JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1970, 14, 373-409 NUMBER 3 (NOVEMBER) PART 2

INHIBITION AND THE STIMULUS CONTROL OF OPERANT BEHAVIOR[1]

ELIOT HEARST, SERENA BESLEY,[2] AND G. WILLIAM FARTHING[3]

UNIVERSITY OF MISSOURI, COLUMBIA

A variety of methods, definitions, and theoretical notions that have been used in the study of inhibitory stimulus control were reviewed and evaluated. Preliminary data from several new operant methods were also described. It was proposed that future workers distinguish clearly between two forms of inhibitory control: (a) the learned power of a *specific stimulus* to reduce behavior, and (b) a *dimensional effect,* in which responding increases as values progressively more distant from the value of that specific stimulus along some dimension are presented (generalization gradient). Conclusions from several important recent studies were shown to be strongly dependent on the individual experimenter's criterion for deciding when a stimulus is inhibitory. The concept of inhibition seems a very valuable one for the field of operant behavior, and it deserves more attention than it has received in the past.

CONTENTS

I. INTRODUCTION

Although the concept of inhibition has assumed an important role in several major behavior theories (Pavlov, 1927; Hull, 1943, 1952; Konorski, 1967; see also Diamond, Balvin, and Diamond, 1963), the experimental analysis of inhibitory phenomena has been neglected by Western psychologists interested in instrumental learning. This neglect may be partially due to a general bias against the use of concepts that have been associated historically with extensive neurophysiological speculation. For example, Pavlov's and Konorski's elaborate neural models of learning stressed the action of unobservable inhibitory processes in the brain, an emphasis that automatically elicits a generalized negative reaction from many American behaviorists. In addition, workers in the field of operant conditioning have been influenced by Skinner's (1938) criticism of the experimental basis and logical status of the concept of inhibition. He argued that the use of the concept is unparsimonious, since inhibition usually refers to nothing more than a reduction in "excitation". However, the greatest

---

[1] We thank Carol Bergfeld, John Flanagan, Dennis Hilker, James McLean, Richard Millar, James Noah, Leigh Shallenberger, Clem Small, Sharron Taus, and especially Minnie Koresko for valuable advice and assistance. The research connected with this article was supported mainly by NIMH Grant MH 12120. Support for the first author was also provided by Biomedical Sciences Support Grant FR-07053 from NIH; by the Department of Psychology and Institute of Human Learning, University of California (Berkeley) during the 1968-69 academic year; and by extended loan of equipment from the NIMH Clinical Neuropharmacology Research Center, St. Elizabeths Hospital, Washington 20, D. C., to the University of Missouri. The many constructive comments of colleagues and students on earlier versions of this article are greatly appreciated. Reprints may be obtained from Eliot Hearst, Dept. of Psychology, Indiana University, Bloomington, Ind. 47401.

[2] NDEA Predoctoral Fellow, 1967-70, at the University of Missouri.

[3] USPHS Predoctoral Fellow (MH 37445) at the University of Missouri during 1967-69. Now at the University of Maine, Orono, Me.

obstacle to the study of inhibitory phenomena has probably been the lack of objective, unambiguous, and relatively direct methods for their detection and measurement.

In this article we intend to review critically the methods that have been used in the past for measuring inhibitory effects of stimuli controlling operant behavior. What experimental conditions must be met for a specific stimulus to be called an inhibitor of behavior? When can we conclude that a given dimension or feature of that stimulus does or does not exercise inhibitory control over a response? An answer to the latter question usually involves the measurement of so-called generalization gradients of inhibition, in which responding along a stimulus dimension is examined as a function of the physical difference between various test stimulus values and a stimulus value at which training that is presumed to be inhibitory has been given. If inhibitory dimensional control is actually present, then responding ought to be minimal at or near the training stimulus and maximal at values far from the training stimulus. This kind of relationship would provide the negative counterpart of the extensively investigated gradient of excitation or reinforcement, in which response strength is inversely related to the distance of a test stimulus from the reinforcement-correlated training stimulus.

Another aim of this article is to introduce several new methods for producing inhibitory gradients that have recently been developed and tested in our laboratory. Past findings by other researchers and our current results with these new methods have convinced us that the concept of inhibition is a very valuable one, particularly as it may be applied to understanding the effects of negative stimuli in animal discrimination learning (Hearst, 1969). Throughout this article, our discussion is confined to stimuli that develop inhibitory properties *as a result of conditioning*. Nonassociative effects (*e.g.*, operant phenomena analogous to Pavlov's "external inhibition", in which the presentation of a novel exteroceptive stimulus weakens ongoing behavior) will not be included, except insofar as they provide valuable control baselines against which to evaluate the effects of stimuli that are presumed to have acquired an inhibitory function.

Furthermore, this monograph focuses on inhibitory effects produced by the presentation of specific external stimuli. It does not cover phenomena that are characterized by certain relatively long-term or progressive temporal changes, such as spontaneous recovery, intrasession fatigue or satiation effects, or the decline in responding during simple extinction. Although certain theorists might classify them as inhibitory effects and discuss them in terms of the buildup or dissipation of inhibition, these phenomena do not normally involve the experimental presentation of discrete external stimuli. For these phenomena, the evaluation of inhibitory effects, free of obvious alternative explanations in terms of the reduction of excitatory effects, is difficult, if not impossible, to achieve. Skinner's (1938) discussion of the general difficulties (see also Gleitman, Nachmias, and Neisser, 1954) remains cogent today.

Many methods to be described are similar to ones used by Pavlov and Konorski in their studies of positive (CS+) and negative (CS−) conditioned stimuli[4] in classical conditioning. Analogous techniques have rarely been employed in studies of operant behavior. Recent elegant demonstrations (Rescorla and Solomon, 1967; Rescorla, 1969*a*, 1969*b*) of the inhibitory effects of stimuli associated with nonreinforcement in a Pavlovian situation suggest that similar effects should exist in operant situations. Isolation and quantitative measurement of these and allied phenomena may provide a good empirical basis for reevaluation of discrimination-learning theories like those of Spence and Hull (Spence, 1936, 1937; Hull, 1943, 1952; see also Jensen, 1961), which posit an interaction of generalization gradients of excitation and inhibition around S+ and S− respectively[4] to account for transposition (Spence, 1937), stimulus intensity dynamism (Perkins, 1953; Logan, 1954), and a variety of other phenomena in operant discrimination learning (Riley, 1968; Hearst, 1969; Hoffman, 1969). Since deficits in "inhibitory processes" are frequently said to characterize the effects of certain brain lesions (*e.g.*, Douglas, 1967; Kimble, 1968) or of certain pharmacological agents (*e.g.*, Carlton, 1963), several of the following methods may also have applied value

---

[4] In conformity with common usage, we will employ the abbreviations CS+ and CS− to refer respectively to positive (paired with reinforcement) and negative (paired with no reinforcement) stimuli in classical conditioning, and S+ and S− to refer to their operant counterparts.

in neuropsychology for measuring lesion-induced or drug-induced changes in the learned inhibitory effects of various stimuli.

II. THE INHIBITORY FUNCTION OF A STIMULUS

Strong differences of opinion have recently been expressed (Terrace, 1966*a*, 1966*b*, 1967; Biederman, 1967; Deutsch, 1967; Bernheim, 1968; Farthing and Hearst, 1968; Bloomfield, 1969; Lyons, 1969*a*, 1969*b*) about findings that purport to demonstrate the inhibitory function of a stimulus associated with nonreinforcement in operant discrimination learning. These disagreements seem to have emerged because at least two different definitions of "inhibitory function" are in current use. Following arguments advanced by Jenkins (1965), Terrace (1966*b*, 1967) suggested that an inhibitory function be assigned to an S— when response strength increases as stimulus values progressively more distant from S— along some *dimension* are presented. An empirical generalization gradient of a certain kind thus serves as Terrace's operational definition of the inhibitory function of a stimulus: "A U-shaped gradient, with a minimum at S—, would indicate that S— was an inhibitory stimulus, while a flat gradient would indicate the absence of any inhibitory function (1966*b*, p. 1678)."

It is noteworthy that in his definition of an inhibitory stimulus Terrace does not imply or include the possible property of this specific stimulus to *reduce* responding that would otherwise have occurred. A stimulus may have this property regardless of whether manipulation of a certain dimension of the stimulus yields a U-shaped gradient. For example, the presentation of a tone S— may reduce the usual rate of operant behavior to a visual S+ that had never before been presented simultaneously with the tone; however, the tone would not qualify as an inhibitory stimulus in Terrace's sense unless one could show that variation of the intensity or frequency of the tone produced a gradient with a minimum at the S— intensity or frequency. Thus, for Terrace, response reduction produced by the addition of a specific stimulus would apparently be an insufficient reason for characterizing a stimulus as inhibitory. Evidence of differential control along at least one dimension of S— seems the indispensable aspect of Terrace's definition.

Another type of definition is different from Terrace's. Hull wrote that "inhibitory potential can be observed only indirectly through the failure to occur of some positive reactions which the antecedent conditions would otherwise produce (1943, p. 289)." Rescorla accepts as the defining attribute of a conditioned inhibitor its control of a "response tendency opposite to conditioned excitation; the question of its other properties is to be decided empirically (1969*a*, p. 67)." In these statements, as in most of the discussions of the concept of inhibition offered by Pavlov (1927) and Konorski (1967), Terrace's idea of differential control along a specific dimension of S— is not included and is clearly not regarded as necessary to the definition of an inhibitory stimulus. The critical point in this second kind of definition is the notion of a directional effect (usually measured as a decrement in behavior) that acts in opposition to the normal level of response strength prevailing under the experimental conditions.

Jenkins' comment that inhibition may be identified with the "development of a response to S— incompatible with the reinforced response (1965, p. 59)" also suggests that a habit is learned that actively opposes the one produced by reinforcement in the situation. Terrace's idea of differential control along some stimulus dimension does not appear, therefore, to be crucial to Jenkins' definition of an inhibitory stimulus, although he definitely did stress that notion in his 1965 article. Jenkins' meaning of the term was clarified and developed further in an interesting study (Brown and Jenkins, 1967) that examined an operant analog to the specific phenomenon Pavlov labelled "conditioned inhibition". Pecking by pigeons on one side (*e.g.*, the left) of a split key was reinforced when one color (A) was present on both halves of the key, and on the other side when a different color (B) was present on both halves. Then a second discrimination was trained involving the presence or absence of a tone. Only Color A was presented on both halves of the key during this phase and correct (left-side) responses to that color were intermittently reinforced when the tone was absent; no responses were reinforced when the tone was present.

In a subsequent phase, Brown and Jenkins made sure that left-side responses still occurred to Color A and right-side responses to Color B, and that left-side responding was greatly reduced during combined presentation of Color A and the tone. Then test trials were given during extinction, in which Color B and the tone were presented together for the first time (interspersed with trials on other compound and component stimuli). The presence of the tone greatly reduced the rate of right-side responses to Color B, but differential responding to the colors was maintained during the test. Control birds, which received tone presentations for the very first time during an extinction test, revealed no detectable effect of the tone. Thus, the reduction of appropriate behavior by the tone in the experimental group was probably not due to some unconditioned effect (akin to Pavlovian "external inhibition").

Brown and Jenkins concluded that the decremental effects of an S− transfer to a different S+. They argued that their experimental design answers many of Skinner's (1938) specific objections to Pavlov's earlier demonstrations of conditioned inhibition in a salivary conditioning paradigm. According to Brown and Jenkins: "the empirical meaning of inhibitory stimulus control is that the presentation of a part or aspect of the nonreinforced stimulus which is *separable* from the positive stimulus produces a decrement in responding (p. 256)." Thus, these experimenters emphasize the production of behavioral decrement as a crucial attribute of an inhibitory stimulus. An auditory gradient with a minimum at or near the tonal S− was never obtained in their experiment. Differential control along a dimension of S− is therefore not necessary to their definition.

Conclusions from specific experiments could certainly be affected by which of these two general types of definition is used. For example, Deutsch (1967) and Terrace (1967) reached virtually opposite conclusions from Terrace's report (1966*b*) that (a) generalization gradients along the S− (wavelength) dimension following discrimination-learning-without-errors (*i.e.*, no responses occur to S− throughout training) are completely flat and are characterized by near-zero responding all along the entire dimension, and (b) gradients similarly obtained following learning-with-errors have minima at S− and well-above-zero responding at values far from S−. Employing a definition based on the presence of differential responding along the S− dimension, Terrace concluded that S− does not acquire an inhibitory function during errorless learning, but does act as an inhibitory stimulus during learning with errors. Employing a definition of inhibition based on the absolute amount of responding along the S− dimension, Deutsch argued that errorless learning, if anything, creates greater inhibition of responding along the S− dimension than does learning with errors. As we will discuss later, there are procedural deficiencies in Terrace's experiment that seem to make both his and Deutsch's conclusion questionable, but it is clear that a large part of the argument between these two investigators arose because of their different criteria for deciding when a stimulus is inhibitory.

We find it useful to keep these definitional problems in mind while evaluating the different empirical methods that have been used or proposed for determining so-called generalization gradients of inhibition. We recommend the reader do the same while these methods are reviewed and evaluated in later sections of this article. Our own tentative definition of an *inhibitory stimulus* in operant conditioning (see also Farthing and Hearst, 1968) distinguishes it from *inhibitory dimensional control,* in analogous fashion to the manner in which one might distinguish between an *excitatory stimulus* and *excitatory dimensional control.* An "excitatory stimulus" is a stimulus (*i.e.*, some physically defined environmental event, presumably of a multidimensional nature) that develops during conditioning (*i.e.*, its repeated presentation in conjunction with certain response-reinforcement dependencies) the capacity to increase response strength above the level occurring when that stimulus is absent. An "inhibitory stimulus" is a stimulus that develops during conditioning the capacity to decrease response strength below the level occurring when that stimulus is absent.

The mere observation of a behavioral decrement would not necessarily mean that inhibitory effects are involved. In our opinion (see also Catania, 1969, especially pp. 741-742, for a somewhat analogous argument), one can convincingly talk about inhibitory effects only when all the conditions that maintain a response (*e.g.*, reinforcement schedule, drive, amount and delay of reinforcement, the presence of S+) are held constant, and then some stimulus is *presented* that leads to much less behavior than would normally occur under those conditions. Decremental effects that are a result of changes in the maintaining variables themselves (*e.g.*, removal of reinforcement during S+, lowered drive, withdrawal or variation of S+) seem more parsimoniously described in terms of manipulations of only one type of factor (excitatory). This argument closely resembles the one Skinner made more than 30 years ago when he listed his objections to conventional uses of the concept of inhibition; the term had been and still is frequently applied to behavioral decrements of all kinds, for example, those brought about by simple extinction, drive manipulations, changes in S+, *etc.* He did note, however, that the concept might prove useful when applied to the case where a second (extraneous) stimulus is added to a behavioral situation: "if it were not for the inhibiting stimulus, a certain amount of activity would be observed (1938, p. 96)" and "The strength of a reflex may be decreased through presentation of a second stimulus which has no other relation to the effector involved. . . . Inhibition refers to a negative change in strength produced by a kind of operation that would yield a positive change under other circumstances (1938, pp. 17-18)."

Therefore, like Rescorla's (1969*a,* 1969*b*) analysis of Pavlovian conditioned stimuli, we conceive inhibitory stimuli in operant conditioning to be controlling a tendency separable from and opposite to that controlled by excitatory stimuli. As in Pavlovian conditioning, major methodological problems arise because the response-reducing capacity of a stimulus often proves difficult to detect, especially when the baseline behavior is at or near zero. Many of the methods to be described in this monograph represent attempts to produce a high enough output of behavior so that one can distinguish inhibitory (reductive) effects from no effect at all.

The term "excitatory dimensional control", in our view, would be applied when new stimulus values that lie at progressively greater distances along a specific dimension from an excitatory stimulus show a graded decremental effect. The terms "inhibitory dimensional control" would be applied when new stimulus values at progressively greater distances from an inhibitory stimulus show a graded incremental effect on the strength of an operant response. It is important to point out that an incremental gradient around some stimulus value is *necessary but not sufficient* for defining inhibitory dimensional control. The specific stimulus at which responding is minimal must also be shown to be inhibitory by some independent test, since it is logically possible that such a stimulus is relatively "neutral" and the other values progressively more excitatory.

Jenkins (1965) implied this problem of defining neutrality in his discussion of two categories of "nonresponding", as have Wilton and Godbout (1970), but an example may be useful here. Suppose that periods of nonreinforcement in the presence of a 1000-Hz tone, interspersed with periods of reinforcement in the presence of a clicking noise, merely convert the tone into a neutral stimulus, in the sense that presentation of a 1000-Hz tone would not increase or decrease behavior from the level occurring in its absence. A generalization test along the tonal frequency dimension might reveal a minimum at 1000 Hz not because that stimulus value was "inhibitory" but merely because the initial excitatory generalization from the clicking noise to the 1000-Hz tone had been reduced by nonreinforcement at that value and not specifically reduced at other Hz values. In order to determine whether the 1000-Hz tone were neutral or inhibitory, some *active* test of its inhibitory properties would have to be arranged, such as (a) pairing it with a new S+ to see if it reduced response strength to that S+ much more than would a 1000-Hz tone in other groups for which, *e.g.*, the tone was novel, or was present during past training but completely uncorrelated with reinforcement and nonreinforcement periods, or (b) determining whether it would prove harder to convert into an S+ than would the novel or previously uncorrelated 1000-Hz tone. Section IV of this article (as well as Konorski [1967], Bignami [1968], and Rescorla [1969b]) presents a brief review of these and several other possible techniques for determining the active inhibitory properties of a specific stimulus and for establishing control groups in which the same stimulus is presumably "neutral".

By our set of definitions, a *particular external stimulus* may be said to be either excitatory (its presentation increases certain measures of operant behavior) or inhibitory (its presentation decreases certain measures of operant behavior) even when (a) no test of dimensional control has been performed, or (b) actual generalization tests reveal horizontal (flat) gradients along dimension(s) of S+ or S− respectively. We would conclude that a certain *dimension* exercised excitatory or inhibitory control only when corresponding decremental or incremental gradients are observed along that dimension, and independent tests show that the specific S+ or S− is actually excitatory or inhibitory relative to some baseline of "neutrality".

This distinction between behavioral control by the presence *vs.* absence of a specific stimulus and control by different values along a specific dimension of that stimulus seems valuable to us because it emphasizes the fact that there is no inevitable relation between the two types of control. Some examples of this possibility have already been given, but another may serve to clarify our argument. Suppose the presentation of a particular stimulus can be shown (via one of the above "active" tests) to reduce operant behavior, *i.e.,* it meets our definition of an inhibitory stimulus, but the experimenter fails to choose a relevant dimension to vary during the determining of generalization gradients. Presentations of a bright, tilted, 550-nm line on a pigeon's response key might meet our definition of an inhibitory stimulus, but if the experimenter varied the intensity or the tilt of the line he might obtain relatively flat gradients. On the other hand, if he varied the wavelength of the line, he might obtain a steep incremental gradient around 550 nm (see, *e.g.,* Newman and Baron, 1965, for an example of strong control by different values along one dimension of a training stimulus and weak control by values along another dimension). The absence of differential responding along one or several dimensions of a stimulus does not at all preclude the possibility that the presence *vs.* absence of the specific stimulus exercises strong excitatory or inhibitory control. In our opinion, this is the reason why the term "stimulus control", very popular now in operant research, is a somewhat confusing concept. One of its proposed meanings (Terrace, 1966*a,* p. 271) stresses control along a dimension ("with respect to the continuum being studied"), and yet many workers use the term mainly to refer to control by the presence *vs.* absence of a specific stimulus. Herrnstein (1966, p. 40) seems to be using the term in the latter way, without reference to dimensional effects, when he states that the "rule of stimulus control" involves the tendency for behavior to be governed by the stimulus in the presence of which it is occurring. Hendry (1969, p. 401) states even more strongly that stimulus control refers to the "degree to which the probability of an operant varies with the presence/absence of a discriminative stimulus." One of the main points of our monograph is the argument that there is value in making a clear distinction between these two meanings of "stimulus control".

Because recent studies concerning the inhibitory effects of stimuli on operant behavior have stressed methods for analyzing dimensional control (generalization gradients) rather than methods for detecting the response-reductive properties of specific stimuli, we will focus on the former area of research in the following review. However, we hope to indicate how experimenters might profitably combine study of both types of effect in future work. Almost all the methods that yield so-called inhibitory gradients around S− do not unambiguously reveal whether S− is inhibitory according to our definition. Several methods for distinguishing among ineffective, irrelevant, and inhibitory stimuli are now being increasingly used in the field of classical conditioning (Rescorla, 1969*a,* 1969*b;* Hammond, 1967, 1968) and ought, in our opinion, to be applied more frequently in the study of operant behavior. Conversely, it would be worthwhile if researchers in classical conditioning were to do more work on the problem of inhibitory dimensional control; for example, the empirical determination of generalization gradients around CS− in the classical conditioning of animals has been very rare.

III. GENERALIZATION GRADIENTS

In past studies of operant behavior, several different experimental methods have been used to assess dimensional control by presumed inhibitory stimuli. The development of new methods and the standardization of old ones ought not only to be valuable for the relatively atheoretical researcher, but should also have theoretical utility. For example, although gradient-interaction theories (*e.g.,* Spence, 1936, 1937; Hull, 1943, 1952) imply that excitatory and inhibitory gradients of a certain shape and slope are established around S+ and S− during discrimination learning, specific tests of these ideas have been based mainly on fragmentary empirical results from studies with large groups of subjects or on more or less arbitrarily constructed hypothetical gradients. Evaluation of theories of this kind would benefit from the development of relatively direct, reliable, and unambiguous empirical procedures for determining inhibitory gradients (IGs)[5] in individual subjects or in small groups of subjects, such as have already been widely developed for studying excitatory gradients (EGs). (See, for example, numerous experiments reported in Mostofsky, 1965). More than 30 years ago Spence implied the need for such techniques and data when he stated: "The selection of the curves of generalization has been more or less arbitrary, as little evidence bearing on the problem is available (1937, p. 434)."

All of the more direct methods for obtaining IGs are based on essentially the same general plan. Equal responding is experimentally established or assumed to exist along a particular stimulus dimension, and presumed inhibitory training (usually extinction of an operant response) is given at one value along that dimension. Then a generalization test follows, which includes several dimensional values differing in their physical distance from the presumed inhibitory stimulus. Every one of these methods can be conceptualized as involving an algebraic summation of excitatory and inhibitory effects. Since excitatory effects are prearranged or presumed to be equal all along the dimension, any obtained gradient with a minimum at or near the training stimulus must, according to this argument, be due to differences in the generalization of inhibitory effects at each point along the dimension. Although the logic of this argument can be criticized for its inability to distinguish between neutral stimuli and inhibitory stimuli, the techniques that have developed from it are very valuable to any experimenter in this area of research. Perhaps the modifications or amplifications of these techniques to include "active" tests for the inhibitory properties of a specific stimulus, as will be suggested at various points in this article and summarized in Section IV, will make them even more valuable.

An appreciable positive level of behavior, usually produced by some reinforcement operation, must be established in order to enable detection of inhibitory effects. That is, an above-zero level of responding obviously has to be present to detect the response decrements caused by presentation of an inhibitory stimulus or to observe the systematic reductions in behavior that characterize inhibitory dimensional control. In the latter case, if values far from the value of the presumed inhibitory stimulus themselves produce zero responding there is no way to measure less responding at the presumed inhibitory stimulus itself.

This requirement typically makes the empirical determination of inhibitory effects more difficult than is the determination of excitatory effects, and often necessitates special pretraining procedures. Special pretraining is not usually necessary when the presence or generalization of excitatory effects is being studied, because in the excitatory case, positive training, involving increases in response strength, can be evaluated against a baseline of no responding or very low operant level. The conditioned excitatory function of a particular stimulus is demonstrated when a response (a) occurs more frequently in the presence of that stimulus than it did before making the stimulus a signal for reinforcement and (b) occurs much less frequently in the absence of that stimulus than in its presence. Excitatory dimensional control is demonstrated by a generalization gradient along some dimension, with maximal responding at or near the excitatory stimulus and less and less responding elsewhere. Although they could conceivably occur, "ceiling effects" have rarely created practical or interpretive problems in past studies of EGs.

Although this point regarding the need for an above-zero baseline, in order to avoid a "floor effect" in studying inhibitory stimuli and inhibitory dimensional control, may appear almost painfully obvious (see also Rescorla's comments on this problem: 1967, 1969*a*, 1969*b*), it has occasionally been overlooked; we will return later to its role in the controversy between Deutsch and Terrace, for example. Moreover, since differences in pre-

---

[5] Throughout the following discussion we will use this abbreviation in a rather loose and theoretically neutral manner, primarily as a convenience in referring to empirical gradients that various experimenters have labelled "inhibitory". We do not mean to suggest that all such gradients actually demonstrate *inhibitory dimensional control*, according to our definition; in fact, one of the points of this monograph is that mere evidence of a nonhorizontal incremental gradient around S— does not unequivocably show that inhibitory control is involved.

training procedures and response baselines inevitably exist when comparing EGs and IGs, experimenters who are interested in a meaningful comparison of the two must search for some way of minimizing these differences.

Each of the methods to be described for producing IGs has inherent strengths and weaknesses. Because of the deficiencies, future workers will probably want to obtain and compare gradients via several different methods in order to avoid misleading conclusions based on a single method. Interpretive problems will undoubtedly arise in future work because different methods will sometimes yield results that are inconsistent with each other, *i.e.*, some will yield evidence of dimensional control in a particular situation, whereas others will not. The experimental and theoretical strategy to be followed in such cases will have to be worked out eventually. For now, however, all we will try to do is to evaluate some advantages and disadvantages of the various methods for obtaining IGs, and the comparability of each with standard methods for obtaining EGs. The inevitable differences between the procedures for obtaining EGs and IGs must be taken into account, we think, whenever one tries to assess their applicability for testing theoretical formulations that posit an interaction between excitatory factors and inhibitory factors during discrimination learning.

A. Intradimensional Methods

The first group of methods for determining IGs involves training with positive and negative stimuli that both lie on the dimension of the generalization test. We will label these techniques "intradimensional", in analogous fashion to Switalski, Lyons, and Thomas' (1966) classification of generalization test methods. Other methods, to be discussed later, involve interdimensional or intermodal training, in which positive and negative stimuli lie either on presumed orthogonal dimensions in the same modality (*e.g.*, tonal frequency *vs.* white noise) or come from two different modalities (*e.g.*, auditory *vs.* visual).

1. Equalization-by-Reinforcement

Until about 10 years ago, only one general method had been employed in experimental attempts to obtain IGs. The first phase of this type of design involves the uniform delivery of reinforcement at several points along a generalization dimension in order to equalize response strengths at all these values before inhibitory training. Following equalization of response strength at these values, one value (S—) is presented alone for a certain number of trials without reinforcement. Gradients around S— are then obtained by retesting response strength at all the original stimulus values, either during complete extinction or while responding is maintained by reinforcement at certain stimulus values. This method is essentially the one used by Pavlov and some of his colleagues (1927, Lecture IX), Bass and Hull (1934), and Hovland (1937) in classical conditioning situations, and by Kling (1952) in a discrete-trial operant study employing a group design.

In a free-operant analog to the above studies, Honig (1961) first trained pigeons to peck for intermittent grain reinforcement (50-sec variable-interval [VI] schedule) at a response key illuminated on different trials by one of 13 possible wavelengths of light (10 nm apart, from 510 to 630 nm). This phase was continued until responding to all 13 stimuli was approximately equal[6] (10 sessions). Then the key pecking of some birds was extinguished at one value (570 nm) for 40 min; for other birds, extinction at 570 nm lasted 80 min. Response output in individual birds was subsequently measured during randomized presentation of all the original training values in extinction, a testing procedure just like that extensively used by Guttman and his colleagues to produce excitatory gradients (Guttman, 1956, 1963). Honig obtained gradients with a clear minimum at 570 nm in eight of his 12 subjects. However, the gradients were "bowl-shaped", much shallower than typical excitatory gradients following intermittent reinforcement at an S+ of 570 nm.

---

[6] One has to consider the possibility that "stimulus superstitions" (Morse and Skinner, 1957; Kieffer, 1965; Herrnstein, 1966) may develop in this phase. *i.e.*, response rates to some specific wavelengths may be higher or lower than to the rest of the wavelengths, even though the scheduled frequency of reinforcement is equal for all wavelengths. Such an outcome is usually attributed to some manifestation of pre-experimental preferences or to accidental correlations between specific stimuli, response rates, and frequency of reinforcement. Fortunately, these effects are often transient and may present no important problems if the equalization phase is continued long enough.

Compared to most prior experiments, which employed this general method in a group design, Honig's procedure had the advantage of yielding within-subject gradients around a stimulus associated with extinction. It was mainly the use of intermittent reinforcement during training that permitted observation of these effects in individuals, since the great resistance to extinction established by intermittent reinforcement ensured that responding would persist during extended presentation of a variety of test stimuli without reinforcement. One of the problems in the earlier studies of Bass and Hull and Hovland with groups of human subjects had been the rapidity of extinction of the CR (galvanic skin response) following continuous reinforcement. Therefore, Bass and Hull alternated periods of reinforcement with periods of extinction and retesting. This feature of their procedure allowed a discrimination to be formed between CS− and reinforcement-correlated stimuli on the same dimension. Thus, there is no strong justification for attributing the obtained incremental "gradients of extinction" in Bass and Hull's study to the operation of purely inhibitory effects. Complex, changing interactions between excitatory generalization and inhibitory generalization along the same dimension were probably involved.

Honig included a valuable control group that was given no extinction to 570 nm before the generalization test. After the equalization-by-reinforcement phase, and without any previous extinction period, these subjects were tested during extinction at all 13 wavelength values. This group yielded flat gradients along the wavelength dimension, a finding which meant that the experimental groups' gradient with a minimum in the middle (570 nm) of the stimulus series was not created during the test itself, *i.e.*, through summation of separate gradients of extinction around each of the 13 nonreinforcement-correlated test stimuli.

Although IGs secured via the equalization-by-reinforcement method[7] do yield evidence of dimensional control, there are several serious methodological and interpretive problems inherent in the technique. First, this method for obtaining IGs is not particularly comparable to usual methods (*e.g.*, Guttman and Kalish, 1956) for obtaining EGs. No initial phase of equal reinforcement for responses to all the test stimuli is normally included in the procedure for determining EGs, and the inclusion of such a phase should weaken any subsequent tendency to respond differentially to the stimuli during generalization testing. As Honig (1969) and others have shown, pretraining with equal reinforcement to only two visual stimuli ("pseudodiscrimination") tends to flatten subsequent EGs, even when these gradients are obtained along visual dimensions different from those used during pseudodiscrimination training. Consequently, the equalization procedure, if applied to the excitatory case, should produce gradients that are flatter than typical EGs obtained without preliminary equalization. It may be justifiable to compare IGs obtained via the equalization method to EGs obtained via an analogous equalization procedure (see some attempts of this kind in Weisman and Palmer, 1969), but it does not seem warranted to compare them to the usual EGs obtained without preliminary equalization. The latter comparison should almost always reveal much flatter IGs than EGs (as Honig [1961] indicated), and such a difference seems better attributed to the presence or absence of a prior equalization phase than to any intrinsic differences in the scope of excitatory *vs.* inhibitory generalization.

Second, equalization training destroys the novelty of the test stimuli to be presented to subjects during subsequent generalization testing. The concept of stimulus generalization has usually been applied to the transfer of learned responses to new stimulus values, different from those present during original training; naturally, if equalization training is given, the test values can no longer be novel for the subjects in that situation. In contrast, subjects undergoing a generalization test following conventional excitatory training procedures (*i.e.*, without equalization, as in Guttman and Kalish's work) are presented with the test values for the first time in their experimental history. Although the role of novelty is not yet clear in accounting for generalization effects, one should bear in mind that the interpretation of differences between EGs and

---

[7] Honig (1966) has also used a very similar procedure to study the stimulus generalization of punishment along a line-tilt dimension. Following equalization of VI behavior at nine stimulus values, electric shock was administered for responses at one line tilt. Subsequently, recovery from punishment was examined under the original VI equalization procedure.

IGs may be complicated by prior procedures that retain or destroy the novelty of the test stimuli. To what extent would any obtained EG *vs.* IG differences be due to the differential novelty of the test stimuli, or to intrinsic differences between excitatory *vs.* inhibitory control, or to both?

Besides the above methodological defects, which seem to preclude meaningful comparisons of IGs with most previously-obtained EGs, the equalization method has other, more serious, weaknesses. Since the nonreinforcement-correlated stimulus value and all the reinforcement-correlated values lie along the same dimension, it is impossible to vary the distance between the nonreinforcement value and particular test values without also changing the distance between specific reinforcement values and these particular test values. Thus, IGs obtained by this procedure are probably a very complex mixture of excitatory and inhibitory dimensional control. Furthermore, a minimum at the nonreinforcement value in an obtained gradient could mean either that extinction had converted that value into a neutral stimulus, with the other test values still retaining varying degrees of excitatory power, or it could mean that the nonreinforcement value was actually inhibitory. Without some definition or baseline of neutrality, we are left with the problem Skinner posed in 1938: are these effects best classified as inhibitory or can they be more parsimoniously viewed as mere reductions in excitation at various points along the generalization dimension? When variations in some property of S— are accompanied by variations in the same property of S+, as intradimensional procedures inevitably involve, this question is particularly difficult to answer. As Honig remarked, the gradients of extinction he obtained by the equalization procedure may not be "primary generalization gradients at all, but result from complex interactions among differentially reduced gradients of acquisition (1961, p. 277)."

2. Maintained Gradients with Prolonged Differential Training

Another intradimensional method that has been used to determine generalization gradients around stimuli associated with extinction is more indirect than the equalization-by-reinforcement procedure and is even harder to analyze in terms of separate excitatory and inhibitory control. Because of its inconclusiveness in this respect, no extensive summary of prior research with this method will be given. We will merely describe some examples of its use and point out the particular problems it leaves unresolved.

Under this procedure, several values along a stimulus dimension are presented to the subject in every session, some of which are consistently associated with intermittent reinforcement and others consistently associated with nonreinforcement. Thus every session is a "test" session, and subjects that are maintained on this steady-state discrimination yield generalization gradients every time they are run. By means of between-subject or within-subject reversals or other manipulations of the reinforcement conditions at each stimulus value, the experimenter attempts to separate generalization effects attributable to reinforcement from those attributable to extinction.

For example, Reynolds (1961) first gave uniform VI reinforcement to pigeons for pecking at 10 different orientations of an isosceles triangle projected on the response key; each orientation was presented twice in every daily session. This initial phase is thus quite similar to the first stage of the equalization method described above. Then, birds were run for 42 sessions on a procedure in which intermittent reinforcement was available only at two adjacent values of triangle-orientation, whereas extiction was in force during the other eight stimulus values. These conditions yielded a gradient with maximal responding in the vicinity of the two S+'s and with generally less responding as the tilt of the triangle was varied further and further from the S+ values. Then, the experimental conditions were reversed, and for 43 sessions subjects were extinguished at the two former reinforcement values with reinforcement at the other eight. On the new procedure, a gradient was obtained with minimal responding in the vicinity of the two extinction values and with generally more responding as the tilt of the triangle was varied further and further from these values.

Reynolds called the first gradient the "reinforcement gradient" and the second gradient the "extinction gradient". Measures of total response output differed between the two gradients, but the shapes of the two gradients were virtually the inverse of each other. Reynolds was particularly interested in the differences in total responding because of their implications for the phenomenon of behavioral contrast, but he did suggest that gradients obtained this way might be used to test certain assumptions and predictions of the Spence-Hull model for discrimination learning.

Other studies have used analogous intradimensional procedures for studying possible interactions between excitatory and inhibitory effects (*e.g.,* along an intensity continuum in Gray, 1965, p. 185, or James and Mostoway, 1969; along a spatial dimension in Catania and Gill, 1964; along a wavelength continuum in Blough, 1969). In these studies, as in that of Reynolds, reinforcement was maintained at certain dimensional values and extinction at other values. An abundance of interesting data, important for our understanding of discrimination learning and behavioral contrast, has resulted from procedures of this kind, but insofar as any clear separation of excitatory and inhibitory generalization is concerned, the method leaves much to be desired. As Catania and Gill pointed out, their most interesting findings could reflect either an inhibitory effect of reinforcement or an excitatory effect of extinction.

Likewise, Reynolds' experiment does not allow us to determine the extent to which the "reinforcement gradient" was due to reinforcement at two values or to extinction at the eight others. Additional data, *e.g.,* gradients obtained after reinforcement at only one or two values without any extinction values, or gradients obtained after extinction at only one or two values instead of eight, might possibly be revealing in this respect. However, such methods of obtaining EGs and IGs (see also Hanson, 1959; Jenkins and Harrison, 1962; Kalish and Haber, 1963; Hearst, 1969) are necessarily very indirect, since they require some kind of mathematical manipulation (*e.g.,* addition or subtraction) of two or more gradients to determine the EGs and IGs; and how does an experimenter decide which numerical transformation of the response data is the most appropriate one? The maintained-gradient procedure, like the equalization procedure, suffers from the fact that S+'s and S−'s along the same dimension cannot be independently varied in terms of their distances from other values along the dimension.

On the equalization procedure, the period of massed extinction at one value was definitely responsible for the minimum at that value in subsequently obtained gradients. This conclusion is warranted because responding to all stimulus values had been demonstrably equal just before the period of massed extinction, and no further reinforcements were given at any value. In the maintained-gradient procedure, however, even if initial equalization of responding is arranged (as in Reynolds, 1961), the later stages that involve (a) reinforcement at one or several values along the dimension, and (b) interspersed extinction at other values, make it impossible to separate the effects of reinforcement from the effects of extinction.

Another problem, which applies to both the maintained-gradient method and the equalization procedure, is created by the use of test stimuli that are not novel for the subjects. Maintained gradients cannot easily be compared to gradients obtained in conventional generalization tests with novel stimuli. In fact, some experimenters prefer to call the former "discrimination gradients"; they would reserve the term "generalization gradient" for those situations in which subjects are given one or very few generalization tests, entirely novel test stimulus values are used, and there is no differential reinforcement delivered during testing.

One other intradimensional procedure ought to be mentioned that does not fit simply into either the equalization or maintained-gradient category, but has aspects similar to both. Thomas and Lanier (1962) used it to compare wavelength gradients in groups of human subjects trained either to respond or not to respond to a particular value (525 nm). Verbal instructions were employed to generate the two different response tendencies. Both groups were instructed to press a telegraph key, but one group was told to release the key whenever 525 nm was presented and to keep pressing it whenever other wavelengths were presented; another group was told to keep pressing the key whenever 525 nm was presented but to release it whenever other wavelengths were presented. Generalization gradients plotting frequency of the release response *vs.* wavelength for the two groups were strikingly similar, except that one was the inverse of the other. Thomas and Lanier suggested that when pre-generalization-test procedures (instructions) are arranged to ensure more or less equivalent tendencies to respond and not-to-respond, generalization gradients would be virtually the same for the two tendencies.

These findings might not be meaningfully related to EGs and IGs obtained in animal studies. Humans have had a long prior history of labelling and discriminating colors. Simultaneous use of positive and negative values along one dimension makes it hard to separate excitatory from inhibitory effects. "Releasing" and "pressing", two specific responses, are probably not appropriate analogs for the opposing response tendencies usually thought to be involved in comparisons of excitatory and inhibitory control. In Jenkins' (1965) analysis of inhibitory effects in animal studies, he pointed out that not-responding may involve a variety of topographically different behaviors; but Thomas and Lanier gave their subjects only two alternatives, releasing or pressing.

Jenkins' argument about the asymmetry of EG and IG methods does imply that if subjects are required to make a specific response in S–, EGs and IGs ought to be much more similar than usual. However, if subjects are literally forced to make a specific different response in S–, as was done in Thomas and Lanier's study, it is legitimate to label the outcome an "IG"? It would seem more appropriate to call it an "EG" for a different response (*cf.* Migler and Millenson, 1969).

B. Interdimensional and Intermodal Methods

All intradimensional methods for obtaining IGs suffer from the fact that an experimenter cannot vary the distance of a test value from S– without also changing the distance of the test value from S+. Thus, any clear or direct separation of excitatory and inhibitory dimensional control is virtually impossible. Moreover, intradimensional methods for obtaining IGs normally involve pretraining procedures that are not very comparable to those typically used for obtaining EGs. To eliminate or minimize these deficiencies, several interdimensional and intermodal methods have recently been developed.

1. Basic Orthogonal Procedure

A new method for obtaining EGs and IGs after highly similar pretraining was independently devised by Jenkins and Harrison (1962; see also Jenkins, 1965), Schwartzbaum and Kellicut (1962), and Honig, Boneau, Burstein, and Pennypacker (1963). It has subsequently been applied and extended by Terrace (1966b), Farthing and Hearst (1968), Hearst (1968, 1969), Marsh (1968), Nevin (1968), Newman and Benefield (1968), Vetter and Hearst (1968), Desiderato (1969), Hoffman (1969), Lyons (1969*a*), Weisman (1969), Weisman and Palmer (1969), and Yarczower (1970). The critical and novel aspect of this procedure is the use of a dimension of S– that is orthogonal to dimensions of S+. Since S+ does not meaningfully lie anywhere on the generalization dimension, different distances from S– on that dimension are all presumed to be at equal distances from S+.

For example, in one of Jenkins and Harrison's experiments, S+ was a white noise and S– was a tone of 1000 Hz. Subsequent variation of S– along the auditory frequency dimension during generalization testing in extinction presumably did not move S– any closer to or farther away from the white noise. Thus, the obtained IGs with a minimum at S– could not have been contaminated by the influence of a nonhorizontal EG along the S– dimension. As Jenkins and Harrison state, by such a procedure "the inhibitory gradient would be an analog to the simple excitatory gradient because in testing for the excitatory gradient we do not move toward or away from a previously nonreinforced stimulus value as we move toward or away from the reinforced stimulus value (1962, p. 435)."

Honig *et al.* (1963) used a similar technique to compare EGs and IGs along a visual dimension. One group of pigeons learned a successive discrimination involving (a) intermittent reinforcement for pecking at a white key on which a black vertical line (S+) was projected, and (b) extinction for pecking at the key when the black line was absent (S–). The EGs were then obtained by varying the orientation of the black line during extinction, as on the standard Guttman-Kalish procedure. A different group of pigeons learned the reverse discrimination: S+ was the blank white key and S– the key with the vertical line on it. The IGs were then obtained from these subjects during extinction, in exactly the same way as for the EG. Since the blank key is presumably no closer to any one line tilt than to any other, EGs and IGs obtained by this method should be free of contamination from nonhorizontal line-tilt IGs and EGs respectively. Although the total number of responses made to line orientations during determinations of IGs in Honig *et al.*'s experiment was appreciably smaller than during determination of EGs, the EGs and IGs based on relative response measures (per cent of total responses) did not differ very much, except for the fact that one gradient was the inverse of the other.

Like the equalization procedure, the simple orthogonal method may be conceptualized as involving an initially horizontal EG along a dimension, which summates with the relatively steep IG produced by subsequent extinction at one value along that dimension. On the equalization procedure, the initially horizontal EG is experimentally established by nondifferential reinforcement at all points along the dimension, whereas on the orthogonal procedure all points along the S– dimension are merely assumed to be the recipients of equal excitatory generalization from S+. According to this argument, nonhorizontal gradients obtained on either procedure could be due only to unequal inhibitory dimensional control, because excitatory effects are uniform all along the dimension.

The orthogonal-training method for obtaining IGs has several major strengths. First, unlike the equalization procedure, it attempts to vary S– along some dimension without also varying the distance of these dimensional values from reinforcement-correlated stimuli. Therefore, in theory at least, it will permit clearer isolation of control by a dimension of S– than does the equalization procedure. Second, very similar pretraining is given all subjects before determination of EGs and IGs; only two stimuli are ever presented before the generalization test, one stimulus associated with reinforcement and the other with extinction. Third, the orthogonal procedure ensures that the various generalization test values will be equally novel for both groups; regardless of whether subjects provide an EG or an IG, before the test they are exposed to only one value along the generalization dimension. Thus, due to the similarity of training and testing conditions, the orthogonal technique ought to permit much more meaningful comparisons of EGs and IGs than is possible with the usual application of the equalization technique. An interesting possibility might be to use a combination of the equalization and orthogonal methods[8] to determine both EGs and IGs (as Weisman and Palmer, 1969, have attempted) and thus to avoid some of the disadvantages of the former method. However, with such a combination, the problem of clearly separating excitatory and inhibitory dimensional control would still remain, since reinforcement and extinction have previously occurred at different points along the same dimension.

Although the orthogonal procedure represents an important methodological advance over previous methods of obtaining IGs, several problems arise in its application and interpretation. First, it is usually difficult, if ever possible, to find an S+ with features or dimensions that can be unequivocably stated to be orthogonal to the dimension that an experimenter wishes to study. If S+ is not actually equidistant from all test values on the S– dimension, then application of the entire method becomes less appealing.

An obvious example of this problem faces the experimenter who wants to determine EGs and IGs along an intensity dimension: what possible stimulus could he choose that would be equidistant from all values along this dimension? Any conceivable "blank" stimulus would have some intensity of its own and therefore it could not fulfill the criterion of equidistance. Usually this problem is much more subtle than in the case of intensity generalization and therefore much more likely to be overlooked.[9] In fact, almost every study that has presumed equidistance between S+ and all values on the S– generalization dimension is open to possible criticisms of this kind.

A few examples should serve to illustrate the problem. Marsh (1968) used a white light on the pigeon's response key as a stimulus value orthogonal to the wavelength dimension, but it is quite conceivable that behavior con-

---

[8] Following equalization training, Kling (1952) continued to provide reinforcement for responding to a presumably orthogonal stimulus (black square) while one value of white-circle size was being extinguished.

[9] Farthing, Koresko, and Hearst (1968) presented strong evidence that the frequency of a clicking sound (pulses per second) is not orthogonal to the absence of a sound; the latter could be conceived of as lying on the click frequency dimension (at 0 pps). In that study, peak shift toward the higher click frequencies was often observed in gradients following discrimination training with S+ (click) and S– (silence), but was not obtained after simple VI training to one click value. Some aspects of these findings resemble the phenomenon labelled stimulus intensity dynamism by Hull (1949).

ditioned in the presence of a white light would occur more to some wavelengths on the key than to others—perhaps because of impurities in the whiteness of the light or because of unknown preferences or actual physical similarities based on saturation or intensity (see also Blough, 1966, pp. 363-364). If this were the case, one would have to utilize psychophysical data and to control or randomly vary the intensity and saturation of the training and test stimuli much more efficiently than is usually possible with the miniature display cells that are frequently used to project stimuli in studies of this sort. Analogously, Terrace's (1966*b*) and Lyons' (1969*a*, 1969*b*) choice of a white vertical line on a black background to serve as a stimulus value orthogonal to the wavelength dimension may not be entirely satisfactory. Various wavelengths may be closer than others to the line stimulus in terms of intensity or saturation.

The vertical line *vs.* blank (no line) visual discrimination used by Honig *et al.* (1963), Hearst (1968, 1969), and many others for obtaining IGs can be similarly criticized. If birds consistently look at a specific area at the top or bottom of the key, near to where the black vertical line is projected on S— trials, then variation of the orientation of the line away from the vertical might make that area on the key more "like" its value during S+ (that is, there is more "blank" space in that area when the line is rotated away from it). Variations in what has usually been presumed to be only a dimension of S— may also have involved correlated changes in some property of S+. Therefore, these experimenters could have been measuring an excitatory rather than (or in combination with) an inhibitory effect.

In their first experiment to determine IGs for auditory frequency, Jenkins and Harrison (1962) observed a slight rise in responding at the higher frequencies. They pointed out that this result could be explained on the assumption that tones of higher frequency are less loud to the pigeon and hence are more similar to S+ (which was the absence of a tone in that experiment). In this case an auditory rather than a visual dimension was involved, but the possibility again illustrates the practical difficulty of discovering an S+ that is equidistant from all values on S— dimension.

All these examples should re-alert workers in this area of research to the related question of whether the nominal, experimenter-defined stimulus dimension is really the functional dimension for the subject. Many prior researchers (see a review in Sutherland, 1961) have demonstrated that animal subjects may attend mainly to very specific features of a complex stimulus in learning a discrimination, or may attend to dimensions other than the one or in addition to the one which the experimenter had intended. Such outcomes would seem more likely with increasingly complex training stimuli. The fact that the "orthogonal" discrimination in Terrace's (1966*b*) and Lyons' (1969*a*, 1969*b*) experiments (*i.e.*, white line on a black background *vs.* a wavelength) could be learned on the basis of color or intensity differences, as well as on the basis of the presence *vs.* absence of a white line, certainly seems to make interpretation of any eventual line-tilt gradients very difficult (see a discussion of this point in the next section on combined-cue procedures.). Beale and Corballis' (1968) tentative explanation ("beak shift") of interocular mirror-image reversal effects in pigeons provides another example of how subjects may possibly be controlled by cues not considered by the experimenter. In studies with retarded boys, Touchette (1969) showed that even tilted lines may themselves be very complex stimuli.

One conceivable way of answering skeptics who may argue in specific cases that S+ is not in fact equidistant from all values along the S— dimension would involve an actual demonstration of the orthogonality of S+ to the generalization dimension of S—. Such empirical tests for orthogonality would be very valuable, but the best way of performing them is not yet known. Prior attempts have not yielded a clear solution.

Jenkins and Harrison (1962), Farthing and Hearst (1968, p. 749), and Lyons (1969*a*) all reported results of such a test. In the Farthing and Hearst report, for example, some birds were first given the usual VI reinforcement to a blank white key for seven sessions; but then, without any further training, a standard "generalization test" was administered that included the blank stimulus and six orientations of a black line bisecting the white key. Unlike subjects in EG or IG groups, these birds had never seen a line on the response key before the generalization test. Although some individual subjects responded more frequently to certain line orientations than to others, the average line-tilt gradient for the group was approximately flat, lending support to the presumption that the blank stimulus was equidistant from values on the line-tilt dimension. Jenkins and Harrison's and Lyons' comparable control data also revealed essentially horizontal gradients along the dimensions (tonal frequency and line-tilt respectively) that they examined for S− generalization in their birds. Such results would satisfy one possible operational definition of orthogonality: a flat, above-zero gradient along some presumably unrelated dimension after reinforcement training to a certain stimulus display.

Unfortunately, however, these empirical tests for orthogonality have serious weaknesses. First, a flat group gradient might merely conceal individual preferences (lack of orthogonality) for various values along the potential S− dimension—idiosyncratic "preferences" which Farthing and Hearst did notice in several subjects and over which the experimenters apparently had little or no control. No one can yet state why a specific bird may peck at or look at a particular area of a blank key, and, as suggested above, variation of the tilt of a line may have different effects, depending on what area of the key a bird actually does attend to. Unless one is interested only in group functions, a method is needed for measuring a subject's preferences along a dimension before actual training with stimuli on that dimension, but this seems impossible to accomplish without obvious contamination of later generalization measures.

Second, the occurrence of a flat gradient does not provide very strong evidence that test values along the potential S− dimension are actually equidistant from S+. As is often the case with flat gradients, the test procedure may simply not have been sensitive enough to reveal differential control by values along the dimension (see, for example, the attention *vs.* cue utilization experiments of Newman and Benefield, 1968, and Honig's review, 1969). Third, and most important, however, is the possibility that a flat gradient obtained on such a test of orthogonality may be useless as an indicator of the degree of orthogonality that actually exists following S+ vs. S− training in the experimental groups. For example, Farthing and Hearst's birds may have begun to look consistently at a particular area of the key *only after* differential reinforcement had been supplied for responding to the blank S+ *vs.* the line S−. Such differential reinforcement is an integral part of the training phase on the orthogonal procedure, but it may destroy "orthogonality". This major problem of discovering an orthogonal S+ and of ensuring that it remains equidistant from values on the generalization dimension of S−, even in the face of subsequent differential training, requires more serious consideration than it has received in the past.

A second major problem with application of the orthogonal procedure arises because generalization of behavior from S+ to the orthogonal S− dimension is often very limited. Thus, even if the first problem were somehow resolved and we could be fairly confident that S+ was equidistant from all values along the S− dimension, presentations of values along that dimension might yield so little overall responding that graded decremental effects would be difficult to detect. After a preliminary phase of VI reinforcement to a blank response key, the response rate of pigeons to the initial presentations of a key with any line on it is only 40 to 60% of response rate to the blank S+ (Farthing and Hearst, 1968; Hearst, 1969). The average response output to various line tilts during later generalization testing ought therefore to be typically lower when line-tilt serves as the S− continuum than when it serves as the S+ continuum and the blank is S−.

In addition to this initial difference in level of responding along the line-tilt dimension, which depends on whether behavior is originally acquired and reinforced to a line-tilt or a blank, the effects of later discrimination training in which responding is extinguished to a line-tilt in the IG group will also serve to make absolute IGs much shallower than EGs obtained for comparison (see the results of Honig et al. [1963], Jenkins and Harrison [1962], and Hearst [1968, 1969]). It is interesting that Spence and Hull assumed that IGs are shallower than EGs, which enabled them to account for learning and maintenance of behavior under partial reinforcement (*i.e.*, reinforcement and extinction at the same stimulus do not cancel each other out). However, the fact that this difference is experimentally obtained with the orthogonal procedure seems more a function of specific aspects of this method than a confirmation of Spence and Hull's assumptions (see also Graham, 1943, p.

365). Jenkins' (1965) incisive discussion of the asymmetry inherent in comparing EGs and IGs by means of the orthogonal procedure is closely related to the above points. His argument also implies that the orthogonal procedure will yield IGs that include appreciably fewer responses than are present in EGs along the same dimension.

The fact that behavior occurring in S+ may be drastically reduced by the presentation of stimulus values along an orthogonal dimension created a serious problem in the interpretation of an important experiment performed by Terrace (1966*b*), which we mentioned briefly earlier in this article. He found that some birds learn the discrimination (S+: white line on a black background; S—: 550 nm of light for one bird, 580 nm for three other birds) with virtually no errors following 14 days of VI training to S+ alone. In other words, these "errorless" subjects showed almost no initial generalization of behavior from S+ to some wavelength. When Terrace subsequently varied wavelength from 490 to 670 nm in 10 to 20-nm steps, he obtained very flat gradients with virtually zero responding all along the wavelength (S—) dimension.

Although Terrace concluded that IGs are flat after such errorless[10] training, this conclusion seems unwarranted to us. When values far from S— produce zero responding there is no way to measure less responding at S—. Subzero response tendencies may actually differ in strength but a "floor effect" prevents detection of any such differences (see also Reynolds [1964], Eckerman [1967], Vetter and Hearst [1968], Johnson, Kinder, and Scarboro [1969], and Farthing and Hearst [1970] for other examples and discussions of this type of problem). Although Terrace may very well be correct that errorless discrimination learning leads to flat IGs, his experiment does not appear to permit this strong conclusion. Some way has to be found to raise the overall level of responding during the generalization test so that dimensional control, if it exists, can be detected. The next two sections of this article describe possible methods of accomplishing this goal.

2. *Combined-Cue Tests*

When little or no responding occurs along the S— dimension during generalization tests after orthogonal pretraining, assessment of dimensional control becomes extremely difficult. The next two methods we discuss ("combined-cues" and "resistance-to-reinforcement") are attempts to handle this problem, primarily by modification of the generalization test procedure itself. Both methods are designed to elevate response output during testing to a level high enough to permit reliable detection of graded decrements along the generalization dimension. If these new ways of testing do achieve this goal, they will solve one of the two major problems we mentioned with the standard orthogonal procedure. Based on results thus far, our current opinion is that resistance-to-reinforcement testing is much superior to combined-cue testing, but we will describe the latter procedure first.

In a combined-cue test, values along the S— dimension are presented simultaneously with a well-trained S+ that is either orthogonal to the S— dimension or comes from another modality. Since the isolated presentation of this S+ produces a great deal of responding, its combination with S— values during generalization testing should increase behavior along the S— dimension to a level well above that to the S— values alone. If combinations of this S+ and a variety of specific values along the S— dimension do yield a reliable gradient with a minimum at or near S—, then dimensional control around S— is thereby demonstrated.

Furthermore, a simpler post-training test may be given in which responding to S+ is compared to responding to the simultaneous presentation of S+ and S—. If this novel combination of S+ and S— produces appreciably less behavior than occurs to S+ alone, then S— is thereby demonstrated to be an inhibitory stimulus—provided of course that appropriate

---

[10] Terrace's procedure of selecting "errorless" subjects on the basis of their actual performance on the original discrimination seems virtually to ensure a zero baseline. These subjects had never responded to the S— wavelength 'during training, and if the wavelength dimension is in fact orthogonal to the S+, then the subjects would be expected to emit almost no responses along the entire wavelength dimension. A less dubious way of arranging for errorless subjects would appear to be through the use of some experimental manipulation, such as "fading", to achieve errorless learning, rather than through the selection of subjects with a strong initial tendency not to respond to wavelengths after S+ (line) training. Bernheim (1968), Farthing and Hearst (1968), and Hearst (1969) have discussed these points in some detail and have remarked on their relevance to the 1967 controversy between Terrace and Deutsch.

control groups (*e.g.*, without prior exposure to or extinction to S−) yield much less or no reduction of behavior when the S− value is added to S+. This control result would eliminate the possibility that the decrements in the experimental group were due to the mere novelty of such a combination or to stimulus change alone. By analyzing generalization gradients along the S− dimension and the effects of simply combining the specific S+ and S−, an experimenter would have a basis for deciding not only whether dimensional control existed around S−, but also whether S− was itself an inhibitory rather than a neutral stimulus (see Section IV).

Within the combined-cue method, S+ serves the vital function of producing a strong response for S− values to reduce. In the standard orthogonal method for obtaining IGs, on the other hand, the experimenter relies on the power of unspecified situational cues (which had been present during the reinforcement of operant behavior), or on generalization from S+ to the S− dimension, to produce enough responding to enable measurement of differential behavior during generalization testing in extinction. Both the combined-cue technique and the standard orthogonal technique for determining IGs can be conceptualized as involving a summation of uniform excitatory influences resulting from the association of reinforcement with S+ or situational cues, and unequal inhibitory influences resulting from the association of nonreinforcement with S−. In his discussion of active tests for Pavlovian inhibition Rescorla (1969*a*, 1969*b*) labels a method that is very similar to our combined-cue procedure the "summation technique". We prefer a different label, since virtually all the methods discussed in this article (and in Rescorla's work) involve summation in some sense.

Because the combined-cue testing procedure requires the simultaneous presentation of positive and negative values, the stimuli employed as S+ or S− during pretraining must necessarily be of a special kind. The S+ and S− must differ in at least two features or dimensions, since otherwise they could not meaningfully be combined. For example, the S+ of no-tone and S− of 1000 Hz used in one of Jenkins and Harrison's experiments (1962) could not themselves be combined without producing S− itself, as would also occur if the blank (S+) and vertical-line (S−) visual stimuli of Honig *et al.*'s study (1963) were employed during initial training. On the other hand, an appropriate training pair for a pigeon might be S+ (red light) *vs.* S− (white line on a dark background), which differ in at least two ways: wavelength and the presence *vs.* absence of a white line on the key. During subsequent generalization testing (which in past research has typically involved variation of line tilt on its normal dark background), different orientations of the white line could be presented for the first time on a red (S+) background. The experimenter presumes orthogonality between the wavelength dimension and the dimension of line tilt, and between the properties of a blank key and a key with a line on it, as he does with the standard orthogonal procedure.

Previous applications of the combined-cue technique have rarely involved the study of generalization gradients. They have concerned the identification of the inhibitory (reductive) power of a specific stimulus, rather than inhibitory dimensional control. Examples of the former type of study, which have often been performed within a classical-conditioning paradigm, can be found in Pavlov (1927, see Leporsky's and Babkin's experiments on pp. 75-77), Rodnick (1937), Cornell and Strub (1965), Brown and Jenkins (1967), Hammond (1967), Konorski (1967, *e.g.*, pp. 320-321), Weiss (1967), Bignami (1968), Bull and Overmier (1968), Reberg and Black (1969), and Rescorla (1969*a*, 1969*b*).

To our knowledge, the only published experiments that have specifically employed the combined-cue method to determine IGs in an appetitive situation[11] are the recent studies of

---

[11] We have concentrated on appetitive tasks in this monograph because almost all studies of IGs in an operant situation have employed such tasks. However, the methods of Desiderato (1969) and Hoffman (1969) in operant aversive situations did yield IGs for Pavlovian CS−s by means of essentially a combined-cue procedure. For example, Desiderato obtained clear EGs and IGs with a Sidman avoidance hurdle-jumping response. The experiment included two different groups of dogs that had received prior Pavlovian orthogonal training either with (a) CS+ (800-Hz tone) and CS− (clicker), or (b) CS+ (clicker) and CS− (800-Hz tone); shock was the US. Discrete presentations of auditory frequencies ranging from 250 to 2600 Hz, or the clicker, were given during avoidance extinction. Thus, the generalization test involved novel combinations of auditory values and the general situational cues that presumably controlled the avoidance response. Avoidance response rate was moderately high, so that both increases and decreases in the baseline could be measured.

Lyons (1969*a*) and Yarczower (1970). In Lyons' experiment, for example, pigeons were trained, either with or without errors, to discriminate between S+ (green light, 555 nm, illuminating the entire response key) and S— (a white vertical line on the otherwise black key). Half the subjects subsequently received generalization tests in which line tilt was varied on its usual black background, whereas the other half received tests in which the tilt was varied on a green (S+) background. By presenting the different S— test lines to one half of the birds on a background consisting of the S+ color, Lyons tried to ensure that at least half of his birds would produce a high response output during testing. As Lyons pointed out, and as we have noted above, Terrace's (1966*b*) finding of flat, near-zero gradients along an S— dimension in subjects trained without errors could have reflected merely the operation of a "floor effect", which prevented the detection of any differential response strengths along the S— dimension. "A test in which each test angle is superposed on the 555-nm light can be expected to induce greater responding in the presence of those stimuli, possibly revealing differences in response strength which Terrace's procedure may have concealed (Lyons, 1969a, p. 491)."

In Lyons' published report he concentrated on measures of relative generalization, which revealed some surprising effects: subjects trained either with or without errors exhibited *maximum* responding at the S— (vertical) line tilt in the combined-cue (superimposition) test groups, whereas irregular gradients ("errorless" group) or gradients with a minimum at S— ("errors" group)[12] were obtained in subjects for which line tilt was varied on its usual black background. Lyons also mentioned that the presence of the S+ color did increase overall generalization-test output for subjects trained with errors, as compared to the output of similarly trained subjects that were tested without the superimposed S+ color. In Lyons' errorless and control ("S+ alone" training) groups, on the other hand, superimposition did not produce any significant response augmentation.

Yarczower (1970) reported that a similar superimposition test did not consistently increase responding in his birds, all of which learned the discrimination with errors. Therefore, in Lyons' and Yarczower's studies the combined-cue technique did not succeed very well in its main purpose, which was to raise the response output during generalization testing.

Lyons concluded that "there is no negativity (in the sense of response suppression) connected with the negative stimulus for the errorless group (1969*a*, p. 491)." Lyons is presumably using the word "negativity" in the sense in which we have defined an "inhibitory stimulus", that is, in terms of the learned capacity of a stimulus to reduce behavior regardless of whether evidence of specific dimensional control is also obtained. However, we do not understand the exact basis for Lyons' specific conclusion, since the appropriate way to determine the negativity of S— would involve a comparison of response output in the presence of S+ (555 nm) and response output in the presence of a combination of S+ and S— (*e.g.*, 555 nm plus the vertical line). To further interpret such a comparison, a single-stimulus control group (which Lyons also included in his experiment) provides indispensable information concerning the extent of the decrement produced by mere stimulus change or by the novelty of the combination of S+ and S—, *i.e.*, after intermittent reinforcement to a 555-nm stimulus only, how much reduction in responding would occur when a novel vertical line is superimposed on the 555-nm stimulus? Unfortunately, a comparison of (a) responding to S+, with (b) responding to the combination of S+ and S—, was not possible in Lyons' study, because during generalization testing he never included trials with S+ alone in any of his groups. Thus, no baseline was established from

---

[12] Robert Kaplan of the University of California (Berkeley) has suggested an ingenious but thus far untested explanation of this paradoxical result in the "errors" group. If Lyon's birds mastered the original discrimination mainly on the basis of color, a salient dimension for pigeons, then the birds may have maintained visual fixation on either the left or right half of the response key—a place that was always either completely green or completely black. Subsequently, as the line was varied from vertical to horizontal on a black background during generalization tests, the "effective stimulus" for *non*responding (*e.g.*, a solid black space on the left half of the key) was changed—it was no longer a completely black area—and responding to it should increase. As the line was varied from vertical to horizontal on a green background, the effective stimulus for responding (*e.g.*, a solid green space on the left) was changed—it was no longer a completely green area—and responding should decrease.

which to measure the decremental effects of S–[13].

Lyons (1969*b*) performed an important counterpart to the above experiment. Three main groups of subjects (errors, errorless, and single-stimulus control) were tested for control along a dimension of S+ following pretraining to the stimuli used in his other experiment (*i.e.*, S+ [555 nm]; S– [white vertical line on a black surround]). Each of the three main groups was subdivided into two subgroups, and visual wavelength was varied during generalization testing either (a) with a superimposed vertical line always present on the key, or (b) with no line ever present on the key.

Steep gradients along the wavelength dimension, with a maximum at 555 nm, occurred in every group. Furthermore, responding in all three groups was decreased by the superimposition of the S– line on the colored background. Interestingly, this reduction was significant only in the errorless group. According to our discussion above, the most appropriate measure of the inhibitory (reductive) effects of the vertical line involves a comparison of those trials on which S+ was presented alone and those on which it was combined with the vertical line. In his published report, Lyons did not present exact data concerning this comparison, but a personal communication from him revealed that the addition of the vertical line reduce responding at 555 nm (S+) by 42% in the control group, 48% in the with-errors group, and 70% in the errorless group. This is the kind of evidence, lacking in Lyons' other experiment, that one needs to determine whether S– is an inhibitory stimulus. Responding at S+ can be compared with responding to a combination of S+ and S–, and a separate control group permits evaluation of the generalization decrement to be expected merely by adding a novel line to the green S+.

Since reduction of S+ responding by the vertical line was significantly greater in the errorless group than in the single-stimulus control group, one has a basis for the conclusion (which Lyons also reached, but on the basis of overall generalization test responding) that his errorless S– was an inhibitory stimulus; it weakened behavior much more than the same stimulus presented after S+-only training. This is virtually the opposite of Terrace's (1966*b*) conclusion that an errorless S– functions as a neutral stimulus. The conclusion from Lyons' work is based on a response-reduction definition of an inhibitory stimulus, which bears obvious similarities to the criterion Deutsch (1967) employed when he argued that Terrace's results demonstrate greater inhibition by S– following errorless discrimination learning than following learning with errors. Problems of this kind were discussed earlier in this article, where we pointed out that Terrace uses a U-shaped empirical gradient along an S– dimension as his criterion for deciding whether or not S– is inhibitory. Some final comments on these definitional problems will be reserved for the last sections of the present article. However, since Terrace's and Lyons' procedures for obtaining learning without errors were so different (*post hoc* selection *vs.* response prevention), it would be worthwhile for someone to repeat Terrace's training procedure and then to attempt a combined-cue test. Until that information is available, the applicability of Lyons' conclusions to Terrace's situation cannot be clearly evaluated.

Since the combined-cue technique has not yet been tested sufficiently, it is really premature to attempt an evaluation of its overall worth for studying stimulus generalization. If, unlike Lyons' and Yarczower's general findings, the technique often does accomplish its main purpose of producing appreciable behavior along the S– dimension, especially in cases where the absence of S+ yields near-zero behavior, then it will probably be a useful addition to an experimenter's list of procedures for determining IGs. It retains many of the advantages of the standard orthogonal procedure, which we enumerated in the last section of this article. Furthermore, data from the technique can provide an "active" test of

---

[13] In recent work with a procedure similar to Terrace's (1966*b*) we have included trials with S+ alone (555 nm) among the series of generalization test stimuli. Every bird, but especially those learning without errors (not defined via a fading procedure, or by the response-prevention method Lyons used, but through *post hoc* selection of the kind Terrace (1966*b*) used), responded much less (usually a decrease close to 100%) when S+ and S– (vertical line) were combined than when S+ was presented alone. Therefore, S– does not appear to be a neutral stimulus following this kind of learning-without-errors. However, all these experiments ought to be repeated with errorless learning induced by a fading procedure (see Footnote 10) and with the inclusion of a comparison group that receives presentations of S+ with reinforcement and no presentations of S– during training (single-stimulus control).

whether S− itself is an inhibitory stimulus, like the summation method Rescorla (1969*a*, 1969*b*) recommends for detecting Pavlovian conditioned inhibitors. Perhaps intermodal combinations (*e.g.*, light and tone: Weiss, 1967; Rescorla, 1969*a*; Wagner, 1969*b*; vom Saal and Jenkins, 1970) will prove more effective than intramodal combinations (Lyons and Yarczower) in the study of dimensional control. Although relatively untested in studies of dimensional control, intermodal combinations have been the method of choice for studying inhibitory control by a specific stimulus via the combined-cue technique.

In spite of possible advantages, we do not feel too hopeful about the ultimate utility of the combined-cue method in studies of stimulus generalization, particularly when intramodal combinations are involved. Several disadvantages of this combined-cue method appear serious. First, orthogonal dimensions are hard to find, and the same comments we made earlier concerning this point apply with even greater force to the combined-cue procedure. Because the pretraining S+ and S− must differ in more than one aspect or dimension in order to enable a meaningful combination of cues, there is an even greater chance that S+ will not be equidistant from all values along the S− dimension. In Lyons' and Terrace's experiments, the S+ and S− were a wavelength of light *vs.* a white line on a dark background. These stimuli differed in color, in the presence of a line, probably in intensity, *etc.* Perhaps this kind of problem could be partially solved by random variation of certain (irrelevant) dimensions, *e.g.*, intensity, during initial discrimination training.

Second, the inevitable cue redundancy in the original discrimination would probably reduce the chances of subjects attending to any single dimension or feature of S−. Gradients around S− may prove relatively flat, not because dimensional control is typically absent following a particular kind of discrimination training, but because another dimension or feature overshadows the generalization test dimension as a probable basis for the original discrimination (see Footnote 12). In that event, variation of the more powerful or salient dimension might have yielded clearer evidence of dimensional control.

Third, the redundancy of the original discrimination may make it so easy that numerous subjects acquire it virtually without errors; many of Terrace's and our recent subjects on a similar procedure have done this. Therefore, gradients obtained after such training may not be particularly applicable to the analysis of ordinary discrimination learning, which typically involves an appreciable number of responses to S−.

To determine whether a given S− is an inhibitory or neutral stimulus, the above-mentioned series of tests comparing behavior in S+ to behavior in the presence of a combination of S+ and S− provide a relatively direct method for isolating the "active" inhibitory function of an S−. Another way of achieving this goal involves a test procedure in which a former S− is converted into an S+ by the introduction of reinforcement during that stimulus. This technique of studying the "retardation of the development of a CR" (Rescorla, 1969*a*, 1969*b*) by a stimulus formerly associated with nonreinforcement is similar in rationale to the one for obtaining IGs to which we now turn.

3. *Resistance to Reinforcement*

A recurrent problem in the application and interpretation of methods for determining IGs has been the small amount of responding that may occur along an S− dimension during generalization tests. Recently, we have developed a new test procedure (Besley and Hearst, 1969) that appears to cope successfully with this problem. The basic idea is extremely simple: instead of the standard method of extinguishing behavior at all values during tests for the IG, we provide VI reinforcement for responding at all the test values. Well-above-zero response rates along the entire S− dimension ought to be ensured by the uniform availability of reinforcement. If some values along the S− dimension are actually more inhibitory than others, then they ought to resist the influence of reinforcement more strongly.

This method is practically the reverse of the popular Guttman-Kalish procedure for obtaining EGs. On the Guttman-Kalish procedure, behavior is first intermittently reinforced at one stimulus value and then resistance to extinction is examined at a variety of different test values along a dimension. On our procedure, however, behavior is first extinguished at one stimulus value and then resistance to reinforcement is examined at a variety of different test values along a dimension. Our initial applications of the resistance-to-reinforcement procedure, to be described shortly, have been successful in producing IGs with a total amount of responding that is *not* lower than in EGs obtained via the resistance-to-extinction technique.

The resistance-to-reinforcement method also bears certain obvious similarities to the equalization procedure we reviewed earlier, in the section describing intradimensional methods. On the equalization procedure (*e.g.,* Honig, 1961), operant responding is first equalized by intermittent reinforcement at various points along a dimension and then one of these values is presented while operant behavior is being extinguished; subsequently, all values are retested to determine the IG. The resistance-to-reinforcement method amounts to a reversal of the first two phases of the equalization procedure, since responding at one dimensional value is first extinguished (with responding usually maintained by reinforcement at an orthogonal S+), and then an equalization phase is run with uniform intermittent reinforcement at all test points along the dimension. This equalization phase is in fact the "generalization test", which can be continued for one long period of time or repeated again and again over many shorter sessions. The method thereby permits observation of changes in IGs that occur as a function of extended testing—a clear advantage over the usual procedure of testing in extinction, where responding along the dimension may reach zero too quickly to make lengthy testing for IGs feasible.

Rescorla (1969*a,* 1969*b*) presented a method for detecting Pavlovian conditioned inhibition that he calls the "retardation of learning" method. This technique is based on the notion that if a stimulus has become inhibitory, its subsequent transfer into a CS+ ought to be slowed up, compared to appropriate control groups. "If conditioned inhibition and excitation are subtractive of one another, then setting up inhibition to a stimulus either prior to or simultaneously with excitatory conditioning should retard the development of an overt CR (1969*b,* p. 82)." The logic of Rescorla's technique is thus extremely similar to the one we have used to justify the resistance-to-reinforcement method for obtaining IGs.

Like the combined-cue method for elevating response output, the resistance-to-reinforcement technique can be conceptualized as involving some kind of algebraic summation of excitatory and inhibitory effects. In the combined-cue procedure, different degrees of inhibitory control along the S− dimension are presumably counteracted, and response output thereby raised, by the simultaneous presentation of some important aspect of an excitatory stimulus (S+). In the resistance-to-reinforcement procedure, on the other hand, different degrees of inhibitory control along the S− dimension are presumably counteracted by the uniform availability of primary reinforcement (food) along the dimension.

Most prior experiments based on a method analogous to our resistance-to-reinforcement technique have employed classical-conditioning procedures and have focused on the measurement of the inhibitory properties of a specific stimulus, rather than on the determination of generalization gradients. Studies of the former kind include those of Pavlov (1927, see especially pp. 196-200 and pp. 302-310), Konorski (1967, see pp. 318-323 and other transfer paradigms in Chap. VII; some of these are also briefly discussed in Bignami, 1968), Hammond (1968, who calls the method "reversal conditioning"), and Rescorla (1969*a,* 1969*b*). Recent studies of "latent inhibition", in which nonreinforcement preexposures of a stimulus hinder its later establishment as a CS+, are also obviously related to the general method under discussion, even though no interspersed trials of reinforcement to a different stimulus occur during the preexposure period (see Lubow, 1965; Carlton and Vogel, 1967). The logic behind numerous studies of discrimination reversal is also quite similar to that of the resistance-to-reinforcement method (*e.g.,* Bower and Grusec, 1964; Mackintosh, 1965; Trapold, 1966; Trapold, Lawton, Dick, and Gross, 1968).

To our knowledge, the only studies of generalization gradients that have utilized a procedure like that of the resistance-to-reinforcement technique have been (a) the "retraining" and "recovery" phases of Honig's (1961, 1966) studies of the generalization of extinction and punishment, (b) an experiment concerned with the generalization of latent inhibition during classical eyelid conditioning of the rabbit (Siegel, 1969), and (c) a wavelength generalization experiment briefly described by Guttman (1965, p. 216), who faced the problem of near-zero responding in the vicinity of S— when he tried to measure "negative peak shift" following intradimensional training. In order to detect response decrements near S—, he provided uniform intermittent reinforcement at all test values on the wavelength continuum. No details of Guttman's results are given, but apparently the method yielded only transient effects in his pigeons, which had all been tested in extinction previously.

Some current data from our laboratory will be used to illustrate the resistance-to-reinforcement technique. Pigeons are trained on a standard orthogonal discrimination (S+: blank white key; S—: a thin black line bisecting the white key) as in Farthing and Hearst (1968), except that (a) a richer VI schedule (30 sec instead of 1 min) is used in S+, and (b) each subject must achieve a 4% discrimination ratio (S— responses/S+ responses) before receiving the generalization test (Farthing and Hearst gave different amounts of training to different groups of birds, but in their 8- and 16-day training groups the 4% criterion had been met by the time IGs were determined). Details of the generalization test along the line-tilt continuum are also essentially the same as in Farthing and Hearst's report, except that the VI 30-sec schedule remains in force throughout the test. Thus, subjects receive the same rate of reinforcement for responding to any of the seven test stimuli.

Figure 1 displays group gradients for five successive daily tests lasting approximately 50 min each (twelve 30-sec presentations of each stimulus, alternated with 7-sec blackouts). For the upper set of gradients, the S— during training had been a vertical line (0°), whereas for the lower set it had been a line tilted 30° clockwise from the vertical. Both groups displayed relatively steep IGs, with a large number of total responses, on the first day of testing. These absolute response levels were not very different from those in typical EGs obtained by the resistance-to-extinction method (see Hearst, 1968, 1969). Evidence of dimensional control persisted for two or three test sessions in the gradients for both S— groups. Although the Test I gradient for S— (0°) was fairly symmetrical around 0°, a mirror-image effect (*i.e.*, minima at both +30° and —30°) occurred reliably in the S— (+30°) group during Test I. In our previous experiments (Hearst, 1968, 1969), which involved testing for the IG in extinction, this effect was seldom observed in empirical gradients. But the resistance-to-reinforcement method has yielded mirror-image effects in a majority of individual subjects.

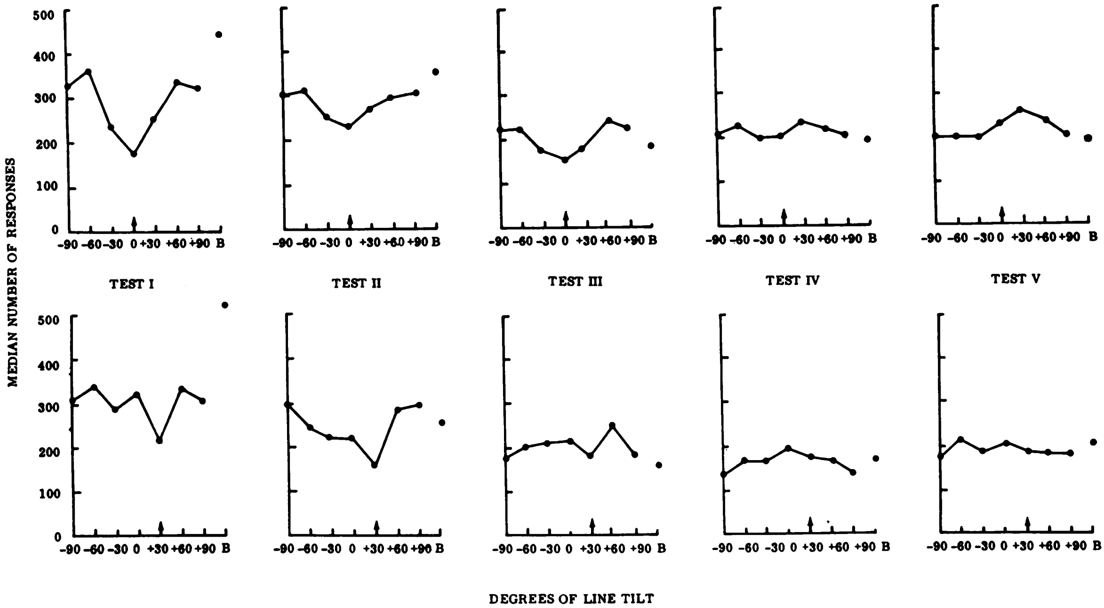

Fig. 1. Group median line-tilt gradients over five successive daily generalization tests with intermittent reinforcement at all stimulus values. The S— during training for the upper group (N = 9) was 0°, for the lower group (N = 8) +30°. The point "B" on the right of the abscissa designates the blank (S+) value.

As testing is continued, the gradients change in several interesting ways. First, the effects of prior nonreinforcement at S− do not disappear very quickly. One might have predicted that differential behavior to the various values would stop soon after subjects had started responding fairly rapidly to all stimuli and were receiving the maximum possible number of VI reinforcements in each stimulus. However, even during Test I the total numbers of reinforcements delivered at each stimulus value were approximately equal in almost every individual subject. Despite such equated reinforcement, differential responding along the dimension was observed during both halves of the first test session, and it frequently persisted for several sessions more in both grouped and individual data.

Second, we have observed unexpected inversions of gradient shape in more than 70% of our individual subjects. As testing is continued for as long as 12 sessions in both S− groups, the entire gradient often "flops over" so that maximal responding occurs to values near S− and minimal responding occurs to values far from S−. This "flop-over" phenomenon does not happen during the same test session for every subject, so that group gradients usually conceal it (but see Test V for the S− [0°] group in Fig. 1).

Figure 2 illustrates this effect for an individual subject that received eight tests. Gradients with maximum responding relatively far from S− were obtained for three consecutive daily sessions from this subject, but during the next session the gradient began to change. By the fifth test session, the entire gradient was inverted and this effect persisted for the remaining sessions. A discussion of the possible significance of this finding (*e.g.*, its relevance for the phenomenon of behavioral contrast) is beyond the scope of the present article. However, the novelty and reliability of the finding seem to provide an additional justification for future study of the resistance-to-reinforcement method.

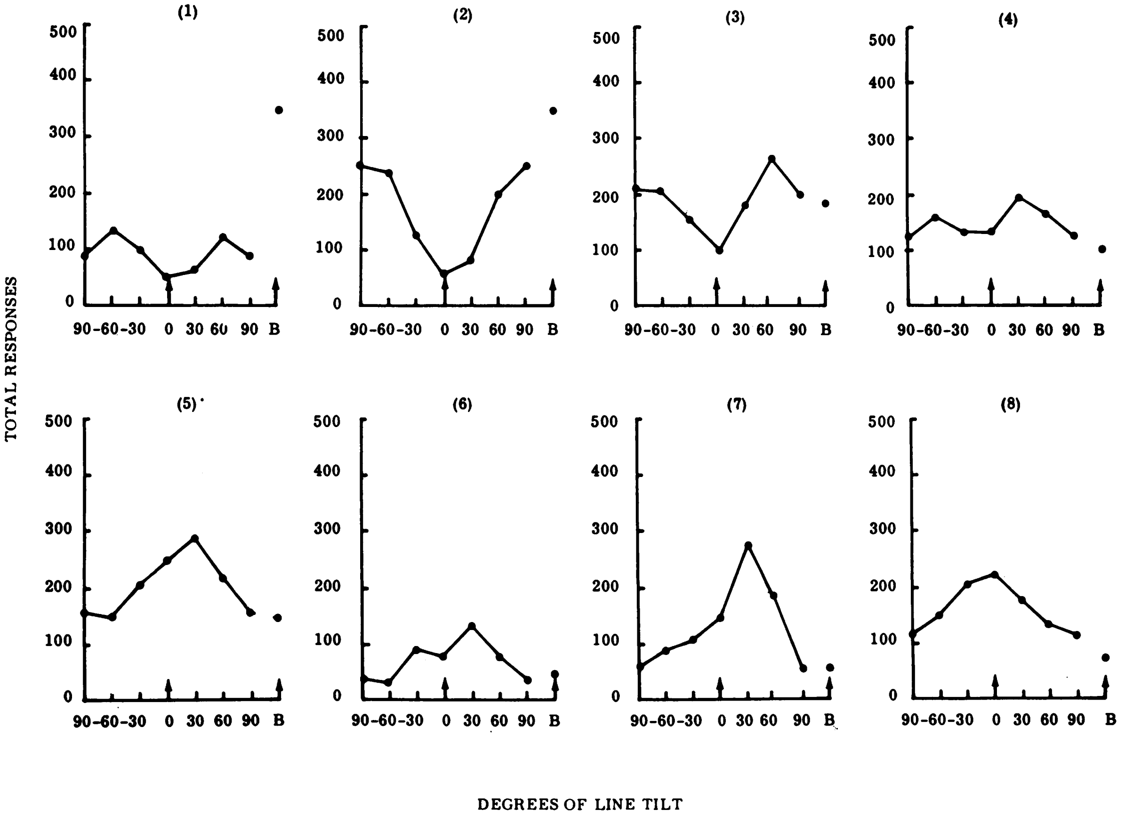

Fig. 2. Line-tilt gradients for Bird No. 1712 over eight successive daily generalization tests with intermittent reinforcement at all stimulus values. The S− during training was 0° for this subject.

These data indicate that the resistance-to-reinforcement technique is both a feasible and sensitive procedure for studying generalization along a dimension of S—. At the present time, insufficient work has been done with the technique to permit a fair evaluation of its worth, but mention of certain potential advantages and disadvantages seems appropriate. We believe that the advantages far outweigh the disadvantages and we hope that the method will receive frequent practical tests in the future.

On the surface at least, there seem to be numerous strengths in the resistance-to-reinforcement method. It will certainly succeed in raising response levels far above those normally obtained during tests in extinction. Moreover, it should yield response rates low enough to prevent ceiling effects and high enough to prevent floor effects in group and individual data. Because the absolute response levels displayed in the IGs of Fig. 1 are not far different from those typical of EGs obtained by the conventional testing-in-extinction procedure, the meaningful comparison and mathematical manipulation of EGs and IGs may be facilitated (see Hearst, 1969). Even more accurate matching of response outputs in EGs and IGs could probably be achieved by appropriate selection of parametric values during training and testing. Obviously, the actual amount of responding during resistance-to-reinforcement testing will depend on the VI schedule, amount of reinforcement, *etc.*

Another important advantage involves the lengthening of the duration of tests for dimensional control. The standard testing-in-extinction procedure typically restricts determination of IGs to one session of less than an hour, because subjects soon stop responding to all values along the S— dimension. The data from the new technique in Fig. 1 and 2 demonstrated that dimensional effects may persist for several sessions, during which time other interesting and rather novel phenomena could be observed.

Generalization tests with nondifferential reinforcement may prove useful not only in the study of inhibitory control, as was stressed above, but also in the detection and analysis of the effects of various specific procedures on EGs or postdiscrimination gradients (PDGs, see Hearst, 1968, 1969). For example, the rapid decline of responding to a near-zero level during tests in extinction may have helped to conceal some potential effects of massed extinction or "preextinction" in previous research (Honig, Thomas, and Guttman, 1959; Hearst and Poppen, 1965); reinforcement testing may permit observation of these effects. In fact, no overwhelming reason exists to support the almost inevitable choice of testing-in-extinction for studies in the field of operant stimulus generalization. In a relevant discussion, Sutherland (1961, pp. 28-29) evaluated the usefulness of two analogous forms of nondifferential testing (*i.e,* uniform reinforcement or extinction) in discussing transfer tests after learning of shape discriminations in animals.

Compared to the combined-cue method, the resistance-to-reinforcement procedure never requires the use of a multiple redundant-cue discrimination during the training phase. We pointed out earlier that discriminations of that kind do not permit very good experimental control over which dimension of S+ or S— the subject will attend to. Therefore, they may greatly limit the usefulness of the combined-cue method for determining dimensional control. In contrast, the reinforcement method can be used after almost any kind of training, including discriminations based on only a single feature or dimension. When used following so-called orthogonal training, the resistance-to-reinforcement method inherits many of that procedure's strengths and weaknesses, which we have already enumerated.

Future experimenters may wish to arrange a variety of different kinds of tests for control by a dimension of S— within individual subjects or for different groups of subjects. One kind of test procedure may prove to be more sensitive in a given bird than are the others. For example, a few of our birds that have shown flat gradients along a dimension of S— during a conventional test in extinction have revealed clear dimensional control during a resistance-to-reinforcement test; on the other hand, one bird yielded a flat gradient during a test with reinforcement but produced an incremental gradient around S— when later tested in extinction. Such results motivated our earlier suggestion that experimenters should continuously attempt to compare a variety of different test procedures for detecting dimensional control, in order to avoid false conclusions that may be drawn from flat gradients obtained on only one type of test. Rescorla similarly comments on the value of constant comparisons between the techniques he labels the "summation" and "retardation of learning" methods.

The resistance-to-reinforcement method is by no means free of potential weaknesses. Although our research does not provide any evidence for the possibility, it is conceivable that certain subjects will quickly discriminate the change to uniform reinforcement of all stimuli and will soon stop responding differentially. In that event, the test would hardly be a sensitive indicator of dimensional control. However, the use of intermittent reinforcement during training and testing ought to make this outcome extremely unlikely.

Another problem might arise if, during testing, a subject's response rate to dimensional values at or near S— remains so low that many fewer reinforcements are actually received at those values than at values far from S—. Such an outcome could act to maintain indefinitely a very low response output to values near S—, an effect analogous to the accidental contingencies that seem to account for many "stimulus superstitions" of the kind reported by Morse and Skinner (1957; see Footnote 6). However, such a result has been extremely infrequent in our research. Even though response rate to values at or near S— is initially low, subjects soon begin to respond to all stimuli; almost every one of our subjects has received approximately the same number of reinforcements at each stimulus value during Test I. This is strong evidence that the shape of our obtained gradients is determined mainly by the nonreinforcement of responding at S— during prior training and not by any inadvertent differential reinforcement received during the test. The choice of a VI schedule, rather than a ratio or other schedule for testing, is certainly the single most important factor in ensuring approximately equal numbers of reinforcements at all stimulus values. Even when response rates are appreciably lower at some stimuli than at others, the actual rate of reinforcement on VI schedules will remain about the same in all stimuli so long as the subject responds at least a few times during each stimulus presentation.

Another potential problem with the resistance-to-reinforcement technique is closely related to the one just mentioned. Owing to the particular range, number, and sequence of interreinforcement intervals constituting the VI schedule, behavior could be temporarily facilitated or suppressed at certain test stimuli by the chance delivery of an unusually large or unusually small number of VI reinforcements during presentations of these stimuli early in the generalization test. These "local" effects could persist long enough to affect gradients over an entire session. Such effects have not been apparent in our current work, but probably they can be completely avoided by modification of the VI schedule so as to arrange, for example, at least one reinforcement availability and no more than *x* reinforcement availabilities during each stimulus presentation. Our VI 30-sec schedule does not permit more than two reinforcements to be obtained during any 30-sec stimulus presentation.

Over the course of our generalization tests, subjects receive many more reinforcements than during their training sessions. Some subjects have gained appreciable weight during the test phase, especially when it continued for 10 to 12 consecutive days. Thus, our specific procedure could create a problem in motivational control. However, by appropriate choice of VI schedule, session length, or amount of reinforcement, such potential complications ought to be preventable.

Results from the resistance-to-reinforcement procedure may conceivably be influenced by summation of separate EGs that develop around each of the different test stimuli during reinforced generalization testing. On a circular dimension such as line tilt this possibility does not seem relevant, but with other dimensions (*e.g.*, wavelength) appropriate control groups probably ought to be included. As noted earlier, Honig (1961) employed an extra group to evaluate potential summation effects under the equalization procedure. The application of similar control procedures in our situation, say by giving some subjects *only* initial training to an orthogonal S+ and then administering a "generalization test" with equal VI reinforcement along the S— dimension, would provide relevant data on this point. Incidentally, the conventional testing-in-extinction procedure of Guttman and Kalish does not escape a similar complication, since one could argue in their case that separate IGs are being established around all the different test values, which summate over the course of testing to complicate interpretation of the obtained EGs.

Finally, evidence of an incremental (U-shaped) gradient around a former S— obtained by the resistance-to-reinforcement method does not in itself permit any strong conclusions as to whether S— is a neutral or an inhibitory stimulus. To answer this question, control groups would have to be added to determine whether responding at S— after the introduction of reinforcement takes longer to reach its eventual asymptote than would be the case for the same stimulus when it is, for example, novel or irrelevant. This is also the essence of Rescorla's (1969*a*, 1969*b*) "retardation of learning" test for Pavlovian conditioned inhibition. He compares the rapidity of CR acquisition for at least two groups, (a) one that has had prior training in which the new CS+ was never followed by US, and (b) another that has had prior training in which the new CS+ was presented independently of US. Discussion of such a procedure leads naturally into the last major section of this monograph, which summarizes some tentative methods for establishing baselines of neutrality from which to measure excitatory and inhibitory effects in operant conditioning.

IV. THE QUEST FOR A ZERO POINT: "NEUTRAL STIMULI"

Throughout this monograph we have argued that the finding of an incremental gradient around S— demonstrates merely that dimensional control exists along the S— continuum; it does not unambiguously indicate whether S— is a neutral stimulus or an inhibitory stimulus. Performance may be relatively low at S— (a) because any response-evoking power originally possessed by S— has simply been counteracted or neutralized by nonreinforcement, or (b) because S— has itself acquired the "active" capacity to reduce behavior. Distinguishing between (a) and (b) presents logical, technical, and theoretical problems that often are connected with the fact that, if a presumably neutral stimulus produces little or no behavior, then on most procedures it will be difficult to measure the decremental effects of a comparison stimulus that is thought to be inhibitory.

Previous and current research on control along a dimension of S— in operant behavior has proceeded quite systematically and efficiently without any real experimental attention being devoted to the question of whether S— is neutral or inhibitory. Many workers think it a meaningless question, others prefer to defer analysis of it because of the inescapable methodological problems it currently presents for solution, whereas still others frequently worry about it as they continue performing significant work on the factors that influence generalization gradients along an S— dimension. We cannot be sure which of these responses of researchers is the most adaptive or profitable one, but a goal of this monograph has been to convince more workers to consider the question carefully.

There do not appear to be any simple general solutions to the problem of distinguishing neutral and inhibitory stimuli, but in this section we would like to review a few potential strategies, most of which have already been implied or suggested at other places in this monograph. Because experiments along these lines have been sparse in the field of operant conditioning, the following discussion will be sketchy and quite tentative. Let us say at the outset that we strongly believe operant research on the topic of S— generalization can continue to be productive and valuable without undue concentration on the problem of determining a zero point and developing specific tests for the active inhibitory properties of a stimulus. In fact, many operant researchers (with good justification) feel that attempts to solve such problems are premature; the field would benefit more, they say, by the collection of additional basic parametric data on the variables affecting dimensional control by S+s and S—s. But, in the long run, we think an analysis of the kind we suggest is inevitable and the sooner the accompanying problems are faced the better.

Although we will be concentrating here on attempts to distinguish between inhibitory and neutral stimuli, such a goal also necessitates the complementary study of excitatory stimuli. Otherwise, we would have no strong justification for our eventual decision as to where a zero point should be located. Since excitatory effects have been extensively studied in the past and seem to present fewer logical difficulties and measurement problems than in the case of inhibitory effects, we will not say very much more about excitatory stimuli (but see Rescorla, 1967, and several relevant portions of our above discussion).

As a direct result of an important series of papers by Rescorla and his colleagues (*e.g.*, Rescorla and LoLordo, 1965; Rescorla, 1967; Rescorla and Solomon, 1967; Rescorla, 1969*a*, 1969*b*), which have obviously had a great influence on the writers of the present article, experimental procedures for distinguishing neutral or ineffective stimuli from inhibitory stimuli are being developed and increasingly used in the field of Pavlovian conditioning. Rescorla sought to distinguish between the simple absence of conditioning to a stimulus and the acquisition of inhibitory properties by a stimulus; if a CR does not occur to the stimulus, there is no obvious way to separate these possibilities. But, as Rescorla remarked, "Intuitively it seems clear that learning that the US does not follow the CS is different from failing to learn that the US follows the CS or learning that the CS is irrelevant to the US (Rescorla, 1967, p. 76)." Rescorla's "truly random" control group, in which CSs and USs are both presented to the subject but entirely randomly and independently so that occurrences of CS predict nothing about subsequent occurrences of US, represents an attempt (see also Prokasy, 1965, and Seligman's criticism, 1969) to establish an appropriate zero point against which both the excitatory effects of a CS+ ("positive contingency", a stimulus predicting the occurrence of US) and the inhibitory effects of a CS− ("negative contingency", a stimulus predicting the nonoccurrence of US) can be evaluated.

These three general kinds of CS-US relationships in Pavlovian conditioning (*i.e.*, positive contingency, negative contingency, and random contingency) yield stimuli that have actually been demonstrated by Rescorla (1966, 1969*a*) to produce increases, decreases, or no effect respectively on baselines of operant (avoidance) behavior. For example, an auditory stimulus that had closely preceded all shocks (US) during prior Pavlovian conditioning sessions was superimposed on a Sidman-avoidance baseline in dogs and it increased response rate; the same stimulus, explicitly not paired with shocks in the Pavlovian sessions, decreased response rate; and the same stimulus (a) randomly presented in relation to shocks in the Pavlovian situation, or (b) merely presented in the Pavlovian situation without any shocks given the subject, or (c) presented for the first time during avoidance responding, had little or no effect on avoidance response rate. In more recent work, Rescorla has stressed *both* his summation method (the superimposition of external signals on Sidman-avoidance behavior could be considered a special case of this technique) and his retardation-of-learning method as useful techniques for differentiating among excitatory, inhibitory, and neutral stimuli. As noted above, these methods are analogous to what we have called the combined-cue and resistance-to-reinforcement procedures for obtaining IGs. Rescorla (1967, p. 78) contends that a major advantage of his approach, which stresses CS-US "contingencies" rather than CS-US "pairings", lies in its potentiality for yielding a manipulable continuum, with a tentative zero point from which both increases and decreases in behavior can be measured.

We are faced with the same kinds of problems and decisions in the isolation of excitatory and inhibitory stimuli in operant conditioning. First, an appropriate procedure must be selected that can be logically and (later) empirically justified as establishing a "neutral" condition, and yet which holds constant as many of the details of training as possible (*e.g.*, number and rate of stimulus presentations, overall density and patterning of reinforcements) in order to enable meaningful comparison with the "experimental" groups, which receive either correlated reinforcement or correlated extinction for responding during the same stimulus. Second, after training is completed, some uniform test procedure must be employed that permits the experimenter to detect both increases and decreases in behavior relative to the behavior of the comparison groups, *i.e.*, ceiling and floor effects must be avoided.

Many potential strategies for attacking these two problems in operant conditioning are analogous to those listed by Rescorla for classical conditioning. Therefore, we can leave implementation of the methodological details, and specific evaluation of these and some other approaches, as an exercise for the interested reader. In the following discussion we will give more than passing mention only to "choice" test procedures, which necessarily involve measures of operant behavior.

As noted above, the first question is: what *training* groups are appropriate as comparison conditions, to provide a possible zero point or standard against which to measure the decremental effects of a presumed inhibitory stimulus, say the S− in a discrimination learned with errors? Some "control" groups of this kind might be (1) a *novel-stimulus* group, which never receives presentations of the stimulus during training, but only during the subsequent test. Such a group has often been used in the past to assess "external inhibition" (*e.g.*, Brown and Jenkins, 1967). But since the stimulus is not novel for the discrimination group, one could argue that a proper comparison stimulus ought not to be novel either. (2) a *habituated-stimulus* group, which receives an equal number of presentations of the stimulus that serves as S− in the discrimination group, but whose subjects are never permitted to emit the specified operant response during the stimulus, and thus reinforcement or extinction for that response never occurs during the stimulus. For example, the same number of presentations of the stimulus could be given before operant training, or stimulus presentations could occur during training periods when the manipulandum is inaccessible. Such procedures are designed to equate stimulus "novelty" before testing in the two groups, but, among other things, the groups are obviously not equalized in terms of the subjects' opportunities to respond to the stimulus. (3) an *uncorrelated stimulus* group, which receives the same number and pattern of stimulus presentations as the discrimination group during training, but these stimulus presentations are not correlated with the periods of reinforcement and extinction of the operant response. Or, such subjects could be *yoked* to subjects in the discrimination group so that reinforcement became available for them at the same times as for the other group, but neither S+ nor S− (though presented the same number of times) would be consistently paired with the availability of reinforcement. The important difference between the treatments of the comparison and discrimination groups would be that reinforcement availabilities are independent of the prevailing stimulus conditions in the former group, whereas they are specifically correlated with external stimuli in the latter group. But one could argue that the uncorrelated condition actually trains subjects not to attend to external signals or "stimulus change" and is not really appropriate as a neutral condition (see Rescorla [1969*a*, 1969*b*] and Honig [1970] for data, discussions and possible tests of this "loss-of-attention" explanation).

The last of the above comparison procedures is obviously an operant counterpart of Rescorla's "truly random" control group. Interestingly, very similar comparison groups have been used in recent studies of selective attention and stimulus validity in rats and pigeons (Wagner, Logan, Haberlandt, and Price, 1968; Farthing, 1969; Honig, 1969, 1970). We think such groups could also provide valuable comparative baselines in studies focused on the inhibitory properties of a particular stimulus, or on inhibitory dimensional control.

These, then, are some comparison treatments that could be logically justified as providing "neutral" stimuli. Currently, no one can state with any great confidence which one, if any, of the treatments is the most appropriate or whether others (*e.g.*, types of extradimensional training [Honig, 1969]) would be equally valuable. But if Rescorla's results (1969*a*) for Pavlovian conditioning hold to any degree for operant conditioning, all three kinds of comparison treatment should yield relatively ineffective stimuli during subsequent testing.

Now that we have settled upon a few possible treatments for producing a presumably "neutral" stimulus, what subsequent behavioral tests can we employ to determine whether these treatments actually do produce ineffective stimuli and whether the same stimulus significantly reduces behavior in the discrimination group? We need to test the stimulus in a new situation that is uniform for all the groups. To repeat, we are interested here in detecting what we have called an *inhibitory stimulus*, not in measuring inhibitory *dimensional* control. Some possibilities in this respect are the following. (1) *Combined-cue tests.* During this kind of test the stimulus is combined for the first time with a new S+ that is known to produce appreciable and fairly equivalent rates of some operant response in all the different groups. In order to qualify as inhibitory, the stimulus in the discrimination group would have to reduce responding to the S+ significantly; and the presentation of the same stimulus in the various comparison groups would have to exert much less or no effect on responding to S+. Thus, our definition of an inhibitory stimulus would be fulfilled: "a stimulus that develops during conditioning the capacity to decrease response strength below the level occurring when that stimulus is absent (above, p. 376)." (2) *Resistance to reinforcement.* In this type of test, reinforcement for some operant response is made possible only in the presence of the former S−, and the acquisition of behavior to that stimulus is studied. We would conclude that the stimulus had become inhibitory in the discrimination group, if learning curves for subjects in that group took longer to reach some common level of performance than did learning curves in the various comparison groups. (3) *Choice (simultaneous) tests.* On this procedure, individual subjects are given a choice between the presumed inhibitory stimulus and another stimulus thought to be neutral. If the stimulus is actually inhibitory it ought to be selected significantly less often than the neutral stimulus. Such a test could be run with or without scheduled primary reinforcement (*i.e.,* with either nondifferential reinforcement or nondifferential extinction of choice responses).

A problem with the third method involves the selection of a neutral comparison stimulus. How do we know that this stimulus is really "neutral" (see Footnote 14)? A logical possibility is that it actually has excitatory effects and then the choice of it over the presumed inhibitory stimulus would also occur if the latter stimulus were merely neutral. Such an alternative explanation could presumably be evaluated by examination of the choice results from the three comparison training groups described above. Or perhaps a neutral stimulus could be selected for the study on the basis of previous preference tests in other (completely naive) subjects that indicate an equal choice of the stimulus to be tested for inhibition and the neutral stimulus.

Some of these problems are exemplified in an experiment of Biederman's (1967). He used a choice procedure to determine whether S− becomes and remains inhibitory during acquisition of a discrete-trial simultaneous discrimination in pigeons. To oversimplify the experimental procedure somewhat, a different compound (color-form) visual stimulus was presented on each of the two keys in a standard Skinner Box. If the subject pecked the correct stimulus (S+), it received 3 sec of grain reinforcement followed by a 12-sec timeout; on the other hand, a response to the incorrect stimulus (S−) produced a 15-sec timeout. Some birds happened to learn this discrimination with very few errors, whereas others made appreciably more errors. Subjects occasionally received probe trials, on which they were presented with a choice between S+ and a "neutral" novel stimulus (S°), or a choice between S− and S°. Subjects learning either with or without errors selected S+ much more often than S° (approximately 80% *vs.* 20%), and selected S° much more often than S− (approximately 70% *vs.* 30%), which led Biederman to conclude that S− is inhibitory following discrimination learning either with or without errors. In terms of our earlier discussion, he is apparently using the word "inhibitory" in the sense of control by a specific stimulus, rather than control by values along some dimension of that stimulus.

Unfortunately, Biederman's experiment is open to several criticisms. His subjects were categorized as "errorless" not on the basis of any experimental manipulation (*e.g.,* "fading techniques", as in Terrace, 1966*a*), but merely because they made very few errors in learning the discrimination. Such selection of subjects undoubtedly meant that errorless birds had a strong initial bias against responding to S− (see also Footnote 10). Therefore, it is difficult to determine the extent to which a subject's choice of S° over S− is due to the development of a *learned* inhibitory function of S− during discrimination training. Biederman performed no direct empirical test of the neutrality of S°; on the basis of his past experience with a variety of similar stimuli, he merely assumed that this novel stimulus (the same for all subjects: a pattern of white dots on a black background) was relatively neutral. It certainly might not have been, since it was selected as much as 20% of the time even when paired with S+. Possibly, white dots on the key resemble pieces of grain. At any rate, his conclusions would have been much stronger if he had obtained the same results with a variety of different kinds of novel stimuli as the S°, or if he had arranged for proper counterbalancing of the stimuli he actually employed (see also Bernheim, 1968, for further comments on Biederman's experiment; and Biederman's reply, 1968*a*).

Despite the apparent flaws in this specific experiment, its basic method of presenting the subject with a choice between positive or negative stimuli and a neutral stimulus seems an excellent one to pursue and develop further (see the interesting and closely related work of Mandler, 1968). Like the combined-cue and resistance-to-reinforcement tests, it can be used to isolate either excitatory or inhibitory stimuli. Furthermore, like the other tests, choice probes given after various amounts of discrimination training could help to determine whether S— becomes more or less inhibitory as training progresses—a question of relevance for D'Amato and Jagoda's (1961, 1962) suggestion (see also Deutsch and Biederman, 1965; Mackintosh, 1965; Biederman, 1968*b*) that the overlearning-discrimination reversal effect (ORE) is due to the nonmonotonicity of S—s negativity over the course of simultaneous discrimination learning. That is, S— presumably becomes less negative or aversive after prolonged training than it is after normal amounts of training, and therefore subjects are more likely to respond to S— during discrimination reversal after overtraining. For most of these authors, "negativity" or "aversiveness" seems to refer to certain motivational properties of what we have called an inhibitory stimulus.

Some of our current research indicates that a choice procedure can be successfully applied in individual subjects to the study of *dimensional control* by S+ and S— in a simultaneous discrimination, with generalization tests that are based on *choice* between two stimulus pairs (*e.g.*, the former S+ or S— and new stimuli) rather than on response *rate*. All previous work on EGs and IGs with the orthogonal method that was described in an earlier section of this article has involved free-operant successive discrimination training before generalization testing and has employed measures of response frequency during testing. To broaden understanding of basic discrimination processes, one would like to be able to compare the development and scope of control along dimensions of S+ and S— for simultaneous *vs.* successive discriminations, with either free-operant or discrete-trial techniques. Our choice procedure may also be worth reviewing because many workers in the field of instrumental learning who prefer to use discrete-trial simultaneous discriminations may be able to adapt it for their own purposes, because Spence's gradient-interaction theory of discrimination learning and transposition was originally developed for the analysis of simultaneous discriminations and may be best tested in that kind of situation, and because experimentation has frequently revealed important differences between simultaneous and successive discriminations (see Honig, 1962, and Riley, 1968).

In our research with simultaneous discriminations, pigeons in a discrete-trial two-key situation learn to choose between a blank white key and a white key bisected by a black vertical line. The keys are 2 in. apart, center to center. For some birds ("line-positive") the line is S+, whereas for other birds ("line-negative") the line is S—. Each session consists of 100 trials, with S+ randomly scheduled on the right or left key, and with grain reinforcement available for 25% of the trials on which a correct choice between the two keys occurs. Each trial is terminated either by a response to one of the keys or by the passage of 10 sec. The intertrial interval (blackout) averages 15 sec. Some subjects ("criterion") are tested in extinction after discrimination training to a level of at least 90% correct choices in one session. Other subjects ("overtrained") are given 600 more trials (*i.e.*, six more sessions) after reaching 90% correct and are then tested in extinction. During tests, the vertical training line (0°) is paired on each trial with one of three other line tilts (30°, 60°, and 90°, tilted in a clockwise direction from the vertical) or with the blank stimulus. The four different test pairs are presented once in each of 20 randomized blocks and the 0° line alternates randomly between the right and left keys. Just as during training, trials are terminated automatically after 10 sec if no response has occurred.

Individual data from the entire generalization tests of the first few subjects we have run on this procedure are shown in Fig. 3. The choice (%) data on the ordinate are of course based only on trials in which a response was made within the 10-sec limit, but all subjects responded on a great majority of the total 20 trials for each pair. There were two birds in each of the four main groups (line-positive, criterion; line-positive, overtrained; line-negative, criterion; line-negative, overtrained). Performance on the line *vs.* blank training pair was maintained at a very high level by all subjects in the four groups during testing in extinction. The line-positive subjects chose the 0° line on a great majority of the trials during the test, whereas the line-negative subjects

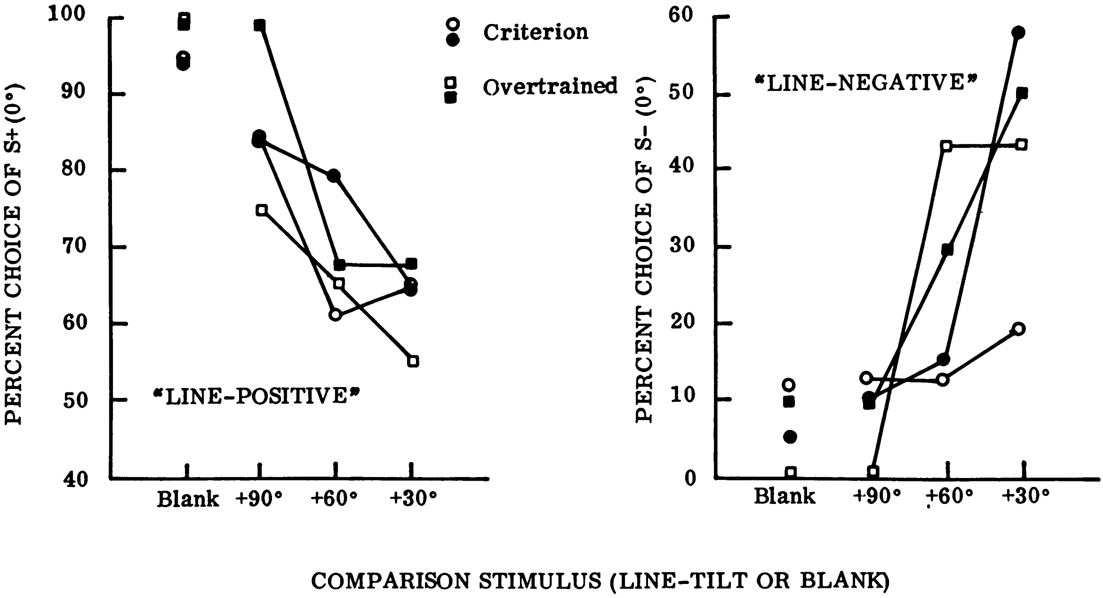

Fig. 3. Generalization gradients obtained from eight individual pigeons during a choice-test after learning of a simultaneous discrimination. Subjects in the "line-positive" condition had been either trained to criterion or overtrained on the discrimination, S+ (vertical line, 0°) *vs.* S— (blank). Subjects in the "line-negative" condition had been either trained to criterion or overtrained on the discrimination, S+ (blank) *vs.* S— (0°).

tended to avoid the 0° line and to choose the other stimulus. The per cent choice of the 0° line in the line-positive and line-negative subjects depended on the degree of physical similarity between the two line-tilts in a test pair. These preliminary data suggest that learning a simultaneous discrimination, like learning a successive discrimination, usually involves control along dimensions of both S+ and S—.[14]

V. OTHER RELATED PHENOMENA

This article has focused on the definition and measurement of (a) specific inhibitory stimuli and (b) inhibitory dimensional control. We think these two kinds of phenomena can be meaningfully distinguished and we have reviewed a number of experimental methods for detecting and analyzing them. There are several other behavioral effects that have traditionally been regarded as close relatives of the phenomena emphasized in this monograph. Sometimes these other effects are even accepted as evidence that "inhibition" is operating in a particular situation (see Blough and Millward, 1965, p. 72; Staddon, 1969, p. 483). A detailed review of research concerning these additional phenomena is beyond the scope of the present monograph,[15] but brief

---

[14] One potential problem in interpreting choice-test data comes from the observation that subjects may respond to stimulus novelty *per se*. Sutherland (1961) discussed this possibility, and Warren and McGonigle (1969) and others have noted that young children, monkeys, and cats often exhibit strong preferences for the more novel stimulus in a choice situation, whereas rats do not. The gradient in the line-negative condition might be explained on the basis of approach to novelty rather than avoidance of S—. However, other data from simultaneous discriminations in our lab suggest that pigeons may not typically respond to novelty *per se*. For example, following training on a line-tilt discrimination they consistently choose the S— line-tilt over a novel color and following training on a color discrimination they choose the S— color over a novel line stimulus.

[15] Other areas of related research not specifically covered in this article include several studies of (a) procedures that apparently do not produce clear or consistent nonhorizontal IGs (*e.g.*, "massed extinction" to a stimulus without interspersed reinforcement at another stimulus: Weisman and Palmer, 1969; mere reduction of response rate to a stimulus: see Weisman, 1970, and Nevin, 1968, for a summary of positive and negative results on this question, and (b) procedural parameters that influence the degree of control along an S— dimension, *e.g.*, greater amounts of training steepen IGs (Farthing and Hearst, 1968). These topics have been omitted, first, because pertinent findings have been somewhat unclear or sparse within the field of operant conditioning and, second, because Rescorla (1969*b*) has already published an excellent review of relevant prior research which, however, has been mainly carried out in a Pavlovian aversive-conditioning situation.

mention of some of them does seem appropriate.

1. Induction and contrast. Pavlov's analysis (1927, especially Lecture XI) of the phenomenon of induction stressed the role of inhibitory processes. For example, positive induction (the intensification of a conditioned response to a CS+ when the CS+ is closely preceded by a CS−) was attributed by Pavlov to specific aftereffects of an inhibitory event. In somewhat analogous fashion, recent findings concerning transient and sustained contrast in operant situations (*e.g.,* see Nevin and Shettleworth, 1966; Terrace, 1968) have often been interpreted in terms of interactions between excitatory and inhibitory effects. Although Terrace (*e.g.,* 1966*b*) has argued forcefully that the development and later maintenance of behavioral contrast are highly correlated with the development and later maintenance of inhibitory properties by S−, other workers (*e.g.,* see Frieman, 1969; D'Amato, 1970, p. 455; Ellis, 1970) have presented evidence and arguments that seem to shed doubt on Terrace's interpretation.

2. Peak shift. A recurrent theme in the operant literature on intradimensional discriminations (*e.g.,* Bloomfield, 1968, 1969; Hearst, 1969; Jenkins, 1965; Terrace, 1966*b*) is the notion that peak shift, the occurrence of maximum response strength not at S+ but at a dimensional value even further away from S−, is a symptom of inhibitory control. Hearst (1969) did take the position that peak shift necessarily involves an incremental gradient around S−. However, it seems too early to state with confidence whether this opinion is correct and whether S− itself must be an inhibitory stimulus (according to our proposed definition), rather than a neutral stimulus in order for peak shift to occur.

3. S− aversiveness (emotionality, negativity, frustration)..The idea that an inhibitory stimulus in an appetitive situation possesses aversive properties such that subjects (a) will work to escape it, or (b) will emit emotional or aggressive behavior in its presence, or (c) will have their ongoing behavior suppressed by its response-dependent presentation, has been the focus of several studies (see Terrace, 1966*a;* Wagner, 1969*a;* Rilling, Askew, Ahlskog, and Kramer, 1969). Some of this research has involved the administration of drugs or the monitoring of physiological measures to determine whether an appetitive S− does function in a manner similar to that of an aversive stimulus like electric shock. This suggestion that an inhibitory stimulus possesses certain negative motivational or emotional properties is an important possibility, which will doubtless receive even further experimental tests in the future.

4. Spontaneous recovery and disinhibition. The study of spontaneous recovery (the return of an extinguished response with the passage of time since original extinction) and disinhibition (the return of an extinguished response upon the presentation of a novel stimulus) is much more common in the field of Pavlovian conditioning than in operant research (but see Brimer, submitted for publication). Rescorla (1969*b*) pointed out that although the occurrence of these phenomena has been accepted by many investigators as evidence that mainly inhibitory processes are at work, this inference is really dependent on several rather dubious assumptions. Nevertheless, we have included the two phenomena here because future research may reveal relationships between (a) the degree to which a stimulus is inhibitory and (b) the susceptibility of behavior in the stimulus to the effects of the passage of time or to the presentation of novel stimuli.

All the above phenomena certainly merit further study in connection with inhibitory effects in operant conditioning. But, as a first step, we think it is necessary to establish precise definitions and standardized methods for isolating what we have called specific inhibitory stimuli and inhibitory dimensional control. Then it should become clear that these other phenomena are operationally different from the two stressed in this monograph. Empirical research seems necessary to determine the nature of the relationships among all these phenomena, which have often been assumed or presumed in the past.

One final point concerns the possibly naive or premature question: what specific mechanisms mediate the response reduction produced by an inhibitory stimulus? Does such reduction involve (a) the direct weakening of the original response, (b) interference produced by the actual acquisition of a new response (or responses) antagonistic to or competing with the original response, or (c) some general suppression of operant behavior? Such questions are related to traditional controversies such as the one between inhibition and interference theories of extinction (see Gleitman *et al.*, 1954; Kimble, 1961; Konorski, 1967). Perhaps some light would be cast on these possibilities by research that seeks to determine the specificity of inhibitory effects with respect to the original conditioned response. Thus, for example, in pigeons combined-cue or resistance-to-reinforcement tests are probably especially useful when a new response is involved (*e.g.*, treadle pressing, shuttling, or lever pressing), one which is topographically different from the key-pecking response that was reinforced and extinguished during prior discrimination training.

If the addition of the presumed inhibitory stimulus were to reduce not only key-pecking behavior in this example, but also a variety of topographically dissimilar operant responses, then it would be unlikely that either possibility (a) or (b) was correct. In that event, one might prefer possibility (c) and posit some central emotional or motivational state (see Rescorla and Solomon, 1967, and Premack, 1969, p. 136) as the basis for the response suppression—a hypothesis also suggested by the analogous general effects of a CER (conditioned suppression) stimulus on operant behavior (see Skinner, 1938, pp. 233-234; Blough and Millward, 1965, p. 72; Grusec, 1968). On the other hand, if the reductive effects of the stimulus were fairly specific to a given response, one might prefer a more peripheralistic interpretation and subsequently undertake research aimed at actual detection and measurement in subjects of a specific motor response (*e.g.*, in S−) that competes with the original conditioned response (*cf.* Migler and Millenson, 1969). These comments and suggestions are very vague and speculative, of course, but we think they imply tangible and significant problems for future research.

VI. CONCLUDING COMMENTS

In this monograph we proposed that a clear distinction be made between control of operant behavior by the presence *vs.* absence of a specific stimulus and control of operant behavior by different values along a dimension of that stimulus. An *inhibitory stimulus* was defined as a stimulus that develops during conditioning the capacity to decrease response strength below the level occurring in its absence. The term *inhibitory dimensional control* was used to refer to an empirical generalization gradient, in which behavior is stronger at test values far from an inhibitory stimulus than at values close by. According to this set of definitions, there is no inevitable relationship between inhibitory control by the presence *vs.* absence of a specific stimulus and control along a given dimension of that stimulus. A specific stimulus may qualify as inhibitory even when (a) no generalization tests have been performed, or (b) empirical gradients along some dimension of that stimulus are flat. Furthermore, dimensional control, demonstrated by an incremental gradient around some training stimulus value, would not be labelled inhibitory unless other tests showed that the specific training stimulus was itself inhibitory, as defined above. This latter provision is necessary because such an incremental gradient could conceivably occur if various dimensional values had differential excitatory effects and the training stimulus was neutral (*i.e.*, responding in its presence was equal to responding in its absence).

A variety of methods for distinguishing between inhibitory and neutral stimuli, and for obtaining generalization gradients around a stimulus associated with the extinction of an operant response, were reviewed and evaluated. Conflicts and contradictions in the interpretation and results of several important experiments in the field of generalization and discrimination seem to have occurred because of the different criteria various investigators used for deciding when a stimulus is inhibitory. We think that such contradictions and controversies may be avoided in the future if our distinction between two types of inhibitory control is more widely accepted.

Many more questions were raised than answered by our review and discussion. We wish to reiterate the tentative nature of many of our suggestions. However, our main goal has been to stimulate research in operant conditioning on these two aspects of inhibitory control. We recommended that in the future a variety of different methods for detecting and measuring inhibitory effects should constantly be compared, so as to avoid faulty conclusions resulting from the inevitable deficiencies of any single technique.

Like Rescorla (1969*b*), we have attempted to employ the concept of inhibition in a purely behavioral sense, devoid of neurological implications. Like Staddon (1969), we believe that inhibitory phenomena have been unjustifiably neglected in the field of operant conditioning. If Pavlov had been a little less excessive in his neurological speculations, and Skinner a little less excessive in his demands for parsimony, the experimental analysis of inhibitory effects might be a more popular area of operant research today. There seems no good reason why neglect of the topic should continue.

REFERENCES

Bass, M. J. and Hull, C. L. Irradiation of a tactile conditioned reflex in man. *Journal of Comparative Psychology,* 1934, **17,** 47-65.

Beale, I. L. and Corballis, M. C. Beak shift: An explanation for interocular mirror-image reversal in pigeons. *Nature,* 1968, **220,** 82-83.

Bernheim, J. W. Comment. *Psychonomic Science,* 1968, **11,** 327.

Besley, S. and Hearst, E. *Inhibitory gradients obtained with nondifferential reinforcement during generalization testing.* Paper presented at the meeting of the Psychonomic Society, St. Louis, 1969.

Biederman, G. B. Stimulus control in errorless discrimination. *Psychonomic Science,* 1967, **9,** 375-376.

Biederman, G. B. Reply to Bernheim. *Psychonomic Science,* 1968, **11,** 328. (*a*)

Biederman, G. B. Stimulus function is simultaneous discrimination. *Journal of the Experimental Analysis of Behavior,* 1968, **11,** 459-463. (*b*)

Bignami, G. Stimulus control: Contrast, peak shift, and Konorskian models. *Psychonomic Science,* 1968, **11,** 249-250.

Bloomfield, T. M. Discrimination learning in animals: an analysis through side-effects. *Nature,* 1968, **217,** 929-930.

Bloomfield, T. M. Behavioral contrast and the peak shift. In R. M. Gilbert and N. S. Sutherland (Eds.), *Animal discrimination learning.* London: Academic Press, 1969. Pp. 215-241.

Blough, D. S. The study of animal sensory processes by operant methods. In W. K. Honig (Ed.), *Operant behavior: areas of research and application.* New York: Appleton-Century-Crofts, 1966. Pp. 345-379.

Blough, D. S. Generalization gradient shape and summation in steady-state tests. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 91-104.

Blough, D. S. and Millward, R. B. Learning: operant conditioning and verbal learning. *Annual Review of Psychology,* 1965, **16,** 63-94.

Bower, G. and Grusec, T. Effect of prior Pavlovian discrimination training upon learning an operant discrimination. *Journal of the Experimental Analysis of Behavior,* 1964, **7,** 401-404.

Brimer, C. J. *Disinhibition of an operant response.* (Submitted for publication).

Brown, P. L. and Jenkins, H. M. Conditioned inhibition and excitation in operant discrimination learning. *Journal of Experimental Psychology,* 1967, **75,** 255-266.

Bull, J. A. and Overmier, J. B. Additive and subtractive properties of excitation and inhibition. *Journal of Comparative and Physiological Psychology,* 1968, **66,** 511-514.

Carlton, P. L. Cholinergic mechanisms in the control of behavior by the brain. *Psychological Review,* 1963, **70,** 19-39.

Carlton, P. L. and Vogel, J. R. Habituation and conditioning. *Journal of Comparative and Physiological Psychology,* 1967, **63,** 348-351.

Catania, A. C. Concurrent performances: inhibition of one response by reinforcement of another. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 731-744.

Catania, A. C. and Gill, C. A. Inhibition and behavioral contrast. *Psychonomic Science,* 1964, **1,** 257-258.

Cornell, J. M. and Strub, H. A technique for demonstrating the inhibitory function of S—. *Psychonomic Science,* 1965, **3,** 25-26.

D'Amato, M. R. *Experimental psychology: methodology, psychophysics, and learning.* New York: McGraw-Hill, 1970.

D'Amato, M. R. and Jagoda, H. Analysis of the role of overlearning in discrimination reversal. *Journal of Experimental Psychology,* 1961, **61,** 45-50.

D'Amato, M. R. and Jagoda, H. Overlearning and position reversal. *Journal of Experimental Psychology,* 1962, **64,** 117-122.

Desiderato, O. Generalization of excitation and inhibition in control of avoidance responding by Pavlovian CSs in dogs. *Journal of Comparative and Physiological Psychology,* 1969, **68,** 611-616.

Deutsch, J. Discrimination learning and inhibition. *Science,* 1967, **156,** 988.

Deutsch, J. A. and Biederman, G. B. The monotonicity of the negative stimulus during learning. *Psychonomic Science,* 1965, **3,** 391-392.

Diamond, S., Balvin, R. S., and Diamond, F. R. *Inhibition and choice.* New York: Harper and Row, 1963.

Douglas, R. J. The hippocampus and behavior. *Psychological Bulletin,* 1967, **67,** 416-442.

Eckerman, D. A. Stimulus control by part of a complex S—. *Psychonomic Science,* 1967, **7,** 299-300.

Ellis, W. R. Role of stimulus sequences in stimulus discrimination and stimulus generalization. *Journal of Experimental Psychology,* 1970, **83,** 155-163.

Farthing, G. W. *Effects of the cue-reinforcement correlation on stimulus control in compound-stimulus learning.* Unpublished doctoral dissertation, University of Missouri, 1969. Pp. 1-102.

Farthing, G. W. and Hearst, E. Generalization gradients of inhibition after different amounts of training. *Journal of the Experimental Analysis of Behavior,* 1968, **11,** 743-752.

Farthing, G. W. and Hearst, E. Attention in the pigeon: testing with compounds or elements. *Learning and Motivation,* 1970, **1,** 65-78.

Farthing, G. W., Koresko, M. B., and Hearst, E. *Generalization, peak shift, and dynamism along a click-frequency dimension in pigeons.* Paper presented at the meeting of the Midwestern Psychological Association, Chicago, 1968.

Frieman, J. *A correlational analysis of successive discrimination learning.* Paper presented at the meeting of the Midwestern Psychological Association, Chicago, 1969.

Gleitman, H., Nachmias, J., and Neisser, U. The S-R reinforcement theory of extinction. *Psychological Review,* 1954, **61,** 23-33.

Graham, F. K. Conditioned inhibition and conditioned excitation in transfer of discrimination. *Journal of Experimental Psychology,* 1943, **33,** 351-368.

Gray, J. A. Stimulus intensity dynamism. *Psychological Bulletin,* 1965, **63,** 180-196.

Grusec, T. The peak shift in stimulus generalization: equivalent effects of errors and noncontingent shock. *Journal of the Experimental Analysis of Behavior,* 1968, **11,** 239-249.

Guttman, N. The pigeon and the spectrum and other perplexities. *Psychological Reports,* 1956, **2,** 449-460.

Guttman, N. Laws of behavior and facts of perception. In S. Koch (Ed.), *Psychology: a study of a science,* Vol. 5. New York: McGraw-Hill, 1963. Pp. 114-178.

Guttman, N. Effects of discrimination formation on generalization measured from a positive-rate baseline. In D. Mostofsky (Ed.), *Stimulus generalization.* Palo Alto: Stanford University Press, 1965. Pp. 210-217.

Guttman, N. and Kalish, H. I. Discriminability and stimulus generalization. *Journal of Experimental Psychology,* 1956, **51,** 79-88.

Hammond, L. J. A traditional demonstration of the active properties of Pavlovian inhibition using differential CER. *Psychonomic Science,* 1967, **9,** 65-66.

Hammond, L. J. Retardation of fear acquisition by a previously inhibitory CS. *Journal of Comparative and Physiological Psychology,* 1968, **66,** 756-759.

Hanson, H. M. Effects of discrimination training on stimulus generalization. *Journal of Experimental Psychology,* 1959, **58,** 321-334.

Hearst, E. Discrimination learning as the summation of excitation and inhibition. *Science,* 1968, **162,** 1303-1306.

Hearst, E. Excitation, inhibition, and discrimination learning. In N. J. Mackintosh and W. K. Honig (Eds.), *Fundamental issues in associative learning.* Halifax, Canada: Dalhousie University Press, 1969. Pp. 1-41.

Hearst, E. and Poppen, R. Steepened generalization gradients after massed extinction to the CS. *Psychonomic Science,* 1965, **2,** 83-84.

Hendry, D. P. Concluding commentary. In D. P. Hendry (Ed.), *Conditioned reinforcement.* Homewood, Illinois: The Dorsey Press, 1969. Pp. 385-402.

Herrnstein, R. J. Superstition: A corollary of the principles of operant conditioning. In W. K. Honig (Ed.), *Operant behavior: areas of research and application.* New York: Appleton-Century-Crofts, 1966. Pp. 33-51.

Hoffman, H. S. Stimulus factors in conditioned suppression. In B. A. Campbell and R. M. Church (Eds.), *Punishment and aversive behavior.* New York: Appleton-Century-Crofts, 1969. Pp. 185-234.

Honig, W. K. Generalization of extinction on the spectral continuum. *Psychological Record,* 1961, **11,** 269-278.

Honig, W. K. Prediction of preference, transposition, and transposition-reversal from the generalization gradient. *Journal of Experimental Psychology,* 1962, **64,** 239-248.

Honig, W. K. The role of discrimination training in the generalization of punishment. *Journal of the Experimental Analysis of Behavior,* 1966, **9,** 377-384.

Honig, W. K. Attentional factors governing the slope of the generalization gradient. In R. M. Gilbert and N. S. Sutherland (Eds.), *Animal discrimination learning.* London: Academic Press, 1969. Pp. 35-62.

Honig, W. K. Attention and the modulation of stimulus control. In D. Mostofsky (Ed.), *Attention: contemporary studies and analyses.* New York: Appleton-Century-Crofts, 1970, *in press.*

Honig, W. K., Boneau, C. A., Burstein, K. R., and Pennypacker, H. S. Positive and negative generalization gradients obtained after equivalent training conditions. *Journal of Comparative and Physiological Psychology,* 1963, **56,** 111-116.

Honig, W. K., Thomas, D. R., and Guttman, N. Differential effects of continuous extinction and discrimination training on the generalization gradient. *Journal of Experimental Psychology,* 1959, **58,** 145-152.

Hovland, C. I. The generalization of conditioned responses: I. The sensory generalization of conditioned responses with varying frequencies of tone. *Journal of General Psychology,* 1937, **17,** 125-148.

Hull, C. L. *Principles of behavior.* New York: D. Appleton-Century Co., 1943.

Hull, C. L. Stimulus intensity dynamism (V) and stimulus generalization. *Psychological Review,* 1949, **56,** 67-76.

Hull, C. L. *A behavior system.* New Haven: Yale University Press, 1952.

James, J. P. and Mostoway, W. W. Stimulus intensity dynamism: effect of nonreinforcement of the highest intensity stimulus. *Canadian Journal of Psychology,* 1969, **23,** 49-55.

Jenkins, H. M. Generalization gradients and the concept of inhibition. In D. Mostofsky (Ed.), *Stimulus generalization.* Palo Alto: Stanford University Press, 1965. Pp. 55-61.

Jenkins, H. M. and Harrison, R. H. Generalization gradients of inhibition following auditory discrimination learning. *Journal of the Experimental Analysis of Behavior,* 1962, **5,** 435-441.

Jensen, A. R. On the reformulation of inhibition in Hull's system. *Psychological Bulletin,* 1961, **58,** 274-298.

Johnson, D. F., Kinder, B. N., and Scarboro, G. N. Determination of selective stimulus control by part of a compound S−. *Psychonomic Science,* 1969, **17,** 7-8.

Kalish, H. I. and Haber, A. Generalization: I. Generalization gradients from single and multiple stimulus points. II. Generalization of inhibition. *Journal of Experimental Psychology,* 1963, **65,** 176-181.

Keiffer, J. D. Differential response rates correlated with the presence of "neutral" stimuli. *Journal of the Experimental Analysis of Behavior,* 1965, **8,** 227-229.

Kimble, D. P. Hippocampus and internal inhibition. *Psychological Bulletin,* 1968, **70,** 285-295.

Kimble, G. *Hilgard and Marquis' conditioning and learning.* 2nd ed.; New York: Appleton-Century-Crofts, 1961.

Kling, J. W. Generalization of extinction of an instrumental response to stimuli varying in the size dimension. *Journal of Experimental Psychology,* 1952, **44,** 339-346.

Konorski, J. *Integrative activity of the brain.* Chicago: University of Chicago Press, 1967.

Logan, F. A note on stimulus intensity dynamism (V). *Psychological Review,* 1954, **61,** 77-80.

Lubow, R. E. Latent inhibition: Effects of frequency of nonreinforced pre-exposure of the CS. *Journal of Comparative and Physiological Psychology,* 1965, **60,** 454-457.

Lyons, J. Stimulus generalization as a function of discrimination learning with and without errors. *Science,* 1969, **163,** 490-491. (*a*)

Lyons, J. Stimulus generalization along the dimension of S+ as a function of discrimination learning with and without error. *Journal of Experimental Psychology,* 1969, **81,** 95-100. (*b*)

Mackintosh, N. J. Selective attention in animal discrimination learning. *Psychological Bulletin,* 1965, **64,** 124-150.

Mandler, J. M. Overtraining and the use of positive and negative stimuli in reversal and transfer. *Journal of Comparative and Physiological Psychology,* 1968, **66,** 110-115.

Marsh, G. *Prediction of the peak shift from gradients of excitation and inhibition.* Paper presented at the meeting of the Psychonomic Society, St. Louis, 1968.

Migler, B. and Millenson, J. R. Analysis of response rates during stimulus generalization. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 81-87.

Morse, W. H. and Skinner, B. F. A second type of superstition in the pigeon. *American Journal of Psychology,* 1957, **70,** 308-311.

Mostofsky, D. (Ed.) *Stimulus generalization.* Palo Alto: Stanford University Press, 1965.

Nevin, J. A. Differential reinforcement and stimulus control of not responding. *Journal of the Experimental Analysis of Behavior,* 1968, **11,** 715-726.

Nevin, J. A. and Shettleworth, S. J. An analysis of contrast effects in multiple schedules. *Journal of the Experimental Analysis of Behavior,* 1966, **9,** 305-315.

Newman, F. L. and Baron, M. R. Stimulus generalization along the dimension of angularity. *Journal of Comparative and Physiological Psychology,* 1965, **60,** 59-63.

Newman, F. L. and Benefield, R. L. Stimulus control, cue utilization, and attention: Effects of discrimination training. *Journal of Comparative and Physiological Psychology,* 1968, **66,** 101-104.

Pavlov, I. P. *Conditioned reflexes.* London: Oxford University Press, 1927.

Perkins, C. C. The relation between conditioned stimulus intensity and response strength. *Journal of Experimental Psychology,* 1953, **46,** 225-231.

Premack, D. On some boundary conditions of contrast. In J. Tapp (Ed.), *Reinforcement and behavior.* New York: Academic Press, 1969. Pp. 120-145.

Prokasy, W. F. Classical eyelid conditioning: Experimenter operations, task demands, and response shaping. In W. F. Prokasy (Ed.), *Classical conditioning.* New York: Appleton-Century-Crofts, 1965. Pp. 208-225.

Reberg, D. and Black, A. H. Compound testing of individually conditioned stimuli as an index of excitatory and inhibitory properties. *Psychonomic Science,* 1969, **17,** 30-31.

Rescorla, R. A. Predictability and number of pairings in Pavlovian fear conditioning. *Psychonomic Science,* 1966, **4,** 383-384.

Rescorla, R. A. Pavlovian conditioning and its proper control procedures. *Psychological Review,* 1967, **74,** 71-80.

Rescorla, R. A. Conditioned inhibition of fear. In N. J. Mackintosh and W. K. Honig (Eds.), *Fundamental issues in associative learning.* Halifax, Canada: Dalhousie University Press, 1969. Pp. 65-89. (*a*)

Rescorla, R. A. Pavlovian conditioned inhibition. *Psychological Bulletin,* 1969, **72,** 77-94. (*b*)

Rescorla, R. A. and Lolordo, V. W. Inhibition of avoidance behavior. *Journal of Comparative and Physiological Psychology,* 1965, **59,** 406-412.

Rescorla, R. A. and Solomon, R. L. Two-process learning theory: Relationships between Pavlovian conditioning and instrumental learning. *Psychological Review,* 1967, **74,** 151-182.

Reynolds, G. S. Contrast, generalization and the process of discrimination. *Journal of the Experimental Analysis of Behavior,* 1961, **4,** 289-294.

Reynolds, G. S. Operant extinction near zero. *Journal of the Experimental Analysis of Behavior,* 1964, **7,** 173-176.

Riley, D. A. *Discrimination learning.* Boston: Allyn and Bacon, 1968.

Rilling, M., Askew, H. R., Ahlskog, J. E., and Kramer, T. J. Aversive properties of the negative stimulus in a successive discrimination. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 917-932.

Rodnick, E. H. Does the interval of delay of conditioned responses possess inhibitory properties? *Journal of Experimental Psychology,* 1937, **20,** 507-527.

Schwartzbaum, J. S. and Kellicut, M. H. Inverted generalization gradients about a nonreinforced stimulus. *Psychological Reports,* 1962, **11,** 791-792.

Seligman, M. E. P. Control group and conditioning: a comment on operationism. *Psychological Review,* 1969, **76,** 484-491.

Siegel, S. Generalization of latent inhibition. *Journal of Comparative and Physiological Psychology,* 1969, **69,** 157-159.

Skinner, B. F. *The behavior of organisms.* New York: D. Appleton-Century Co., 1938.

Spence, K. W. The nature of discrimination learning in animals. *Psychological Review,* 1936, **43,** 427-449.

Spence, K. W. The differential response in animals to stimuli varying within a single dimension. *Psychological Review,* 1937, **44,** 430-444.

Staddon, J. E. R. Inhibition and the operant. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 481-487.

Sutherland, N. S. Shape discrimination by animals. *Experimental Psychology Society Monographs,* 1961, **1,** 1-68.

Switalski, R. W., Lyons, J., and Thomas, D. R. Effects of interdimensional training on stimulus generalization. *Journal of Experimental Psychology,* 1966, **72,** 661-666.

Terrace, H. S. Stimulus control. In W. K. Honig (Ed.), *Operant behavior: areas of research and ap-*

plication. New York: Appleton-Century-Crofts, 1966. Pp. 271-344. (*a*)

Terrace, H. S. Discrimination learning and inhibition. *Science,* 1966, **154,** 1677-1680. (*b*)

Terrace, H. S. Reply to Deutsch. *Science,* 1967, **156,** 988-989.

Terrace, H. S. Discrimination learning, the peak shift, and behavioral contrast. *Journal of the Experimental Analysis of Behavior,* 1968, **11,** 727-741.

Thomas, D. R. and Lanier, W. G. A comparison of stimulus generalization of tendencies to respond and not to respond. *Psychological Record,* 1962, **12,** 61-65.

Touchette, P. E. Tilted lines as complex stimuli. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 211-214.

Trapold, M. A. Reversal of an operant discrimination by noncontingent discrimination reversal training. *Psychonomic Science,* 1966, **4,** 247-248.

Trapold, M. A., Lawton, G. A., Dick, R. A., and Gross, D. M. Transfer of training from differential classical to differential instrumental conditioning. *Journal of Experimental Psychology,* 1968, **76,** 568-573.

Vetter, G. and Hearst, E. Generalization and discrimination of shape orientation in the pigeon. *Journal of the Experimental Analysis of Behavior,* 1968, **11,** 753-765.

vom Saal, W. and Jenkins, H. M. Blocking the development of stimulus control. *Learning and Motivation,* 1970, **1,** 52-64.

Wagner, A. R. Frustrative-nonreward: a variety of punishment? In B. A. Campbell and R. M. Church (Eds.), *Punishment and aversive behavior.* New York: Appleton-Century-Crofts, 1969. Pp. 157-181. (*a*)

Wagner, A. R. Incidental stimuli and discrimination learning. In R. M. Gilbert and N. S. Sutherland (Eds.), *Animal discrimination learning.* London: Academic Press, 1969. Pp. 83-111. (*b*)

Wagner, A. R., Logan, F. A., Haberlandt, K., and Price, T. Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology,* 1968, **76,** 171-180.

Warren, J. M. and McGonigle, B. Attention theory and discrimination learning. In R. M. Gilbert and N. S. Sutherland (Eds.), *Animal discrimination learning.* London: Academic Press, 1969. Pp. 113-136.

Weisman, R. G. Some determinants of inhibitory stimulus control. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 443-450.

Weisman, R. G. Determinants of inhibitory stimulus control. In J. H. Reynierse (Ed.), *Current issues in animal learning.* Lincoln: University of Nebraska Press, 1970, *in press.*

Weisman, R. G. and Palmer, J. A. Factors influencing inhibitory stimulus control: discrimination training and prior non-differential reinforcement. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 229-237.

Weiss, S. J. Free-operant compounding of variable-interval and low rate discriminative stimuli. *Journal of the Experimental Analysis of Behavior,* 1967, **10,** 535-540.

Wilton, R. N. and Godbout, R. C. Stimulus control in discrimination learning. *British Journal of Psychology,* 1970, **61,** 109-114.

Yarczower, M. Behavioral contrast and inhibitive stimulus control. *Psychonomic Science,* 1970, **18,** 1-3.

Received 29 December 1969.