DEMO1: Introduction to Perceptual Control Theory

A TEACHER'S TUTORIAL AND GUIDE, PART I

Demo1 is intended as a self-paced introduction to the phenomenon of control. In a series of instructional screens the user is guided through the basics and is called upon to exert increasingly informative types of control. First there are simple compensatory and pursuit tracking, followed by extension of the same principles to controlling orientation, size, shape, the pitch of sound, and other kinds of variables. Some basic principles are illustrated, among them being the principle that living control systems, like all others, control their own perceptions by varying their actions in answer to environmental disturbances. They do not control their actions.

Some preliminaries

To help me fix the audience in mind as I write, I'm going to assume that you're a person who is professionally interested in human behavior—a psychologist or sociologist, perhaps, and a teacher. To make my job a little easier, I'm going to assume that you don't know very much about control theory, but that you're willing to learn. I will also assume that you have your own ideas about what is wrong with the present state of theory about human behavior and that you're moderately to very dissatisfied with it. Therefore, while I may drop a remark here and there about conventional ideas, I won't dwell on what you already know is wrong and I'll teach control theory as if no other theories existed. Perhaps by the time you have learned all the material you will agree that this approach is appropriate. You can do your own straining of babies from bathwater.

It will be helpful to run the program as you read this manual. I'll be looking over your shoulder, as you will do with students, offering amplifications and tips that will help you go beyond the terse descriptions one can cram into a small display.

This first program can easily serve as the basis for a one-quarter or one-semester course. As you go through it, preparing to teach the material, you will probably see how each point that is made touches on central problems in behavioral theory—it would not be difficult to spend a week on each screen after the first few introductory ones. Those who teach from these programs will probably put together different structures for the course, but all, I hope, will end up as the students will—wondering how on earth these strange phenomena could be explained by any theory. That is the point of Part 1. The theoretical model is built up in Part 2.

When the program comes up, a menu screen appears. The highlight points to the first item: select it by pressing the Enter key (also called the Return key, as it acts like a carriage-return on a typewriter). Each time you finish a screen and return to the main menu, the highlight will be on the next item. Just press Enter again. You can also type the letter corresponding to the item. You can, however, select items in any order. I recommend going in sequence the first time.

So let's begin.

Step A: Introduction

The purpose of Part 1 is to explore the phenomenon called control, not to explain it. Control is the process by which organisms (or machines designed to imitate organisms) act on their environments to create and maintain specific effects that matter to the organisms.

The term "behavior" refers to the visible acts of organisms and to consequences visible to an external observer that those acts produce. Before the process of control was understood, behavior was thought to end with those acts and consequences. The scientific approach to behavior, given that definition, quite logically focused on looking for its causes: the circumstances and antecedent events that seemed systematically related to behavior.

The goal of studying behavior scientifically was not just to understand it: it was to predict it, and ultimately to control it. Most behavioral scientists adopted that goal—after all, if sciences like biology and psychology were to have any value to society, they had to produce social benefits such as the curing of illnesses and the prevention of crime.

In its most general form, control results in some aspect of the outside world being brought to a specific state by the actions of the controlling system. Because the world is variable and full of unexpected disturbances, there is (outside the laboratory) no one action that will always produce the required result. What marks the process of control as something special is that fact that a successful control system can produce whatever action is required to produce a specified result.

An excellent example of a scientist behaving this way was given by the behaviorist B.F. Skinner, in his description of how a behaviorist "shapes" an animal's behavior. It is, of course, necessary to know beforehand what the final behavior of the animal is to be, and preferably to state it (lest anyone think that one is just accepting whatever happens and claiming to have caused it). But as Skinner pointed out with some care, there is no particular act that a scientist can perform that will result in shaping the behavior of an animal. Instead, the scientist must watch the present behavior, and whenever it seems to include a move in the direction toward the desired final form, reward the animal to reinforce that part of the behavior. One has to improvise, said Skinner, and be ready to change course at any time. While a scientist can shape the behavior of one animal after another to create the same final behavior, such as walking in a figure-eight pattern, the actions of the scientist during this process probably are never the same for any two animals.

That is the essence of control. One must have an internal picture of the final outcome that is desired. One must be ready to vary one's actions so as to keep the current state of affairs moving toward a match with the internal picture. The same result may occur over and over, but the means of producing it do not normally repeat. Under highly controlled conditions, repeating an action might result in repeating its consequences. But under ordinary conditions there is too much variability in the world where action takes place to allow that. Nevertheless, control does happen, regularly and reliably, in that variable world.

In observing the process of control, we do not see any "internal picture" in the controlling system (unless we are the system). Such ideas belong in theories. Our goal for now is just to look at this process called control, to explore its characteristics in ways ranging from the obvious to the surprising. There are many details to notice, many relationships that can escape one's attention. And there is a special way of looking at behavior that is required before these details can be grasped and appreciated.

Steps B And C: Measuring action

It has long been the custom in most of the behavioral sciences to observe, record, and analyze behavior in terms of "events." A person answers a question, picks up a wallet, chooses an alternative, responds to a stimulus, does some named thing when some other named thing happens.

In this view, nature is divided into packets that either occur or do not occur. What happens inside the packets, or in the intervals between then, is ignored. The processes that lead from one recognizable packet to another are ignored. Even the question of how, physically, one event leads to another is ignored. All that is recorded is the fact that one event occurred, and then another event occurred or did not occur.

To make this kind of analysis look more scientific, scientists employ statistics and search for events that seem to follow systematically from others. This does not mean that a causal event is always followed by the effect-event. It means only that in the long run, over many trials with many subjects, there is a tendency for the effect to follow the cause. By its very nature, this approach does not predict individual occurrences. It only tells us how to bet, assuming that we will have many opportunities to bet and won't go broke before the odds have their effects.

More important, this method of looking at statistical regularities among named events can't tell us why those regularities occur. Even if the odds are so good that the bet is a sure thing, we still don't understand what is going on.

The theory of control systems is expressed in terms of continuous relationships among variables. To prepare for understanding that theory, we must begin observing behavior in those terms.

In Steps B and C, we start making the transition to thinking in continuous and quantitative terms. On the screen, in step C, is shown a number. This number changes as you move the control handle. The point is to show that there is always a number corresponding to handle position. There is no time "between handle movements." The handle always has a position that is representable as a number, whether it is stationary or moving. You can create "events" by patterning the movements: up-down-up, for example. The number simply follows. The parceling of handle behaviors into events is entirely a product of the observer's way of looking at them. What is actually going on is that the computer is continuously converting handle position into an internal number, which it displays on the screen 30 times per second (even when the handle is stationary and the same number is shown over and over).

Step D: Effect of handle on cursor

In Step C, you move a control handle to cause a cursor to move. The cursor position is also shown as a number that always represents cursor position (here it's the same as the handle number). When you move the handle, the cursor moves in exactly the same way, and the numbers representing cursor and handle position vary together.

We could describe what happens in terms of events: "When the handle moves up, the cursor moves up." Instead, we will look on the relationship of handle position to cursor position as continuous, with no natural planes of cleavage to break one event apart from the next event. We will see that the positions of the handle and of the cursor vary together in a specific relationship. The same relationship is present no matter where the handle and the cursor are, whether they are stationary or moving, whether they are moving fast or slowly. The cursor position is always proportional (by a factor of 1, for now) to the handle's measured position, relative to the center of the screen. That relationship of proportionality is timeless. No "sequence" is involved: handle and cursor move together.

This demonstration involves a real handle and a real cursor; you could see them whether or not they were represented by numbers. At the bottom of the screen, however, we have both handle and cursor shown as numbers. The numbers vary as the handle and cursor move. We are taking a step toward modeling behavior, by seeing how symbols can represent physical variables. Here what matters are the quantities associated with the physical variables; the numbers associated with handle and cursor. The numbers and the variables are continuously tied to each other; by watching the numbers you could estimate where the physical variables are in their range of variation. The numbers don't represent everything about the handle and the cursor—their color, for example. But they represent the aspects of handle and cursor with which we are concerned here: their positions.

Perhaps the most important fact about this exceedingly simple demonstration is that we understand everything about it (except how the human being is moving the handle). There is nothing mysterious or metaphorical in what we see. The numbers are created by a process we can understand in as much detail as we care to know. We might have to ask a programmer how this result is brought about, but we can be perfectly confident that there is an explanation. We can be sure that the explanation will not involve any mysterious statistical "tendencies," or that it occurs only because that it how it has always occurred before.

As you come to understand the process of control, you will see that it is always exactly this understandable.

Step E: Disturbances

The text screen for this step asks you to select a disturbance number from 1 to 10. There are 10 pre-calculated tables of disturbances in the folder, kept in a file called MAINDIST. If you want a new selection of disturbances, delete the file MAINDIST before starting the program. The program will then create a new master table of disturbances.

On the second screen of step E, we see the same thing we saw before: a cursor that you can move with the handle. Now, however, we find that the cursor moves even when the handle is stationary. The reason is that a disturbance is being applied to the cursor independently of the effect of handle position. The cursor position is now determined by the sum of the disturbance and the handle position. The magnitude of the disturbance is shown as a number at the bottom of the screen, along with the cursor and handle numbers. As you can verify by starting and stopping the action (using "p" for pause), the cursor number is always the sum of the handle number and the disturbance number. The cursor position still corresponds exactly to the cursor number.

In the previous steps, you could tell where the handle was by looking at the cursor, and you could tell where the cursor was by looking at the handle. If the handle was up, the cursor was high on the screen. If the handle was moving downward, the cursor was moving downward (if you have a game joystick or mouse, "up" means "forward," and so on). That is no longer true. If you move the handle slowly up and down, you will see that the cursor may or may not behave the same way. The handle might be moving up while the cursor is moving down. With the cursor in the center of the screen, the handle might be anywhere in its range of movement. The position of the cursor is no longer dependent on handle position alone. It now depends partly on the magnitude of the disturbance. The disturbance, in fact, can have just as much effect on the cursor as a full-scale change in handle position has.

When we speak of "the disturbance," it should be clear that we mean the invisible cause inside the computer, not the effect on the cursor. If the disturbance number is increasing while the handle number is decreasing at the same rate, the cursor will not move; the effect that the disturbance would have when acting alone is being canceled by the changing effect of handle position.

The disturbance is being created by the program: it is permanently invisible. We can see, on this screen, a number telling us its magnitude, but we cannot see the disturbance itself, which exists as electrical voltages in silicon chips. To get a proper feel for the process of control, you have to recognize the difference between a disturbance as an independent influence and the effects of such a disturbance. When you see tree branches waving back and forth, you might casually say to yourself that a wind is blowing. You are seeing, however, the effect of the wind, not the wind. The wind is invisible. If the tree branches were wired invisibly in place, the wind would still be exerting the same varying forces on the branches, so the disturbance would still be present—but it wouldn't be having any effect. Assume that you can sense the wind in other ways, such as by reading a number on an anemometer. If you understood that normally a wind ought to move the branches of the tree, you might deduce that the wires, or something equivalent, must be present—something is acting on those limbs to create forces opposite to the wind forces. Otherwise the limbs would be moving.

So what we might normally mean by "a disturbance" should be taken to mean "the effect of an unopposed disturbance," and whenever we use that word in a technical sense, as we will do here, you should hear it as meaning the physical variable that is behind the effect. This means, of course, that we must also understand that there are mechanisms connecting the cause to its effect. The wind, just by existing, doesn't affect the branches. It's the friction of the air moving past the limb that produces the effect. How much effect there is depends on laws of aerodynamics, the shape of the limb, and so on.

In Part 2, when we model the effects of disturbances, we always put a box between the cause and the effect. The box expresses the physical link, which might be an abstract law or a series of concrete intermediate linkages.

We will usually speak of "the" disturbance, but in natural situations there might be many independent disturbing influences acting at the same time. A moment's thought, however, will show you that it doesn't matter how many disturbances there really are. All that matters here is their net effect on the cursor. The disturbance number at the bottom of the screen measures that net effect. Also, even a single disturbance may have effects on many things. But we are only interested here in one of those effects—the effect on the cursor's position. It wouldn't matter if the computer contained six sources of disturbance all acting at once on the cursor. It wouldn't matter if that same disturbance were making some integrated circuits heat up. We're interested only in the net effect on the cursor, which we can attribute to a single equivalent disturbing influence. That is why we need only one disturbance to represent all possible disturbances.

The main point to understand now is that the cursor position no longer depends reliably on handle position, and that the reason for this is that an invisible and variable disturbance is contributing to the cursor position.

Step F: Compensatory tracking

If you have ever driven a car in a gusty crosswind or past a speeding truck, you know intuitively the difference between a disturbance and its effect. The whole point of steering the car is to keep the wind-disturbance from having an effect on the path of the car. When the wind is blowing from the right, you twist the steering wheel to the right. Your steering effort applies one force to the car; the wind applies another. In order for the car to go straight, these forces must always be equal and opposite. That's a simple fact of nature known since the 16^{th} Century.

In Step F, we find a similar situation. There are two stationary target bars in the right center of the screen. Between them is the cursor, which can move up and down as you move the handle—and as the invisible disturbance varies. The task is like that of steering a car in a straight line. You are trying to keep the cursor between the target marks. The disturbance is always tending to push the cursor up or down by varying amounts. To do the task you have to move the handle up and down just enough to have an equal and opposite influence on cursor position.

When you actually do this, it turns out to be easy. The disturbance doesn't change magnitude very fast. You can keep the cursor quite close to the target. After the experimental run, which you start by pressing the space bar (read the instructions on the screens!), the results are plotted and the program computes some very interesting statistical results.

To read the plot (not everyone finds this obvious), scan from left to right, following one trace. Imagine that you're dragging an index card slowly along, watching the place where the trace comes out from under it (you can actually do this). The point where the trace becomes visible moves up and down, representing the rise and fall of that variable as it actually occurred. The heavy trace shows what the handle did; the symmetrical light trace shows the disturbance variations. The trace that wiggles up and down near the middle shows how the cursor moved.

The disturbance trace shows something that you couldn't see during the run, the variations in the cause of cursor deviations. Imagine for a moment that this trace is missing from the plot. What is left, then, is the cursor trace showing how the cursor moved, and the handle trace showing how the handle moved.

You could, of course, see the cursor. In conventional terms, the movements of the cursor would be called a "visual stimulus," and the handle movements would be interpreted as a "response" to that stimulus. Just looking at these two traces, however, can you see any relationship between them? Does the handle rise with every rise in cursor position? Is the handle trace always above center when the cursor trace is above center? Can you even see any trend that is the same in both traces, ignoring their different sizes?

We can find out more objectively what relationship there is between the supposed "stimulus" and the "response." All we have to do is calculate the correlation of handle position with cursor position over the whole run (900 data points, of which only about 1 in 10 is shown). If you read the text on the screen you will see that correlation filled in after the run; it is computed from the actual data. It will be somewhere around 0.2 or less. That amount of correlation (1.00 is perfect) would be interpreted by any statistician as no relationship. The more you practice this task and the better you get at it, the closer this correlation will come to zero.

This is the first large surprise that comes out of looking carefully at control phenomena. We have a situation in which a response is unrelated to the only visible stimulus on which it seems to depend. Of course mere lack of a correlation would not be so surprising by itself, but there is another relationship to consider. Also filled in in the text at the top of the screen is the correlation between the handle position and the magnitude of the invisible disturbance. That correlation will be around –0.95, higher with practice. That is a nearly perfect negative correlation, meaning that the handle position varied in a way almost exactly equal and opposite to the disturbance variations.

The only way you could get any information about the magnitude of the disturbance was by watching the cursor, because the disturbance's presence is indicated only by its effects on the cursor. But as you can see, there is no relationship between the cursor movements and the disturbance variations (that correlation, not shown, is about the same as the correlation between cursor and handle—0.2 or less). So we have the handle movements showing no relationship to the cursor movements, yet showing an extremely high negative correlation with the disturbance variations, which can be known only through their effects on the cursor.

Under any conventional understanding of how behavior works, this is impossible.

But do the run again. We are looking at facts. By varying your action, you maintain an external variable in a state that you want for it: between the target marks. In doing this you automatically cancel the effect of a completely invisible disturbance on the variable you are controlling.

Step G:

In Step G, a second disturbance makes the target move. Now to control the position of the cursor relative to the target you must move the handle to counteract the effects of two independent disturbances which act on the visible relationship in different ways. The results are the same.

Step H: Beyond tracking

Tracking experiments have been done for over 40 years, without anyone's appearing to have noticed the odd facts we have just been exploring. But tracking is just one example of control behavior.

In Step H you can choose any of five different variables to control: the size of a geometric figure, the orientation of a figure, the shape of a figure, the pitch of a sound, and a purely numerical display (trying to keep the displayed number at 50). One point is to show that position is only one of many variables that can be controlled. The second point is to show that the action by which control is exerted is of almost no importance. In fact, the same action is used in all these examples: moving the control handle just as you did before. Furthermore, the same disturbance is used in all the examples, so control actually involves exactly the same physical movements of the control handle, nearly to the last detail.

To satisfy purists, you can quit this step, choose it again, and repeat it with a different disturbance, choosing the tasks in a different sequence. This should prove to you that no memorization of the disturbance pattern (of any significance) is going on. You control just as well with a new disturbance as with one you've experienced before. The correlations prove it.

After each experimental run, the results are plotted as before, and the same correlations are calculated. We are looking at the same phenomenon despite the fact that different variables are being controlled—different not only in form but in basic perceptual type.

This step should begin to show that control is more a matter of perception than of action. The actions we use to control all the myriad variables of experience are of few types: basically they boil down to push, pull, twist, and squeeze. What makes the difference between controlling the pitch of a sound and the shape of an odd line drawing is not the movements we make with our limbs or which muscles we tense, but the perceived effects of those acts—and the inner knowledge of what effect we prefer to experience. In the example in which the pitch of a sound is controlled, the only reference is given by the tone at the start, which goes away once the run begins. As you keep the pitch near that initial pitch, what are you using to judge whether the presently-heard pitch is too high or too low? All that is available is the memory of the initial pitch. You are comparing what you now perceive with what you first perceived: a memory. This tells us something, but takes us into theory, and we are still looking at phenomena.

Step I: Accidental vs. Intentional consequences

One of the puzzles of behavioral theory is that actions may have many consequences, but only some of those would be considered intentional. When you make your car veer, it may simultaneously avoid a pothole and scrape a fender against a stop sign. Both avoiding the pothole and scraping the fender are outcomes of the same action; either outcome could be called a "behavior" in conventional terms. We might say, "That person is always scraping fenders against stop signs," defining the actions by their outcomes. But few drivers would claim (especially to their insurance companies) that they scraped the fender on purpose, while all of them would probably agree that they intended to avoid the pothole.

The question has always been, how is intentional behavior different from unintentional behavior? Those who claim to take the purely scientific point of view maintain that there is no difference—that intentions, being "mental" if they exist at all, are at best side-effects of responses to stimuli, and have no force before the fact, no causal capacities. In short, the usual claim is that all outcomes of action are equally valid measures of behavior. At least no scientist operating from the conventional base of interpretation has found any way to distinguish accidental outcomes from intentional outcomes.

We will now see that with the process of control in mind, rather than reaction to stimuli, not only can we distinguish reliably between intentional and accidental effects of action, but a computer can do the same thing.

In this step you will see not one but three cursors that move up and down between the target marks. As you can verify by moving the handle before starting the run, the handle has exactly the same amount and direction of effect on all three cursors. This remains true throughout the run: any movement of the handle has the same effect on all three cursors.

The task is simple: pick any one of the cursors, and hold it level with the target marks. This is just as easy to do as when there was only one cursor.

Each cursor, however, is subject to a different pattern of disturbance. There are now three independent disturbances tending to move the three cursors. This means that a handle position that would bring one cursor to the center (by canceling the effect of its disturbance) would not necessarily bring either of the other two to the center. In fact, most of the time it won't, except in passing and by chance.

After you have chosen a cursor to control and have started and completed the experimental run, the results are shown for all three cursors. One of the plots shows the symmetrical relationship between handle and disturbance that is by now familiar, with the cursor remaining nearly centered. The other plots show exactly the same pattern of handle movement, but very different behavior of the disturbances and the cursors. The computer calculates the correlation between handle and cursor movements for all three cursors, then chooses the lowest correlation as indicating which cursor you were controlling. It tells you that you were controlling LEFT, MIDDLE, or RIGHT. No hedging.

It always chooses the correct cursor.

Do the run several times, picking different cursors to control (or the same one if you like). You will see symmetry only for the cursor you chose to control. The computer will show that this is the cursor with the lowest correlation with handle movements. It is, of course, also the one for which handle movements showed the highest negative correlation with the disturbance, but we are assuming that the observer can't see the disturbances any better than the participant can. Using only the information visible on the screen, we must search for the lowest correlation of handle with cursor—and that is always the key to choosing which cursor was under intentional control.

Richard Marken (1989) has shown that this same principle will distinguish accidental from intentional consequences of action even when the intended consequence is a pattern of movement of the cursor (or other object) chosen by the subject and incomprehensible to any other observer.

If by the term "behavior" we mean the actions of the participant, it is clear that intentional behavior is exactly the same as accidental behavior. The actions are just actions and their consequences are just consequences. There is no extra flavor of intentionality to be found in either one. In the present experiment, we know for certain that there is only one action, and that its effects on the three potential outcomes are identical for each cursor. Whatever it is that makes one cursor's movements intentional and another's not is not to be found in the action or any of its immediate effects.

The difference is only to be found in the participant's deciding that one cursor is to remain between the target marks. We might guess, and the participant would most likely agree, that the participant looked at one cursor and selected for it the target condition "between the marks." The actions were then based on deviations of that cursor alone from the specified position. The other cursors were ignored.

To anticipate the final screen, it is also perfectly possible for the participant to decide to hold one of the cursors some fixed distance above or below the target marks. The computer will pick that up, too, as correlations don't depend on constant offsets. Or the participant could decide that one of the cursors is to move up and down in some regular way. If the participant is very clear about that pattern and really tries to maintain it, the computer will select the correct cursor again—not every time, because the correlations from cursor to cursor will be less different, but definitely over a series of runs. To maintain even an arbitrary but predetermined pattern of cursor movement, the participant must systematically oppose the effect of the disturbance that is acting on that cursor. That systematic opposition reduces the effect of the disturbance on the cursor, and lowers the correlation of cursor movement with handle movement, too. Given enough time to compute valid correlations, the computer can always find out which effect of action was intended.

When we understand behavior as control, there is no longer any mystery about intentionality. Intentions are an integral part of the process of control.

Step J: Changing the feedback factor

The demonstration in this step can be seen just as an oddity of control, or as having far-reaching significance—that depends on what you believed before. All that is involved here is altering the amount of effect that the handle has on the cursor. Through all the previous demonstrations, one unit of handle movement (as measured by the computer) had one unit of effect on cursor position. Now you are allowed to change that factor.

This demonstration is worth going through because it refutes once and for all the idea that we are looking at a one-way causal chain. The conventional one-way analysis must treat cursor movements as consequences of handle movements. It really can't account for the handle movements by treating the cursor as a stimulus (handle movements don't correlate with cursor movements); it must accept the disturbance as the real cause of the actions, because the actions correlate almost perfectly with the disturbance. If we keep the disturbance the same the handle movements should remain the same; if we double the effect of handle movements on the cursor, the cursor should move twice as much. And of course if we halve the effect of handle on cursor, the cursor should move half as much.

After choosing a factor between 50 and 200 percent, and after doing the experimental run, you will see a plot of the results that shows the exact opposite of what the conventional view would predict. When you double the effect of handle on cursor, the handle moves half as much. When you halve the effect of handle on cursor, the handle moves twice as much. The cursor stays put.

If you consider why this is so—in control-system terms—you will see that this result is obvious. With the handle effect doubled, if the handle moved the same amount as before, the cursor would be moving radically up and down—just as far as the disturbance alone would move it, but the other way. The participant would be overcompensating by 100 percent. What actually happens is that the participant keeps the cursor stationary as usual. If the handle has twice the former effect, this stability of the cursor can be and is achieved with only half the handle movement. If the handle has half the former effect, the handle has to move twice as much to keep the cursor from moving. That is what happens.

Step K: Control of remote effect

This demonstration is in the same vein as the previous one. It shows that the conventional view of cause and effect in behavior is wrong.

The conventional view sees motor behavior as the outcome of a series of prior processes. S-R psychologists trace the prior processes all the way back to sensory stimuli; cognitive psychologists, some of them, trace back only to the level of cognitive decisions. Neurologists see behavior as originating as neural activity in higher centers of the brain. In any of these views, if there is some regular remote consequence of motor actions, it follows that everything between the motor actions and the consequence must have remained the same. This is not true.

In this demonstration the "motor action" is the positioning of a cursor, as previously. But the cursor has attached to it the end of a "string" that passes down, around a "pulley," up and over another "pulley", and back down to a free end. As the handle and cursor move up and down, the free end of the string moves up and down. The task is to keep the free end aligned with a mark on the screen.

The catch is that both pulleys, when the run starts, begin to wander up and down, independently. They therefore disrupt the regular connection between the behavioral action and its remote result, the position of the free end of the string. When you carry out the task, you will find that it is trivially easy if you just watch the free end of the string and ignore the pulleys. If some theory tells you that you must watch the pulleys and calculate their effects on the string, and move the handle so as to compensate for those effects, you will fail miserably.

Many famous theoreticians, including cyberneticians among them, have proposed that remote effects are stabilized by sensing the causes of disturbances, calculating their effects on the remote consequence, and computing an action that will have just the opposite effect when it gets to the end of the line. Obviously none of them ever tested that idea. We can easily test it here. All you have to do is cut out a small rectangle of cardboard and tape it on the screen to hide the cursor and the pulleys, so all you can see is the free end of the string. If you were actually paying attention to the pulleys and calculating the response that would compensate for their movements, this should completely prevent you from keeping the end of the string where it is supposed to be. Of course nothing like that happens. You control the end of the string just as easily as before. Maybe more easily—you won't be tempted to look at the pulleys.

Control relies not on computing outputs, but on sensing whatever is to be controlled and comparing what is sensed with an inner standard that defines what is to be sensed. Control is organized around perception, not around action. Control systems do not compute output. They control sensory input.

Step L: Feedforward vs feedback

Engineers use the term feedforward to mean reacting to the cause of a disturbance in a way that roughly compensates for it, thus making the task of control somewhat easier (and in some circumstances, faster). An example would be connecting an outside thermometer to a sensor that would turn the furnace on for a few minutes when there was a sharp drop in outside air temperature, or off when then there was a sharp rise. With careful adjustment of this effect, the regulation of inside temperature might be somewhat better than it would be if the thermostat waited for the inside air temperature to change, especially in parts of the house far from the thermostat's sensor. You can probably see that there would be some problems with this method—the system would at least have to know whether the outside air was cooler than the inside air, which implies sensing the inside temperature anyway. Feedforward is seldom used in engineering control systems, and never by itself when control has to be even moderately accurate. It is found mostly in pre-20th-Century devices such as temperature-compensated pendulums or barometers.

But the engineering uses of this term cannot account for the way it has been used in psychology and biology. The most likely explanation for its appearance in the literature is a sense of verbal symmetry: if there is feedback, why shouldn't there also be feedforward—or for that matter, feedsideways?

The idea that behavior is organized around feedforward is, once you experience feedforward for what it is, impractical. To work with feedforward means sensing a disturbance and reacting to it in exactly the right amount and in the right direction to stabilize some remote effect that you can't see. If you can see it, it isn't feedforward but feedback that is at work. So you are guessing how much action would be called for to prevent a change in a variable that you aren't sensing, which is basically impossible.

In this step, you can go back and forth between feedback and feedforward. In the feedback version you keep the cursor between the target marks—this is just compensatory tracking. You can start paying attention to how the handle moves, getting an idea of its range of movement.

In the feedforward phase, the "cursor" on the screen now represents the magnitude of the disturbance; the real cursor is not shown. You are asked to watch the disturbance and adjust the handle so that the real cursor would remain between the target marks. As you know by now, this means keeping the handle displacement from center equal and opposite to the disturbance, now shown as a displacement of the cursor from center.

As you go back and forth between feedback and feedforward, you gradually learn how much to move the handle, and you learn to make the handle pass zero just as the disturbance passes zero (but going the other way). Looking at the plots that are shown each time tells you whether you are over or under-reacting to the disturbance in the feedforward phase. With a few rounds of practice you can learn to compensate quite well with feedforward. But you never learn to control as well as you can do with feedback, when you can see the cursor. Controlling the position of the cursor is best done by watching the cursor, not by controlling a symmetry on which cursor position indirectly depends. You can't judge symmetry as well as you can judge position.

On the screen are shown, after each run, various measures of how good the control was. The root-mean-square deviations of the cursor are much larger when you can't see it. The handle position isn't nearly as symmetrical with the disturbance when you can see the disturbance but can't see the cursor.

What is happening here is not feedforward at all, but a higher level of feedback. What you're learning to do in the feedforward phase is to control the relationship between sensed handle position and sensed position of the mark that represents the disturbance. Your're adjusting the scale factors involved in perceiving handle position as "equal and opposite" (or better, proportional and opposite) to disturbance/cursor position. You're also finding where the handle should be when the disturbance crosses zero.

This adjustment would be impossible if you never saw how the real cursor behaved. The plots of results and the feedback phase are essential for adjusting the relationship-control process to sense and control the right relationship. In every workable feedforward system, someone must watch the effects on the controlled variable and adjust the parameters of feedforward until the compensation cancels the effect of the disturbance. That process of adjustment is a feedback process.

Feedforward is just a bad way to speak about higher levels of control.

Part I: Reference level of controlled variable

The final step does not bring in any new demonstrations. It points out that in all the experiments done so far, the participant was accepting a certain condition of the display as being the "right" one—cursor between target marks, one figure matching another, a tone at a given pitch, a number at a given value. In every case, the participant could have chosen a different condition as the "right" one.

By returning to previous experiments and choosing different conditions as targets, the participant can experience the unique ability of an organism to manipulate the state of the perceived environment. Once control becomes skillful, it is easy to counteract reasonable disturbances. Then it is just as easy to pick new target levels for the controlled variable, and achieve them.

In fact achieving them is so easy that we tend to ignore the control process itself and simply "will" that the variable change to a new state. During compensatory tracking, after you're bored with keeping the cursor centered, you can "move it to a new position" just by wanting it to be somewhere else. Then you can hold it in the new position, continuing to counteract the disturbance without having to pay much attention to it.

The target state toward which your actions continually urge the controlled variable is called the "reference level" of the controlled variable. Most of what we call behavior is really defined in terms of reference levels for controlled variables. When you claim that you are steering a car in a straight line, you're not describing what the car is actually doing; you're describing what you want it to do—you're describing a reference level of straightness for the path. The actual path, as any unkind passenger could point out, is only approximately straight, especially when there are bumps and tilts and gusty winds and passing trucks. When you say "I'm eating lunch" you may actually be standing in line at a cafeteria, not yet having eaten a bite, or you may be drinking coffee with an empty plate in front of you. No matter what you say you're doing, at the instant you say it you're probably actually doing something else.

What you mean is that you have set a reference level for what your actions are to accomplish, and you're in the middle of a control process that keeps the perceived state of affairs as close as possible to the reference state. You say "I'm keeping the cursor between the target marks" when it's obviously well above or below them. But "cursor between target marks" is the perception you want to experience. It's the reference level that explains your actions—that, together with the disturbances that make it necessary to act.

By choosing different reference levels for the controlled variables in earlier steps, you can see the effects on the plots. You will see how action continues to oppose disturbance, but also maintain an offset that is needed to keep the controlled variable at the new reference level. And you may well begin to see that this ability to vary reference levels needs explanation if we are to understand how behavior works.

Part 2 (DEMO2: Modeling control) develops that explanation in the form of a working control system model.