. When you pull the lever, pull it repeatedly, that is, towards you—back to neutral—towards you, and so on.

Everytime you pull the lever you will cause the yellow light in front of you to flash momentarily. Occasionally you will also cause the green light to flash and, at the same time, one unit will be added to the counter. At the end of the session you will be given as many dimes as are shown on the counter. Remember that you will have to respond many times in order to earn a dime. You will not, however, be able to predict which responses will produce a dime. Thus, your best strategy is to respond at a steady rate at all times when the square is presented.

At the end of each half minute you will see the stimulus on the screen shut off momentarily. After a pause of two seconds the square will reappear at which time the lever will again be operative.

This is the first of a series of sessions of varying length. The first few sessions will be shorter than the subsequent sessions. At times, particularly in the longer sessions, you may find the task of pulling the lever boring. Just the same it is necessary to concentrate on the stimuli.

One final instruction: Please keep your hand on the lever at all times—this is true even if you are not operating it.

To make a response, it was necessary to move the stick towards the subject from a "neutral" position in a manner similar to that required to engage a floor-mounted gearshift lever of a car with a standard transmission. In its upright position, the top handle of the joystick was 30 cm above the floor. At the end of its excursion, the joystick made contact with a microswitch and produced a flash of light on a panel under the stimulus screen. In order to make a second response, it was necessary to return to the neutral position. Thus, as is the case with microswitches housed in pigeon keys and rat bars (*cf.* Ferster and Skinner, 1957), it was necessary to activate and release the microswitch operated by the joystick before a second response could be detected. Only the top 9 cm of the joystick could be seen by the subjects. The bottom portion of the joystick, including the areas that struck the microswitch and a transducer, were hidden from view in a wooden box.

In the neutral position, the joystick rested against a rigid mechanical stop that prevented it from travelling in the opposite direction. Unknown to the subject, the mechanical stop was the external edge of the piston of a force transducer. The transducer (Statham Laboratory model #104-D8a-5a-120) produced a voltage proportional to the force applied to it. The relationship between applied force and the resultant voltage was linear over a range from zero to five pounds. This arrangement allowed the experimenter to measure any tendency to push the joystick in a direction opposite to that required to operate the microswitch. This tendency will be referred to as antagonistic responding.

The output of the transducer was fed into a voltage-to-frequency converter that produced a digital response for each 1.4 pounds (6.2 N) of force applied to the transducer. The arbitrary choice of one antagonistic response per 1.4 pounds of force was based upon the results of a pilot experiment. This value provided similar ranges of microswitch and antagonistic responding during S—. While the microswitch and the antagonistic responses differ slightly with respect to force requirement and feedback (none was provided for antagonistic responding), there appears to be no *a priori* reason to compare their frequencies directly.

Two groups of four subjects were each trained to discriminate two types of squares composed of random arrays of white dots (on a black background) generated by a computer.² These squares were generated as follows. The probability of the dot appearing in each cell in the outer portion of a square matrix was 0.5. S+ refers to those squares in which the probability of a dot appearing in a cell in the inner portion of the matrix was 0.12. S— refers to those squares in which the value of this probability was 0.88. As shown in Figure 1, the inner matrix of S+ was darker than the inner matrix of S—. The discriminative stimuli shown in this figure are randomly chosen examples of two series of 160 S+s and 160 S—s.

The experiment consisted of three phases: during five sessions of non-differential training, only S+ was presented. During seven subsequent sessions of discrimination training, S+ alternated irregularly with S—. The last phase of training consisted of three non-differential sessions. Sessions were conducted on consecutive days. The scheduled duration of each presentation of S+ was 30 sec, after which the screen was darkened for 2 sec. Responding to S+ was reinforced following intervals of various length whose average was 30 sec (VI 30-sec). Reinforcement was indicated by a brief flash of a light, and by incrementing a large digital counter located under the stimulus screen. Each session was terminated after the subject had earned 30 reinforcers (\$3.00). Non-differential (VI) sessions were approximately 15 min long. The duration of discrimi-

²The stimuli were prepared by Dr. B. Green of the Department of Psychology, Johns Hopkins University, Baltimore, Maryland, and used with his kind permission.

DISCRIMINATION LEARNING: TRADITIONAL TRAINING

Fig. 1. Samples of S+ and S- used during traditional training procedure. See text for additional details.

nation (VI EXT) sessions was at least 30 min. As described below, the duration of the discrimination sessions was in part determined by the subject's behavior during S-. During each session, the only source of illumination in the experimental room was provided by the stimulus and by the apparatus displaying how many reinforcements had been earned.

Before the first discrimination session began, each subject was informed of a procedural change during that and the six subsequent sessions. They were told that these sessions would be longer and that they would see a new type of square whose significance they would have to determine for themselves. The instructions read to each subject were as follows:

As you know, today's session will be longer than the earlier sessions. Starting with today's session you may see some new squares. You may also note that on some occasions it will not be possible to earn dimes. Which squares these are should become clear during the course of the experiment.

REMEMBER: Keep your hand on the lever at all times, even when not responding.

Throughout discrimination training, responding to S+ was reinforced on a VI 30-sec schedule. Responding during S- was never reinforced. If responding occurred during the presentation of a particular S-, the duration of that trial was prolonged until 15 sec without a response had elapsed.

One group of four subjects was trained to discriminate S+ and S- by a traditional pro-

cedure whereby S- was introduced as shown in the upper portion of Figure 1. A second group was trained to form the same discrimination by a fading procedure. Representative examples of the stimuli used during the fading series are shown in Figure 2. At the start of discrimination training, S- was a homogeneous array of dots. That is, $P(\text{of a dot appearing in the edge of the stimulus}) = P(\text{of the dot appearing in the center of the stimulus}) = 0.5$. The initial duration of S- was 2 sec rather than 30 sec. During the course of successive presentations of S-, its duration was progressively prolonged until it reached 30 sec, the duration of S+. During the second phase of the fading procedure, the duration of S- was held constant at 5 sec, while the probability of a dot appearing in the center was progressively increased until it reached its terminal value of 0.88. During the third and final phase of the fading procedure, the duration of S- was gradually increased from 5 to 30 sec. All of the changes called for by the fading procedure were made during the first discrimination session. Aside from fading, all aspects of discrimination training coincided with those followed under the traditional procedure.

RESULTS AND DISCUSSION

The frequencies of microswitch and antagonistic responding of each subject during all phases of the experiment are shown in Figure 3 (traditional method) and Figure 4 (fading method). The functions showing frequency of responding to S- during traditional training

DISCRIMINATION LEARNING: FADING

Fig. 2. Samples of S+ and S- used during fading procedure. See text for additional details.

are quite similar to those obtained from a wide variety of experiments on discrimination training that used animal subjects. As can be seen in the solid function in the top half of Figure 3, there is an orderly decrease in the number of

errors during the course of discrimination training. The range of errors of the four subjects trained by the traditional procedure was 3892 to 12,914. The lower portion of Figure 3 shows that each subject exhibited behavioral

Fig. 3. Frequency of microswitch responding, during S+ and S-, and the frequency of antagonistic responding, during S+ and S-, of each subject trained by the traditional method.

Fig. 4. Frequency of microswitch responding, during S+ and S-, and the frequency of antagonistic responding, during S+ and S-, of each subject trained by the fading method.

contrast (*cf.* Reynolds, 1961; Terrace, 1966). For example, Subject R.B.'s average baseline rate of responding during the five pre-discrimination sessions was 87.8 responses per minute. During discrimination training, this subject's rate of responding increased to an average value of 137.8 responses per minute.

The main datum of interest in Figure 3 is the frequency of responding in a direction opposite to that required to operate the microswitch (antagonistic responding). During both series of non-differential reinforcement sessions, the frequency of antagonistic responding of subjects of both the traditional and the fading groups was zero or near zero. During discrimination training, however, each subject who made errors made many antagonistic responses (Figure 3, upper portion, dashed function).

In each case, the function showing the frequency of antagonistic responding lagged behind the function showing the frequency of microswitch responding. This indicates that, during S-, the microswitch was replaced by a response of opposite topography. As discrimination training continued, however, the fre-

quency of behavior antagonistic to the reinforced response gradually approached zero.

The data obtained from those subjects trained by the fading procedure (Figure 4) differed from the data obtained from those trained by the traditional method in three respects. These subjects made virtually no errors. Also absent in each case were behavioral contrast and responses antagonistic to the microswitch response.

A direct comparison of the number of errors and antagonistic response emitted by the subjects trained by the traditional and by the fading methods is not possible. Because of the consequences of the correction procedure, the actual durations of S- were considerably greater for subjects of the traditional group than they were for subjects of the fading group. The design of the present experiment does not allow one to compare the frequencies of responding to S- and of antagonistic responses during equivalent periods of time. The results shown in Figures 3 and 4 do, however, indicate that a sufficient condition for the occurrence of antagonistic responding is the occurrence of errors.

The amount of contrast observed in the group that learned with errors was a monotonic function of the amount of antagonistic responding. This can be seen in Figure 5, which shows the per cent increase in the rate of responding to S+ of each subject during discrimination training as a function of the number of errors and of the number of antagonistic responses. The per cent increase in the rate of responding to S+ was calculated by comparing each subject's average rate of responding during the five baseline (VI) sessions with his or her average rate of responding during the seven discrimination (VI EXT) sessions. Whereas the amount of contrast increases monotonically with the amount of antagonistic responding (Figure 5, left), no systematic relationship can be seen between the magnitude of the contrast effect and the number of errors that occurred during discrimination training (Figure 5, right). These data clarify the lack of relationship between contrast and the number of errors, a finding consistent with a similar observation based upon a comparison of more than 100 pairs of error and contrast functions obtained from experiments on pigeons in my laboratory (*cf.* Terrace, 1966). Contrast does not occur following discrimination learning without errors. Why then is the amount of contrast that does occur following discrimination learning with errors unrelated to the number of errors? The present results suggest that how forcefully a subject "holds back" in not making

the microswitch response during S- is a good predictor of his rate of responding when S+ is again presented.

The occurrence of antagonistic responding only after discrimination learning with errors is consistent with other differences in performance following discrimination learning with and without errors. It also appears that both the occurrence and the ultimate disappearance of antagonistic responding is related to the aversive function of S-. As Amsel (1958, 1962) suggested, non-reinforcement can result in a drive state that he has referred to appropriately as frustration. One way to avoid the aversive and frustrative effects of non-reinforcement is to engage in behavior specifically incompatible with the response that results in frustration. However, the occurrence of such antagonistic responses ensures that frustration will no longer occur. Accordingly, one would expect a reduction in the reinforcing value of an antagonistic response. This could account for the decline in the frequency of such responses. Support for this interpretation can be seen in the temporal separation between the error and the antagonistic response functions. As shown in Figure 3, antagonistic responding in each case declines *after* microswitch responses (errors) have declined.

In this experiment, the definition of an antagonistic response was seriously constrained by the nature of the apparatus. Only one antagonistic response was possible. Ultimately,

Fig. 5. Per cent increase in rate of responding to S+ during VI EXT training of each subject of the group trained by the traditional method as a function of the number of antagonistic responses; right panel, Per cent increase in rate of responding to S+ during VI EXT training of each subject of the group trained by the traditional method as a function of the microswitch responses during S-.

the concept of an "antagonistic" response appears to be best defined by its stereotypy and by its level of activity, rather than as behavior that is "opposite" to the reinforced response. The difficulty with referring to the opposite of a response is, of course, the absence of any universal standard for defining opposite. In the present experiment, the only antagonistic response that could be measured was one whose direction was 180° opposite to that of the reinforced response. An isometric stick, with strain gauges suitably mounted for measuring responding in a 360° field (rather than only along the Y axis, as in the present experiment) would provide a more meaningful measure of the stereotypy of responding than could be obtained with a joystick free to move in but two directions. Whereas one would not expect all antagonistic responses to occur in a direction 180° opposite to that which defines the positive response, one would require a high degree of stereotypy (in whatever direction) before referring to a response as antagonistic.

Examples of other responses that might be related antagonistically are a human subject moving a finger from a central position to the right, as opposed to moving it to the left, and a rat running clockwise in a running wheel, as opposed to running counter-clockwise. In these, as well as in other examples, there does not appear to be any *a priori* criterion of antagonism. It should also be recognized that the discovery of antagonistic responses of high stereotypy and high amplitude does not ensure that such responses are the *only* antagonistic responses occurring during S—. It may very well be the case that other antagonistic responses also occur, responses that the experimenter is not equipped to measure.

Despite the restricted range of the antagonistic response studied in the present experiment, the results clearly support two conclusions: (1) the tendency to make an antagonistic response differs following discrimination learning with and without errors, and (2) extinction is not a passive process. The present results also show that the extinction of responding to S—in discrimination learning with errors consists of two phases. The first is an active phase, during which a new response, incompatible with the reinforced response is conditioned. That such new behavior can be conditioned where the response in question cannot increase the frequency of S+ has been shown in this

author's study of escape from S— (Terrace, 1971). The active phase of extinction is followed by a passive phase, during which one can observe a wide variety of behavior bearing no systematic relationship to the reinforced response. During discrimination learning without errors, it appears as if non-responding is of the passive variety throughout discrimination training.

The present results may also prove relevant to formulations of conditioned inhibition, a concept that has had a somewhat checkered history in psychology (*cf.* Kimble, 1961, pp. 302-317; Wendt, 1936; Hearst, 1972). These results, along with those of a wide variety of other experiments in which non-responding is the datum of interest (*cf.* Laties *et al.*, 1965), indicate a need for distinguishing between active and passive inhibition. Active inhibition would appear to apply to those situations in which a particular stimulus occasions a specific response that is incompatible with the reinforced response. This usage of inhibition is analogous to the classical usage of inhibition by physiologists as formulated by Sherrington (1906). The frequency of one response is decreased by virtue of eliciting a second response that evokes a response incompatible with the first. Passive inhibition, on the other hand, would appear to apply to those situations where numerous stimuli, often those directly manipulated by the experimenter, occasion a wide variety of non-responding. The topography of passive non-responding bears no systematic relationship to that of the reinforced response. Given the dearth of empirical data on the nature of non-responding, the distinction between active and passive inhibition is admittedly speculative. The present results are consistent with this distinction. At the very least, they define an aspect of non-responding that falls outside the realm of contemporary accounts of conditioned inhibition.

REFERENCES

- Amsel, A. The role of frustrative non-reward in non-continuous reward situations. *Psychological Bulletin*, 1958, **55**, 102-119.
- Amsel, A. Frustrative non-reward in partial reinforcement and discrimination learning: Some recent history and a theoretical extension. *Psychological Review*, 1962, **69**, 306-328.

- Amsel, A. Inhibition and mediation in classical Pavlovian and instrumental conditioning. In R. A. Boakes and M. S. Halliday (Eds.), *Inhibition and learning*. London and New York: Academic Press, 1972. Pp. 275-299.
- Azrin, N., Hutchinson, R. R., and Hake, D. F. Extinction-induced aggression. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 191-204.
- Brogden, W. J., Lipman, E. A., and Culler, E. The role of incentive in conditioning and extinction. *American Journal of Psychology*, 1938, 51, 109-117.
- Ferster, C. and Skinner, B. F. *Schedules of reinforcement*. New York: Appleton-Century-Crofts, 1957.
- Hearst, E. Some persistent problems in the analysis of conditioned inhibition. In R. A. Boakes and M. S. Halliday (Eds.), *Inhibition and learning*. New York: Academic Press, 1972. Pp. 5-39.
- Kimble, G. A. *Hilgard and Marquis' conditioning and learning*. New York: Appleton-Century-Crofts, Inc., 1961. Pp. 302-317.
- Latties, V. G., Weiss, B., Clark, R. L., and Reynolds, M. D. Overt "mediating" behavior during temporally spaced responding. *Journal of the Experimental Analysis of Behavior*, 1965, 8, 107-116.
- Pavlov, I. P. *Conditioned reflexes*. (Translated by G. V. Anrep) London: Oxford University Press, 1927.
- Reynolds, G. S. Behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1961, 4, 57-71.
- Sherrington, C. S. *The integrative action of the nervous system*. New York: Scribner & Sons, Inc., 1906.
- Terrace, H. S. Discrimination learning with and without errors. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 1-27.
- Terrace, H. S. Stimulus control. In W. K. Honig (Ed.), *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. Pp. 271-344.
- Terrace, H. S. Escape from S-. *Learning and Motivation*, 1971, 2, 148-163.
- Terrace, H. S. By-products of discrimination learning. In G. Bower and J. Spence (Eds.), *The psychology of learning and motivation*. New York: Academic Press, 1972. Pp. 195-265. (a)
- Terrace, H. S. Inhibition and discrimination. In R. A. Boakes and M. S. Halliday (Eds.), *Inhibition and learning*. New York: Academic Press, 1972. Pp. 99-119. (b)
- Wagner, A. R. and Rescorla, R. A. Inhibition in Pavlovian conditioning: application of a theory. In R. A. Boakes and M. S. Halliday (Eds.), *Inhibition and learning*. New York: Academic Press, 1972. Pp. 301-336.
- Wendt, G. R. An interpretation of inhibition of conditioned reflexes as competition between reaction systems, *Psychological Review*, 1936, 43, 258-281.

Received 26 October 1973.

(Final Acceptance 17 January 1974.)