Why Mathematical Models?

ROBERT A. BJORK *University of Michigan* [1]

If you are a member of the set of people to whom this article is addressed, that is, if you are an undergraduate student in psychology, you certainly possess some or all of the following characteristics: (*a*) you are interested in understanding behavior; (*b*) you are concerned about a number of environmental, social, and political issues; (*c*) you find the topics on the "soft" side of psychology somewhat more interesting than the topics on the "hard" side of psychology; (*d*) you know little, if anything, about mathematical psychology (and are somewhat suspicious of the whole endeavor); and (*e*) you do not have the time to do everything that you want to do.

Given an audience with those characteristics, this article has relatively limited goals. It says a little about what mathematical models are, gives two specific examples, and indicates why you should care to know something about mathematical models.

There is a real question as to whether or not undergraduate students in psychology should spend a chunk of their limited time on mathematical psychology. For some of you the question does not exist because no such course exists among the undergraduate courses available to you. The absence of such a course is sometimes by design—some of those in charge of undergraduate curricula in psychology do not feel that such a course is appropriate for undergraduates—and is sometimes necessitated by there being no faculty member who is (or feels he is) qualified to teach such a course. Teaching such a course is not easy; at least it is not easy to teach it well. In fact, for those of you at schools where there is a course of sorts in mathematical psychology, the way it is taught is sufficient reason for you to avoid it. The fact that undergraduate courses in mathematical psychology typically have small enrollments is not altogether the result of undergraduates' initial predisposition against such courses. It *is* difficult to compete with current high-demand courses such as The Drug-Crazed Deviant Individual, but, when taught well, courses in mathematical models can be and are well received and even semipopular.

You may feel that you do not have the mathematical training needed to profit from a course in mathematical psychology, and that may be true. Today, however, the minimal training for such a course is certainly not more than the background in mathematics and statistics required to be sophisticated in any area of psychology. You may also feel, given your particular interests, that mathematical models have little relevance to you because few, if any, efforts at quantitative theorizing have been attempted in your area of interest within psychology. That again may be true, but the basic techniques of formulating and testing models are being applied ever more widely in psychology; only a few corners of psychology remain free of mathematical (or computer) models, and they are not likely to remain completely free for very long. There are some good reasons for some of you to expend the effort required to learn something of mathematical models. Those reasons are discussed later in this article.

Why Mathematical Psychology?

Before attempting to define mathematical models and to characterize their properties and virtues, some comments on the role of mathematics in psychology are in order. We tend to think of mathematical psychology as a recent development within psychology, starting somewhere around 1950, and it is, in the sense of being pervasive, identifiable, and generating books, journals, and graduate programs of study, but the history of mathematical

---

[1] This article is the text of an invited address sponsored by Psi Chi at the annual meeting of the American Psychological Association, Washington, D.C., September 1971.
Requests for reprints should be sent to Robert A. Bjork, Human Performance Center, 330 Packard Road, Ann Arbor, Michigan 48104.

theorizing in psychology is long and substantial. As Miller (1964) said, "Psychology has had a long and not always happy affair with mathematics." The drive to quantify has been as strong in psychology as in other sciences. Any number of examples from the last century and the first half of this century of efforts to characterize behavior precisely and quantitatively could be cited, for example, Weber's and Fechner's efforts to specify the relationship between the intensity of a physical stimulus and the subjective intensity of the resulting sensation, Ebbinghaus' efforts to specify the forgetting function for learned verbal items, and Hull's attempts to postulate a general theory of behavior.

It is standard practice, of course, that we report our observations of behavior in numerical terms. The advantages of quantifying data are by now so obvious that we take for granted, whatever the area of psychology, that our observations should be recorded as precisely as possible. Thus, we record response probabilities, amplitudes, speed, frequency, and so forth.

When one examines the history of the sciences that are more developed than psychology, the importance of mathematical theories in scientific progress is very apparent. Not every application of mathematics in science has proven fruitful, but, in general, as knowledge has grown in a field of study, theories have become more mathematical; and as theories have become more mathematical, knowledge has grown. Few people committed to scientific psychology would question the ultimate need to phrase theories in quantitative form. The extent to which current theories in psychology are not quantitative reflects both the complexity of behavior and the relative immaturity of psychology as a science. It is possible, of course, that current attempts to work toward quantitative theories are premature, or, more fundamentally, that behavior is characterized by an intrinsic lack of lawfulness, which makes quantitative theories ill advised.

An important point to make about mathematical psychology as an area of study is that it is not an area of study. Mathematical psychology does not denote a particular content area—not in the way that terms such as *learning, perception,* or *personality* denote particular content areas. Rather, mathematical psychology refers to an attitude, a way of going about the study of various topics within psychology. In fact, to know that someone is a mathematical psychologist working in a particular area of study says nothing of his theoretical biases in that area, except that he is an adherent of the mathematical psychology approach. He may have one or more of a variety of different biases or inclinations with respect to his area of study: the tools of mathematics are available to anyone; they are theoretically neutral.

What Are Mathematical Models?

The precise definition of a mathematical model is not trivial to state and is something that philosophers of science worry about, but the customary usage in psychology is fairly straightforward: "A mathematical model is a set of assumptions together with implications drawn from them by mathematical reasoning [Neimark & Estes, 1967, p. v]."

The most important aspects of a model are the theoretical notions it embodies. A model is a formal system, an abstract representation, and it can be stated or phrased in different ways. The assumptions of the model might be stated in axiomatic form, used to generate a program in a computer, or provide the basis for constructing a mechanical model. The adequacy of a model is judged by how well its behavior corresponds to the behavior of the system it was designed to represent. What is tested in any such judgment are the formal properties of the model, not the specific format in which the model is stated. Thus, we might postulate that the human attentional system can be represented as a clerk (processor) waiting on customers (stimulus inputs) arriving at a counter (receptors). The clerk might be assumed to have particular properties (e.g., he can only handle one customer at a time, and it takes him a certain period of time to handle a customer), and the customers might also be assumed to have particular properties (e.g., they may arrive and choose among one of several counter locations, corresponding to different sense modalities, and if they have to wait too long in line they may leave before being served). If such a representation of the human attentional system were fully and precisely specified, its predictions ("behavior") could be compared to the behavior of a real human attentional system in some experimental situation, say one involving the classification of rapid inputs presented visually and auditorily.

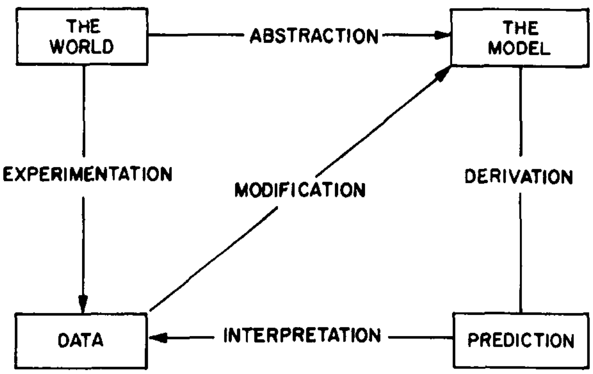

FIG. 1. Schematic illustration of the modeling process (with permission, after Coombs, Dawes, & Tversky, 1970).

The point of the preceding example is that any test of the model's adequacy is a test of its formal properties, not a test of its specific representational format. We can state with some confidence that there are no clerks or counters in the human head, but we cannot reject the model on that basis. There may be a quite close correspondence between the assumed properties of the clerk and the properties of some central processor in the human information-processing system.

You might wonder why one would care to postulate an adequate model of a particular behavioral system. The answer to that question is quite simple: We want to understand behavior. Suppose we were able to construct a mechanical rat that learned to run mazes in exactly the way a genuine rat learned to run mazes; that is, the learning behavior of each had precisely the same characteristics. In an important sense, if we know what we did to make our mechanical rat learn—what learning and motivating mechanisms were assumed—we can say that we understand how a real rat learns.

Figure 1 attempts to clarify some of the points stated above. A model is an abstract representation, a caricature, of some behavioral system. The model consists of a set of assumptions and some correspondence rules. By "correspondence rules" is meant the specification of what in the model corresponds to what in the world, that is, to what in the behavioral system and situation of interest.

In practice, quantitative theorizing is a dynamic process. Typically, a simple model is postulated to account for behavior of a specified type in a well-defined experimental situation. The model is then tested, found to be inadequate to some degree, modified in some fashion designed to correct the inadequacy, tested again, modified again, and so forth. The process constitutes a kind of struggle by the theorist to understand the behavior he is studying, a continuing effort to increase the adequacy and generality of his theory. Much more often than not, some particular discrepancy between the observed and predicted behavior is so compelling that the model is clearly beyond repair. In other cases, the total of small but consistent discrepancies comes to be overwhelming.

It is worth commenting that a model can be rejected or falsified, but it cannot be proven. Coombs, Dawes, and Tversky (1970) illustrated that point with a good example. Suppose a particular model of some learning process predicts that Teaching Method A should be better than Teaching Method B. Finding that Method B is better than Method A rejects the model, whereas finding that Method A is better than Method B only offers inductive support for the model, but does not prove the model. There may be other quite different models that are consistent with the result as well. Thus, we may gain ever increasing confidence in a model. We are, for example, so confident of certain models in physics that we think of them as the quite literal truth rather than as models. However, we can never prove that the model is a completely accurate representation.

Mathematical Models versus Verbal Models

It should be apparent from the foregoing discussion that the psychological assumptions underlying a model are the critical characteristics of the model as a theory of some behavioral process. A model may involve very elegant and sophisticated mathematics and still be a very bad model, and a model stated in a verbal, nonmathematical form may be a quite good model. There are, however, some clear-cut virtues in attempting to phrase a model in a mathematical form.

A MATHEMATICAL MODEL IS MORE READILY FALSIFIABLE

A long-standing problem in psychology is how to get rid of old, inadequate theories. People have a way of becoming committed to theories, especially their own theories, and they tend to interpret experimental results as being consistent with their theory, however strained such an interpretation might be. In contrast to the fate of old army generals, old theories are not only not permitted to die, they are not permitted to fade away either. The problem is a problem in direct proportion to the imprecision of our theories. If a theory is not clearly and unambiguously rejectable, it is of little usefulness. Mathematical models, because they yield numerical predictions, are much more falsifiable than are verbal models, which yield, at most, predictions of ordinal relationships. Thus, mathematical models predict not only whether something is greater or lesser, or faster or slower, than something else, but also how much greater or lesser, or faster or slower.

MATHEMATICAL MODELS FORCE A THEORIST TO BE PRECISE

In order to derive quantitative predictions from a model, the assumptions on which the model is based must be specified completely and precisely. Attempting such a specification is often a useful exercise for a psychologist. The exercise forces one to clarify one's notions and to choose among alternative notions; it can become painfully obvious at times that one is simply not clear as to exactly what he thinks—or that certain hidden assumptions reside in one's opinions.

THE CONSEQUENCES OF ASSUMPTIONS CAN BE DERIVED

A related virtue of mathematical models is that implications can be derived from the assumptions of a model via mathematical logic, some of which may be far from apparent. Such unexpected implications, which would be unobtainable if the model were phrased in a nonmathematical form, often provide a new basis for testing the model and frequently suggest new experiments.

MATHEMATICAL MODELS IMPROVE DATA ANALYSIS

The implications and numerical predictions derivable from a mathematical model often suggest new analyses of experimental data. Also, the relative value of different data analyses in terms of characterizing the behavior being studied can be determined by the power of those analyses to decide between alternative mathematical models. In short, mathematical models contribute to our being aware of the richness in our data.

MATHEMATICAL MODELS HAVE MORE PRACTICAL APPLICATIONS

An even fairly satisfactory mathematical model can have a number of practical applications. For example, we might use a learning model to optimize the programmed instruction of a child in reading or arithmetic; a decision-process model might facilitate medical diagnosis; the format of printed or lighted displays might be improved on the basis of a model of pattern recognition; a model of attitude change might serve the good or bad purposes of governments and other agencies in influencing the public; games designed to teach cooperative behavior to children might be derived from models of two-person interactions; and so forth.

Some Examples

It is likely at this point that you are possessed of one or both of two dissatisfactions. You may feel that the discussion thus far is too general and abstract, that everything is too vague. Or you may feel that not enough has been said about how the actual steps in model building—specifying assumptions, establishing correspondence rules, deriving implications and numerical predictions, and judging the extent to which the model is an adequate representation of the behavior it purports to model—are actually carried out. In an attempt to alleviate the first dissatisfaction, two simple models are formulated and discussed below. Attempting to alleviate the second dissatisfaction, however, is nontrivial if done seriously and is beyond the goals of this article. There is a great deal to say about each of the steps in the model-building process; it would require an article much longer than this one to do justice to any one of them. Some of the considerations are very interesting, some of them are very complex, and should you be interested, there are places for you to find out about them.

Nothing seems to clarify what mathematical models are all about as well as simply examining the assumptions and implications of a variety of different models, and at this point in time there is a large variety from which to choose. There are models of sensation, perception, memory, learning, problem solving, avoidance conditioning, multi-person interactions, choice behavior, attitude change, language learning, perceptual-motor skills,

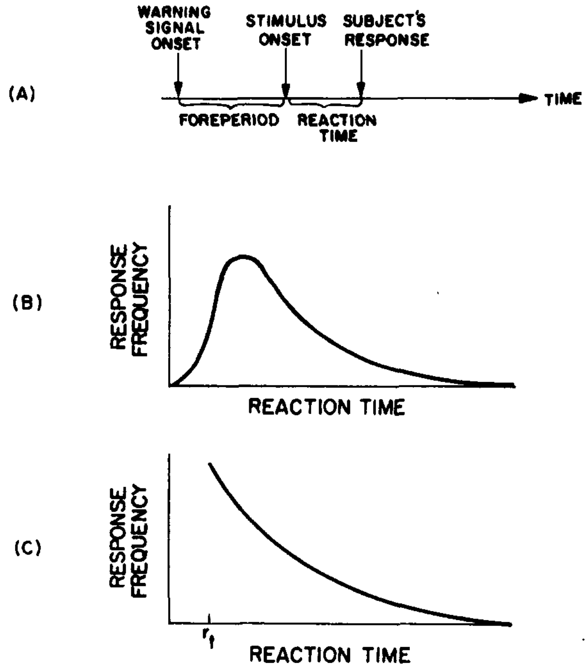

FIG. 2. (A) A trial in a simple reaction time experiment; (B) a typical distribution of simple reaction times; (C) predicted distribution of reaction times.

and other behavioral processes as well. The two models discussed below were chosen arbitrarily to illustrate several different points about mathematical models.

SIMPLE REACTION TIME

Assume, for a minute, the following situation. A person is sitting at a table. On a vertical board facing him there is mounted a small light, and on the table in front of him there is a response key. He is sitting with his right hand on the response key, and he is wearing headphones.

Now assume that we simply require the person to press the response key as rapidly as he can when the light goes on. This elemental task is the simple reaction time task. An actual experiment might consist of a very large number of trials as shown at the top of Figure 2. A trial begins with a warning buzzer presented through the headphones, there is a fixed foreperiod (usually a few seconds) from the warning signal until the light comes on, and the subject's reaction time is measured as the time from the onset of the light until he presses the response key.

The simple reaction time task is, in fact, a simple task, and the subjects' reaction times are very short. A subject's behavior in the task is not all that simple to understand, however. One problem is the inherent variability of simple reaction times. A given subject can be characterized as having a certain average reaction time in the task, but if we record his reaction times from a large number of trials, they will constitute a distribution varying over a considerable range on either side of his average reaction time. An example of such a distribution is shown in Figure 2B.

What produces the variability in a subject's simple reaction time? One might propose a number of different attentional, perceptual, and response mechanisms in an attempt to account for the observed reaction time behavior. Consider the model generated by the following assumptions.

1. A subject must detect that the light is on, and he must execute the response. Both processes take time, and the subject's reaction time on any trial is the sum of the times required for each process. That is,

$$RT = d_t + r_t,$$

where d_t is the time taken to detect the light on a trial, and r_t is the time taken to execute the response.

2. The response execution time does not vary from trial to trial; that is, it is equal to some fixed number of milliseconds.

3. The subject's detection time varies from trial to trial through the following process. During any one instant of time (Δt) following the onset of the light, the subject is either attending to the light or attending to something else. The subject need only attend to the light during some one instant of time following its onset to detect the light. During any one Δt following the light's onset, the subject attends to the light with probability p or attends to something else with probability $1 - p$.

Assumption 3 of the model generates a probabilistic distribution of detection times. It is as though the subject's attention is governed by a switching mechanism characterized by a fundamental intermittency; that is, switching can occur only at Δt intervals of time. Thus, the subject's detection time on any trial might take on any of the values, Δt, $2\Delta t$, $3\Delta t$, $4\Delta t$, . . ., depending on when the subject first detects the presence of the light. The probabilities that the subject's detection time will be equal to any of the possible times are shown below.

DETECTION TIME	PROBABILITY
Δt	p
$2\Delta t$	$(1-p)p$
$3\Delta t$	$(1-p)^2 p$
$\vdots$	$\vdots$
$n\Delta t$	$(1-p)^{n-1}p$
$\vdots$	$\vdots$

For any particular values of r_t, p, and Δt, the model generates an actual predicted distribution of reaction times. The exact form of the distribution will depend on the particular values chosen, but all of the distributions will have a form similar to that shown in Figure 2C. Assume, for example, that $p = .30$. The values of p, $(1-p)p$, $(1-p)^2p$, and so on, will then be .30, .21, .15, . . ., respectively; and any other value of p will produce steadily decreasing values as well.

Suppose now that we have obtained an actual distribution of simple reaction times from a subject. What we would like to decide is how well the distribution predicted by our model corresponds to the subject's distribution. Deciding the matter involves (*a*) estimating what values of r_t, p, and Δt will yield a predicted distribution closest to the obtained distribution, and (*b*) conducting a goodness-of-fit test on the two distributions to measure the likelihood that the subject's distribution could have been generated by the processes assumed by the model.

In this particular case it would not be necessary to actually carry out *a* and *b* above, because one can see by inspection of Figure 2 that the model predicts distributions with certain general properties that do not correspond to the properties of actual distributions. The model predicts a distribution with a geometric (monotonically decreasing) shape. That is, the most likely reaction time is the fastest possible reaction time ($r_t + \Delta t$), whereas observed distributions do not have that property.

Some of you may be thinking that it was a mistake to assume that response execution time does not vary, that if we modified the model to include variation in response times it might do a better job—maybe so. It is just such speculations that seduce theorists into the modify, test, modify, test, etc., struggle referred to earlier.

If one were to attempt seriously to postulate a model for simple reaction time, there are additional phenomena one would need to worry about. For example, reaction time varies with the intensity of the stimulus, becoming faster as the intensity of the stimulus increases. In the model specified above, stimulus intensity might be expected to affect p, the probability of attending to the stimulus in any instant of time. Simple reaction times also vary with the duration of the foreperiod, becoming longer with longer foreperiods. There is no immediately obvious way to reflect such an effect in the model.

Attempting to extend a model to account for additional phenomena in a particular behavioral situation is one way to test the correspondence rules embodied in the model. If the model is an adequate representation of a behavioral situation, then changes in the situation usually can be represented by certain natural changes in the model. If making the natural changes in the model affects the predictions of the model in a different way than the changes in the situation affect behavior, the adequacy of the model as a representation of the behavioral situation is questionable. When, as in the foreperiod example above, there seems no obvious way to interpret a certain experimental manipulation in the structure of a model, one wonders whether the structure of the model is complete enough to be adequate.

AVOIDANCE CONDITIONING

The second illustration is a model for avoidance conditioning postulated and tested by Bower and Theios (1964). This model deals with a very different behavior by a different experimental animal than that represented by the first model, and, in clear contrast to the first model, this model seems quite promising.

The experimental situation is as follows. A white rat is placed in an apparatus that consists of two compartments. The compartments are separated by a vertical partition with a door in the middle that can be raised, permitting the rat to go from one compartment to the other.

An avoidance-conditioning trial consists of the following series of events: the rat is in one of the compartments minding his own affairs; a light goes on in the rat's compartment, and, simultaneous with the onset of the light, a buzzer sounds and the door between the compartments is raised; three seconds after the onset of the light and buzzer, an electric shock is administered through the metal grid on the floor of the rat's compartment. If the rat runs to the other compartment before the shock is administered, he avoids the shock. If the rat does not make the avoidance response in time, he is shocked, and the electric current stays on until the rat manages to escape it by running to the other compartment. The avoidance trials continue until the rat learns to make the avoidance response consistently.

Bower and Theios assumed that the rat, at any point during the avoidance training, is in exactly one of three states of learning. The rat starts in an unconditioned (U) state in which he is naive and is certain to be shocked. As a result of the rat's being shocked, there is some probability c that a fear response to the light and buzzer stimulus is conditioned in the rat. In the fear-conditioned (F) state, there is some probability p that the rat successfully avoids the shock; the rat's fear in response to the light and buzzer motivates activity such as running and jumping, which may lead to his getting through the door prior to the shock. In the event that a rat in State F does not avoid the shock, he must escape it, and escaping the shock results in his learning the avoidance response with probability c. Once the rat has learned the avoidance response, he is in State L and he will successfully avoid the shock thereafter.

These assumptions are represented by the matrix below. The matrix consists of the state-to-state transition probabilities assumed by the model.

$$
\begin{array}{ll}
 & \text{State on Trial } n+1 \\
 & \begin{array}{ccc} \quad\text{L} & \qquad\text{F} & \qquad\text{U} \end{array} \\
\begin{array}{r} \text{State on} \\ \text{Trial } n \end{array}
\begin{array}{c} \text{L} \\ \text{F} \\ \text{U} \end{array}
& \begin{bmatrix} 1 & 0 & 0 \\ (1-p)c & 1-(1-p)c & 0 \\ 0 & c & 1-c \end{bmatrix}
\end{array}
$$

For example, the matrix says that an animal in State F on Trial n will, as a result of the events on Trial n, make the transition to State L on Trial $n+1$ with the probability $(1-p)c$. The model is of a familiar type about which a great deal is known: it is a three-state Markov model. From the model a great many aspects of the trial-by-trial avoidance learning of the rat can be predicted; for example, once values are assigned to c and p, the model predicts the probability of an avoidance response on any trial, the average trial of the last shock, the number of trials prior to the first successful avoidance, and many other statistics of performance.

This particular model has been found to account quite well for avoidance learning in the situation described. As is often the case, some of the predictions of the model that can be derived mathematically are not at all obvious. For example, the model predicts that if we look at a rat's performance on the trials that fall between the trial on which the rat first avoided the shock and the trial on which the rat was last shocked, there will be no improvement in his performance. Bower and Theios, in analyzing actual avoidance-learning data, found strong support for that far from obvious prediction of the model.

Why Should You Be Interested?

For some of you, there is an obvious reason why you should be interested: you might want to get involved. Some of you have the skills, interests, and inclination to contribute to mathematical psychology and, thereby, the study of behavior. Although this article has focused on mathematical models, the arguments apply equally well to computer models, and some of you may have the abilities and motivation to immerse yourself in the study of behavior by means of such models. Newell and Simon's (1972) recent book provides an in-depth coverage of the computer simulation of human problem solving.

Others of you may simply want to know enough about mathematical or computer models to read articles that involve models and to discuss issues involving models with other psychologists. An increasingly higher percentage of important theoretical papers in psychology involve quantitative theorizing. In the study of human cognitive processes, for example, nothing less than a kind of revolution has taken place in the last decade, and much of the progress has been based on quantitative models (for a discussion of these developments, see Greeno & Bjork, 1973). Books by Atkinson, Bower, and Crothers (1965), Coombs et al. (1970), and Restle and Greeno (1970) comprise a good treatment of general topics in mathematical psychology.

REFERENCES

ATKINSON, R. C., BOWER, G. H., & CROTHERS, E. J. *An introduction to mathematical learning theory.* New York: Wiley, 1965.

BOWER, G. H., & THEIOS, J. A learning model for discrete performance levels. In R. C. Atkinson (Ed.), *Studies in mathematical psychology.* Stanford: Stanford University Press, 1964.

COOMBS, C. H., DAWES, R. M., & TVERSKY, A. *Mathematical psychology.* Englewood Cliffs, N.J.: Prentice-Hall, 1970.

GREENO, J. G., & BJORK, R. A. Mathematical learning theory and the new "mental forestry." In P. H. Mussen & M. R. Rosenzweig (Eds.), *Annual review of psychology.* Palo Alto, Calif.: Annual Reviews, 1973.

MILLER, G. A. *Mathematics and psychology.* New York: Wiley, 1964.

NEIMARK, E. D., & ESTES, W. K. *Stimulus sampling theory.* San Francisco: Holden-Day, 1967.

NEWELL, A., & SIMON, H. A. *Human problem solving.* Englewood Cliffs, N.J.: Prentice-Hall, 1972.

RESTLE, F., & GREENO, J. G. *Introduction to mathematical psychology.* Reading, Mass.: Addison-Wesley, 1970.

Directions for Submission of Manuscripts to the *American Psychologist*

Because manuscripts submitted to the *American Psychologist* are often sent to outside reviewers, they should now be submitted in triplicate in order to be considered for publication. All manuscripts must be accompanied by an abstract of 100–120 words.

In line with our current policy of blind reviewing, the author's name and institutional affiliation should appear only on a cover sheet. The first page of the manuscript, however, should include the title. Footnotes containing information that might reveal the identity of the author or his institutional affiliation should be on a separate page.