

# 5



Dear Dr. Marlowe:

*Quick question. I don't want to harp on my cooperating teacher's style too much, but she thinks kids actually learn academic skills better when they get rewarded for it. Everything is a competition for her affection, or for points, or stickers, or more recess time, or getting your name on the "Student of the Day" chart. Last week students were asked to write a short paragraph that finishes the prompt, "On a windy day, I . . ." When the kids were done writing, she asked for volunteers. After each kid who volunteered read their piece, she instructed them to pick up a ticket (kids turn these in for prizes at the end of the week) and also to add two stars to the behavior chart. It's dizzying. There are rewards for everything!*

*When I told the cooperating teacher that my education courses raised questions about using rewards for learning, she said something like, "Psychology classes are interesting, but this is the real world. Kids need immediate feedback to learn. When they get something they like, they know they did what you asked correctly. That's how you learned when you were an infant and your mom praised you for eating your food, or gave you an M&M or praise later on when you were toilet trained, or for being nice to your siblings, or riding a bike. That's just how we all learn. We try something, and if we get good feedback we know we did it right. That's why we give grades, so children know how they are doing. Otherwise, they just don't understand when they're doing something correctly and when they need to improve."*

*Dr. Marlowe, any advice you could offer would be great.*

*Michael Lopes*

## 46 HOW DOES LEARNING OCCUR?

## ❖ HOW WOULD YOU RESPOND?

Does competition promote learning? How about rewards for good behavior or time on task? Michael's teacher believes that kids learn best when the expectations are clear and when teacher feedback is immediate and positive. Can learning occur another way, or is reinforcement always necessary? Think about student behavior in the classroom. Will students sit still, do their work, and act respectfully in the absence of rewards? And, what about academic behaviors, like reading or learning to solve complex mathematics problems? Must teachers provide incentives for students to stay engaged? Can critical thinking and abstract reasoning be developed more easily if students are praised or given other positive feedback when they are moving in the right direction toward meeting these objectives? Are grades simply a more mature version of M&Ms? What kinds of evidence could help answer these difficult questions? Is there a way you could test out the hypothesis that student learning is more efficient when it is rewarded? Keep these questions in mind as you read "The Science of Learning and the Art of Teaching" by B. F. Skinner. What questions do *you* have about learning and rewards? How can you help extend the discussion of these ideas in class? Finally, how would you respond to Michael Lopes?

---

❖ THE SCIENCE OF LEARNING AND  
THE ART OF TEACHING

*B. F. Skinner*

Some promising advances have recently been made in the field of learning. Special techniques have been designed to arrange what are called "contingencies of reinforcement"—the relations which prevail between behavior on the one hand and the consequences of that behavior on the other—with the result that a much more effective control of behavior has been achieved. It has long been argued that an organism learns mainly by producing changes in its environment, but it is only recently that these changes have been carefully manipulated. In traditional devices for the study of learning—in the serial maze, for example, or in the T-maze, the problem box, or the familiar discrimination apparatus—the effects produced by the organism's behavior are left of many

fluctuating circumstances. There is many a slip between the turn-to-the-right and the food-cup at the end of the alley. It is not surprising that techniques of this sort have yielded only very rough data from which the uniformities demanded by an experimental science can be extracted only by averaging many cases. In none of this work has the behavior of the individual organism been predicted in more than a statistical sense. The learning processes which are the presumed object of such research are reached only through a series of inferences. Current preoccupation with deductive systems reflects this state of the science.

Recent improvements in the conditions which control behavior in the field of learning are of two principal sorts. The Law of Effect has been taken seriously; we have made sure that effects *do* occur and that they occur under conditions which are optimal for producing the changes called learning. Once we have arranged the particular type of consequence called a reinforcement, our techniques permit us to shape up the behavior of an organism almost at will. It has become a routine exercise to demonstrate this in classes in elementary psychology by conditioning such an organism as a pigeon. Simply by presenting food to a hungry pigeon at the right time, it is possible to shape up three or four well-defined responses in a single demonstration period—such responses as turning around, pacing the floor in the pattern of a figure-8, standing still in a corner of the demonstration apparatus, stretching the neck, or stamping the foot. Extremely complex performances may be reached through successive stages in the shaping process, the contingencies of reinforcement being changed progressively in the direction of the required behavior. The results are often quite dramatic. In such a demonstration one can *see* learning take place. A significant change in behavior is often obvious as the result of a single reinforcement.

A second important advance in technique permits us to maintain behavior in given states of strength for long periods of time. Reinforcements continue to be important, of course, long after an organism has learned *how* to do something, long after it has acquired behavior. They are necessary to maintain the behavior in strength. Of special interest is the effect of various schedules of intermittent reinforcement. Most important types of schedules have now been investigated, and the effects of schedules in general have been reduced to a few principles. On the theoretical side we now have a fairly good idea of why a given schedule produces its appropriate performance. On the practical side we have learned how to maintain any given level of activity for daily periods limited only by the physical exhaustion of the organism

## 48 HOW DOES LEARNING OCCUR?

and from day to day without substantial change throughout its life. Many of these effects would be traditionally assigned to the field of motivation, although the principal operation is simply the arrangement of contingencies of reinforcement.

These new methods of shaping behavior and of maintaining it in strength are a great improvement over the traditional practices of professional animal trainers, and it is not surprising that our laboratory results are already being applied to the production of performing animals for commercial purposes. In a more academic environment they have been used for demonstration purposes which extend far beyond an interest in learning as such. For example, it is not too difficult to arrange the complex contingencies which produce many types of social behavior. Competition is exemplified by two pigeons playing a modified game of ping-pong. The pigeons drive the ball back and forth across a small table by pecking at it. When the ball gets by one pigeon, the other is reinforced. The task of constructing such a "social relation" is probably completely out of reach of the traditional animal trainer. It requires a carefully designed program of gradually changing contingencies and the skillful use of schedules to maintain the behavior in strength. Each pigeon is separately prepared for its part in the total performance, and the "social relation" is then arbitrarily constructed. The sequence of events leading up to this stable state are excellent material for the study of the factors important in nonsynthetic social behavior. It is instructive to consider how a similar series of contingencies could arise in the case of the human organism through the evolution of cultural patterns.

Co-operation can also be set up, perhaps more easily than competition. We have trained two pigeons to co-ordinate their behavior in a co-operative endeavor with a precision which equals that of the most skilled human dances. In a more serious vein these techniques have permitted us to explore the complexities of the individual organism and to analyze some of the serial or co-ordinate behaviors involved in attention, problem solving, various types of self-control, and the subsidiary system of responses within a single organism called "personalities." Some of these are exemplified in what we call multiple schedules of reinforcement. In general, a given schedule has an effect upon the rate at which a response is emitted. Changes in the rate from moment to moment show a pattern typical of the schedule. The pattern may be as simple as a constant rate of responding at a given value, it may be a gradually accelerating rate between certain extremes, it may be an abrupt change from not responding at all to a given stable high rate, and

so on. It has been shown that the performance characteristic of a given schedule can be brought under the control of a particular stimulus and that different performances can be brought under the control of different stimuli in the same organism. At a recent meeting of the American Psychological Association, C. B. Ferster and I demonstrated a pigeon whose behavior showed the pattern typical of "fixed-interval" reinforcement in the presence of one stimulus and, alternately, the pattern typical of the very different schedule called "fixed ratio" in the presence of a second stimulus. In the laboratory we have been able to obtain performances appropriate to *nine* different schedules in the presence of appropriate stimuli in random alternation. When Stimulus 1 is present, the pigeon executes the performance appropriate to Schedule 1. When Stimulus 2 is present, the pigeon executes the performance appropriate to Schedule 2. And so on. This result is important because it makes the extrapolation of our laboratory results to daily life much more plausible. We are all constantly shifting from schedule to schedule as our immediate environment changes, but the dynamics of the control exercised by reinforcement remain essentially unchanged.

It is also possible to construct very complex *sequences* of schedules. It is not easy to describe these in a few words, but two or three examples may be mentioned. In one experiment the pigeon generates a performance appropriate to Schedule A where the reinforcement is simply the production of the stimulus characteristic B, to which the pigeon then responds appropriately. Under a third stimulus, the bird yields a performance appropriate to Schedule C where the reinforcement in this case is simple the production of the stimulus characteristic of Schedule D, to which the bird then responds appropriately. In a special case, first investigated by L. B. Wyckoff, Jr., the organism responds to one stimulus where the reinforcement consists of the *clarification* of the stimulus controlling another response. The first response becomes, so to speak, an objective form of "paying attention" to the second stimulus. In one important version of this experiment, as yet unpublished, we could say that the pigeon is telling us whether it is "paying attention" to the *shape* of a spot of light or to its *color*.

One of the most dramatic applications of these techniques has recently been made in the Harvard Psychological Laboratories by Floyd Ratliff and Donald S. Blogh, who have skillfully used multiple and serial schedules of reinforcement to study complex perceptual processes in the infrahuman organism. They have achieved a sort of psychophysics without verbal instruction. In a recent experiment by

## 50 HOW DOES LEARNING OCCUR?

Blogh, for example, a pigeon draws a detailed dark-adaptation curve showing the characteristic breaks of rod and cone vision. The curve is recorded continuously in a single experimental period and is quite comparable with the curves of human subjects. The pigeon behaves in a way which, in the human case, we would not hesitate to describe by saying that it adjusts a very faint patch of light until it can just be seen.

In all this work, the species of the organism has made surprisingly little difference. It is true that the organisms studied have all been vertebrates, but they still cover a wide range. Comparable results have been obtained with pigeons, rats, dogs, monkeys, human children, and most recently, by the author in collaboration with Ogden R. Lindsley, human psychotic subjects. In spite of great phylogenetic differences, all these organisms show amazingly similar properties of the learning process. It should be emphasized that this has been achieved by analyzing the effects of reinforcement and by designing techniques which manipulate reinforcement and by designing techniques which manipulate reinforcement with considerable precision. Only in this way can the behavior of the individual organism be brought under such precise control. It is also important to note that through a gradual advance to complex interrelations among responses, the same degree of rigor is being extended to behavior which would usually be assigned to such fields as perception, thinking, and personality dynamics.

From this exciting prospect of an advancing science of learning, it is a great shock to turn to that branch of technology which is most directly concerned with the learning process—education. Let us consider, for example, the teaching of arithmetic in the lower grades. The school is concerned with imparting to the child a large number of responses of a special sort. The responses are all verbal. They consist of speaking and writing certain words, figures and signs which, to put it roughly, refer to numbers and to arithmetic operations. The first task is to shape up these responses—to get the child to pronounce and to write responses correctly, but the principal task is to bring this behavior under many sorts of stimulus control. This is what happened when the child learns to count, to recite tables, to count while ticking off the items in an assemblage of objects, to respond to spoken or written numbers by saying “odd,” “even,” “prime,” and so on. Over and above this elaborate repertoire of numerical behavior, most of which is often dismissed as the product of rote learning, the teaching of arithmetic looks forward to these complex serial arrangements of responses involved in transposing, clearing fractions, and so on, which modify

the order or pattern of the original material so that the response called a solution is eventually made possible.

Now, how is this extremely complicated verbal response set up? In the first place, what reinforcements are used? Fifty years ago the answer would have been clear. At that time educational control was still frankly aversive. The child read numbers, copied numbers, memorized tables, and performed operations upon numbers to escape the threat of the birch rod or cane. Some positive reinforcements were perhaps eventually derived from the increased efficiency of the child in the field of arithmetic, and in rare cases some automatic reinforcement may have resulted from the sheer manipulation of the medium—from the solution of problems or the discovery of the intricacies of the number system. But for the immediate purposes of education the child acted to avoid or escape punishment. It was part of the reform movement known as progressive education to make the positive consequences more immediately effective, but anyone who visits to lower grades of the average school today will observe that a change has been made, not from aversive to positive control, but from one form of aversive stimulation to another. The child at his desk, filling in his workbook, is behaving primarily to escape from the threat of a series of minor aversive events—the teacher's displeasure, the criticism or ridicule of his classmates, an ignominious showing in a competition, low marks, a trip to the office "to be talked to" by the principal, or a word to the parent who may still resort to the birch rod. In this welter of aversive consequences, getting the right answer is in itself an insignificant event, any effect of which is lost amid the anxieties, the boredom, and the aggressions which are the inevitable by-products of aversive control.

Secondly, we have to ask how the contingencies of reinforcement are arranged. When is a numerical operation reinforced as "right"? Eventually, of course, the pupil may be able to check his own answers and achieve some sort of automatic reinforcement, but in the early stages the reinforcement of being right is usually accorded by the teacher. The contingencies she provides are far from optimal. It can easily be demonstrated that, unless explicit mediating behavior has been set up, the lapse of only a few seconds between response and reinforcement destroys most of the effect. In a typical classroom, nevertheless, long periods of time customarily elapse. The teacher may walk up and down the aisle, for example, while the class is working on a sheet of problems, pausing here and there to say right or wrong. Many seconds or minutes intervene between the child's response and the



## 52 HOW DOES LEARNING OCCUR?

teacher's reinforcement. In many cases—for example, when papers are taken home to be corrected—as much as 24 hours may intervene. It is surprising that this system has any effect whatsoever.

A third notable shortcoming is the lack of a skillful program which moves forward through a series of progressive approximations to the final complex behavior desired. A long series of contingencies is necessary to bring the organism into the possession of mathematical behavior more efficiently. But the teacher is seldom able to reinforce at each step in such a series because she cannot deal with the pupil's responses one at a time. It is usually necessary to reinforce the behavior in blocks of responses—as in correcting a worksheet or page from a workbook. The responses within such a block must not be interrelated. The answer to one problem must not depend upon the answer to another. The number of stages through which one may progressively approach a complex pattern of behavior is therefore small, and the task so much the more difficult. Even the most modern workbook in beginning arithmetic is far from exemplifying an efficient program for shaping up mathematical behavior.

Perhaps the most serious criticism of the current classroom is the relative infrequency of reinforcement. Since the pupil is usually dependent upon the teacher for being right, and since many pupils are usually dependent upon the same teacher, the total number of contingencies which may be arranged during, say, the first four years is of the order of only a few thousand. But a very rough estimate suggests that efficient mathematical behavior at this level requires something of the order of 25,000 contingencies. We may suppose that even in the brighter student a given contingency must be arranged several times to place the behavior well in hand. The responses to be set up are not simply the various items in tables of addition, subtraction, multiplication and division; we have also to consider the alternative forms in which each item may be stated. To the learning of such material we should add hundreds of responses concerned with factoring, identifying primes, memorizing series, using short-cut techniques of calculation, constructing and using geometric representations or number forms, and so on. Over and above all this, the whole mathematical repertoire must be brought under the control of concrete problems of considerable variety. Perhaps 50,000 contingencies is a more conservative estimate. In this frame of reference the daily assignment in arithmetic seems pitifully meagre.

The result of all this is, of course, well known. Even our best schools are under criticism for their inefficiency in the teaching of drill



subjects such as arithmetic. The condition in the average school is a matter of widespread national concern. Modern children simply do not learn arithmetic quickly or well. Nor is the result simply incompetence. The very subjects in which modern techniques are weakest are those in which failure is most conspicuous, and in the wake of an ever-growing incompetence come the anxieties, uncertainties, and aggressions which in their turn present other problems to the school. Most pupils soon claim the asylum of not being "ready" for arithmetic at a given level or, eventually, of not having a mathematical mind. Such explanations are readily seized upon by defensive teachers and parents. Few pupils ever reach the stage at which automatic reinforcements follow as the natural consequences of mathematical behavior. On the contrary, the figures and symbols of mathematics have become standard emotional stimuli. The glimpse of a column of figures, not to say an algebraic symbol or an integral sign, is likely to set off—not mathematical behavior—but a reaction of anxiety, guilt, or fear.

The teacher is usually no happier about this than the pupil. Denied the opportunity to control via the birch rod, quite at sea as to the mode of operation of the few techniques at her disposal, she spends as little time as possible on drill subjects and eagerly subscribes to philosophies of education which emphasize material of greater inherent interest. A confession of weakness is her extraordinary concern lest the child be taught something unnecessary. The repertoire to be imparted is carefully reduced to an essential minimum. In the field of spelling, for example, a great deal of time and energy has gone into discovering just those words which the young child is going to use, as if it were a crime to waste one's educational power in teaching an unnecessary word. Eventually, weakness of technique emerges in the disguise of a reformulation of the aims of education. Skills are minimized in favor of vague achievements—educating for democracy, educating the whole child, educating for life, and so on. And there the matter ends; for, unfortunately, these philosophies do not in turn suggest improvements in techniques. They offer little or no help in the design of better classroom practices.

There would be no point in urging these objections if improvement were impossible. But the advances which have recently been made in our control of the learning process suggest a thorough revision of classroom practices, and, fortunately, they tell us how the revision can be brought about. This is not, of course, the first time that the results of an experimental science have been brought to bear upon the practical problems of education. The modern classroom does not, however, offer much evidence that research in the field of learning has been respected

## 54 HOW DOES LEARNING OCCUR?

or used. This condition is no doubt partly due to the limitations of earlier research. But it has been encouraged by a too hasty conclusion that the laboratory study of learning is inherently limited because it cannot take into account the realities of the classroom. In the light of our increasing knowledge of the learning process we should, instead, insist upon dealing with those realities and forcing a substantial change in them. Education is perhaps the most important branch of scientific technology. It deeply affects the lives of all of us. We can no longer allow the exigencies of a practical situation to suppress the tremendous improvements which are within reach. The practical situation must be changed.

There are certain questions which have to be answered in turning to the study of any new organism. What behavior is to be set up? What reinforcers are at hand? What responses are available in embarking upon a program of progressive approximation which will lead to the final form of the behavior? How can reinforcements be most efficiently scheduled to maintain the behavior in strength? These questions are all relevant in considering the problem of the child in the lower grades.

In the first place, what reinforcements are available? What does the school have in its possession which will reinforce a child? We may look first to the material to be learned, for it is possible that this will provide considerable automatic reinforcement. Children play for hours with mechanical toys, paints, scissors and paper, noise-makers, puzzles—in short, with almost anything which feeds back significant changes in the environment and is reasonably free of aversive properties. The sheer control of nature is itself reinforcing. This effect is not evident in the modern school because it is masked by the emotional responses generated by aversive control. It is true that automatic reinforcement from the manipulation of the environment is probably only a mild reinforcer and may need to be carefully husbanded, but one of the most striking principles to emerge from recent research is that the new amount of reinforcement is of little significance. A very slight reinforcement may be tremendously effective in controlling behavior if it is wisely used.

If the natural reinforcements inherent in the subject matter is not enough, other reinforcers must be employed. Even in school the child is occasionally permitted to do "what he wants to do," and access to reinforcements of the behavior to be established. Those who advocate competition as a useful social motive may wish to use the reinforcement which follow from excelling others, although there is the difficulty that in this case the reinforcement of one child is necessarily

aversive to another, and only when that has failed need we turn to the use of aversive stimulation.

In the second place, how are these reinforcements to be made contingent upon the desired behavior? There are two considerations here—the gradual elaboration of extremely complex patterns of behavior and the maintenance of the behavior in strengths at each stage. The whole process of becoming competent in any field must be divided into a very large number of very small steps, and reinforcement must be contingent upon the accomplishment of each step. This solution to the problem of creating a complex repertoire of behavior also solves the problem of maintaining the behavior in strength. We could, of course, resort to the techniques of scheduling already developed in the study of other organisms but in the present state of our knowledge of educational practices, scheduling appears to be most efficiently arranged through the design of the material to be learned. By making each successive step as small as possible, the frequency of reinforcement can be raised to a maximum, while the possibly aversive consequences of being wrong are reduced to a minimum. Other ways of designing material would yield other programs of reinforcement. Any supplementary reinforcement would probably have to be scheduled in the more traditional way.

These requirements are not excessive, but they are probably incompatible with the current realities of the classroom. In the experimental study of learning it has been found that the contingencies of reinforcement which are most efficient in controlling the organism cannot be arranged through the personal meditation of the experimenter. An organism is affected by subtle details of contingencies which are beyond the capacity of the human organism to arrange. Mechanical and electrical devices must be used. Mechanical help is also demanded by the sheer number of contingencies which may be used efficiently in a single experimental session. We have recorded many millions of responses from a single organism during thousands of experimental hours. Personal arrangement of the contingencies and personal observation of the results are quite unthinkable. Now, the human organism is, if anything, more sensitive to precise contingencies than the other organisms we have studied. We have every reason to expect, therefore, that the most effective control of human learning will require instrumental aid. The simple fact is that, as a mere reinforcing mechanism, the teacher is out of date.