

THEORIES OF CREATIVITY

Edited by
Mark A. Runco
Robert S. Albert



SAGE PUBLICATIONS
The International Professional Publishers
Newbury Park London New Delhi

FIRST PRINTING, 1990

6

Generativity Theory and Creativity

ROBERT EPSTEIN

Novel human performances have recently been predicted with some success by equations and a computer algorithm, instances of a theory called "generativity theory" (Epstein, 1985a). This theory asserts that ongoing behavior is generated as the probabilities of a large number of behaviors are continuously subjected to a number of simple transformation functions, which are presumed to have physical reality in the nervous system. It treats behavior as novel, fluid, and probabilistic, rather than as stereotypic and repetitious, and hence the theory is a departure from many conventional, learning-oriented theories of behavior. Generativity theory grew out of a series of studies in which novel, complex performances were constructed with pigeons.

The Columban Simulations

In 1978, at Harvard University, B. F. Skinner and I began a project called the Columban Simulation project, after *Columba livia*, the taxonomic name for the pigeon. We set about trying to get pigeons to do some of the complex

AUTHOR'S NOTE: This chapter is based on a talk given at a conference on creativity at Pitzer College in November 1988. Portions were prepared while the author was in residence at the Neurosciences Institute at Rockefeller University. Direct comments or reprint requests to Robert Epstein, Cambridge Center for Behavioral Studies, 11 Waterhouse Street, Cambridge, MA 02138.

and mysterious things people do, and we were successful in several instances. The simulations are significant in the present context in how they differ from most previous research on animal behavior. Previous research was concerned largely with training or conditioning—with procedures that brought about the *acquisition*, or, in some cases, the *maintenance* of behavior.

Pavlov's dogs, for example, acquired a new way of responding to a bell, and Thorndike's cats learned, rather clumsily, to escape from a box; they acquired, through reinforcement, the behavior they needed to open a door. In the extensive research on schedules of reinforcement that was systematized by Ferster and Skinner in 1957, we learned how behavior is maintained by the pace at which reinforcement is delivered. Particular patterns and rates of reinforcement, it was determined, produce particular patterns and rates of responding. Many researchers are still concerned with these relationships.

The first of the simulations was in this vein. Two pigeons were taught to engage in an exchange that suggested "symbolic communication." With the birds on either side of a clear partition; one "informed" the other about a hidden color, the latter "thanked" the informer, and so on (Figure 6.1). The birds engaged in the exchange repeatedly, and they did fairly little that had not been taught (Epstein, Lanza, & Skinner, 1980).

A second simulation, a follow-up of the first, was more illuminating. In the previous study, one bird had learned to peck the color corresponding with each of three letters, and the other pigeon had learned to peck the letter corresponding with each of three colors. We switched the positions of the birds and trained each in the other's role; now each knew both sequences. Then, without providing further training, we removed the partition and gave a single bird access to both of the response panels at once.

Over a period of 15 minutes or so—without our intervention—the two sequences came together to produce a new one. The bird pecked behind a curtain at a hidden color, then pecked and thus illuminated a corresponding letter, and, finally, pecked the corresponding color on the second response panel. By successfully matching the hidden color, the bird received food, and thus it repeated the sequence many times. It appeared that the bird was pecking the letter key "in order to help it remember" the color it had seen behind the curtain. The same sequence emerged when we tested the other bird, and control procedures suggested that the birds were indeed using the letter keys as memoranda (Epstein & Skinner, 1980).

We had witnessed the spontaneous interconnection of two separate repertoires of behavior. We had this opportunity because—for a few minutes, anyway—we were not concerned with teaching the bird anything; rather, we allowed the bird to do something on its own. We were concerned not with




Figure 6.1. Typical Communication Sequence

NOTE: (A) Jack pecks (and thus illuminates) the WHAT COLOR? key. (B) Jill thrusts her head through the curtain and pecks the color illuminated there (red, green, or yellow). (C) Jill pecks the corresponding letter (in this case, G for green), as Jack looks on. (D) Jack pecks THANK YOU, which operates Jill's feeder, as Jill looks on. (E) Jack pecks the corresponding color (in this case, green), which operates his feeder. (F) Both birds eat. The color keys below the WHAT COLOR? key are yellow, red, and green, respectively. The symbol keys are black on white.

acquisition or maintenance but with what some linguists and Gestalt psychologists call "production" or "generation."

A third study also resulted in significant behavior that was not trained. A bird that had been taught (a) to use a mirror to locate spots on the walls and floor of its chamber and (b) to peck blue spots on visible parts of its body proved able, without further training, to use a mirror to locate spots on its body that it could see only in a mirror (Epstein, Lanza, & Skinner, 1981). Some psychologists claim that the ability to use a mirror in this fashion is a

sign of "self-awareness" (for further discussion, see Epstein, 1986; Epstein & Koerner, 1986).

Questions began to take shape: On those occasions when behavior appeared but had not been trained, where had it come from? And did tools exist with which one could predict it? The answers were discouraging. Still, the Harvard studies led to others—conducted at Simmons College, the Foundation for Research on the Nervous System in Boston, and elsewhere—in which more dramatic cases of novel behavior were generated that eventually also led to tools for analyzing, predicting, and constructing such performances in both humans and nonhumans.

Analyzing Novel Behavior

The analysis of novel, ongoing behavior involves two separate tasks. First, one must assess the contribution made by an organism's history (see Birch, 1945; Köhler, 1925; Schiller, 1952; Shurcliff, Brown, & Stollnitz, 1971). What an organism will do, moment to moment in time, either in an old situation or in a new one, depends on what has happened to it in the past. Ideally, we would assess the contribution of both ontogenic and phylogenic histories.

The contribution of previous learning is studied easily enough, at least with laboratory animals. One simply provides or fails to provide certain experiences and then places the animal in the situation of interest. By systematically varying histories in this fashion, one can assess the contributions that they make to the performances that emerge in that situation.

But the behavior that appears will, almost certainly, be different from the behavior that has already been established. Even in old situations, organisms do new things. With the same teeth and the same toothbrush, one never brushes one's teeth the same way twice. And even in familiar situations, every sentence one speaks or writes is new in some way. In new situations, especially ones in which old behavior is ineffective, dramatic new behavior can occur—behavior people sometimes label "creative" or "insightful."

Old situation or new, a second, more difficult task is suggested. We must discover a set of principles that will allow us to predict how previously established behavior is transformed into new behavior in given situations—a set of "transformation functions" (see Bingham, 1929; Chomsky, 1965; Hull, 1935; Maier & Schneirla, 1935; Sternberg, 1988; Wertheimer, 1945). I will elaborate these tasks by exploring three classic cases of complex behavior.

The first two involve novel performances in pigeons and the third, problem-solving behavior in humans.

“Insight”

In 1917 the German psychologist Wolfgang Köhler showed that chimpanzees could solve a variety of mechanical problems in a remarkably humanlike way. A wide range of performances were reported. As is the case with humans, the animals failed most of the time, and most of the successful performances were haphazard and clumsy. But, occasionally, something remarkable happened. An animal looked thoughtful for a while and then, quite suddenly, solved the problem rapidly and completely. It did everything but shout “Eureka.” Köhler (1925) called this kind of performance “insightful,” and its occurrence was said to show that the mechanistic analysis of behavior was inadequate.

Perhaps Köhler’s most famous case of “insight” involved the box-and-banana problem: Six chimpanzees were placed in a large room in which a banana was suspended out of reach. A wooden box was available on the floor. After a number of fruitless attempts to get at the banana by jumping, most of the chimps lost interest. But one, Sultan, looked pensive. He looked back and forth from the box to the banana, and then, after about five minutes, suddenly moved the box less than a meter away from the position on the floor beneath the banana and, wrote Köhler, “springing upwards with all his force, tore down the banana” (Köhler, 1925, p. 38).

Epstein, Kirshnit, Lanza, and Rubin (1984) showed that pigeons could solve the box-and-banana problem in an equally dramatic fashion and could do so the first time it was presented. Each pigeon was confronted with the following situation: A small facsimile of a banana was suspended out of the pigeon’s reach in a position selected by a random number, and a box was placed elsewhere in the chamber. The pigeon had received food in the past for pecking the banana, and it had also learned that making contact with the banana by jumping and flying did not bring food. (See Figure 6.2).

Each of three pigeons that had acquired relevant skills before the test—skills that chimpanzees and children acquire when they are very young—solved the problem in about one minute: At first the pigeon looked confused. It stretched toward the banana, looked back and forth from the banana to the box, and so on. Then, quite suddenly, it began to push the box toward the banana, sighting the banana as it pushed. Each pigeon stopped pushing when the box was beneath the banana and then immediately climbed and pecked.




Figure 6.2. “Insight” in the pigeon. (A, B) The bird looks back and forth from banana to box. (C) it pushes the box toward the banana. (D) It climbs and pecks.

Trying to understand this performance, we first investigated the role that previous experience played in its emergence. The three birds who were successful had learned the following:

- (1) *Directional pushing.* The pigeons had received food for pushing the small box toward a green spot placed at random positions along the base of the chamber. It takes weeks or months to train a pigeon to push toward targets in this way, and many steps are involved. At first the pigeon earns food simply by pushing, then by pecking the green spot, then by orienting toward the spot and pushing toward it slightly, and so on.
- (2) *Climbing and pecking.* The pigeon received food for stepping onto a series of progressively taller boxes. Then a box was fixed in position on the floor, and the banana was suspended over it. The pigeon received food for climbing onto the box and pecking the banana. The position of the box and banana was varied.
- (3) *Extinction of jumping and flying.* The banana was suspended out of the reach of the bird, no box was available, and no food was presented. The bird was left in this situation until the behavior of jumping and flying toward the banana disappeared.

Different training histories produced systematically different outcomes. With one bird we established both the climbing and the pushing repertoires, but we did not extinguish jumping and flying. Like Sultan, this bird jumped (and, unlike Sultan, it also flew) toward the banana when given the test. After about five minutes, jumping and flying disappeared, and the bird solved the problem within the following two minutes.

We gave two other birds food for pecking the banana when it was within reach, but we did not teach climbing. When the banana was suspended over the box, each bird stretched repeatedly toward the banana, but neither successfully climbed and pecked. Apparently, without a history of climbing, a pigeon will have trouble with the box-and-banana problem. With two other birds we taught climbing and pecking, and we also extinguished jumping and flying, but we did not teach pushing. Neither bird pushed the box when given the test (Figure 6.3, panel A).

We taught two birds simply to push the box for long periods of time. We never taught them to push toward a target. We also established climbing and pecking, and extinguished jumping and flying. When given the test, each bird pushed the box aimlessly for long periods of time. Each looked up only rarely. One bird, after 14 minutes of pushing, looked up when the box was beneath the banana; it immediately climbed and pecked (Figure 6.3, panel B). In contrast, the three birds that had learned directional pushing pushed smoothly toward the banana (Figure 6.3, panel C). It seems that novel performances vary systematically with training history, but that only raises more difficult questions. How, moment to moment in time, do previously established repertoires become interconnected? Where, in short, did the new performances come from?

For convenience, I will divide the performance into four parts, although these divisions are somewhat arbitrary and the analysis of the performance will be oversimplified because of them. I will consider (a) the period of apparent confusion, (b) the first push, (c) the pushes that bring the box to the banana, and (d) the cessation of pushing.

Apparent confusion. The first, apparently chaotic, responses we see can be understood in terms of a phenomenon called *stimulus matching*. The bird's training has made two stimulus configurations meaningful: banana-over-box, which is the occasion on which pecking the banana has paid off, and box-with-spot, which is the occasion on which pushing has paid off. These two stimuli can be considered ends of a continuum of stimuli in which the test configuration is an intermediate case. In other words, the test situation contains elements of two stimuli that control two different behaviors. It contains what might be called "multiple controlling stimuli." When an organism is first confronted with more than one controlling stimulus, each




Figure 6.3. Different Training Histories and Different Outcomes

NOTE: Birds that had been trained to climb and peck but never to push did not push the box in the test situation (a). Birds that had been trained (i) to climb and peck and (ii) to push the box aimlessly for long periods of time pushed the box over much of the floor space of the chamber. The birds rarely looked up while pushing. One of the birds stopped pushing in the appropriate place and climbed and pecked the banana after having pushed for more than 14 minutes (b). Birds that had been trained (i) to climb and peck and (ii) to push the box toward a green spot placed at random positions along the base of the chamber solved the problem efficiently and in a manner suggestive of human problem-solving behavior (c). Other controls are described in the text. The times given are in minutes and seconds. A boxed time is the time to solution.

of the responses controlled by the separate stimuli tends to occur. For example, a new driver who is approaching a stoplight on which both red and green are illuminated will feel confused and will tend to stop *and* go—liter-

ally to move the right foot back and forth between the accelerator and brake pedals. The competition is unstable; at the intersection, other factors come into play, and the driver eventually proceeds.

In the laboratory, the effects of multiple controlling stimuli or intermediate stimuli can be studied in an experiment of this sort: A pigeon is placed in an enclosed chamber in which are located a standard automatic feeder and two standard plastic keys. Over the two keys is a row of 12 small lights. When the first of the lights is illuminated, the pigeon receives food for pecking the left key; when the light on the other end of the row is illuminated, the pigeon receives food for pecking the right key. Soon, the pigeon reliably pecks left or right, according to which of the two small lights is illuminated. Light 1 now "controls" left key pecks, and light 12 now "controls" right key pecks. When an intermediate light is illuminated, the pigeon pecks *both* keys, and the pecks distribute themselves roughly equally across the two keys. Lights 2 or 11, on the other hand, produce pecks almost exclusively on the left or right keys, respectively. In other words, the distribution of pecks roughly matches the position of the illuminated light.

In the box-and-banana test, then, we should at first expect our subject to show signs of confusion: to behave with respect to both the banana and the box, to stretch toward one, turn, orient toward the other, and so on, which is indeed what we observed.

First push. As was true with our driver, this competition should be unstable; it should change over time. Recall that, before the test, the bird had seen the banana alone and out of reach; jumping and flying had been extinguished in this situation. Thus, in the competition between the two behaviors, behavior with respect to the box should quickly triumph. (It follows that, if jumping and flying have not been extinguished, we should predict a great deal of behavior with respect to the banana before the other repertoire wins out. As noted above, we achieved such a result.)

In the dynamic interplay between the behaviors, the bird should come to face the box more and more directly. It thus comes to face a close approximation of the stimulus that controls pushing, hence, it begins to push.

Pushes toward banana. Why the bird pushes toward the banana is a more complicated matter. We cannot yet give a definitive account of this behavior, although we can offer various suggestions.

One possibility—which we will reject—is that, to the pigeon, the banana, raised 41 centimeters in the air, looks like the round green spot at ground level. If so, we would have a case of what is called "stimulus generalization"—a spread of effect from one stimulus to another because of common physical characteristics. The particular characteristics of the stimuli (which, after all, don't seem very similar) might not even be important in this instance.

Perhaps the bird treats the banana like the green spot because, during the test, the banana is the only conspicuous stimulus in the chamber (the spot is absent). Perhaps the bird has learned not to push toward the spot but to push toward a conspicuous stimulus.

We tested this possibility by training two birds who had never seen the banana to push toward the green spot. Then the box was placed in the center of a cylindrical chamber, and the banana was suspended at an edge of the chamber at a position determined by a random number. Neither bird confronted with this situation appeared even to orient toward the banana, and neither bird pushed the box beneath the banana during three two-minute trials.

But then, with the box absent from the chamber, we lowered the banana and taught each bird to peck it. Now, confronted with the test situation, each bird oriented toward the banana repeatedly, and each pushed the box beneath the banana on two of three test trials. Because climbing had not been trained, neither bird then climbed onto the box.

Thus stimulus generalization seems not to be involved in successful performances. The birds push toward the banana, it seems, only when pecking the banana has been reinforced. This is encouraging, for we can now say that the bird is pushing toward the banana for roughly the same reasons that an intelligent child might do so: because it has learned to push toward things, and, loosely speaking, because it thinks the banana is important.

Pushing toward the banana suggests another kind of generalization that has been called "functional generalization" (Bruner, Goodnow, & Austin, 1961). We often treat two things alike not because they have common physical characteristics but because they have common functions (for example, "things to sit on") or because we have common histories with respect to them (for example, "things that burn"). Consider a child faced with a variant of the box-and-banana problem: A cookie jar is out of reach on a kitchen shelf. After some reaching, the child clumsily moves a chair beneath the shelf, climbs, and brings down the jar. Will she do the same to retrieve a roll of toilet paper that resembles the cookie jar? Probably not. But she will probably retrieve a small toy, even though the latter looks nothing like the jar. Again, stimuli that have little physical resemblance can produce similar behaviors.

Various accounts may be given of the directional pushes as they occur moment-to-moment in time. N. E. Miller (personal communication) suggested that directional pushes win out over other ones because the bird is inclined both to push *and* to orient toward the banana. The pushes that triumph are the sum of these two responses. M. Branch (personal communication) suggested that, because of the bird's history of pecking the banana, the banana is a conditioned reinforcer. Animals will work to clarify stimuli

of this sort. By pushing the box toward the banana, the pigeon brings the banana closer and thus clarifies it.

Cessation of pushing. The bird stops pushing in the correct place because of a simple but important phenomenon called "automatic chaining." In the laboratory, we make responses occur in sequence by using the stimulus that controls one response as the reinforcer for another response. Because we usually establish the last response first, the second-to-the-last response second, and so on, this procedure is often called "backward chaining."

In the real world, an organism often generates its own chains of behavior, because its own behavior often produces a stimulus that controls other behavior. In other words, an organism's behavior changes its environment, and a changed environment, in turn, changes behavior. A movement as simple as turning one's head can have a profound effect. You turn your head toward your calendar and are reminded of an important appointment. You open your bedroom shade and see someone stealing your car. Indeed, one can hardly do anything at all without changing the probability of subsequent behavior.

As the pigeon pushes, it arranges for itself closer and closer approximations to a stimulus that it has seen before: box-under-banana, the stimulus that controls climbing and pecking. Thus it reaches the banana, climbs, and pecks.

In summary: The period of apparent confusion can be understood as an effect of multiple controlling stimuli. The bird starts to push because of the dynamics of the competition between behavior with respect to the banana and behavior with respect to the box. The bird pushes toward the banana because of its history of directional pushing and its history of pecking the banana (how these histories manifest themselves moment-to-moment in time is not yet clear). And the bird stops pushing because of automatic chaining: Its own behavior produces a stimulus that controls other behavior.

Tool Use

Epstein and Medalie (1983) presented a pigeon with another classic problem—variously called the "marble-under-the-couch" problem or the "stick" or "rake" problem (Hobhouse, 1901; Köhler, 1925; Shurcliff, Brown, & Stollnitz, 1971). A young boy is playing with a marble, which rolls under a couch, just out of his reach. He stretches repeatedly toward the marble, to no avail. Emotional behavior appears; the child whines, and he strikes the couch with his fist. After a minute or two, both the reaching and the emotional behavior subside. The child's eyes fall on a magazine on the floor beside him.

He clutches the magazine and clumsily thrusts it under the couch. He strikes the marble, which, unfortunately, moves the wrong way.

Even though he was not entirely successful, the child has done something remarkable: He has used an object as a tool to extend his reach. How can we account for such behavior?

A disclaimer is necessary. Usually, the first time children do such things, they do so for relatively trivial reasons, or for many reasons, some of which are relatively trivial (see Epstein, 1986). For example, they have almost certainly seen cartoon characters or siblings or parents solve such problems before. If so, the behavior is imitated, or imitation may at least contribute to its occurrence.

If the children can follow instructions, then, almost certainly, language also plays a role. A baby-sitter or parent says, "If you want to get the marble, hit it with the magazine." Current or past instructions of this sort probably always contribute to the emergence of novel behavior in humans. After all, most of the new dishes one prepares originate with recipes. But a question remains: What determines novel behavior that is not controlled by models or instructions? How might we account for a child's success in the marble-under-the-couch problem in the unlikely event that the child has never seen anyone solve such a problem and has never been told how to solve such a problem?

Some experiments with pigeons are suggestive. Medalie and I first provided a pigeon with a skill it would need to solve a certain problem: The pigeon was trained to push a flat hexagonal box toward green spots placed at random positions around the base of a cylindrical chamber. Then a marble of sorts—a target—was created: A Plexiglas partition was added to the pigeon's chamber, which had a 6-centimeter gap along its base. In the center of this gap we placed a small, square metal plate, and, with the hexagonal box absent from the chamber, the bird received food for pecking this plate. Each depression of the plate produced a brief, high-pitched tone, but the feeder was operated only occasionally after a tone.

Over a period of days, the plate was gradually moved to a position 10 centimeters behind the Plexiglas wall, so the bird had to stretch its neck beneath the wall to peck the plate. Then the hexagonal box was placed in the center of the chamber on the bird's side of the Plexiglas wall, and pushing the box was extinguished while pecking the plate continued to be reinforced. In other words, with the wall and plate in the chamber, the bird was taught not to push the box.

A small change in the situation turned it into a problem for the bird: We moved the metal plate to a point 16.5 centimeters behind the wall, so that the

plate was just beyond the bird's reach. We now had a situation similar to the marble-under-the-couch problem. The bird had a potential tool (the box) available to it that it knew how to use. And it had a target (the plate), which, presumably, it "wanted to reach." With the box, the bird could presumably reach the plate. Note, however, that the bird had never learned to push the box under anything and that it had never learned to push the box toward the plate.

At first the pigeon stretched repeatedly and forcefully toward the plate. It also behaved "emotionally": It turned, raised its wings, and scraped its feet on the floor. As behavior with respect to the plate weakened, behavior with respect to the box began to appear. At 30 seconds into the session, the bird pecked weakly at the box but did not move it. At 92 seconds into the session, the bird suddenly began to push the box directly toward the wall. When the box crossed beneath the wall, the bird pushed it back and forth erratically several times. Then it stretched again toward the plate and quickly pushed the box against it, thus producing the high-pitched tone. The bird began now to peck the box repeatedly and thus maintained the tone continuously.

A simple principle, called the "principle of resurgence," sheds light on the appearance of the bird's pushes and perhaps on the appearance of comparable human behavior: When, in a given situation, behavior that was recently successful is no longer successful, behavior that was previously successful in similar situations tends to recur (see Enkema, Slavin, Spaeth, & Neuringer, 1972; Epstein, 1983, 1985b; Epstein & Skinner, 1980; Estes, 1955; Lindblom & Jenkins, 1981; Mowrer, 1940; O'Kelly, 1940; Sears, 1941).

The dynamics of resurgence can be demonstrated experimentally. A pigeon is placed in a standard three-key chamber, and, at first, pecks on the right key are occasionally reinforced with food (Figure 6.4, line 1). The pigeon pecks this key exclusively. Then reinforcement is shifted to the center key: Many thousands of pecks occur on the right key over a period of weeks, but they eventually disappear, and the bird comes to peck the center key exclusively (Figure 6.4, line 2). Finally, all reinforcement is withheld (Figure 6.4, line 3).

Figure 6.5 shows an actual record of responding on each of the three keys during the first hour in which reinforcement was withheld for one bird. Nearly 2,000 pecks occurred on the center key during the hour, and the rate of pecking this key fell off dramatically during the last 20 minutes. Responding on the left key, where the bird had no history of reinforcement, was negligible. But responding on the right key reappeared. The bird did not peck it at all for 40 minutes, but pecking reappeared on this key just as the rate of pecking the center key began to decline. The bird pecked the right key nearly 900 times during the last 20 minutes.

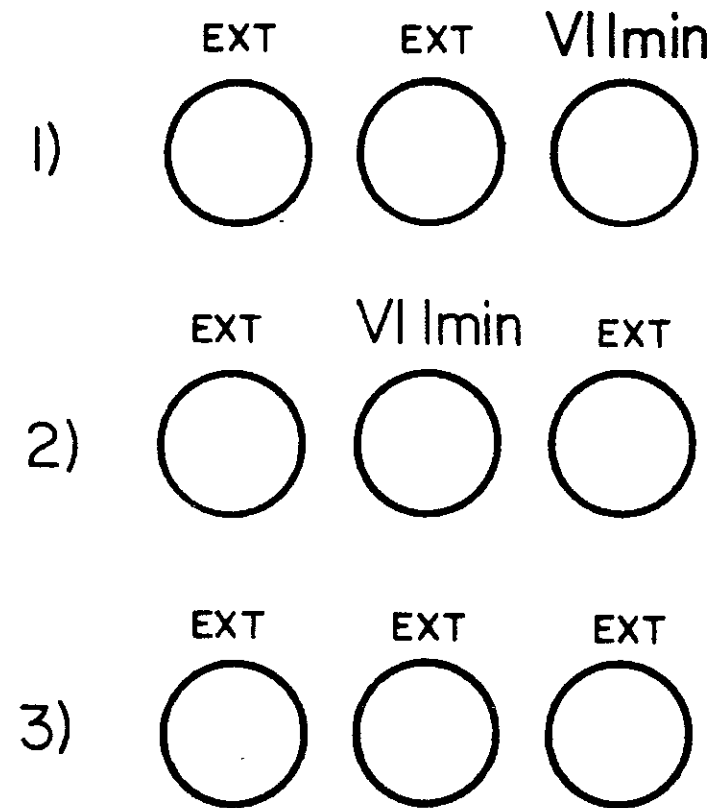


Figure 6.4. Procedure for a Simple Experiment on Resurgence

NOTE: At first pecks on the right key are reinforced with food an average of once per minute. Then pecking is reinforced only on the center key. Finally, in the third condition of the experiment, reinforcement is withheld entirely.

The recurrence of previously reinforced behavior is a common phenomenon, reports of which have appeared many times in both the experimental and the clinical literatures of psychology. If you are turning a doorknob that has always turned easily, for example, and it fails to turn, any and perhaps all of the behaviors that have ever gotten you through doors are likely to appear: You may turn harder, pull up on the knob, kick the door, shout for help, and so on. Freud's concept of regression could be considered a special case of resurgence in which the behavior that recurs is infantile.




Figure 6.5. Cumulative Record Segments from the First Session in Which Reinforcement Was Withheld on All Keys for Bird 13YP
 NOTE: Each of the three segments shows responding on one key. Responding on the center key had been recently reinforced according to a variable-interval one-minute schedule. During the first half of the session, a high steady rate of responding was maintained on this key. It became less stable after the first 1,000 responses. A smooth deceleration is apparent during the last 20 minutes. (Note that the pen resets downward after 500 responses.) The lower line corresponds to the left key, upon which pecking had never been reinforced. The pigeon pecked this key only seven times during the hour. The shaded line corresponds to the right key where there was a history of VI one-minute reinforcement. No responses occurred on this key while responding on the center key was strong, but the pigeon began pecking it at a high rate at about 40 minutes into the session. It pecked the key nearly 900 times during the last 20 minutes.

It would seem that our tool-using pigeon becomes increasingly interested in the box because its behavior with respect to the plate is unsuccessful. As behavior with respect to the box weakens, behavior with respect to the plate should appear. Why the bird pushes toward the plate is another matter, which was discussed earlier.

Note that, in the latter experiment, multiple repertoires of behavior were made available by multiple controlling stimuli, whereas in the tool-use experiment, multiple repertoires were made available by resurgence.

Extrapolating to Human Behavior

There are two ways in which we might demonstrate a relationship between nonhuman animal behavior and human behavior. First, we might use a nonhuman animal to simulate the human case. Rather than providing the animal with an arbitrary set of experiences, as a circus trainer might do, we might first study human subjects to try to identify the experiences they had that allowed them to emit some interesting behavior in a new situation. Then we would provide the animal with comparable experiences. If it subsequently emitted humanlike behavior in the new situation, we would have taken a step toward showing the importance of those experiences in the emergence of the behavior. Moreover, our conjecture that those experiences were responsible for comparable human behavior will have received some support. With laboratory animals, we might then demonstrate that these experiences are necessary for the emission of the behavior; with humans, unfortunately, our conjecture will likely remain a conjecture.

For example, in the insight experiment, pigeons that had learned (a) to push directionally and (b) to climb and reach solved the problem in a humanlike fashion. Did humans and chimpanzees learn these things before they solved comparable problems? The answer seems to be yes, although controlled experiments with children probably cannot be performed. However, Birch (1945) showed that laboratory-reared chimpanzees who had never been allowed to handle sticks could not solve Köhler-type stick problems. After the chimps were given the opportunity to handle sticks, they solved the problems readily.

The second way to show a relationship between nonhuman animal behavior and human behavior is also indirect. One can cast the transformation principles that one has developed with nonhuman animals into formal terms—into equations or a computer algorithm—and then see how well the formal apparatus can predict human behavior.

Note that *neither of these methods tells you what you really want to know*. Neither the experiences nor the principles you have uncovered with nonhuman animal subjects is necessarily responsible for human behavior—even if they successfully predict it. Unfortunately, as is true in neurology, astronomy, meteorology, evolutionary biology, and other disciplines, we must rely on indirect methods to further our understanding of the subject matter. To carry out a rigorous analysis with human subjects would require experimentation that is unethical.

Generativity Theory and the Two-String Problem

In the 1920s N. R. F. Maier devised what has become a classic and frequently studied problem, called the “two-string” or “pendulum” problem (Maier, 1931). Maier confronted adult human subjects with the following situation: Two long strings were suspended from a high ceiling, and various objects were placed on a table near the strings. A subject was told, “Your task is to tie the ends of these strings together. If necessary, you may use one these objects to help you.” In one variation of the experiment, the critical object was a pair of pliers.

Almost invariably, the first thing subjects do is to pull one string toward the other, only to find that the second string is too far away to reach. Having failed with the first string, and geometry notwithstanding, many subjects then try to pull the second string toward the first. Eventually, a subject picks up the available objects. With the pliers, many subjects repeatedly try to reach the second string by holding the first string in one hand and the pliers in the other; the pliers extend the reach a few inches, but not nearly enough to reach the other string.

The solution is to use the pliers (or some other heavy object) to construct a pendulum. One ties the object to the end of one string, sets the string in motion, brings the first string to the center position, and catches the second string when it swings to the center. Many subjects have trouble with this problem. Verbal or nonverbal hints often help. A particularly effective hint, Maier (1931) found, was to set one string in motion slightly by brushing against it. According to our current understanding, such an act would work for two reasons: First by setting the string in motion, the experimenter has modeled some of the behavior the subject must emit to solve the problem. Second, the moving string itself is a “discriminative stimulus” for behavior with respect to pendulums; that is, it makes such behavior more likely.

I presented 30 college students, divided into two groups of 15, with the two-string problem to determine whether the principles I have described

earlier could be used to predict human performances in this situation. Two observers monitored nine behaviors continuously during each session. Observers recorded every instance that they saw of behaviors such as “pulls one string toward the other,” “picks up object,” “ties object to string,” “reaches with object” “sets string in motion from below shoulder line,” and so on. The behaviors were defined precisely and in such a way as to make them easily discriminable. While they recorded their observations, the observers listened, through earphones connected to a common tape recorder, to instructions that allowed them to synchronize their observations in successive 15-second intervals.

The instructions and setup were identical for each group, but the first group had access to a short object, and the second to a long object. These objects were topological distortions of each other, members of a series of objects in which a rod got increasingly longer and thicker, a hook got increasingly larger, and the hook opening got increasingly larger. We conjectured that the short object would produce rapid solutions to the problem, because it might strengthen behavior with respect to weights and pendulums, and that the long object might interfere with a solution, because it might induce subjects to try to extend their reach, as many subjects had done with Maier’s pliers. The long object was not long enough to allow the subject to solve the problem by reaching.

In a baseline study with 148 students, none of whom was a subject in the experiment, our supposition that the short object would produce more rapid solutions was reinforced. Students were shown one object or the other and asked, “What would you do with this object?” Of the students who were shown the short object, 80% indicated they would use it as some sort of weight (“use it as a paperweight,” “make a pendulum,” “a weight,” “a weight on a balance,” and so on). No one indicated that it might be used to extend one’s reach.

In contrast, more than 80% of the students who were shown the long object indicated that they would use it to extend their reach (“pull down shade,” “fish hook,” “open high window,” “ice hook,” and so on), and not one of these subjects indicated that the object might be used as a weight.

Simulation and predictions. A computer simulation of the two-string problem was constructed as follows: Four linear equations were used to represent some of the phenomena I discussed earlier. In each equation, the probability of some behavior during one cycle of the program was determined by the probability of that behavior on the previous cycle, minus some fraction of the latter probability or plus some fraction of 1 minus the latter probability (Figure 6.6). Thus the probabilities always fell between 0 and 1. In each cycle of the program, each of the equations operated on each of the

- (1) *Extinction:* $y_{n+1} = y_n - y_n * \epsilon$
- (2) *Reinforcement:* $y_{n+1} = y_n + (1 - y_n) * \alpha$
- (3) *Resurgence:* for $\lambda_{yy'} < 0$ and $y'_n - y'_{n-1} < 0$,
 $y_{n+1} = y_n + (1 - y_n) * (-\lambda_{yy'}) * y'_n$
- (4) *Automatic Chaining:* for $\lambda_{yy'} > 0$ and $y'_n - y'_{n-1} > 0$,
 $y_{n+1} = y_n + (1 - y_n) * \lambda_{yy'} * y'_n$

Figure 6.6. Equations Used to Generate the Probability Profile

NOTE: y_n is the probability of behavior y at cycle n of the algorithm, y'_n is the probability of behavior y' at cycle n of the algorithm, ϵ is a constant for extinction (it determines the rate at which the probability of behavior y decreases over cycles of the algorithm), α is a constant for reinforcement (it determines the rate at which the probability of behavior y increases over cycles of the algorithm as a result of certain environmental events), and $\lambda_{yy'}$ is the constant of interaction between behaviors y and y' .

behaviors for which an initial probability was specified. Seven behaviors were specified in the case that will be described.

The first two equations represented extinction and reinforcement, respectively. For the two-string problem, it was assumed that no reinforcers were available, and, therefore, alpha was set at 0. Thus every time the program cycled, every behavior was assumed to decrease in probability according to equation 1. The continual and gradual weakening of all possible responses would seem to characterize all problem situations, because by definition, reinforcement is unavailable in such situations.

Equations 3 and 4 represented the phenomena we have labeled "resurgence" and "automatic chaining," respectively. They each made use of a matrix of values that specified the manner in which every possible behavior was assumed to interact with every other possible behavior in the situation. Negative values specified the resurgence relation: As one behavior decreased in probability, another behavior increased in probability. Positive values specified the automatic chaining relation: As one behavior increased in probability, the environment was changed in such a way that another behavior increased in probability. Thus a -0.2 relating "pulls string" to "picks up object" suggests that, when pulling the string is unsuccessful, picking up the object will become somewhat more likely. A 0.4 relating "picks up object" to "ties object to string" suggests that, after one picks up the object, one is then more likely to tie the object to the string.

The interaction matrix is, in effect, a numerical summary of that part of the organism's history that is pertinent to the problem, including the instructions that may have been given to a human subject. For example, if the subject had never learned to tie, 0s would appear at every position in the "tie" row

and "tie" column of the matrix, and, presumably, a solution would not appear. If the task were simply "to select one string and pull it as hard as possible," the behavior "connect the two strings" would probably not appear in the matrix.

The matrix also specifies relationships that are possible between the behaviors in the problem situation. Picking up an object has no obvious relationship to pulling one string toward another, but it does directly change the environment in a way that should make tying more likely.

In future applications of this model, all of the free parameters could, in theory, be obtained in baseline studies. In the present application, only the initial probabilities were obtained from data, and other parameters were estimated. The choice of values was not especially critical in this case. It appears that, as is the case in catastrophe theory, the outcome of the simulation depends little on the value of any one parameter or small subset of parameters. The solution "explodes" over a wide range of values.

The model produces probability profiles of the behaviors that have been specified in it—overlapping curves show the manner in which the probabilities of the various behaviors change over time (Figure 6.7). With just three linear equations (1, 3, and 4) that describe empirically-established behavioral processes, the model generated a solution to the two-string problem. The dynamics of the behaviors seemed reasonable.

Though verbal descriptions in terms of the principles we discussed earlier oversimplify these dynamics (as they did in our discussion of earlier problems), some interpretations can be offered: The probability of "pulling one string to the other" started out high because of the subject's instructions, but, because the behavior was not successful (it produced neither a solution nor stimuli that occasioned other behavior), its probability decreased over time. The extinction of pulling led to the resurgence of other possible behaviors, including "picking up the object" (not shown), which, in turn, increased the probability of "tying object to string." That, in turn, increased the probability of setting the string in motion, which increased the probability of "catching swinging object" (not shown). As the probability of catching increased, the probability of connecting the strings quickly exploded, driven by the long progression of related behaviors. With the long object, the solution still appeared, but it was impeded by the appearance of reaching, which occurred early in the session and then gradually disappeared.

This approach could, in principle, be used to predict the dynamic interactions of any number of behaviors in any situation. To extend the model to other situations, one would estimate the equation constants, the initial probabilities of each of the behaviors, and the interaction values. The success of the prediction would depend on the accuracy of one's estimates, the applicability of the equations to that organism, and the stability of the situation.




Figure 6.7. A Probability Profile Generated by the Transformation Functions Described in the Text, Shown for Five Behaviors Relevant to Maier's Two-String Problem

NOTE: The profile was generated with parameters for the short object (object 1), which produced rapid solutions to the problem and no irrelevant reaching. Note that pulling one string to the other decreases steadily in probability and that other behaviors increase in probability in an orderly sequence. Tying the object to the string makes swinging more likely, which, in turn, makes connecting the strings more likely.

Note that the probabilities generated in a probability profile do not necessarily sum to 1 at any given point in time. Far from being illogical, this is a fundamental requirement of the model. The nervous system can presumably support many behaviors simultaneously, far more than the behaving individual may be aware of at any point in time. When the system is highly active, so that several incompatible behaviors are each highly probable, the individual presumably feels confused and stressed. The system may also be relatively inactive and behavior weak across the board. In any case, the generative dynamics presumably do not stop and may be identical even at very different levels of activity. How probability curves sum at a given point in time—or, more generally, how one decides what behavior one will actually see—requires yet another level of analysis, beyond the scope of the current discussion.

Results. In a unique, novel performance, the probability of a particular response cannot be computed by observing a single subject, just as the

heritability of a trait cannot be estimated by looking at a single phenotype. Therefore, even though the computer model could be used to estimate the probability of ongoing behavior in an individual, it cannot be validated on an individual. Hence, we pooled our data across subjects in each of the two groups.

As we predicted, subjects solved the problem faster and more easily with the short object. All of the subjects who had the short object solved the problem within the allotted 15 minutes, and the average solution time was 2.75 minutes. Only 11 of the 15 subjects who had the long object solved the problem within 15 minutes. If we grant the four unsuccessful subjects 15-minute solution times, the average solution time for the long object proves to be 7.25 minutes.

We also examined transitional probabilities, or, more precisely, the proportion of cases in which one behavior followed another within the same 15-second observation interval or during the next 15-second observation interval. In some cases these proportions were revealing. For example, the probability of tying within an interval of picking up the object was high for each object: 0.59 for object 1 and 0.48 for object 5. But the probability of reaching with the object within an interval of having picked it up or having tied it to a string was 0.00 for the short object and 0.21 for the long object.

Both transitional probabilities and an analysis of films revealed cases of what appear to be automatic chaining. With the short object, for example, the probability of setting the string in motion from above the shoulders (which suggests a deliberate swing) within an interval of having set it in motion from below the shoulders (which suggests simply putting the object down or letting it go) was extremely high—0.47.

When Maier (1931) set the string in motion for his subjects, they solved the problem easily. And, indeed, when one looks closely at successful performances, one often finds that a subject has “accidentally” set the string and weight in motion just before a successful swing. In one case the long object swung back and forth in front of a subject's face while she tied it to the string; she swung it almost immediately after it was secured. In several cases, with the object tied to the string, subjects dropped the object or let it go in a way that suggested that they had “given up.” But a heavy object at the end of a long string invariably swings. The sight of a swinging string soon led to a solution for most subjects.

In order to assess the value of the probability profiles that had been generated by the computer, we constructed frequency profiles for each group. We tallied the number of times each behavior was observed during each tenth of a session. The dynamics of the interactions between the behaviors (treating the group as a single subject) looked similar to the dynamics predicted by the

simulation. As predicted, in each group the frequency with which pulling was observed started out high and gradually decreased over the session. For the group that had the long object, reaching appeared early in the session and then gradually disappeared. The frequency with which swinging was observed increased toward the end of the session in each group, and, of course, the curve for connecting the strings rose rapidly at the end of the session.

Thus the equations accomplished two things: They generated a solution to the two-string problem, and they predicted some of the dynamics of real solutions.

Four Repertoires

The approach I have outlined has proved fruitful in allowing us to construct complex, novel performances in laboratory animals. For example, we recently achieved the spontaneous interconnection of four repertoires of behavior in a pigeon (see Epstein, 1985c; 1987). The pigeon had learned (a) to open a Plexiglas door, (b) to move a box toward targets, (c) to climb (but not to peck anything overhead), (d) to peck the banana when the bird was on the floor and the banana was within reach above it, and (e) not to jump or fly when the banana was out of reach. The bird never saw any two of the training objects (box, banana, or portable doorway) together at the same time.

It was then presented with the following problem: The banana was placed out of the bird's reach, and the box was placed behind the door. In just under four minutes, the bird managed to retrieve the box from behind the door, push it to the right place, climb, and peck the banana. A formal analysis of this performance predicts many aspects of it, including the floor path along which the bird pushed the box.

Generativity and Creativity

It appears that previously established behavior manifests itself in new situations in orderly ways. New sequences, new topographies, or behaviors that have new functions may appear. The manner in which such behavior emerges is describable by a set of transformation functions, each of which operates on every possible behavior that can occur in the situation. The functions predict continuous and probabilistic changes in behavior.

This approach to understanding the emergence of ongoing, novel behavior has proved useful in the prediction and engineering of increasingly complex, intelligent performances in nonhuman animals. It has also led to the devel-

opment of an effective problem-solving algorithm derived from empirically based principles of behavior, and it has allowed us to predict, with reasonable success, ongoing, novel, intelligent performances in human subjects, most recently in studies with retarded children (Epstein, 1988).

Note that I have avoided using the language of creativity in discussing generative phenomena. Such language is heavily value laden (Cziksztentmihályi, this volume), and hence it may obscure an understanding of the generative phenomena on which it depends. Behavior called "creative" by one group might be harshly judged by another. Moreover, the language of creativity is often reserved for the *product* of behavior, not for the behavior itself. The product is of necessity a poor index of the creative process. The product is continuously edited and often rejected by the behaving individual as he or she, having been creative one minute, steps back during the next and *judges* the product—acting, in effect, as an agent for a larger cultural entity. Action is generative, but *reaction* is often corrective and inhibiting. The language of creativity obscures these distinctions and hence should be used with great caution in an analysis of generative phenomena.

Although I have not explicitly discussed thinking and perception in this chapter, it should be apparent that the dynamics I have described are largely covert—that is, they are not necessarily expressed in observable behavior. Generativity theory may in fact be applicable to an explicit analysis of thinking, reasoning, and perceptual processes.

References

- Bingham, H. (1929). Chimpanzee translocation by means of boxes. *Comparative Psychology Monographs*, no. 3.
- Birch, H. (1945). The relation of previous experience to insightful problem solving. *Journal of Comparative Psychology*, 38, 367-383.
- Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1961). *A study of thinking*. New York: John Wiley.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge: MIT Press.
- Enkema, S., Slavin, R., Spaeth, C., & Neuringer, A. (1972). Extinction in the presence of free food. *Psychonomic Science*, 26, 267-269.
- Epstein, R. (1983). Resurgence of previously reinforced behavior during extinction. *Behavior Analysis Letters*, 3, 391-397.
- Epstein, R. (1985a). Animal cognition as the praxist views it. *Neuroscience and Biobehavioral Reviews*, 9, 623-630.
- Epstein, R. (1985b). Extinction-induced resurgence: Preliminary investigations and possible implications. *Psychological Record*, 35, 143-153.
- Epstein, R. (1985c). The spontaneous interconnection of three repertoires. *Psychological Record*, 35, 131-141.
- Epstein, R. (1986). Bringing cognition and creativity into the behavioral laboratory. In T. J. Knapp & L. Robertson (Eds.), *Approaches to cognition: Contrasts and controversies* (pp. 91-109). Hillsdale, NJ: Lawrence Erlbaum.

- Epstein, R. (1987). The spontaneous interconnection of four repertoires of behavior in a pigeon (*Columba livia*). *Journal of Comparative Psychology*, *101*, 197-201.
- Epstein, R. (1988, May). *Generativity theory*. Paper presented at the 14th Annual Convention of the Association for Behavior Analysis, Philadelphia.
- Epstein, R., Kirshnit, C., Lanza, R. P., & Rubin, L. (1984). "Insight" in the pigeon: Antecedents and determinants of an intelligent performance. *Nature*, *308*, 61-62.
- Epstein, R., & Koerner, J. (1986). The self-concept and other daemons. In J. Suls & A. Greenwald (Eds.), *Psychological perspectives on the self* (Vol. 3, pp. 27-53). Hillsdale, NJ: Lawrence Erlbaum.
- Epstein, R., Lanza, R. P., & Skinner, B. F. (1980). Symbolic communication between pigeons (*Columba livia domestica*). *Science*, *207*, 543-545.
- Epstein, R., Lanza, R. P., & Skinner, B. F. (1981). "Self-awareness" in the pigeon. *Science*, *212*, 695-696.
- Epstein, R., & Medalie, S. (1983). The spontaneous use of a tool by a pigeon. *Behaviour Analysis Letters*, *3*, 241-247.
- Epstein, R., & Skinner, B. F. (1980). Resurgence of responding during the cessation of response-independent reinforcement. *Proceedings of the National Academy of Sciences U.S.A.*, *77*, 6251-6253.
- Epstein, R., & Skinner, B. F. (1981). The spontaneous use of memoranda by pigeons. *Behaviour Analysis Letters*, *1*, 241-246.
- Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, *62*, 145-154.
- Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. Englewood Cliffs, NJ: Prentice-Hall.
- Hobhouse, L. T. (1901). *Mind in evolution*. New York: Macmillan.
- Hull, C. (1935). The mechanism of the assembly of behavior segments in novel combinations suitable for problem solution. *Psychological Review*, *42*, 219-245.
- Köhler, W. (1925). *The mentality of apes*. London: Routledge & Kegan Paul.
- Lindblom, L. L., & Jenkins, H. M. (1981). Responses eliminated by noncontingent or negatively contingent reinforcement recovery in extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, *7*, 175-190.
- Maier, N. R. F. (1931). Reasoning in humans. II. The solution of a problem and its appearance in consciousness. *Journal of Comparative Psychology*, *12*, 181-194.
- Maier, N. R. F., & Schneirla, T. C. (1935). *Principles of animal psychology*. New York: McGraw-Hill.
- Mowrer, O. H. (1940). An experimental analogue of "regression" with incidental observations on "reaction-formation." *Journal of Abnormal Psychology*, *35*, 56-87.
- O'Kelly, L. I. (1940). An experimental study of regression. I. The behavioral characteristics of the regressive response. *Journal of Comparative Psychology*, *30*, 41-53.
- Schiller, P. (1952). Innate constituents of complex responses in primates. *Psychological Review*, *59*, 177-191.
- Shurcliff, A., Brown, D., & Stollnitz, F. (1971). Specificity of training required for solution of a stick problem by rhesus monkeys (*Macaca mulatta*). *Learning and Motivation*, *2*, 255-270.
- Sears, R. (1941). Non-aggressive reactions to frustration. *Psychological Review*, *48*, 343-346.
- Sternberg, R. J. (1988). *The nature of creativity: Contemporary psychological perspectives*. Cambridge: Cambridge University Press.
- Wertheimer, M. (1945). *Productive thinking*. New York: Holt, Rinehart.