

SKINNER .MIS
Skinner's Mistake. Essay by Bill Powers

Presented at the annual CSG meeting in Pennsylvania, 1990.

PREPARATORY NOTE: B. F. Skinner died during the last CSG meeting. I wrote this piece as a salute from the opposition, trying to put control theory and behaviorism into the context of the progression of scientific ideas rather than describing Skinner's ideas as merely an aberration. I hope I have not made too many mistakes in trying to show how a reasonable man could have been led astray by adoption of the wrong premises.

SKINNER'S MISTAKE
W. T. Powers
The Control Systems Group

In over fifty years as a psychologist, B. F. Skinner made only one major mistake. He made it when he first observed operant behavior taking place in apparatus of his own invention, and he continued to make it until the day he died. The mistake was to believe that he could see the apparatus -the environment -- shaping and then maintaining the behavior of an experimental animal.

That is not what he saw in those first now-classical operant conditioning experiments. What he did see, what anyone could have seen, was an animal first nosing about a cage, then accidentally pressing a lever, and finally pressing it over and over, repeatedly causing the apparatus to release an occasional bit of food which the animal ate. The mistake lay not in what he observed but in what he imagined. He imagined that he could see something about the apparatus, or the food, or the combination of the two, that was *causing* the animal to change its pattern of behavior.

THE DEVELOPMENT OF 'SELECTION BY CONSEQUENCES'

The reason Skinner thought he could see this causation was that he thought he *should* see it. According to all he had learned, the behavior of organisms is caused by influences in the environment. Spontaneous behavior, he had learned, would be capricious. It would require some kind of magical power to be at work inside the organism. Organisms, as far as he knew, wouldn't act at all unless something stimulated them into action. He had to change his mind about that point, but he never changed his mind about the controlling role of the environment.

The important point is that Skinner knew what he should see before he saw it. What he actually saw was an animal acting through an apparatus to cause food to appear where it could be eaten. What he imagined was that the apparatus and the food were somehow causing the animal to press the lever as it did.

It soon became evident to Skinner that the story was not quite that simple. An animal might employ many different detailed acts to get the lever to go down - - press with a front or rear paw, sit on it, nose it down, or anything else that would work. It might approach the lever from either side or the front, from close to it or from far away. This was clearly not a mere matter of reflexive muscle movements. If that were the story one would expect stereotyped and often futile movements, but those did not appear. Behavior, Skinner decided, could not be described as reflexive responses to stimuli. It could only be defined in terms of classes of detailed responses. The obvious way to form a class was in terms of a common final effect of the behaviors: detailed behaviors with the same consequences are classed as the same behaviors. In the original Skinner box that common consequence was the appearance of a bit of food.

The same consideration applied to the idea of a stimulus. Because the animal might be in many different positions and orientations at the moment the food was delivered, the effective stimulus must be highly variable. Furthermore, changing the size or color of the food pellets did not alter the final

behavior -- many kinds of changes that one would expect to change the stimulus failed to have an effect on the final result. The only consistent way to define the stimulus was in terms of its effect on behavior: again, only a class could be defined, now the class of stimuli that have the same effect or behavior.

The only factor that could be defined non-recursively was the setting of the apparatus: changing the number of lever presses needed to produce each pellet was immediately and predictably followed by a change in the rate of lever pressing.

Finally, there was the matter of how behavior came to be appropriate in the first place. Animals do not normally feed themselves by pressing levers. Initially, the animals did not do so. But as soon as their random wanderings resulted in a lever press and delivery of food, there was a strong increase in the probability that the lever would get pressed again. To account for the initial behavior, Skinner had to introduce the idea of *emitted* behavior -- that is, behavior produced spontaneously, or at least not as the result of any specific or relevant stimulus. At this point, it is said, he ceased to think of himself as a stimulus-response psychologist and -- because that represented a break with the traditions from which he came -- started calling himself a radical behaviorist.

He found that the technique of "deprivation," borrowed from older conditioning studies, revealed a basic effect. Animals maintained at 80 per cent of their free-feeding body weight, a standard experimental condition, were exposed to the apparatus with the ratio set to various numbers. When a large number of presses was needed to produce each pellet, the animals would still press the lever, but at a low rate. As the ratio of pellets to presses was increased, 50 the previous rate of pressing produced more pellets, the animals would begin to press faster. From this observation, Skinner deduced that an increase in the rate of food reward caused the rate of pressing to increase. This relationship continued to hold as long as the behavior did not produce so much food that satiation would begin to appear.

From this, Skinner deduced his principle of operant conditioning in essentially its final form. The consequences of behavior, he generalized, can tend to reinforce the behavior that produces them. "Reinforcing" behavior means increasing its probability of occurrence or its rate of occurrence. Continued behavior is maintained as long as it produces continued reinforcing consequences. If there is more reinforcement, the behavior is maintained at a higher rate.

Skinner then found that altering the setting of the apparatus in a systematic way could be used to "shape" behavior -- lead it gradually from one form to another. In the simplest version of shaping this effect was very dramatic. Starting with an "easy" schedule under which each press produced a food pellet, he could gradually raise the ratio of required presses to pellet deliveries. With each small change, the rate of pressing would rise. Using pigeons pecking at keys, he was able to get the animals to peck thousands of times for each food pellet, over long enough periods to wear their beaks down to stubs. They would do this even though they were getting only a small fraction of the reinforcements initially obtained. It was thus that he showed the power of *intermittent* reinforcement.

Eventually Skinner formulated his proposition this way. Organisms naturally emit behaviors. These behaviors have consequences. Some of the consequences are reinforcing, in the presence of discriminative stimuli. The result is that the initial aimless behaviors are selected by their own environmental consequences, until only those that are systematically reinforced remain. Thus the environment selects, modifies, and ultimately controls the form of behavior. Subjective phenomena such as thoughts, feelings, intentions, and so on arise during this process as inner behaviors, but are results of environmental shaping, not causes in their own right.

Having established this basic explanation of how behavior works (or is made to work), Skinner then began applying these principles to all areas of behavior. He remained consistent with what he believed: he repeatedly admonished readers that the behaviorist must always find ways to express the causes of behavior

so that the environment is ultimately responsible. This was the conviction with which he began his career, even before any evidence was in, and with which he ended in his final speech.

CRACKS IN THE MONOLITH

Skinner's mistake and his steadfast commitment to it led him to overlook a highly significant disparity in his own findings. The basic principle he deduced from the deprivation studies was that an increase of reinforcement increases behavior. But in the shaping studies where very high rates of behavior were induced, exactly the opposite relationship was seen: the rate of reinforcement went down just as dramatically as the rate of behavior went up. He never brought these two phenomena together, because in each case he was looking at what he thought of as manipulations that controlled the behavior of the animal. In the deprivation studies, he was looking at the way increasing the amount of reward per press would increase the behavior rate. In the shaping studies, he expressed the relationship the other way: requiring more presses per reward raised the rate of pressing. If he had expressed this second relationship the same way he expressed the first, the problem might have been more obvious: in shaping, decreasing the ratio of reinforcements to behaviors (so the previous rate of pressing results in fewer reinforcements) causes the rate of pressing to increase.

Only long after Skinner's early years did other workers go back to the original experiments and extend the conditions past the point where so-called "satiation" began. They found that if the schedule were varied between very hard (many presses per pellet) and very easy (at least one pellet for every press), the relationship became an inverted U. When the amount of reinforcement received was low enough, the relationship Skinner expected appeared: more reinforcement, more behavior. But then the behavior rate ceased to increase with increasing reinforcement (the satiation point where the original experiments were ended), and soon it turned around and began sharply dropping with further increases in reinforcement. This was not just "satiation:" it was a complete reversal. The most telling single feature of these data is the fact that the turnaround occurs at levels of (food) reinforcement where the animal would just barely be able to stay alive on that food input. Once that point is reached, further increases in the reinforcement/behavior ratio uniformly cause behavior to decrease. In short, under conditions in which an animal could survive, the real relationship between reinforcement rate and behavior rate is exactly the opposite of what Skinner thought he had found. His focus on controlling the animal through altering its environment led him to cast the shaping phenomenon in a way that concealed the contradiction.

That, however, is only the beginning of the problem. When any fixed idea is carried to extremes it is bound to develop contradictions -- no one idea suffices to explain everything. No doubt if the contradiction had been pointed out clearly, Skinner could have found a way around it. When the desired conclusion is known in advance of the data, there is always a reinterpretation that will make them seem still to be correct. Skinner engaged in a great deal of reinterpretation.

The real problem goes back to Skinner's original observations where he made the mistake of believing as a fact something that he imagined: the causal links from environment to behavior. He thought he had the best of justifications for doing this -- science itself, as he saw it, demanded that all control remain in the environment. He later wrote a book in which he said that radical behaviorism was as much a philosophy as a science. It was the philosophy, not the science, that made what he imagined seem as real as what he observed.

The question that remains is whether any other interpretation of behavior exists that does not involve Skinner's mistake, but still allows us to take advantage of the many novel phenomena that Skinner discovered.

AN ALTERNATIVE EXPLANATION OF BEHAVIOR

Running through Skinner's analysis of operant conditioning phenomena are factual statements that do not depend on an imagined effect of the environment

on behavior. His propositions about reinforcement and operant conditioning in general were all attempts to resolve apparent problems with older explanations of behavior as it is actually observed. In his apparatus, these problems became more evident than in traditional procedures.

Consider the basic observation behind the term "operant." There is no problem in the observation that behavior has consequences. Where the problem arises in terms of traditional logic is in the same finding that William James noted 30 years before Skinner began his work: *variable* behavior has *repeatable* consequences. The appearance is that the consequence, being repeated, is intended by the organism and that the organism simply finds whatever action will create the desired consequence. It was, in fact, specifically this interpretation that Skinner felt he had to refute in the name of science. If the environment is in control, then we cannot have, at the same time, an organism's intentions or desires being in control (aside from the fact that such things are unobservable).

But just speaking factually, it is clear that the focal point of behavior is not the muscle actions that are involved in it, or even the detailed limb movements, but the consequences that keep reappearing. For decades there had been controversy about intentional behavior, but few factual observations ever crept into the arguments. Skinner may have been the first behaviorist ever to face squarely the phenomenon around which the controversies revolved. Behavior *does* vary from one circumstance to another, and the result very often *is* the production of an invariant consequence. The evidence, in other words, actually favors the purposive explanation.

At the time in the 1930s when Skinner was working out a way around these problems, only one kind of mechanistic explanation of behavior was known in the life sciences (although alternatives had been proposed, mainly in biology, since the time of Claude Bernard and before). This explanation was basically the stimulus-response or reflexological view of behavior. Like the machines of the 18th and 19th centuries, organisms were made to behave by prior causes. Muscles were made to tense by prior nerve signals, which were caused by prior stimuli. A chain of cause and effect proceeded from environmental events to stimulus events to neural events to motor events to visible consequences that we call behavioral events. The only alternatives to this picture were cast in abstract and basically untestable forms, invoking invisible factors such as traits and tendencies. Behaviorists scorned such explanations and Skinner was a behaviorist.

When Skinner tried to apply the standard cause-effect analysis to his operant conditioning phenomena, he realized that it did not work. It could not account for the stability of consequences in the face of the variability of behavioral acts. No other mechanistic theory being known to him, he decided to abandon all attempts to guess what went on inside the organism, and following John Watson's credo, rely only on what could be observed from outside the organism's boundaries. He stuck to this approach faithfully until later in his career, when mounting criticisms forced him to speak about "internal behavior" (he never said how one manages to observe that, especially in another organism).

His decision to speak of behavioral classes rather than detailed behavior was his way to put aside the problem. He simply accepted the fact that regular consequences are in fact brought about by variable actions, and classified the actions in terms of the regular consequences. Thus a "barpressing" behavior would be any act that made the bar drop, and an "operant" behavior would be any behavior whose consequences end in a reinforcing event.

It was here that his mistake came into play. Even though a consequence of behavior is clearly a dependent variable in relation to any measure of behavior, even though every step of the causal chain can be observed in the environment, Skinner elected to treat the consequence as an independent variable -- a cause, with behavior as its effect. To be precise, his mistake was to substitute an imaginary causal chain running from the consequence to behavior for the observable causal chain running from behavior to the consequence. behavior, he said, is determined by its consequences. That was a pure act of assertion, going directly against observation, invoked not as a necessary explanation of the observations but as a premise needed to support a predetermined philosophical position.

It may be that Skinner's sympathies with the philosophical stance of behaviorism would have been too strong to allow him to consider any alternative that would go against that stance. But one can speculate what might have happened had Skinner known what was going on in the field of electrical engineering at the same time he was building and testing his first Skinner Box. During the 1930s (starting with H. S. Black's critical insight in 1927) electrical engineers discovered how to build devices in which causation ran in a circle. Moreover, these devices were built in order to perform tasks that formerly required the guidance of a human operator. They were not stimulus-response machines. They could, in fact, vary their output actions in any way required by environmental disturbances, so that some external variable would be brought to a specific condition or be maintained in that condition against perturbations.

If Skinner had observed one of these early control systems in a suitable environment, he might have noticed some striking parallels with the animals in his experiments -- and some striking differences. The most obvious similarity would have been that the actuators of these devices were in constant movement, apparently random, while the variable they affected -- a position, a direction, a temperature, a voltage, a fluid level -- remained almost stationary. Yet by the turn of a knob, the operator of such a machine could make it suddenly bring the variable to a new level, after which it would keep the variable in the new state by further variations in its output. Skinner would have seen that the detailed actions of these devices could not be seen as "reflexive" -- they would have to be classified in terms of their common consequence, which was the stabilized state of the controlled variable.

The most striking difference between these machines and Skinner's animals would have been the machine's lack of ability to alter its behavior in any basic way. It did not start by producing random consequences and then gradually produce more and more regular consequences. Either it produced regular consequences right away or it didn't work at all. It couldn't learn.

But this would have taught Skinner something: a highly significant difference between performance and learning. A functioning control system, which cannot learn, can nevertheless change the apparent character of its behavior when external links from its actuators to the controlled consequence are altered. If, for example, the amount of effect a given actuator movement has on the controlled variable is *halved* (Skinner would have called that a change in the schedule), the actuator movements will *double,* and the controlled variable will remain almost undisturbed. That is essentially the relationship Skinner found in his less-extreme shaping demonstrations. This would have been informative, because the machine clearly did not alter its way of responding to inputs at all (it couldn't), yet its behavior changed radically. An organism showing similar effects therefore might not actually be changing its basic behavior, either.

On the other hand, if the connection from the actuator to the controlled consequence is *reversed,* the control system will instantly fail: its action will drive the controlled variable away from the stable state until a limit is hit or something breaks. The organism will start to do that -but will quickly reverse the sense of its responses and once again produce the same stable consequence. Something has clearly changed in the organism in a way that the machine is not designed to change.

The organism can, in effect, rewire itself -- that is learning. Both the organism and the control system can adapt instantly to a large variety of alterations in the external world without changing the internal wiring at all; both can quickly resist unpredictable disturbances: that is performance. Skinner did not make that distinction: he treated all behavioral changes as if they were part of a single phenomenon. But perhaps he would have seen the difference, had real control systems been available for his inspection.

Without changing his behavioristic philosophy at all, Skinner could have found a great variety of suggestive phenomena in the behavior of control systems, had he known about them. But if he had pursued such studies and learned how control systems actually work, he would have been faced with a philosophical dilemma. The only effect that external circumstances have on the operation of

a control system enters the system as a sensory report on the state of the external controlled variable. If organisms are like control systems, then the only effect a reinforcing consequence can have on the organism is to change the state of its perceptions (or physiological condition, which is also sensed). If any *additional* sort of change is to take place, the agent of change must exist inside the organism, not in its environment. If the artificial control system is to alter its basic way of acting in response to radical changes in the environment, it must be equipped with internal means of assessing its own capacity to control, and altering its own parameters accordingly. Nothing would have prevented Skinner from proposing that changes in the external environment were the cause of those alterations (in a hypothetical adaptive control system, which was not to exist until much later than this imaginary scenario would have happened). But the engineers could easily have shown that no such causal link existed -- or was needed. Adaptation in an adaptive control system is made *necessary* by changes in the environment: but it is made *possible* by capacities inside the system itself.

Perhaps, had he been able to undergo this experience with control systems, Skinner might have remembered why it was that he abandoned attempts to explain behavior through theories about the internal organization of organisms. He did so because none of the theories he could find worked. Either they were highly specific and specifically wrong (as in S-R theory), or they were vague and untestable (personality theory and so on). But along with control systems there comes another kind of theory, control theory. This theory explains how control systems behave the way they do. It is a highly specific theory suitable for quantitative testing, and it works. The behavioral properties of control systems mimic the behavioral phenomena that Skinner discovered; his operant conditioning experiments are set up exactly as a control-system engineer would construct an experiment to measure the properties of a control system concerned with control of the same variable -- what Skinner called the reinforcer.

He would not have found all the answers he wanted in control theory. But control theory was the outcome of an effort to imitate human behavior: specifically, the kind of behavior that is called control. As far as it goes, it imitates this behavior well enough to have brought on the automation revolution: the replacement of human controllers with automatic controllers. To extend it far enough to suit the purposes of a student of living organisms, it must be developed further, into realms of complexity that do not interest practical engineers. Simple adaptive control systems must be developed into systems that can alter their own organizations in far more fundamental ways than any artificial system now can, before they can match the behavior of a naive rat in a Skinner Box.

By the time the attempts of cyberneticists came along, and the first engineering psychologists, and the early behavioral control-system theoreticians, Skinner had gone far down a different road, following out the implications of his initial assumptions. Control theory turned from a potential source of enlightenment into a rival view and a threat. Radical behaviorism came under general attack, and in defending against that attack, Skinner and his adherents did not single out control theory as a different sort of deviation from what Skinner proposed. Perhaps it was inevitable that Skinner would have to play out the whole scene, bring his structure of thought to its own logical conclusion, and never see that control theory would vindicate his earliest observations -- even if it showed that he chose the wrong conception of causation. In some alternate universe, perhaps it came out differently. That B. F. Skinner would probably be having a wonderful time as a control theorist.