Action always equal and opposite disturbance

Unedited posts from archives of CSG-L (see INTROCSG.NET):

"Coach" - Equal and opposite!

Date: Fri Apr 17, 1992 2:05 pm PST

[From Bill Powers (920417.1100)]

On testing models:

Part 1.

Some time ago I remarked that the most common model in psychology is a cause-effect model in the form of a regression equation. The hypothesis is that the effect depends on the cause linearly, as in y = ax + b. To test this model, you'd take the values of a and b determined from a formal study, and try to predict new values of y from new observations of values of x.

David Goldstein commented that this concept of using a model for predictions is not the way such findings are used in psychology. Once the regression line is drawn through the data points, that's the end of it. The model equation describes the data, but isn't then used for predictions.

On thinking this over, I agree that no formal use is generally made of the regression equation, but the findings are certainly used to predict individual behavior. Suppose the dependent variable y is a clinical measure of depression, and the independent variable x is a depression- factor score on a personality test. In computing the correlation between the test score and the clinical measure (in a study of many people), a regression equation of the form y = ax + b is the basic premise behind the correlation calculation. If the correlation is positive and statistically significant, the conclusion drawn is that depression is predicted by the test score. Then the test is administered to a new individual (presumably from the same population), and if the depression- factor score is high, the person is diagnosed as depressed.

This isn't a formal application of the regression equation: you don't say that a test score of exactly 7 predicts a depression of exactly 25 units on the clinical scale, even if that's what the regression equation says. But a person who measures 15 on the test score would be judged as more depressed than a person who measures only 3. So while the slope and intercept coefficients aren't explicitly used, the general trend is implicitly used, and there are semi-quantitative judgements made.

The scatter in data of this kind is so great, of course, that literal application of the regression equation would be silly. The prediction for any individual when the correlation is as low as 0.8 would be seriously wrong most of the time, often even getting the sign of the relationship wrong for one person. The only correct way to make a prediction would be to begin with another equally large sample of the population and do the whole study again. You would predict that the same regression coefficients would be found.

But there is an urge to predict for individuals, and the form of the urge follows the regression line: a higher clinical score ought to predict a more severe depression. While it is folly to give in to this urge when the data are so bad, the motive behind doing so is consistent with the principle of modeling.

If the principle of modeling were followed through formally, the regression line would indeed be used to predict behavior. If the line has the equation y = 3x + 5, and the depression-factor test score for a new individual is 4, the model predicts that a clinical evaluation of depression will come up with 17 on the clinical scale for that person. To follow the test through, one would then submit the person to the same clinical evaluation as used in setting up the model, and see what number actually results.

Suppose the actual depression measure is 12 on the clinical scale. This is a deviation of -5 units from the value of 17 predicted from the test score, for an error of -29 percent. Is that good, or is that bad? The answer depends on how important it is to get the evaluation exactly right.

Of course in this case we know the clinical measure of depression, and if we believe it we can just ignore the test score and the prediction. But what if we want to make the diagnosis on the basis of the test score alone? Now the generally expected error for an individual prediction becomes relevant. If you're going to prescribe electroshock therapy that will most likely severely disturb the person's life for many years, maybe even permanently, you might decide that a 29 percent error is too large to allow. Perhaps even an error of 5 percent would be too large if the person is a borderline case. On the other hand, if you're going to prescribe a tranquilizer that won't do any permanent harm even if the person isn't really depressed, then perhaps you can allow errors as large as 29 percent.

I've gone through this to illustrate that prediction errors can't be judged as good or bad without taking the context into account. But what if the context is that of testing a general model of behavior? Now the actions taken as the result of a diagnosis are no longer in the picture. All we want to know is which theory is better. Now the errors of prediction under different models are judged not against practical standards, but against each other. The smaller the expected error, the better.

I have also tried to show that even in standard approaches, the method of modeling is there just beneath the surface. It's probably not mentioned much because the predictions made from literal application of the model -- the regression equation -- are so poor. But the model is there. It's that model that we have to compare against the control-theory model, and the way we do the comparison is through making quantitative predictions using the actual form of the model.

Let's look at the rubber-band experiment. Suppose we just measure the position of the experimenter's end of the rubber bands and of the subject's end, designating the positions as e and s. Let's confine the experiment to a line, so we consider only one dimension. The zero point on the line can be chosen arbitrarily, with all measurements made relative to that zero.

If we now measure the positions e and s over a long series of movements by the experimenter, we will obtain a data set consisting of pairs of values of e and s. We can do a correlation between e and s. From the normal calculations, we can derive a regression line.

The regression line will have the form s = ae + b. The position of the subject's end will depend on the position of the experimenter's end. If the rubber-bands are identical, the coefficient a will be very close to -1. Half of the intercept b will correspond to a position on the line. That position will be the average position of the ends of the rubber bands: with a = -1, we will have (s + e) = b, or (s+e)/2 = b/2.

In fact, half of the intercept b will turn out to be a position nearly underneath the knot where the rubber bands are connected. The knot, as it will turn out, remains very nearly at the position b/2 all during the experiment.

There's a moral to this story, but it's not quite obvious yet. The first part of it is that when you do an SR experiment in the usual way, to get a regression coefficient, you can SOMETIMES translate it directly into a control-system experiment. If you find that the intercept b corresponds to something in the experimental situation that's remaining nearly constant at that value, you've found a controlled variable -- actually, by finding its reference level first.

The second part of the story concerns the accuracy of the prediction. The SR prediction will be accurate only if the two rubber bands have identical characteristics, or strictly proportional characteristics. If their characteristics are different, the correlation coefficient you derive from the data corrected for the different rubber-band properties will be very much higher than the one derived from the model s = ae + b, which assumes identical rubber bands.

Part 2.

In testing the control-system model, the basic procedure is to assume that all behavior without exception is control behavior, predict behavior on that basis, compare the prediction with the appropriate data, and let the match or mismatch decide the issue. You can never prove that a particular control-system model is the only correct one, but you can show that it is incorrect.

Considering the low correlations that are found in S-R experiments, it might seem hopeless to substitute a PCT model for the linear regression model. When the data are that noisy, how can any clear decision be made? This objection, however, assumes that the SR experiment has correctly represented the data. While we can't prove that ALL SR experiments could be translated into relatively noise-free PCT experiments, there are excellent reasons to think that this can be done in a significant number of instances, maybe even most instances. To do this, however, can require some changes in viewpoint that may be hard to achieve.

An SR "fact" is expressed as an effect of a cause. Doing something to a person results in that person's doing something else. If the relationship expressed in this "fact" isn't clearcut and quantitative, then the control theorist has to start asking questions about the data.

The basic question is, what is it that was affected by the "stimulus" that was also affected by the "response?" If you utter encouraging words to someone, and that someone then shows added efforts to achieve something, you have an SR relationship. Now you have to try to guess: what did the encouraging words affect that was affected EQUALLY AND OPPOSITELY by the increased efforts?

Equally AND OPPOSITELY? There's the rub. You would like to think that there is something you said that helped this person do better. But control theory says that if your words of encouragement had some regular effect on the person's behavior (apparently), that behavior was aimed at COUNTERACTING your influence. If this is true, then you don't have the control over the person's behavior that you thought you had, even for the good. You are seeing yourself as helping the other person to do better. The other person, however, is seeing the situation differently: you're disturbing something, and the other person is acting to cancel the effect of the disturbance.

This may not be true, but if you're going to test the PCT model honestly, you have to pretend it's true and try to make sense of it. You can't test a model if you don't follow its logic faithfully and literally as far as you can. You can't look ahead and think "If PCT is right, then I haven't been helping people the way I thought I was -- so PCT must be wrong." You have to be prepared to change your ideas about anything at all. Otherwise your reasoning is just a sham.

Let me give you a real example from my high-school days. We had a coach, named Coach, who was tremendously popular, a great guy. We all loved him and wanted his approval above anything else. Coach would say "You can do better than that, I know it -- just give it one more try and you'll make it." And by golly, we'd give it one more try and we'd make it, sometimes.

Now it would seem that his encouragement and belief in us caused us to try a little harder than we thought we could, so we achieved something we couldn't do before (sometimes). I suppose that Coach looked at it that way, as any reasonable person would. But I can tell you that from inside at least one person (and at the time I guessed this was true of a lot of the others), it wasn't all that nice.

The basic problem was that Coach went around all the time saying to people, "What you're doing isn't good enough to please me." That's what "You can do better" says. I was already doing better than I thought I could, in number of pushups, speed of climbing a rope, time in the 40-yard dash, or whatever. And

I was damned tired and hurting, and not necessarily interested in doing any better. I liked physics a lot better than physical education. But here's Coach telling me that he doesn't like what I'm doing. That mattered to me. So I got myself together and made it REALLY hurt, and I felt great -- because now Coach wasn't displeased with me. Not because I'd achieved something I wanted, but because I'd done something to counteract his disapproval.

From Coach's point of view, he had helped me put out that extra bit of effort to surpass my previous achievements. No doubt if I had continued to go along with this, worked out, built up a lot of strength, learned the football playbook by heart, and all of that, satisfying the coach more and more all the time, I might have achieved even more. I might have been a college football star; I might even have become a professional football player and ended up as a coach myself, by now. I might be bold, aggressive, commanding, and rich. But I certainly wouldn't be writing this. I also wouldn't be the Bill Powers you know.

What actually happened was that many of us simply gave up on pleasing Coach because we didn't buy the goal. It wasn't pleasant to do that -- to decide we were trying as hard as we cared to try toward that particular end, and that we would simply endure the disapproval. We still loved Coach, and we tried to fend off his disapproval by seeming to try harder. But the price was too high to really do it. When Coach was called into the Navy and left in 1944, there was a huge tearful farewell ceremony for him, and I'm sure that amid the sorrowful participants there were many hearts filled with relief.

To apply the PCT model, this is the sort of thing you have to think about. It's especially difficult when the hoped-for effect on a person is beneficial. There's an almost-inescapable tendency to suppose that what you think of as beneficial is also considered beneficial by the other person; that what you consider harmful is also thought harmful by the other. Coach would have been completely baffled by the present discussion. He would have said "Well, you did try harder, didn't you? And you did do something you thought you couldn't do, didn't you? What's so bad about that?"

The SR viewpoint encourages this sort of naive projection of one's own goals onto the behavior of others. I shouldn't even call it the "SR" viewpoint. It's really this viewpoint, adopted innocently by well-meaning people who have never heard of stimuli and responses, that led naturally into SR theory.

To test the PCT model in real life, you have to be prepared to follow its logic all the way. Forget about whether the "response" is good or bad. The question is how to find the controlled variable, the thing that is disturbed by what is done to the person, and is protected against more disturbance by the action that the person takes. If you find such a controlled variable, you will understand that person far better than you did before. If you want to help that person, you might even find out what he or she really wants and figure out ways that person could get there.

It's possible that you won't find any such controlled variable in a given circumstance. But if you don't look for one, you will certainly not find one even if it's there staring you in the face.

The basic message here is that to test PCT, you have to make predictions from it and from nothing else. You have to follow out the logic even when it seems to say things you don't believe. Then you have to look carefully to see whether, in fact, the prediction holds true. This requires being consciously open-minded and willing to take a chance. You simply have to trust that if the theory does predict correctly, you'll be better off knowing what it predicts than not knowing, letting the chips fall where they may.

Best to all Bill P.

Date: Sat Apr 18, 1992 2:34 pm PST
Subject: Coach; conflict

[From Bill Powers (920418.1500)]

Martin Taylor (920418.1340) --

You've sort of taken off at right angles to the line of thought I was developing. The "Coach" example was meant to illustrate how an apparent SR relationship (encouragement --> doing better) can lead to quite a different interpretation when explored from the viewpoint of control theory. I wasn't trying to generalize from the particular way I and probably others dealt with Coach's urging us to overachieve. With another person or in another circumstance, a similar encouraging remark leading to improved performance could work in a different way. But it will never be a cause-effect way. My point was that to test control theory you have to think of possibilities other than the surface appearances.

Since I'm into high school stories, I remember another instance with a mathematics teacher. I didn't much like or dislike this teacher -- he knew his stuff but wasn't strong on making things clear. The class was doing an exercise, each person trying to prove a trigonometric identity. I was stuck -- something was wrong and I didn't know if I was even getting close. The teacher was going around the room seeing how everyone was doing. When he got to me, he said "That's fine, you're almost there."

This told me that I hadn't made any mistakes so far and was headed in the right direction. So I stopped worrying and went ahead and finished the proof, my first one. That felt nice. The 60th proof didn't feel so nice.

Apparent SR relationship: he said what he said, I then went ahead to reach the goal. Cause and effect? No. Information. I wanted to know if I'd made some stupid mistake, and he told me (in effect) that I hadn't. With that information, I could stop looking for a mistake and devote my efforts to something more productive. I didn't finish because I liked the teacher or in order to please him. I finished because I wanted to be able to prove the identity. His remark wasn't a disturbance of something I was trying to control; it provided a missing perception so I could get unstuck from looking for a nonexistent error.

My Coach example was one in which the apparent stimulus actually did disturb something I was controlling for, and my response opposed the effect of the disturbance. The result was to put a very different light on what seemed like a simple S->R chain. That's all I was trying to show -- not that there's something inherently bad about encouragement or that being as pushy as Coach was necessarily leads to resentment and bad feelings. In fact I never resented Coach; not many did. He was a nice guy. I just resisted him. I regretted not wanting to live up to his expectations, but not enough to change my mind.

Re: your comments on conflict.

Conflict doesn't "lead to" anything in particular. What it leads to depends on how you resolve it, or fail to resolve it. Most conflicts are unimportant; we just shrug and turn to something else, or go into a little fit of reorganizing and think of a different way out. This happens all the time; we have natural machinery for resolving inner conflicts and it usually works very well.

The degrees-of-freedom problem doesn't normally cause conflict because we've learned to use only those control systems that are compatible when working at the same time. The balancing of reference signals contributed by many higher-level systems isn't a conflict unless one of the higher systems is unable to keep its own error reasonably small because of the interference of other systems at the same level. The usual case is that all active higher-level systems keep their errors small despite the fact that no one lower-order system's reference signal is the exclusive property of one higher order system. The systems just find the analog solution of the simultaneous equations and they all are successful.

When opposing muscles are used to control limb position, there's no conflict. In fact there are two controlled variables that are independently adjustable: for the tendon reflex, one is the difference between the tensions in the two muscles, the other is the sum. The sum-of-tensions signal is controlled to produce a specific muscle tone. The difference signal controls the net applied force. Because the muscle is highly nonlinear, the sum (muscle tone) signal effectively alters the spring constant of the combined muscles near the zero-error condition, thus adjusting the static loop gain of the tension control system (and also the stretch control system).

Conflict is a problem only when it concerns some variable important to the organism, is severe, and goes unresolved for a long time. That's what brings the clients to the therapist or counsellor. Serious conflict destroys control or reduces its effective range to the point where it's not sufficient for the purposes normally served by the control systems.

A control system that keeps its error very small isn't likely to be "placid and content." It's able to keep the error small because it has a very high loop gain. This means that even the smallest disturbance will evoke an opposing effort, and that opposition will keep the controlled variable nailed to its reference condition. When you're driving a car along a mountain road with a washout on the cliff side, you tighten up that control system so the car stays precisely on the path you've picked to squeeze past the danger point. I don't think that "placid and content" describes that control system. But it's not in conflict, either: if it is, you have a problem because you won't be able to move the wheel as much as if there weren't any conflict.

There's a problem with your suggestion that "a system with tension and conflict will be more robust than one that is placidly content." The problem is that reorganization will start because of the chronic conflict. As a result, precise control will become impossible: the parameters of the control systems are going to be changing at random. What you get is a jittery and unpredictable control system that could literally do ANYTHING without warning.

Just because of neural response curves, I can believe that some slight amount of tension would help with rapidity of response to disturbances, because near zero signal the slopes of the functions will be very low and the loop gain will be low. But this is relevant only when the control point is set to zero and there are no disturbances. Most reference signals specify values of perceptual signals that are far from zero -- somewhere in the normal range between zero and maximum. And there's normally some amount of disturbance to raise the error signals above zero, if only gravity. In those cases, there's no advantage to conflict because conflict won't raise the sensitivity or speed of the system and will only reduce its range of control. I think that the best state to be in for possible action is one of alertness and calm. You should feel just a little zingy, but you certainly shouldn't be in white-knuckle conflict with yourself. You want everything working in the same direction.

So I guess I agree with your concluding remark: tension, conflict, and uncorrectable disturbance are good, but not in excess. I would figure something like 5 percent of the range of control. The rest of your reserve you would want to save for affecting the environment.

Uncalled-for remarks on social conflict.

In the background I suspect is an idea that competition is good for us (if not in your mind, then in others). Up to a point, while it's fun, I agree. We like to set problems for ourselves and solve them, and get better at solving them. But competition as a way of life doesn't work that way, except for a few winners. A social system based on serious competition is just a step from violence (in the US, a very short step). The losers vastly outnumber the winners: we end up with a society of losers, winners being an anomaly. In situations where the terms of the game determine that only a few can win, chronic losers can get very nasty; in fact, they tend to abandon whatever social principles there might be that make civilization better than life in the jungle. I don't think that the price is right. Competition -- interpersonal conflict -- is the lowest level of social intelligence. I don't like to admit that even a little conflict can be a good thing, because we've accepted a HUGE amount of conflict as good and natural for far too long. It's time to get smarter.

Best Bill P.