# 2

# *Characterization and Classification of Choice Experiments*[1,2]

Robert R. Bush
*University of Pennsylvania*

Eugene Galanter
*University of Washington*

R. Duncan Luce
*University of Pennsylvania*

*1. This work was supported in part by grants NSF G-8864, G-14839, and G-17637 from the National Science Foundation to the University of Pennsylvania and by the Office of Naval Research, Contract NONR-551(37). Reproduction in whole or part is permitted for any purpose of the United States Government.*

*2. We wish to thank Ward Edwards, Francis W. Irwin, Elizabeth F. Shipley, and Richard L. Solomon for their extremely helpful comments and criticisms of drafts of this material and our students whose puzzlement over various paragraphs or sections revealed some fuzziness in our thinking, which, hopefully, is less apparent now.*

# Contents

# *Characterization and Classification of Choice Experiments*

A large part of the naturalistic analysis and an even larger part of the successful experimental and mathematical analysis of human and animal behavior involves the notion of choice. Whenever an organism, either in its natural environment or in the laboratory, is confronted with a set of two or more mutually exclusive courses of action—responses—and it selects or performs one of them, we say that a choice has occurred. If the word "choice" suggests a conscious decision by the organism, that is unfortunate. To be sure, some choices may reasonably be called "conscious"; for example, in a psychophysical discrimination experiment a person says "louder" after having "consciously" decided that the second tone is louder than the first. But decisions are not usually called "conscious" when a rat "chooses" to turn right or left in a T-maze, and we want to call that a choice. The term is strictly behavioristic. What an experimenter elects to observe determines, in part, whether or not choices are made. The scientific significance of these observations, however, probably rests on the wisdom of the experimenter's choices.

By a choice experiment, then, we mean one in which there is a set of two or more empirically defined alternative responses from which the subject chooses just one whenever he is given an opportunity to do so. We call these opportunities "trials." The set of response alternatives is usually finite, often having only two or three elements, but occasionally infinite sets are employed.

In all theories about experimental choice behavior a measure function is imposed over the response set; most commonly it is a probability measure, but other measures have been used as well (e.g., Luce, 1959; Audley, 1960). But neither the kinds of measures occurring in a theory nor the types of behavioral observations actually made (response frequency or response time, for example) have anything to do with classifying an experiment as a choice design.

In psychophysics and learning the vast majority of experiments performed are choice ones, as we are using the term. Signal detection, discrimination, T-maze learning, and binary prediction studies are obvious examples. Nevertheless, there are some notable exceptions. Simple reaction time experiments are not of the choice variety because only a single experimentally defined response is made on every trial. For the same reason, animal-runway and many operant conditioning experiments are not choice designs. Some analyses of reaction time and runway studies (McGill, Chapter 6; Estes, 1950; Bush & Mosteller, 1955, Chapter 14) postulate unobservable responses in order that existing choice models can be used to analyze the phenomena, but this does not alter our classification of the experiments; the experimentally identified response set contains only one element, and so they are not choice experiments.

At the other extreme the obvious identification of responses in some experiments leads to extremely large response sets. For example, in magnitude estimation studies the response set is well defined but nondenumerable, for the subject is restricted only to choosing nonnegative real numbers. It seems clear that only much smaller, finite, but unknown subsets are actually available to most subjects. Similarly, in verbal conditioning experiments in which the subject is asked to emit any word that comes to mind and the experimenter reinforces a certain class of words, such as plural nouns, the set of possible responses is finite and well defined but large. Again, however, the set of all words is much larger than the real but unknown response set which at most is the subject's vocabulary.

It is very difficult to see how to exclude either of these studies formally from the class of choice experiments, yet few theoreticians feel that there is much hope in trying to analyze them exactly as we do choice experiments with small response sets except possibly when some appreciable redefinition of the responses can be made. For example, in verbal conditioning we might define as one response the class of plural nouns and as the other all remaining words, in which case we then view the experiment as an example of two-response learning. Admittedly, these two-response classes are most heterogeneous, but that is in no way unique to this experiment. We have little assurance—indeed, there is evidence to the contrary—that responses such as left turns in a T-maze are homogeneous. The conclusion, then, is that any experiment that can be viewed as a choice experiment can always be so viewed in several different ways. For many experiments, however, the identification of the response set seems so obvious and unambiguous that there is little to be decided. For example, no one has suggested any radical reduction in the responses for a magnitude estimation experiment, and theorists have been obliged to work with continuous response theories rather than with the simpler discrete theories postulated for most other psychophysical experiments.

Another troublesome class of experiments contains those for which the response set is well defined and small, but a trial is terminated either when one of these responses occurs or when a fixed time has elapsed following the stimulus presentation. Thus no prescribed response need occur on a trial.

Examples are (1) avoidance training in which an animal avoids or escapes a shock by a particular response or, failing escape in a certain time, the shock is automatically terminated; (2) free-recall verbal learning in which a subject either recalls a word or not in a specified time; and (3) classical Pavlovian conditioning in which the (conditioned) response may or may not occur, following the presentation of the conditioned stimulus, within the observation period of a trial.

If the prescribed responses are taken to form the response set, then these are not choice experiments because a response need not always occur during a trial. If, however, the response set is augmented by including what may be called the "null" response, namely, any behavior not previously prescribed as a response, then formally it is a choice experiment. Some theorists are uneasy about treating the null response as a response because they do not feel that a choice is made in the same sense in the two cases. But this is little more than a feeling, and the ultimate decision about the appropriate identifications of responses probably can be made only when we know whether these experiments demand inherently different theories from those used to account for behavior in what are clearly choice experiments.

In branches of psychology other than psychophysics, learning, and preference studies the choice experiment paradigm has been little used. Although mathematics is employed extensively in psychometrics, for example, the concern there has not been primarily with behavior as such, but rather behavioral data are used to establish "cultural" scales of social variables. Similarly, modern psychological studies of language are less concerned with the behavior of the speaker and listener than with the formal structure of language itself (see Chapters 11, 12, and 13). Some social psychological experiments, such as those of small group behavior, are designed with an implicit or explicit choice paradigm in mind, but many others are not (see Chapter 14).

Nonetheless, a sizable fraction of all the work done in mathematical psychology is currently concerned with choice experiments, and a considerable unity and systematization has evolved over the years. For this reason, a chapter devoted to a fairly careful characterization and classification of choice experiments seemed useful to us. What we have to say is directly relevant to the expositions found in Chapters 3, 4, 5, 8, 9, 10; it is largely irrelevant to most of the other chapters in this work.

# 1. THE ABSTRACT STRUCTURE OF A CHOICE EXPERIMENT

## 1.1 The Classical Stimulus-Response Paradigm

For several decades stimulus-response psychology has been based mostly on a single skeletal design which identifies three significant events that occur in experimental trials. In brief, it is

$$\text{stimulus} \rightarrow \text{response} \rightarrow \text{outcome.}$$

A stimulus—be it called that, an environmental situation, signal, collection of cues, or some other similar term—is presented by the experimenter. Next, the subject whose behavior is under study makes a response or performs an act which may be motor, verbal, or physiological (e.g., GSR). This is followed by experimenter-determined events, whatever they may be, which we call outcomes. Outcomes are sometimes termed rewards and punishments, environmental events, or payoffs; the first two terms are, however, also used as theoretical notions in some theories. Outcomes may be the presentation of tangible objects or substances, or they may be signals that convey information. Sometimes they are omitted from the design, as in much of classical psychophysics, in many modern studies on the scaling of stimuli, and in classical conditioning.

The paradigm is general, and so its identifications are bound to be ambiguous, but little violence is done to the basic ideas if we say stimuli are particular aspects of the environment during the period of the trial that precedes the response and outcomes are particular aspects during the period of the trial that follows the response. Stimulus-response behavior theories are concerned with how the subject sets up a "connection" between his responses and the stimuli, but we need not deal with this problem here. Our present goal is to characterize the experiment, not to describe or explain the subject's behavior. Of course, later chapters in this work are primarily concerned with behavior theory as such.

The S-R-O framework has had a controlling influence on essentially all mathematical learning models developed in the last decade. Estes (1950) in one of the founding papers made the parts of the paradigm quite explicit. Bush and Mosteller (1951a, 1951b, 1955) and Bush, Mosteller, and Thompson (1954) made them more or less explicit, depending upon the problem discussed. The psychophysical literature, both experimental and mathematical, seems to have taken the paradigm for granted, except for the traditional lack of concern about outcomes. Galanter and Miller (1960) pointed out and objected to this strong S-R influence on mathematical psychology. Rather than describe or discuss their objections here, we shall attempt instead to formalize and clarify the S-R-O paradigm and to present a somewhat new classification of experiments. This necessarily leads to some new nomenclature.

Because we want a scheme that applies to several different substantive areas, each of which has its own conventional notations for the same notions, we are forced to create a compromise notation that is not entirely consistent with any of the existing ones. In fact, we have departed more than is strictly necessary in order to satisfy certain conventions that seem useful: insofar as possible we denote sets by italic capital Latin letters and their elements by the corresponding letters in lower case; we use as the names of sets the first letter of the word that describes the elements; and we denote functions by the lower-case Greek letters corresponding to the symbol for the set of elements constituting the range of the function. If this or any other consistent notation were accepted and used in the several fields, certain advantages of communication would accrue.

## 1.2 Stimulus Presentations

The flow diagram of Fig. 1, which is explained fully in the following pages, gives much of our scheme for a choice experiment. A more detailed summary appears in Table 1, p. 96.

In all of the experiments we shall consider, the events that occur are partitioned into a sequence of *trials*. The sequence of trials may therefore

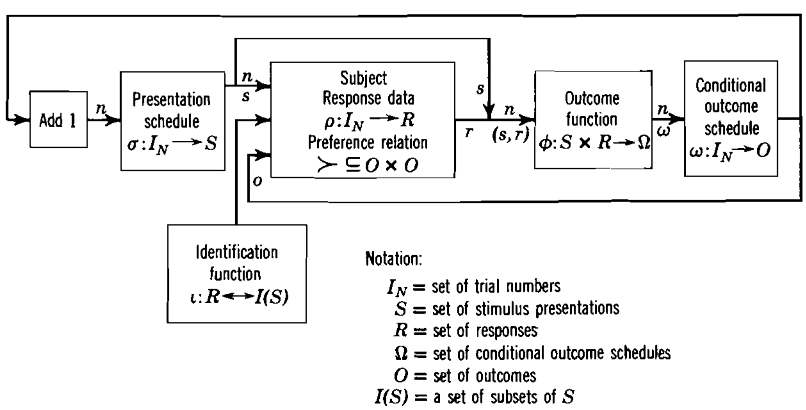


Fig. 1. Flow diagram of a choice experiment.

be identified with the sequence $I_N = \langle 1, 2, \ldots, n, \ldots, N \rangle$ of the first $N$ integers, where $N$ is the total number of trials in the experimental run. Although in some experiments $N$ is a random variable, sometimes dependent upon the subject's behavior, it is often fixed in advance.

We enter the flow diagram at the left box which issues the trial number $n$. Given this, the experimenter enters a table or schedule, represented by the next box, which he has prepared in advance of the experiment, to find out which of several possible stimulus presentations is actually to be presented to the subject. This schedule may or may not depend upon the previous behavior of the subject.

Let $S$ denote the set of possible *stimulus presentations*, its typical element being denoted $s$, or sometimes $s_i$, possibly with a prime or superscript. For example, in a simple discrimination study one of the presentations might be a 100-ms, 60-db, 1000-cps tone followed in 20 ms by a 100-ms, 62-db, 1000-cps tone; in such a design the set $S$ consists of all the ordered pairs of tones that are presented to the subject during the experimental run. Later, in Sec. 2.1, we discuss more fully the structure of $S$. In a simple learning experiment the set $S$ typically has only a single element; the same stimulus presentation occurs on every trial. (The reader should not confuse the elements of our set $S$ with the "stimulus elements" of the stimulus-sampling models described in Chapter 10.) In so-called "discrimination learning" experiments $S$ usually has two elements, each of which is presented on half the trials.

The *presentation schedule*, then, is simply a function

$$\sigma : I_N \rightarrow S.$$

In many experiments $\sigma$ is decided upon by some random device with the property that

$$P(s) = \Pr\,[\sigma(n) = s]$$

is a function of $s \in S$ but not of the trial number $n$. We call these *simple random presentation schedules* and refer to $P(s)$ as the *presentation probability* of $s$. Of course,

$$\sum_{s \in S} P(s) = 1.$$

Although many experiments use simple random schedules, some do not. For example, the schedules used in studying the perception of periodic sequences as well as those having probabilities conditional upon the response are not simple random.

### 1.3 Responses

Following each presentation, the subject is required to choose one response alternative from a given set $R$ of two or more possible *responses*. Typically, we use $r$ or $r_j$ for elements of $R$. By his choice, the subject assigns a response alternative, say $\rho(n) \in R$, to the trial number $n$. This is to say, his responses generate a function $\rho$, where

$$\rho : I_N \to R.$$

This function we call the *response data* of the experiment. For example, in a discrimination experiment the subject may be asked to state whether the second of two tones forming the presentation is louder or softer than the first, in which case $R = \{\text{louder, softer}\}$, and $\rho$ is nothing more than the abstract representation of the data sheet with its assignment of "louder" or "softer" to each of the trial numbers.

In many contemporary models a probability mechanism is assumed to underlie the generation of the response data. Moreover, in spite of some evidence to the contrary, most analyses of psychophysical data and most psychophysical models assume that the responses are independent in the sense that they depend directly upon only the immediately preceding stimulus presentation. Thus the postulated probabilities are

$$p_n(r \mid s) = \Pr[\rho(n)=r \mid \sigma(n)=s],$$

where

$$\sum_{r \in R} p_n(r \mid s) = 1, \qquad (s \in S,\ n \in I_N).$$

### 1.4 Outcome Structure

For a time, let us ignore the preference relation located in the box in Fig. 1 marked "subject" and the identification function feeding in from below and turn to the two boxes to the right. These are intended to represent the mechanism for feeding back information and payoffs, if any, to the subject. In many psychophysical experiments today, and in almost all before 1950, this structure simply is absent, but for reasons that will become apparent later many psychophysicists now feel that it should be an integral part of the design of many experiments, as it is in most learning experiments.

The set $O$, which we call the set of possible experimental *outcomes* (in learning theory certain outcomes are called reinforcers), consists of the direct, immediate, experimenter-controlled consequences to the subject which depend in part upon his behavior. We let $o$ or $o_k$ denote typical elements of $O$. If there is only feedback about the correctness of the responses, then $O = \{\text{correct, incorrect}\}$; if there are payoffs as well, such as 5¢ for a correct response and −5¢ for an incorrect one, then $O = \{5¢ \text{ and correct}, -5¢ \text{ and incorrect}\}$. This last set is usually just written $\{5¢, -5¢\}$, on the assumption that the sign of the payoff indicates which response is correct. This is not, however, a necessary correlation (see Sec. 3.3).

Although in psychophysics it has been usual for the outcomes, when they are used at all, to be determined uniquely by the presentation-response pair, in learning and preference studies matters have not been so simple. The more general scheme could and some day may very well be used in psychophysics also. Instead of selecting an outcome directly, the presentation $s$ and the response $r$ select a function, denoted as $\omega$ or $\omega_l$, from a set $\Omega$ of such functions. In turn, this function assigns an outcome to the trial number, that is, if $\omega \in \Omega$, then

$$\omega : I_N \rightarrow O.$$

Such a function we call a *conditional outcome* (or reward) *schedule*. Usually $\omega$ is determined by some sort of random device; if so and if

$$\pi(o \mid \omega) = \Pr\,[\omega(n) = o]$$

is independent of $n$, then we say that it is a *simple random conditional schedule*. Of course,

$$\sum_{o \in O} \pi(o \mid \omega) = 1.$$

The function $\phi$ that selects the conditional schedule to be used,

$$\phi : S \times R \rightarrow \Omega,$$

we call the *outcome function*.

In the mathematical learning literature an outcome function is said to be noncontingent if and only if it is independent of the response, that is,

$$\phi(s, r) = \phi(s, r')$$

for all $s \in S$ and all $r, r' \in R$. Otherwise, it is called contingent.

The sequence of events, then, is that a presentation $s$ on trial $n$ elicits a response $r$, and, given these, the outcome function $\phi$ selects a conditional outcome schedule $\omega = \phi(r, s)$, which in turn prescribes the outcome $o = \omega(n)$ to be administered to the subject.

As we have said, in contemporary psychophysics the conditional outcome schedules, if they are used at all, are random (independent of $n$) with the very special property that $\pi(o \mid \omega) = 0$ or 1. In these cases we can by-pass $\Omega$ entirely and think of $\phi$ as a function from $S \times R$ into $O$ by defining

$$\phi(s, r) = o \quad \text{if and only if} \quad \phi(s, r) = \omega \quad \text{and} \quad \pi(o \mid \omega) = 1.$$

When we do this, we say $\Omega$ does not exist and refer to $\phi$ as a *payoff function.*

In simple learning experiments in which there is only one stimulus presentation, it is usual to suppress explicit reference to the outcome function $\phi$ and simply to subscript the conditional outcome schedules by the corresponding response symbol under $\phi$. When this is done, it is usual to speak of the schedule rather than the outcome function as contingent or noncontingent, even though strictly this does not make sense.

In certain studies having simple random conditional outcome schedules it is necessary to refer to or to describe the random mechanisms that generate the schedules. For example, we may ask a subject to choose which of two schedules is to be used when all he knows is the nature of the devices that generate the schedules. We let $\boldsymbol{\omega}$ denote the device, or a description of it, as the case may be, that generates $\omega$ and $\boldsymbol{\Omega}$ the set of devices or descriptions of them, corresponding to the set $\Omega$ of schedules. It is, of course, perfectly possible for the same device to generate two different schedules, that is, for $\boldsymbol{\omega} = \boldsymbol{\omega}'$ even though $\omega \neq \omega'$.

It is generally assumed that subjects have preferences among the elements of $O$. Usually, these are assumed to be representable by an asymmetric (therefore, irreflexive), transitive binary (preference) relation $\succ$. In principle, a separate experiment must be performed to discover $\succ$ (see Irwin, 1958). The form of this experiment is (1) $S \subseteq O \times O$, (2) $R = \{1, 2\}$, (3) $\Omega$ does not exist, and (4) if $s = \langle o^1, o^2 \rangle \in S$, then $\phi(s, r) = o^r$. The assumption that $\succ$ is a relation is equivalent to the assumption that $p_n(r \mid s)$ has the value 0 or 1, independent of $n$. Should this prove false, then a more complicated preference structure over $O$ must be postulated. In practice such experiments are rarely carried out because the results are presumed to be known: a subject is thought to prefer being correct to incorrect, a larger to a smaller sum of money, no shock to a shock, etc. It is clear, however, that if we come to perform experiments with conflicting components to the outcomes—for example, 5¢ and a shock if correct versus $-3$¢ if incorrect—then these preference subexperiments cannot be by-passed.

Following the selection and administration of the outcome to the subject or, when there is no outcome structure, following his response, the system returns the trial number $n$ to the box labeled "Add 1" which then generates the next trial number, $n + 1$, unless $n = N$, in which case the experimental run is terminated.

In the course of developing this description of the class of choice experiments, we have omitted several topics. We turn to these now.

## 2. CLASSIFICATION OF THE ENVIRONMENT

### 2.1 Stimuli

The presentation set $S$ is often constructed from a simpler set $\mathscr{S}$, which we may call the *stimulus set*. For example, in a psychophysical discrimination experiment, $S$ may consist of pairs of tones, in which case it is reasonable to say that $\mathscr{S}$ is the set of these tones and that

$$S \subseteq \mathscr{S} \times \mathscr{S} = \{\langle \mathscr{s}^1, \mathscr{s}^2 \rangle \mid \mathscr{s}^1, \mathscr{s}^2 \in \mathscr{S}\}.$$

Normally we use $\mathscr{s}$ and $\mathscr{s}_i$ to denote typical elements of $\mathscr{S}$, but when a Cartesian product is involved we use superscripts to indicate the several components. Thus $\langle \mathscr{s}^1, \mathscr{s}^2, \ldots, \mathscr{s}^k \rangle$ and $\langle \mathscr{s}_i^1, \mathscr{s}_i^2, \ldots, \mathscr{s}_i^k \rangle$ are both typical elements of $\mathscr{S} \times \mathscr{S} \times \ldots \times \mathscr{S}$ ($k$ times), which is abbreviated $\mathscr{S}^k$.

In most psychophysical studies it is not very difficult to decide which physical events should be identified as the elements of $S$ and $\mathscr{S}$, even though in principle they are not always uniquely defined. Often, $\mathscr{S}$ is a fairly homogeneous set in the sense that its elements are identical except on one or at most a very few simple physical dimensions, such as amplitude, intensity, frequency, or mass. The most important condition that the stimuli must satisfy is reproducibility: we must be able to generate an occurrence of a particular stimulus at will within an error tolerance that is small compared with the subject's ability to discriminate. For some purposes, but not all, it is also important to characterize the stimuli in terms of well-known physical measures, so that, among other things, other laboratories can reproduce them.

In the psychophysical experiments considered in Chapters 3, 4, and 5, it is always possible to define $\mathscr{S}$ in a natural way so that $S \subseteq \mathscr{S}^k$, and the order of the Cartesian product corresponds either to the time order of presentation or, if the stimuli comprising a presentation occur simultaneously, to their spatial location. For example, in the psychophysical discrimination experiment we have mentioned previously $S \subseteq \mathscr{S}^2$, so $s = \langle \mathscr{s}^1, \mathscr{s}^2 \rangle$ is a typical presentation in which $\mathscr{s}^1$ is the first tone presented and $\mathscr{s}^2$ is the second. If an experiment involves visual displays of sets of three objects arranged in some fixed pattern and $\mathscr{S}$ is the set of all objects used, then $S \subseteq \mathscr{S}^3$, where each $\mathscr{S}$ is associated with one of the locations.

In so-called "discrimination learning" experiments one often feels that he can identify the set $\mathscr{S}$ as well as $S$. Indeed, the now conventional distinction between "successive" and "simultaneous" discrimination is based upon just such identifications (Spence, 1960). Suppose we have two "stimuli," a black card and a white card. If the black card is presented on some trials and the white on the rest, then it is called "successive discrimination." It is natural to say $S = \mathscr{S} = \{\text{black card, white card}\}$. If both cards are present on each trial but on some the black is to the left of white and on others to the right, then it is called "simultaneous discrimination." If $\mathscr{S}$ continues to be defined as above, then $S \subseteq \mathscr{S}^2$. If, however, $\mathscr{S}$ is defined as a black-white and a white-black card, as might seem natural if the presentations were two cards, one of which was black on the left and white on the right and the other just the reverse, then $S = \mathscr{S}$. Thus, the distinction between successive and simultaneous procedures rests upon the experimenter's identification of $\mathscr{S}$. It is easy to invent pairs of stimulus presentations which are difficult to decompose into separate stimuli. For example, one element of $S$ might consist of a rectangular grid of black lines, and the other, a number of concentric circles. In such a case, it is anyone's guess what the stimuli (elements of $\mathscr{S}$) are, but all would agree that there are two distinct and readily identified stimulus presentations. We conclude, then, that when it is possible and useful to identify the elements of $\mathscr{S}$, the set $S$ can be generated as ordered $k$-tuples from it. If not, we simply identify the elements of $S$ directly. Our position on this point is essentially atheoretical; our goal here is to characterize and classify experimental procedures, not to discuss substantive questions such as the influence on behavior of the geometric and physical properties of the elements of $S$.

## 2.2 Background and Residual Environment

The remainder of the subject's environment we divide into two classes of events, the background and the residual environment. *The background* consists of experimenter-controlled constant stimulation that often is relevant in some direct way to the choice being made. The background is usually measured in physical terms by the experimenter and reported in the description of the experiment, and it may very well be altered as an experimental parameter in different experimental runs. For example, in many signal detection experiments a background of white noise is present throughout the run and, at regular intervals, stimuli, such as short bursts of a tone, are introduced into the background. In the quantal experiments a fixed energy level of a tone is present at all times, except periodically, when the level is raised slightly for a fraction of a second. These energy increments are considered the stimuli, and the always present energy, the background. Like the definition of $\mathscr{S}$, the defining characteristics of the background are subject to debate, and probably there is no way of fully defining it in the abstract, but again in practice there is usually little difficulty in gaining agreement about what constitutes the background.

All remaining stimulation not included in the presentation or background we refer to as the *residual environment*. Little attempt is made to characterize this except in the most general terms: "The subject sat at a desk in a small, sound-attenuating room" or "The subject's eye was kept 30 cm from the target, and the room was totally dark." Attempts are usually made to control the residual environment, but very little of it is measured. As it is sometimes impossible to control relevant extraneous stimulation adequately, a well-controlled background may be used to mask it.

## 3. INSTRUCTIONS, PRETRAINING, AND THE IDENTIFICATION FUNCTION

### 3.1 Instructions and Pretraining

We have not yet discussed one feature of every human experiment, namely, the instructions. These have at least the following two roles. First, they inform the subject about the nature of the stimulus presentations and background (usually by example); about the response set $R$ and how responses are to be made; and about the outcome set $O$, the generation of the conditional outcome schedules $\omega \in \Omega$, and the outcome function $\phi$. Second, they attempt to convey to him the judgment it is desired he make or, what is the same thing, what significance his responses will have to the experimenter.

Verbal or written instructions pose, shall we say, technical difficulties in animal experiments, and so various kinds of pretraining procedures are substituted. Prior to an experimental run of "reward training," the animal is generally partially deprived of food, water, or whatever is to be used as a reward, and in pretraining he finds out what sorts of behavior can possibly lead to reward. In escape and avoidance training studies a known "noxious" stimulus is used to motivate responding. These procedures have much the same purpose as telling a human subject that "your task is to do . . . in order to be correct." Where the human subject can be told what elements compose the stimulus, response, and outcome sets, the animal must find them out through experience. This learning is sometimes effected during preliminary trials before the choice experiment begins, as, for example, by forced trials in a pretraining phase of a T-maze experiment. This preliminary process, interesting as it is, has not been much studied.

## 3.2 Identification Functions

Returning to human experiments, three examples of instructions are the following:

> Two tones will be presented, one right after the other. One will always be louder than the other, and you are to report whether the second is louder or softer than the first. If you think that the second is louder, push the button marked "louder;" if you think that the second is softer, push the "softer" button.
>
> At regular intervals this light will come on for one second. During that time a tone may or may not be introduced into the noise. If you think the tone is there, push the Yes button; if not, push the No button.
>
> Three different tones will be presented successively. You are to decide whether the first is more similar to the second or the third. If you think it is more similar to the second, push the button labeled "2;" if it is more similar to the third, push the button labeled "3."

Although these instructions are extremely parallel, the third is really quite different from the first two because in those the subject knows that the experimenter has an unambiguous physical criterion to decide whether or not a response is correct. The experimenter knows which tone is more intense or whether a tone is present, but he has no nonarbitrary criterion to decide which of two tones is more similar to a third (when all three are different).

Let us consider for the moment situations in which the experimenter has an unambiguous criterion to decide for each stimulus presentation the response or subset of responses that is "correct." In other words, we assume a relation of "correctness" on $R \times S$. If $\langle r_j, s_i \rangle$ is an element of this subset, we say that response $r_j$ is correct for presentation $s_i$ and, by making that response, the subject has correctly identified the stimulus presentation. For this reason we call the "correct" subset of $R \times S$ an *identification relation.* It is possible to design experiments such that one or more responses are not correct for any presentation or such that for one or more presentations no response is correct. Because these designs are not common, we shall not consider them further, that is, we shall deal only with identification relations in which every response maps into (is correct for) at least one presentation and every presentation is assigned to at least one response.

In terms of the identification relation, we can define an *identification function* $\iota$ by letting $\iota(r)$ denote the subset of presentations for which $r$ is a correct response. The range of the function $\iota$, which is a subset of the power set of $S$ (i.e., of the set of subsets of $S$), is denoted by

$$I(S) = \{\iota(r) \mid r \in R\}.$$

If each stimulus has a unique correct response, then $I(S)$ is simply a partition of $S$.

Part of the role of the instructions and pretraining in many human experiments is to convey the identification function $\iota$ to the subject. To the degree that this has been successful, we feel entitled to interpret a response of $r$ to a stimulus presentation as meaning that the subject believes the presentation to have been an element of $\iota(r)$, whether or not it was in fact.

When the experimenter has no objective criterion of correctness but uses such terms as "similar," "equally spaced categories," "half-way between," and the like, in instructing human subjects, it is tacitly assumed that these words induce in the subject something analogous to an identification function. The purpose of such experiments usually is to discover this induced criterion, which we assume is at least partially revealed by his behavior during the course of the experiment. Of course it is possible, and may sometimes be useful, to impose an arbitrary identification function and to look for effects it may have on behavior.

In animal studies in which one response is always rewarded for a particular stimulus presentation and the others are not, it is natural enough to say that the rewarded response is "correct." This does not mean that the identification and outcome functions are identical, because a variety of different outcomes can be consistent with a single identification function, but they seem to amount to nearly the same thing. When the responses of animals are partially reinforced (with fixed probabilities), it is not so clear that it is useful to think of an identification function as existing at all because it is not evident what the correct responses are. The response having the largest probability of reward is a possible candidate, but this definition leads to difficulties when the maximum reward probability is not the same for all presentations. For example, suppose the probability of reward is 1 for $s_1$ and 0.7 for $s_2$. It seems reasonable to view $s_1$'s 100 per cent rewarded response as somehow "more correct" than $s_2$'s 70 per cent rewarded one.

In spite of the fact that identification functions may be of less importance for animal experiments than for human ones (see the next section for further comments on this point), they are nonetheless sometimes useful in classifying animal experiments. For example, in what has conventionally been called "discrimination learning," one normally presents stimuli that are so clearly different from one another that there can be little doubt about the animal's ability to discriminate them perfectly in the usual psychophysical sense. Yet learning occurs. If the animal is not learning to discriminate the stimuli, what is it learning? It seems evident that it is discovering the experimentally prescribed identification function, for, in spite of pretraining, the animal can obtain "information" about that function only through experience in the choice experiment itself. This point leads us to discard the term "discrimination learning" as seriously misleading; instead, we propose "identification learning." True discrimination experiments in the psychophysical sense are rarely performed with animal subjects; nevertheless, the word discrimination should be reserved for them and not be used in other ways.

## 3.3 Compatibility of Payoffs with the Identification Function

Although the notion of an identification function has not, to our knowledge, been formally discussed in the literature, there seems to have been some tendency to act as if correctness is synonymous with the outcome structure. This is clearly not so. In classical psychophysics an identification function was often defined, but there was no experimental outcome structure. But even when both are defined they need not be compatible with one another; however, it has been generally felt that they should be coordinated in some way.

This is done through the binary preference relation which the subject is assumed to have over the elements of $O$. With $\succ$ known and $\Omega$ nonexistent, we say that a payoff function $\phi$ and an identification function $\iota$ are *compatible* if

1. $s, s' \in \iota(r)$ and $r \in R$ imply $\phi(s, r') = \phi(s', r')$ for all $r' \in R$;
2. $\iota(r) = \iota(r')$, $r, r' \in R$ imply $\phi(s, r) = \phi(s, r')$ for all $s \in S$;
3. $s \in \iota(r)$ and $s \notin \iota(r')$, $r, r' \in R$ imply $\phi(s, r) \succ \phi(s, r')$.

In words, the first condition states that if two presentations have the same correct response then they have the same outcome pattern over all responses. The second, that if two responses designate the same set of stimuli then they have the same outcome pattern for all stimuli. And the last, that the outcome of a correct response to a presentation is preferred to the outcomes resulting from incorrect ones.

When the payoff and identification functions are compatible, it is often possible to describe the payoff function as a square matrix. The columns are identified with sets of equivalent responses, the rows with the subsets $\iota(r)$ of stimulus presentations, and the entries with the outcomes. By condition iii, the most preferred entry of each row is in the main diagonal, assuming the usual coordination of columns with rows.

It is not at all clear, as we noted earlier, what it means for an outcome structure that is not a payoff function (i.e., partial reinforcement) to be compatible with an identification function. Indeed, in animal experiments an identification function cannot exist without there being a compatible payoff function because we have no choice but to use the payoffs to "teach" the animal the identification function. Of course, having done the teaching, we can place the animal in a new experiment with a different outcome function, but we would probably interpret the results by saying that the animal extinguished on the old identification and learned a new one. In general, the distinction between identification and outcome functions is less clear in animal studies than it is in human ones. People can be told what is correct and at the same time be rewarded for being wrong, but this is not easily arranged with animals.

This remark, however, suggests a way to think about identification functions which may reduce the apparent differences between animal and human experiments. If we treat the identification function in human experiments as a very special kind of outcome structure, namely a payoff function in which the outcomes are necessarily the concepts of being correct and incorrect, then we can say that in these human experiments there are two distinct outcome structures. When they are compatible, as is usually the case, they can be treated as one, but the fact that they can be put into conflict if we choose shows that they are distinct. Viewed this way, the analogous animal experiment must also have two independent outcome structures, and these may or may not be compatible. One might use food outcomes for the one and shock for the other. When the first, say, is a payoff function and the second is not, then the first can be treated as inducing an identification function and the second as the outcome structure. Although experiments of this kind are rare in the animal literature, they have been performed to demonstrate the acquisition of "moral" behavior in animals (R. L. Solomon, in preparation). In these experiments punishment is used to induce an identification function that is incompatible with the natural preference ordering for two different foodstuffs that serve as the outcomes of choice. Thus a formal parallel with the human experiment exists, but it cannot be considered more than formal until it is shown to have the same sort of special properties that human identification functions seem to have. Little has yet been done to develop mathematical theories for behavior in the presence of conflicting outcomes, but such research seems potentially interesting and important (see introduction to Chapter 5).

## 4. DEFINITION AND CLASSIFICATION OF CHOICE EXPERIMENTS

### 4.1 Definition of Choice Experiments

Much of what we have said so far is summarized for convenient reference in Table 1. We make the following assumptions. The response data function $\rho$ and, in principle, the preference relation $\succ$ are generated by the subject and observed by the experimenter. The presentation schedule $\sigma$ is generated in some manner by the experimenter, and it is not usually revealed to the subject. Psychologists generally feel that if $\sigma$ is revealed to the subject then his responses may be biased by that knowledge and, at least at this stage of development, that this is undesirable. There are exceptions, such as psychophysical quantal experiments in which $\sigma$ is known to the subject, and that aspect of their design has been one major criticism of them. The presentation set $S$ and the set $\mathscr{S}$, if it is defined, are chosen by the experimenter; sometimes they are completely described to human subjects (e.g., in recognition experiments) and at other times they are only partially described (e.g., in discrimination experiments in which the subjects are only informed that pairs of tones will be presented, but not the specific pairs). The outcome function $\phi$, the set $\Omega$ of conditional outcome schedules, and the set of outcomes $O$ are selected by the experimenter, and they are more or less completely described or revealed in pretraining to the subject. Often the specific conditional outcome schedules are not revealed, but the means of generating them may be.

Any experiment in which these assumptions are met and in which the response set $R$ is well defined and contains two or more elements is called a *choice experiment*. Psychophysical experiments are almost always choice experiments, but open-ended designs are not because at least the experimenter has no real idea what $R$ is.

### 4.2 Identification Experiments

The class of choice experiments can be divided into those for which an identification function exists and those for which it does not. If that function exists, the choice experiment is called an *identification experiment*.

Table 1 A Summary of the Concepts Involved in a Choice Experiment

| Name | Mathematical Status | Symbol | Relation to Other Symbols | Remarks |
|---|---|---|---|---|
| Trials | Sequence | $I_N$ | $I_N = \langle 1, 2, \ldots, n, \ldots N \rangle$. | Total number of trials, $N$, may be a random variable. |
| Stimulus set | Set | $\mathscr{S}$ | $\mathscr{s}, \mathscr{s}_i, \mathscr{s}^h \in \mathscr{S}$. | Consists of the "elementary" stimuli defined by the experimenter. |
| Stimulus presentation set | Set | $S$ | $S \subseteq \mathscr{S}^k$, $s, s_i \in S$; image of $\sigma$; part of domain of $\phi$. | Consists of the stimuli that are presented on a trial to the subject, which when $\mathscr{S} \neq S$ are ordered $k$-tuples of elementary stimuli. |
| Range of $\iota$ | Set of sets | $I(S)$ | A set of subsets of $S$ in 1:1 correspondence with $R$ under $\iota$. | |
| Response set | Set | $R$ | $r, r_j \in R$; in 1:1 correspondence with $I(S)$ under $\iota$; part of domain of $\phi$. | Consists of two or more responses prescribed by the experimenter. |
| Set of conditional outcome schedules | Set | $\Omega$ | $\Omega$ is the set of the functions $\omega$; image of $\phi$. | The set of random devices generating the $\omega$'s in $\Omega$ is denoted $\mathbf{\Omega}$. |
| Outcome set | Set | $O$ | $o, o_k \in O$; range of $\omega$. | Consists of the immediate information feedback and outcomes prescribed by the experimenter. |
| Preference relation | Relation | $\succ$ | $\succ \subseteq O \times O$; $\succ$ is assumed to be asymmetric and transitive. | Although $\succ$ is often presumed known, it is a property of the subject and is therefore discovered experimentally. |
| Presentation schedule | Function | $\sigma$ | $\sigma : I_N \rightarrow S$ (onto) | It is said to be simple random if $\Pr[\sigma(n) = s]$ is independent of $n$. |

Table 1 (continued)

| Name | Mathematical Status | Symbol | Relation to Other Symbols | Remarks |
|---|---|---|---|---|
| Response data | Function | $\rho$ | $\rho: I_N \to R$ (into) | In many theories $p_n(r \mid s) = \Pr[\rho(n) = r \mid \sigma(n) = s]$ is assumed to exist, and the basic problem is to account for these probabilities. |
| Outcome function | Function | $\phi$ | $\phi: S \times R \to \Omega$ (onto) | It is called a payoff function if $\Pr[\omega(n) = o] = 0$ or 1 independent of $n$. Such functions are usually compatible with $\succ$ and $\iota$. |
| Conditional outcome schedule | Function | $\omega$ | $\omega: I_N \to O$ | It is simple random if $\pi(o \mid \omega) = \Pr[\omega(n) = o]$ is independent of $n$. The random device generating $\omega$ is denoted $\boldsymbol{\omega}$. |
| Identification function | Function | $\iota$ | $\iota: R \leftrightarrow I(S)$ (1:1) | Prescribed by the experimenter if it exists at all. |

It seems useful to us to partition such experiments still further in terms of the type of identification that exists; one possible and reasonable partition is described in the following four paragraphs.

ONE:ONE. Suppose the identification function is a one-to-one correspondence between the sets $R$ and $S$; that is, each response is correct for one and only one presentation and each presentation has precisely one correct response. In this case it is called a *complete identification experiment* because the subject completely and uniquely identifies each presentation (correctly or incorrectly) when he makes a response. Clearly, the number of elements in $R$ must equal the number of elements in $S$. One example is a simple detection experiment in which there is a noise background and either nothing or a particular tone is presented on each trial and the subject responds either "yes" or "no." In another detection design a tone is presented in the noise in precisely one of $k$ time intervals, and the subject picks the interval. Still another example is the classical paired-associates learning design, provided that the experimentally defined associate symbols are considered to be the response set $R$.

ONE:MANY. An experiment in which some response is correct for two or more presentations, but only one response is correct for each, is called *partial identification.* Each element of $R$ may map into several elements of $S$, but each element of $S$ maps into only one element of $R$. Put another way, $I(S)$ is a partition of $S$ that is not identical to $S$. Clearly, $S$ must have more elements than $R$ because $I(S)$ is in one-to-one correspondence with $R$. An example is a psychophysical discrimination experiment in which numerous pairs of tones are presented on different trials, and the subject responds either "louder" or "softer" to each. The identification is partial because the subject does not uniquely identify each pair of tones; he only assigns the pair (correctly or not) to one of two classes. Similarly, the usual concept formation experiments involve partial identification because the several instances of the concept require the same response.

MANY:ONE. The logical counterpart of the preceding class of designs is that in which two or more responses are correct for at least one presentation but in which each response is correct for only one presentation. We call this *optional identification* because there is at least one presentation for which the subject has an option of which response to make. The identification relation defines a partition of $R$ that is in one-to-one correspondence with $S$, hence $R$ must contain more elements than $S$. Although it is not difficult to invent such a design, we know of no classical ones that fit this description. In most standard experiments in psychophysics and learning the various correct responses for a particular stimulus presentation are not distinguished and so they are collapsed into a single response element. For example, the responses "yes," "yep," and "uh huh" are treated as equivalent in a detection experiment. The only purpose that we can see in performing an optional identification experiment would be to study the effects of some variable, for example, the amount of work required, that differentiates the optional correct responses. Although an extensive literature exists on the work involved in responses (Solomon, 1948), none of the studies used choice designs.

MANY:MANY. When at least one response is correct for more than one presentation and at least one presentation has more than one correct response, we call the experiment *ambiguous identification.* As an example, we could have $r_1$ correct for $s_1$ and $s_2$, $r_2$ correct for $s_2$ and $s_3$, etc. There is no restriction on the relative numbers of elements in $R$ and $S$. Again, we know of no standard experiment that falls into this category. Presumably, it is so complex that it is not of any real research interest at the moment; however, the prevalence of such complexity in everyday language suggests that these experiments may ultimately be important. Many different vehicles (stimulus presentations) are partially identified by the optional terms (responses) "automobiles," "autos," and "cars."

The simplest identification experiment of all is the two-response, two-presentation case in which $r_1$ is correct when and only when $s_1$ is presented and $r_2$ is correct when and only when $s_2$ is presented. This experiment is a special problem. In our scheme it is classed as the simplest example of a complete identification experiment. It is equally easy, however, to think of it as, and to revamp the definition so that it is classed as, the simplest example of a two-response partial identification design. Our choice, therefore, may be misleading from either an empirical or theoretical viewpoint or both. For example, the limiting case of the standard psychophysical discrimination design involves two stimuli presented in the two possible orders. Whether the data are best described by a model for partial identification discrimination in which there are two or more pairs of stimuli but only two responses or by a model for complete identification in which the stimulus presentation and response sets are in one-to-one correspondence is an open question. Perhaps the answer depends upon the instructions used, whether the subject is asked to say which tone is louder or merely asked to identify each presentation by one of two neutral labels, or perhaps they are equivalent, in which case it does not matter how we class it. It is a question of fact, and the information apparently is not available.

## 4.3 Asymptotic and Nonasymptotic Behavior

Our second mode of classifying experiments is of a different sort, for it rests upon assumptions about the nature of the behavior. It is thus a behavioral division rather than a methodological one. It is simply the question of whether the behavior is, or is assumed to be, statistically unchanging.

In psychophysics we frequently assume that the response probability $p_n(r \mid s)$ is independent of $n$, in which case we say that the behavior is *strictly asymptotic*. In an attempt to satisfy this assumption, pretraining is usually carried out to get the subject beyond the learning phase, and, to check it, simple statistical tests are often made. This strong assumption is surely wrong, however, if the behavior is thought to be the end product of a learning process in which the individual response probabilities continue to fluctuate even after stochastic equilibrium has been reached. Current learning models specify a branching process, so one must deal with a distribution of $p_n(r \mid s)$ on a particular trial. If that distribution is independent of $n$, we say that the behavior is *stochastically asymptotic*. We will sometimes speak simply of *asymptotic* behavior when we do not wish to say whether it is strict or stochastic.

In most models one can estimate the mean of the asymptotic distribution either by averaging over identical subjects on a single trial or by averaging over many trials for a single subject.

Most work in psychophysics—both experimental and theoretical—is devoted to asymptotic choice behavior. Possibly more attention should be paid to the preasymptotic behavior of subjects in psychophysical experiments and to the asymptotic fluctuations predicted by learning models. On the other hand, learning experiments of the choice variety usually focus on the transient preasymptotic behavior and seldom provide adequate information about asymptotic behavior, the main exceptions being some over-learning experiments with partial reinforcement. Although most operant conditioning studies do not use choice designs, it is worth noting that the interest is mostly in asymptotic behavior and the changes in asymptotes as various experimental factors are manipulated.

## 4.4 Summary of Classifications

Using the existence or nonexistence of an outcome structure as a third important mode of classification, Tables 2 and 3 summarize the distinctions we have made. Where standard examples are known, they are listed in the appropriate cell.

Most of the examples mentioned in the tables are familiar learning and psychophysical experiments, and they fall naturally within our scheme. It is less apparent how the experiments that have been performed to study preferences among outcomes are to be characterized. These are the ones that have come to be known as gambling experiments and that are associated with theoretical studies of the notion of utility. They can be partially described by (1) $\mathscr{S} = \mathbf{\Omega}$, where it will be recalled $\mathbf{\Omega}$ is the set of random devices which generate the conditional outcome schedules, (2) $S \subseteq \mathbf{\Omega}^k$ (usually $k = 2$), (3) $R = \{1, 2, \ldots, k\}$, (4) $\iota$ does not exist, and (5) if $s = \langle \boldsymbol{\omega}^1, \boldsymbol{\omega}^2, \ldots, \boldsymbol{\omega}^k \rangle \in S$, then $\phi(s, r) = \omega^r$. It is easy to see that this is simply a generalization of the previously described experiment used to determine the preference relation $\succ$ over $O$. Such experiments are of considerable interest, and formally they are intriguing because of the intimate connection between the stimulus presentation set and the outcome structure which does not exist in other experiments.

A notable feature of Tables 2 and 3 is the pattern of omissions. First, there seem to be no standard psychological experiments we could classify as optional or ambiguous identification. It is not clear to us whether this represents a serious gap in experimentation or whether such experiments are considered to be of little interest. It is evident that most complete identification experiments could be viewed as optional ones with responses within a class being treated alike by the experimenter. It is well known that animals often develop idiosyncratic, stereotyped, or "superstitious" behavior in situations in which several response patterns are equally "functional," but this has seldom been studied systematically.

Table 2 Identification Experiments

| | Complete Identification (one:one) $N(R) = N(S)$ | Partial Identification (one:many) $N(R) < N(S)$ | Optional Identification (many:one) $N(R) > N(S)$ | Ambiguous Identification (many:many) |
|---|---|---|---|---|
| Nonasymptotic | Simple animal "discrimination learning"<br>Classical paired-associates learning | Concept formation | | |
| Asymptotic | Simple detection<br>Simple recognition<br>$k$-alternative forced-choice detection | Psychophysical discrimination<br>Detection of unknown stimulus<br>Method of single stimuli | | |

Second, within the nonempty cells of Tables 2 and 3 there is a remarkable complementary relation between the asymptotic and nonasymptotic rows. It appears that standard designs in learning are nonexistent in psychophysics and vice versa. This seems unfortunate because independently developed theories may well be inconsistent. A single theory should predict both learning and asymptotic behavior, and both sets of predictions should be tested experimentally. Few such tests have been carried out.

Table 3 Nonidentification Choice Experiments

| | | $N(S) > 1$ | |
|---|---|---|---|
| | $N(S) = 1$ | Outcome Structure Exists | No Outcome Structure Exists |
| Nonasymptotic | Simple learning | | |
| Asymptotic | Overlearning | Gambling experiments | Similarity experiments<br>Category judgments<br>Attitude scaling<br>Magnitude estimation |

## *References*

Audley, R. J. A stochastic model for individual choice behavior. *Psychol. Rev.*, 1960, **67**, 1–15.

Bush, R. R. & Mosteller, F. A mathematical model for simple learning. *Psychol. Rev.*, 1951, **58**, 313–323(a).

Bush, R. R. & Mosteller, F. A model for stimulus generalization and discrimination. *Psychol. Rev.*, 1951, **58**, 413–423(b).

Bush, R. R. & Mosteller, F. *Stochastic models for learning*. New York: Wiley, 1955.

Bush, R. R., Mosteller, F. & Thompson, G. L. A formal structure for multiple choice situations. In R. M. Thrall, C. H. Coombs, and R. L. Davis (Eds.), *Decision processes*. New York: Wiley, 1954. Pp. 99–126.

Estes, W. K. Toward a statistical theory of learning. *Psychol. Rev.*, 1950, **57**, 94–107.

Galanter, E. & Miller, G. A. Some comments on stochastic models and psychological theories. In K. J. Arrow, S. Karlin, and P. Suppes (Eds.), *Mathematical methods in the social sciences*, 1959. Stanford: Stanford Univer. Press, 1960. Pp. 277–297.

Irwin, F. W. An analysis of the concepts of *discrimination* and *preference*. *Amer. J. Psychol.*, 1958, **71**, 152–165.

Luce, R. D. *Individual choice behavior: a theoretical analysis*. New York: Wiley, 1959.

Solomon, R. L. The influence of work on behavior. *Psychol. Bul.*, 1948, **45**, 1–40.

Spence, K. W. *Behavior theory and learning*. Englewood Cliffs, N.J.: Prentice-Hall, 1960.