Perceptual Control Theory: several demonstrations

Chapter 6 from *People as Living Things* by Philip J. Runkel

Do it yourself

I have put before you a lot of pages of theory and argument. It is time to give you relief from stretching your imagination and let you stretch something with your hands. I will describe a few games you can play with a friend. I urge you to do them. The games will give you some experience with (a) consciously observing yourself controlling, and (b) observing another person controlling. You will get an understanding of the basic principles that words alone cannot convey. My description here follows very closely Powers's Chapter 5 in his 1998 book, even to using many of his sentences (for which thanks).

THE RUBBER BANDS

Get two rubber bands just alike, three or four inches long. Knot them as shown in Figure 6–1 by passing one through the other and pulling them tight. You will also want a table where you can sit across from your friend or side by side. And you will need a mark on the table between the two of you. You could put a mark on a piece of paper and lay the paper between you. Or use a dent or mark already on the table. (You can do this exercise without a table, but a table is comfortable.) Each person now hooks a finger through an end of the rubber bands, stretching them horizontally an inch or so above the paper. If you sit side by side, use your outside hands to avoid bumping into each other.

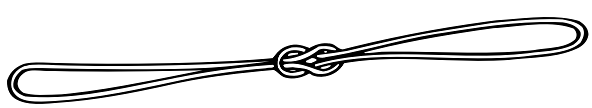

Figure 6–1. The Rubber Bands

Designate one person as **E**xperimenter and the other as **C**ontroller. (Change roles from time to time so that both people can see what's going on from both viewpoints.)

The task of the **C**ontroller (**C**) is simply to keep the knot that joins the rubber bands exactly over the mark. The *internal standard* that **C** must adopt to perform this task is the relation between the knot and the dot—namely, the knot holding directly over the dot.

The **E**xperimenter (**E**) uses **E**'s end of the rubber bands to disturb the position of the knot. **E** can do that by moving the finger forward or back, left or right—in any horizontal direction (not up toward the sky or down toward the earth). **E** should understand that the object of this experiment is *not* to prevent **C** from controlling the position of the knot. You cannot keep the knot stationary (exercise control) if the other player moves faster than your natural reaction time can compensate. Move smoothly, not too fast. The lessons to be learned will be much more obvious to both of you if **C** is able to keep the knot always close to the mark. Of course, after the basic observations are made, **E** can try all sorts of things to see what control looks like under difficult conditions. But especially at first, we want to keep the conditions easy by letting **C** learn to get good control of the knot. **E** moves the disturbing end of the rubber bands around in any kind of slow pattern, while **C** concentrates on keeping the knot accurately over the dot. A few minutes' practice should be enough.

You will notice very soon that every motion of **E**'s finger is reflected exactly by a motion of **C**'s finger. When **E** pulls back, **C** pulls back. When **E** moves inward, **C** moves inward. When **E** circles left, **C** circles left. **C** must do that, of course, to keep the knot stationary. Discounting small control errors, at every moment **C**'s hand is exactly as far from the dot as **E**'s hand (if the rubber bands are identical). The action illustrates very plainly the phenomenon of control—that we act in opposition to a disturbance.

If a third observer happened on this scene, what would the first impression of these actions be? It would be that **C** is mirroring the movements of **E** symmetrically around the dot. It would not be obvious which person is putting in the disturbances and which one is counteracting them. Even if **E** confessed to being the disturber, it would still not be obvious that control is happening. Much more likely, the third observer would see **E** doing things and **C** reacting to them: stimulus and response. The third observer would say that what **E** does causes the acts of **C**. The third observer might not notice that the knot stays over the dot.

This interpretation, based on a quick judgment, would be reasonable. The third observer might well lose interest at this point, and leave with the impression that control theory is just the same old stimulus-and-response idea that's been around since great-grandfather's day. But a quick glance is not enough to grasp that control is going on.

Remember the basic organization proposed by PCT: perception, comparison of the perception with an internal standard, detection of error, and conversion of error into an action that affects the perception. **C** is perceiving the present position of the knot relative to the dot. The perceived relationship is compared with an internal standard—knot over dot. The difference (the perceived horizontal distance of the knot from the dot) is converted into an action (a motion of **C**'s end of the rubber bands) that will bring the perception of the knot-to-dot distance to the distance required by the internal standard—zero.

How could we test whether the PCT model is right, or whether the stimulus-response interpretation is just as good? According to PCT, what is being controlled is a perception of the knot and dot. The stimulus-response interpretation (in one form) says that **C** is responding to movements of **E**'s hand. So the two theories are actually claiming that **C** is responding to different perceptions of the situation, and we ought to be able to decide which claim is right.

An easy test would be to get a piece of cardboard and use it to keep **C** from seeing first **E**'s hand, and then the position of the knot. If **C** has been responding to movements of **E**'s hand, then blocking the view of **E**'s hand while still allowing the knot to be seen should greatly modify **C**'s behavior. On the other hand, if **C** is perceiving the relationship of knot to dot, blocking the view of **E**'s hand should have no effect on **C**'s actions, while blocking the view of the knot and dot should make control much worse, if not destroy it. If you want to be sure what would happen, you can get a piece of cardboard and actually do those two things, though it would be easier simply to ask **C**, "Are you watching **E**'s hand or the knot?" **C** will deny paying attention to **E**'s hand.

Doing this test more formally, using instrumentation and computers, shows that control of the knot-to-dot distance depends critically on the controller's being able to see the knot and the dot, and not at all on the ability to see the cause of disturbances of the knot. I will show several examples of this fact, demonstrated by the use of computers, in the next chapter. Recognition of this fact is one of the crucial differences between PCT and other psychological theories. Other theories try to explain how it comes about that people perform particular acts—such as moving the end of a rubber band in a particular direction. PCT tries to explain how it can come about that people maintain a particular *perception*—such as the relation between a knot and a dot. Recognizing the fact makes a huge difference in the success of the explanation.

As well as using a piece of cardboard to hide the knot, there is another way to test for control. The idea here is simply to find out whether the knot is doing what it would be doing under solely physical effects. Let **C**, for a moment, hold **C**'s end of the rubber bands stationary. Let **E** start with the rubber bands almost slack, and then pull directly away from the dot by about six inches. Watch the knot. The knot will move half as far as **E**'s end of the rubber bands moves. This shows us the effect on the knot that **E**'s disturbance has when **C** does nothing. **E** could figure this out without any help from **C** at all. **E** wouldn't need **C**'s finger to hold one end of the rubber bands in place. **C** could go to lunch, and **E** could use a dowel in the table to hold **C**'s end in one position, and **E** could watch the knot move half as far as **E**'s finger moved.

But with **C**'s finger hooked into a rubber band and with **C** acting to control the position of the knot, **E** can now apply exactly the same disturbance as before and observe what the knot does. Now, of course, pulling back by a calibrated amount will have essentially no effect on the position of the knot. The knot will move only a tiny fraction of the amount that it moved when there was no control system attached to the other end. This failure of the disturbance to have the physically predicted effect is a strong clue that there is a control system acting. It is not infallible as a proof that control exists, because you still have to rule out simpler explanations for the lack of effect, but it is infallible in the other direction. If the amount of movement of the knot is exactly what you would predict under the assumption that there is no control system, then you have ruled out the existence of a control system. This test can eliminate wrong guesses very quickly, which is almost as helpful as being told what the right guess would be. Indeed, these two tests—cutting off **C**'s sight of the knot and cutting off **C**'s control of the knot—are essential parts of the procedure known in PCT lore as The Test for the Controlled Quantity, which is the core of experimental method in PCT. You can see that this method is eminently suitable to examining control on the part of an individual. I will say more about The Test in Chapter 7.

I have mentioned in earlier chapters that PCT includes multiple levels of feedback loops, though I have not yet explained much about that. We can, however, illustrate two levels of control with the rubber-band game. To do so, let **C** make the knot move very slowly and uniformly around the dot in a circle, with a radius of about one inch. The knot should take at least ten seconds to go once around the circle. **E**, of course, continues to move the other end of the rubber bands in big, smooth, slow, random patterns. If **E** sees that **C** is having trouble, **E** should slow down the disturbances. We want to see the controller succeeding, not failing.

Obviously, the internal standard is no longer "knot on dot." Perhaps, as many theoreticians in this field have done, you unconsciously assumed that the dot was specifying the internal standard—that the knot was the controlled perception, and it was brought to the standard set by the dot. Now, however, we can see that the controlled variable was really the *relationship* between the knot and the dot. Now the knot is being maintained in an ever-changing relationship to the dot. And if you still think the dot is not simply part of the controlled perception, we can let **E** choose to move the piece of paper as well as the rubber band—the two simultaneously. **C** is controlling a relationship between two perceptions, one of the dot and the other of the knot, and keeping this relationship in a match with an internal standard that now involves continuous motion.

If you are only reading this description, this won't be obvious, but if you are actually doing the experiment, you will realize that the experimenter, all this time, has been moving the disturbing end of the rubber bands around in big continuous patterns. You may have been thinking that to make the knot move in a circle, **C** has to make the hand holding the rubber band move around in a circle—bigger than the knot's circle, but a circle. Actually, if **C** were to hold a marking pen through the loop in the rubber band so as to leave a record of hand movements on the paper (this is worth trying), the trace would show not circular movements but a random mess.

In the movements of the knot relative to the dot, we are seeing the internal standard that **C** has chosen. The internal standard determines what the controlled perception will do. But in the movements of **C**'s hand, we see a composite of the effect of the internal standard and the even larger effect of the disturbances. The hand movements correspond neither to the internal standard nor to the disturbance; they represent what has to be done to maintain control as the disturbance changes.

Let **C** now stop the motion of the knot at a point one inch to the left of the dot while **E** continues to apply disturbances. Now we are back to the original case where **C**'s hand movements are symmetrical with those of **E**—but **C** is now maintaining the knot in a different and now stationary relationship to the dot. The control process is just like the first one, but with a different internal standard. We can call this one level of control.

The second level of control is the one that perceives continuous change. When the internal standard for this kind of change is the perceptual equivalent of "one revolution every 10 seconds," the knot moves in a circle because the *internal standard for knot position is being changed* so as to maintain that perceived circular movement. The first level of control, which is concerned with maintaining a particular, relative position of the knot and dot, is being used as the output of the second level of control, which is being used to maintain a perception of circular movement. The position control system is being used as part of a motion or trajectory control system. **C** could use a different trajectory control system, and make the knot write **C**'s name. Many different higher-level control processes could be carried out using this same position-control system (although not at the same time).

Many more variations are possible, involving various internal standards, simultaneous control of more than one perception, more than two people, and multiple rubber bands. They are fun to explore. You can also do these experiments against paper on an easel, so that an audience can watch. Here, I will describe two further uses of the rubber bands.

Two Controllers

This is a demonstration of conflict. On the piece of paper, add a second dot about ¼ inch to one side of the dot that's already there. Now **E** disappears and becomes another controller, **C2**; we have **C1** and **C2** controlling the same knot. The experiment is simple. **C1** aims to hold the knot exactly over the old dot, and **C2** aims to hold it exactly over the new dot. Their internal standards differ by 1/4 inch. If both controllers insist on keeping the knot over the "right" dot, there's only one possible outcome. A rubber band will break.

This seemingly innocent situation exemplifies the most serious problem that can arise between control systems, whether they are in different people or inside one person—conflict. PCT explains how conflict works and how it can cause immense difficulties; I will return to this topic in later chapters, especially 9, 23, 28, 29 and 33.

Four Controllers

This game can demonstrate cooperation, too. Neither paper nor pencil is needed. It is convenient to do it standing. Get eight rubber bands. Connect four of them in a circle, and attach the other four to the four knots. Find four obliging people. Ask each to take hold of one of the four rubber bands attached at the knots. Tell them, "Make a square" (of the first four bands). They will quickly do so, without needing to talk about it. Think for a moment about all the ways that the other three people can disturb the corner that is held by any one person. Despite the fact that any motion by one person will to some extent disturb the positions of all the other corners, the four people, without consultation, will somehow move into positions that result in a reasonably accurate square!

You can try various experiments with this layout. You can have someone give instructions about how to go about making the square. Will that square be made faster or better? You might have one group of four do it as described in the previous paragraph and another group (who have not watched the first group) do it after discussing the task and agreeing on how to do it. How would the performances differ? You might hook two rubber bands at each of the joints and use eight people. What would the additional people do? You can think of more variations.

Try to imagine for a moment all the sorts of little motions the four people around the rubber-band circle might make while bringing the bands into a square. Many people who design artificial intelligence for robots believe that a robot (or a person) cannot act without having inside itself, before it acts, a detailed "map" of the environment in which it is going to act. Can you imagine each of the people with the rubber bands trying to anticipate what each of the other three might do next? Each small motion by any person changes the environment for the other three, and all do that simultaneously and continually. Any map would become out of date the moment that anyone made any motion whatever. Yet people do this task of squaring the rubber bands with very little difficulty. I will say more about the idea of making an internal map in Chapter 24 under the heading "Model-Based Control." And I will say more about trying to anticipate specific future acts in Chapter 36 on planning.

You can find other descriptions of the game in Powers (1973, pp. 235–236 and 241–244), in Robertson and Powers (1990, Chapter 4), in Runkel (1990, pp. 105–108), and in Cziko (2000, pp. 87–89).

THE COIN GAME

Get four coins, a flat surface (a table-top or a patch of sand at the beach), and a friend. The four playing-pieces need not, actually, be coins. They could be checkers, or chess pieces, or little shells. They can be alike or different, as you choose. Charles Tucker, who teaches PCT at the University of South Carolina, prefers to use paper disks or poker chips, all alike. But here I'll suppose you will be using coins. As before, let one person be the **E**xperimenter and the other the **C**ontroller. Let **C** arrange the coins on the table in any pattern **C** chooses. **C** might choose to have all the coins in a straight line. (That would be pretty easy for **E** to discern.) **C** might choose to have three of the coins in a cluster while leaving one isolated. Or **C** might choose to have the imaginary line joining one pair of coins always crossing through the imaginary line joining the other pair. You can think of a dozen other patterns, some obvious, some subtle.

The task of **E** is to discover, *without any discussion about it*, the pattern (internal standard) that **C** is exemplifying in the way **C** has laid out the coins. **C** should write down a description or definition of the pattern the coins are exemplifying. Now **E** can begin probing to discover **C**'s pattern. **E** pushes a coin (or more than one) to a new position. If the result changes the pattern away from **C**'s internal standard for the pattern, **C** must *correct the error*—that is, push the coin back to its previous position or to some position that corrects the error. If **E**'s push of the coin does not take the pattern away from **C**'s internal standard, **C** can merely wait or can say, "No error." ("No error" means "You have not caused me to feel that the pattern is now in error.") This process continues until **E** becomes certain of being able to make three moves that will bring corrective moves from **C** and three moves that will bring only a "no error" response. If **C** corroborates **E**'s certainty, **E** and **C** compare their definitions.

Typically, **E** will begin the game feeling reasonably confident of eliciting a correction from **C**, and will be surprised when **C** says, "No error." Playing this game, it becomes very obvious how easy it is to think up explanations of "what **C** is doing" and how easy it is to be wrong about it. The game demonstrates, too, the relation between doing and talking. The three correction-eliciting moves and the three "no error" moves demonstrate that **E** can now *do* what **C** was doing, but **E**'s oral description of the pattern may not sound very much like what **C** wrote down at the beginning of the game. **C** might have written down "Large to small," and **E** might have called it "a string of drops of water." An observer might say, "You are doing the same thing; I don't care what you call it." Or, after the three correcting moves and the three "no error" moves, **E** might say, "You were making a Z." And **C** might say, "No, it was an N." And an observer might say, "I thought it was a zig-zag."

Playing this game as **E**, you come not rarely to the point where you are sure of the pattern the other person has been controlling only to discover that the pattern was something very different. You might have settled on a geometric pattern when **C** was actually keeping the coins in order by date, or by size, or alphabetically by name: dime, nickel, penny, quarter.

This procedure, which is a variant of The Test for the Controlled Quantity, can make it easy for you to understand what it means to say, "You cannot tell what people are doing just by watching what they are doing." But I will phrase that more transparently: You cannot guess very accurately what people's purposes are just by watching their actions. That sounds reasonable, but most of us most of the time, I think, are too ready to believe we can descry the purposes of others. The coin game will help you to look at your own belief.

In psychological experimenting, as in other domains of social life, the pitfalls of language leave us very uncertain whether we have arrived at the condition we sought or have gathered the facts we envisioned. I devoted Chapter 6 of my 1990 book to the weakness of language, and in Chapter 11 there I told about some researches that were carried through with a minimum of language. The Test for the Controlled Quantity can often be carried out with no talking (or writing) at all; the coin game, after you have agreed with the other person on the procedure, can be played that way—silently. Saying "no error" speeds the game, but it is not necessary; you can just say nothing and let **E** conclude that your perception of the pattern is not disturbed, because you have not pushed a coin. You can see that The Test is not limited to humans; it can be used with any sort of creature.

When you play the coin game, remember that you are using it to see how control on the part of another person can be discerned. If you are playing the part of **C**, you want to see how **E** can discover the pattern you have in mind. Sometimes, maybe out of habit with games, a player seems to want to "win" the game by choosing a pattern that will be impossible for **E** to guess. If you do that, you will lose your chance to learn about control.

Do actually try these games. They yield insights you will never get by trying to imagine what the words here mean. The games will help you to discern control and non-control in everyday life. It is fun, too, to make up your own variations of these games. If rubber bands or coins seem beneath your dignity, remember that Galileo Galilei (1564–1642) discovered the shape of gravity by rolling little balls down a slanted piece of wood.

Still another thing you can do without having a laboratory or a budget is to run tutorials, demonstrations and simulations on your computer.

The publisher's website

http://www.livingcontrolsystems.com,.

features DOS and Windows programs, introductory explanations and articles as well as links to other resources that illustrate PCT in various ways. Be sure to check the section on Perceptual Control Theory (PCT).

Richard Marken's demonstrations of several features of PCT at http://www.mindreadings.com/demos.htm are programmed in Java and can be run using a browser on any kind of computer.

And you can get a DVD video and script entitled "Rubber-Band Demonstration" by Dag Forssell (1993), based on an outline by William T. Powers.