Explanations, theories, models: harping on a theme

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

Date: Sun May 08, 1994 8:46 pm PST
Subject: Harping on a theme

[From Bill Powers (940508.1800 MST)]

On explanations, theories, and models: harping on a theme.

Theories in psychology and other similar pursuits are commonly expressed as a statement to the effect that y depends on x. The death rate from breast cancer depends on the degree of poverty, the point average of a graduate depends on the scores achieved on entrance examinations. The "theory" in each case is simply the assertion that the one measure should be found to depend on the other. To uphold the theory means to show that the proposed dependency does exist, with a certain level of confidence. The most common way of testing theories of this sort is to do an experiment in which the independent variable is varied, and its relation to the dependent variable is subjected to a statistical analysis.

When a statistical relationship between an independent and a dependent variable is measured, two numbers come out of the analysis. One is the correlation, or degree of relatedness. The other is the regression line, the average relationship expressed as average slope of output per unit input, or average amount of effect per unit cause. The correlation and associated measures of uncertainty are used to determine the degree of relatedness and test the theory. But, as David Goldstein once pointed out to me, the regression line is seldom used for anything.

In most statistical tests of theories, the only conclusion to come out of it is that y is or is not dependent on x. It is as though the number scale had been reduced to 1 (related) and 0 (not related). This turns what started out as a quantitative theory into a qualitative one. This also explains the either-or nature of theoretical statements that come out of statistical analysis: students who are taught by computer either do or do not remember the material better. People who take one aspirin per day either are or are not protected against heart attacks. There is no room for equivocation in statistical results: the quality of the results generally would not support any statements about the degree of an effect, other than the statement that it exists or does not exist.

This, furthermore, is why, as David G. said, the regression line is seldom used. When you have a correlation of 0.5 between the IV and the DV, the regression line might have a slope of 0.2. But you can't predict y from knowing that y = 0.2x (plus some constant, the intercept). The reason is that for any given value of x, a prediction of the value of y would be seriously wrong in most cases, either far too great or far too small (or even of opposite sign). The regression line only describes an equivalent average relationship. So the slope of the regression line fitted to the relationship of y to x is an unusable fact, in most circumstances.

A major consequence of the unusability of the regression line as a predictive model is that explanation and prediction become estranged from each other. You can explain male chauvinism in terms of an unresolved oedipal conflict, but given the fact that an unresolved oedipal conflict is independently known to exist, there is no way to predict whether an individual will become a chauvinist. A certain proportion of ten thousand people, perhaps; one person, no. If a person graduates near the top of a class, and it is found that this person did very well on the entrance examinations, you can say "There, you see? People who do better on entrance exams rank higher in their graduating classes." But given the score on an entrance examination, it would be a mistake to try to predict any particular individual's ultimate class rank.

There follows from the estrangement between explanation and prediction a custom of interpreting positive instances of a theoretical relationship as supporting the theory, but not the custom of interpreting negative instances as counting against the theory. In fact, neither positive nor negative instances are relevant to a statistically-justified theory; only the whole data set can support or deny the theory. But it seems to be irresistibly tempting to bolster statistical findings by pointing to individual examples that support it, and of course to avoid pointing to examples that run counter to the theory. So in defense of statistical theory, we have the anecdote which illustrates how the theory works when it seems to work, accompanied by a blank lack of interest in counterexamples.

The next stage in the estrangement is complete divorce of explanation from prediction. Once the custom of supporting a theory by referring to positive instances of it is well-established, the concept of predicting can quietly be dropped. Now what matters is whether the explanation seems plausible, appealing, beautiful, ingenious, or mathematically sound. All that matters now is whether the theory offers an attractive and compact description of relationships seen after the fact. The theoretical explanation takes on an independent life, the only criterion for acceptance being whether it is thought convincing and whether it is internally consistent. The thought of actually using a theory to predict an unknown y from its proposed dependence on a known x recedes into the background. The theory is no longer considered to be at risk; it is simply expected to work, with exceptions being blamed on unforseen circumstances.

Accompanying this divorce is a change in attitude toward the world. With factual testing of theories by requiring them to make predictions put aside, the concept of a physical world with properties of its own fades away. Now the mere occurrence of one event can lead to the occurrence of another event. It is not necessary to verify that there is an actual physical link between the events. If there is a mathematical or logical relationship between the events that can be stated, that in itself constitutes a sufficient explanation. If events prove to be related as the theory says, all is well. If they do not, one simply switches to a different explanation, a theory designed to handle the aberrant cases. Sometimes theory A works best, and sometimes theory B works best. You use whichever one works.

What I am describing here is a regression from the methods of science that have developed over the past three centuries to the state of thought that existed before, approximately, Galileo's time. The scientific revolution brought into being the concept of a nature that has properties independent of the observer. This viewpoint has some obvious problems, but along with it came the principle that pure thought could not establish the existence of a property of matter and energy. The only way to defend a proposition about nature was to cast a theory in a form from which a prediction could be made, and then do an experiment to test the prediction.

The theories thus became models: propositions about entities and their relationships to each other below the level of observation. These models had specific characteristics, because they were stated in quantitative form. Given those characteristics, and mathematical or logical reasoning, we could state what the model would _necessarily_ do when subjected to certain experimental manipulations. The behavior of the model, generated out of its own characteristics, could then be compared with observations. And it was possible for observations to contradict the predictions. That is what made all the difference.

It is also what kept sciences like psychology from joining the revolution. In order to develop testable models, one must have models that work well enough to generate predictions. Description and explanation alone are not enough. Plausibility and beauty are not enough. A preponderance of positive examples over counterexamples is not enough. And most specially, the regression line obtained from statistical analyses in which correlations of 0.25 or so are customary is not enough.

Explanation and prediction have to be put back together. Any explanation of a phenomenon should be followed immediately by a prediction: if the independent variable is manipulated in a certain way, the dependent variable should behave in a certain way. And if it does not, the prediction should be accepted as wrong, and the model behind it as wrong.

As we look toward developing PCT in the realms of higher levels of organization, we must renounce the custom of substituting a plausible explanation for a test of a prediction. It is not enough to analyze the data and show that a certain control organization would, if it existed, suffice to explain what we see. The essential next step is to use that explanation, the "regression line," to predict new outcomes under new conditions.

The future of PCT depends on doing this, and on the standards we apply while doing it. The only standard that will, in the long run, result in knowledge worth knowing is that our predictions should be accurate within the limits of measurement of the variables. If we can detect any error at all between our predictions and what actually happens, we should take the error as information for improving the model.

This may seem an impossible requirement: either we predict behavior with the accuracy of a physical experiment, or we remain mute and say nothing. But that is not the only choice.

The other alternative is to make more modest claims, commensurate with our ability to measure phenomena. If the variables we measure are simple enough, and our predictions are simple enough, we can achieve the degree of predictivity that physics demonstrates. Remember Dick Robertson and David Goldstein and their series of experiments for showing that self-concepts are controlled variables. At first they started with a complicated experiment using complicated methods of assessing self-concepts. The results they got were quite consistent with the results in conventional investigations of personality: miserable and unreproducible. After several tries, they decided to do a more modest experiment using simpler methods and measures: when a person makes a self-referential statement, contradict it. If the person opposes that contradiction, take that as an indication that a disturbance has been opposed. "I'm a very shy person," says the subject. The experimenter says "Not at all, I think you're very outgoing and bold." And the subject says "That's ridiculous, that's not at all how I am." Score one correct prediction.

By doing this, they got 25 confirmations of the control hypothesis, one equivocal response, and no counterexamples. Within the crude measure of resistance to disturbance employed, no error in the hypothesis was detectable. Of course it would be hard to get a conventional personality theorist to take such a simple hypothesis and such a simple prediction seriously. The result is too obvious. We learn nothing of deep serious import about human nature (or so it seems).

But look at what the personality theorists get by way of confirmation of more complex hypotheses: correlations ranging from 0.5 on down. The errors of prediction are not only easily detectable, they have a range that is many times the values of the variables being measured. Prediction is simply impossible.

Which is better, to make simple predictions that are true on every test, or to make complex predictions that are wrong almost every time they are made? We would _like_ to be able to make complex predictions that are always right, but we can't do that. We can't even come close. The sciences of human behavior are simply not advanced enough to hope to succeed with anything above the kindergarten level.

So we should go to kindergarten and learn how to make simple predictions that are always borne out by observation. There is no short-cut. If we want to make a science of behavior, there is no alternative. All apparent alternatives are illusions born of wishful thinking.

A science can be built only on predictions that are essentially always born out by observation. By learning to make simple predictions that are always right according to crude measures, we can learn to make more complex predictions that are always right according to more discriminative measures. This is how physics and chemistry began: not with lasers and DNA, but with pendulums timed against the human pulse, with judgements as to whether an object weighed more or less after burning. We still honor the discoverer and formalizer of the fact that the elongation of a spring is proportional to the applied force: Hooke's Law.

The gross mistake made by conventional behavioral scientists was to think that any piece of knowledge, however uncertain, is better than no knowledge at all. It is not; it is far worse than total ignorance. If we know we are ignorant, we will keep trying to better our theories and methods. But if we think that in discerning a faint trace of regularity we have reached the goal of understanding living systems, we will settle for less than a science, declare our methods to be good enough, and think we have become wise. There could be no better assurance that our ability to predict will never improve any father.

I've been chided for telling people that they have wasted a life's work. But I never tell anyone that. I merely present my analysis of methods and knowledge as I am doing here. If I have made my case clearly, then people will be able to draw their own conclusions about their own life's work, both what it has been and what it will be in the future.

Best to all, Bill P.

Date: Mon May 09, 1994 10:57 am PST
Subject: Re: More harping

[From Bill Powers (940509.0815 MDT)]

Hans Blom (940509 --

> The reliability of the outcome of the experiment is thus much more reduced than seekers for 'hard' quantitative truths demand. We have recently seen such problems with AIDS-related drugs. This is the fate of many experiments in medicine (which I am more familiar with than psychology), but I assume that the argument is valid more generally: the more elaborate -- and thus costly -- the experiment, the more certainty is obtained. If cost (or time; time is money) is to be reckoned with, less than absolute certainty will often be an acceptable compromise solution.

Acceptable to whom? To the patient who gets exactly one chance to be cured, or to the doctor who has thousands of chances to cure different patients, and can prove that more people are cured than killed by the treatment?

You're quite right, of course, in terms of practical considerations in the world as it is now. We always have to do the best we can with what we know, and this always means we settle for less than full understanding. But should we then decide that since this is the best we can do, we should focus on defending the way we do things instead of trying to do better?

> What about humans as control systems? I am often operating under conditions of less-than-certain knowledge. Yet I may have to act with some urgency. What to do? Perform experiments first in order to collect additional knowledge? But while I do that, I cannot control, and that may be costly.

When you feel you must act on insufficient knowledge, then you must act. This, however, doesn't mean that acting would necessarily be better than not acting, unless you can count on having many chances to succeed. If you don't bear the costs of your own failures, it's easier to act under the policy that some action is better than no action.

> That may be costly as well. We have a real dilemma here. In psychology, this dilemma is visible in the people we call perfectionists. In extreme cases, their search for a "perfect" solution takes so much time that they are doomed to a life of total inertia. To some degree this style of behavior shows up in a great many scientists.

This is the positive side of your case. The negative side appears when the research money is handed out. Where should we apply our efforts? To doing practical things with what we already know, or to extending and improving what we know? It's easy to label a scientist who wants to improve knowledge as a "perfectionist," but that can too easily turn into a simple defense of the status quo.

> In medicine, one bit often suffices: Do I take drug X or not? Do I take drug X or drug Y? I see nothing wrong with this approach if the data supply not enough information to know more. You would like to know _how much_ of drug X to take (correlation, regression line), and so would I, but medicine -- due to the difficulty of experiments -- must often be satisfied with standard dosages (per kg of body weight, as custom dictates). Sometimes more is known about the response of a drug, yet most frequently your family doctor will prescribe either one, two, or three pills a day or some similarly coarse alternative. The accuracy of medical prescriptions is often in the order of 20 to 50 percent, I would guess, corresponding with 1 to 3 bits of information per decision.

This is a good description of the dilemmas facing us when our knowledge is primarily statistical. But it is a description completely within the universe of statistical thinking. What I would hope to see in medicine is the development of understanding of the biochemical and behavioral systems, so we can not only predict what the effect of a given drug will be, but explain how that drug perturbs all the interconnected systems in the body, and how the response of the whole system to that perturbation generates the changes of operation that we call "the effect" of the drug.

My feeling is that when we finally do develop a systems analysis of the organism, in 100 or 1000 or however many years it will take, we will look back on today's drug treatments with horror.

> Remember that statistics -- and all theories that are based on it, such as optimal control theory -- started out as the theory of gambling. Much of behavior can be seen as gambling: do I take path X or path Y? PCT deals only with situations where behavior is both unconstrained (except maybe by extremes) and immediately goal-directed: the path to the goal must be there to perceive.

In PCT we don't believe that it's necessary to perceive "the path to the goal." The goal is not something that exists outside the organism or in the future. It's a reference signal that specifies, in present time, the desired end-state of perceptions, not the means of achieving that state. Since the current difference between what is being perceived and the state defined by the reference signal is always calculable, it is possible to base present actions on that error so the actions tend to make the error smaller. Constraints on behavior are simply disturbances, that are overcome by adjustments of action which keep the error decreasing. To an outside observer this looks like "finding a path to the goal," but inside the control system something quite different is happening.

As to the relation between gambling and statistics, I fully agree. Acting on the basis of statistical perceptions is always a gamble, and as all gamblers do, one tends to remember the wins more vividly than the losses.

> In poker, the choice is between playing a card and forcing a showdown in a situation where lying (mis-information) is explicitly allowed. PCT is not the theory for these uncertain and/or one-out-of-a-few choices, but other theories are (more or less).

Why would PCT not be the theory for uncertain situations? You have a desired outcome, and a perception of the actual or anticipated outcome, and act on the basis of the discrepancy. The only difference between that kind of control and the kind based on accurate representations of the world is that the statistical approach provides much poorer control of what is perceived. You may justify controlling on the basis of statistical perceptions by saying that this is all we have available, but that does not make the quality of control any better. If your perception is uncertain by 20%, your errors will amount to at least 20% of the goal specification, and you can't do any better than that.

> On the other hand, those other theories cannot express what PCT can. Conflicts between theories? Sure. Is that a problem? No. Unless you pretend that PCT is the one-and-only theory to d describe _all_ behavior.

I really don't see why these "other theories" don't fall under PCT. Are they not concerned with making perceptions match some desired state?

> I maintain that _every_ theory is a description only. An 'explanation' is a description in terms of already known concepts. I see a theory as just a concise summary of observations or experimental data.

Yes, I've heard that idea before. I think that idea is strongly influenced by the statistical approach to nature. It doesn't however, account for theories that propose underlying mechanisms that are not directly observable. "Already known concepts" have to come from somewhere; if one took your definition literally, there could be no descriptions because there would be a grossly insufficient supply of already known concepts and no way of generating more. Where did the first "already known concept" come from?

Concepts have to be _invented_, and a very fruitful way of inventing them is to propose models containing imaginary underlying mechanisms. Models which start out as ad-hoc explanations of phenomena often end up predicting new phenomena, because their components can be imagined to interact in new ways according to their postulated properties, and these new ways often lead to new observations. When the modeling process is subjected to the discipline of experimental test, models can become uncannily useful; we often find, at the end of this process, that advances in technology reveal to us the very components we had guessed must exist. This is not just description.

> Prediction becomes possible only if you have the additional knowledge that the future will be like the past, that the relationships between variables will remain unchanged over time. This is particularly untrue for systems that learn (i.e. that do things differently at different times), such as adaptive control systems and (most :-) humans.

When predictions can't be made in detail, sometimes they can be made on a more abstract level. We can say, for example, that whatever results from the learning process, it will be a control system. Often, as in PCT, we can make accurate predictions about the _relationships_ between variables without being able to predict the _values_ of the variables. So we can say that actions will oppose the effects of disturbances on controlled variables, even though we can't predict what disturbances will occur and therefore what actions will occur. The better we understand the control hierarchy, the more of these relationships we will be able to predict, even though predicting the actual values of controlled variables becomes impossible.

In studying an adaptive or self-reorganizing organism, prediction centers on identifying the criteria of reorganization: the intrinsic reference states that explain when and how rapidly reorganization will occur, and under what conditions it will stop. We can't predict how the rat will cause the lever to depress in the Skinner Box, but we can predict that the method of pressing will cease changing when the result is that the animal receives enough reinforcers to satisfy its needs. If we happen to know what behavior is required to achieve that end, we can even predict what behavior will then be occurring.

>> So in defense of statistical theory, we have the anecdote which illustrates how the theory works when it seems to work, accompanied by a blank lack of interest in counterexamples.

> Actually the opposite is true in science, as Popper has indicated: one counterexample destroys a theory, whereas an infinity of positives cannot make it true. Regrettably, if a theory has a noise term, the situation becomes much more complex, because in such cases the theory says something like "in X percent plus or minus Y percent of the cases Z is true". Testing such propositions requires some education in statistics and may be too costly -- see above.

I was speaking of the way statistical findings are actually used and described in the literature, not the "scientifically correct" way of dealing with them. If statistical findings were reported in a way that is scientifically impeccable, we would never hear that taking one aspirin per day will protect against heart attacks. We would hear that in a population of individuals who take one aspirin per day, there is, for no known reason, a very small difference in the rate of heart attacks in comparison with a population that does not take one aspirin per day. We would also hear that the chances of this regime helping any particular person are so small that the benefits probably do not exceed the cost of going to a store and buying the aspirin.

One counterexample can indeed destroy a theory, or at least demand that it be revised. But too much is made of the other logical conclusion, that no number of positive examples can prove the truth of the theory. This is often used as a reason for accepting confidence levels that are ridiculously low. Why should we accept $p < 0.05$, when we might achieve confidence levels of $p < 0.000001$? There is a qualitative difference between such confidence levels, the difference between having an isolated fact and a fact that can be used in extended reasoning with many other equally reliable facts. It is the difference between superstition and science.

> You seem to vacillate about the point whether prediction is possible at all (as I understand some of your previous postings) and the high predictive power that a theory must have. Have I lost you here?

Yes, I meant to deal with that and forgot to. Prediction is possible under conditions where we can specify disturbances and reference signals. That is what must done to test any model. It turns out that we can do this in the laboratory well enough to predict behaviors with an accuracy that is quite good. But the point of doing this is not to predict behavior. It is to establish that we have the relationships in the model right. When a person is outside the laboratory situation we can no longer predict what the important disturbances will be, or how the reference signals will be set. But we can still test the model after the fact by seeing whether it explains the relationships among variables that we are observing.

The more levels of control we can test in the laboratory, the more situations there will be where the model continues to explain behavior "in the wild." Higher-level reference signals change more slowly than lower-level ones, so we can see behavior at the higher levels under conditions of approximately constant reference signals, and by using the Test get some idea of the controlled variable and its reference level. Then we can explain actions in some detail by seeing what disturbances are acting.

There must be _some_ way to challenge a theory through experimental testing, to make it meet the requirement of correctly predicting outcomes. Without that, no theory should be let loose into the world. Laboratory tests give us the opportunity to make the model handle difficult cases and apparent exceptions, if it can.

>> Sometimes theory A works best, and sometimes theory B works best. You use whichever one works.

> Yes, that is the currently most prevalent view of what theories are. Theories are _tools_, and just as we have only a limited set of tools in our toolbox, words in our language, muscles in our body, reflexes in our motor system, the task is to accept those constraints and come to an "optimal" solution within the limitations of those constraints. Those limitations are what makes modelling a twelve degrees of freedom arm so difficult, yet they are real, so we'd better take them into account.

Why does this sound like a cop-out to me? If you have a collection of phenomena, and a collection of different theories about them, then no matter what happens you will be able to find one theory that would have predicted the observation. Now, however, the problem of prediction becomes the problem of knowing which theory to use. If you can predict in advance that theory A should be used, and not B, C, or D, then you have a theory that encompasses A, B, C, and D, and should use it instead. If you have no such encompassing theory, then you might as well flip a coin. Not knowing in advance which theory to use, you remain completely unable to make a prediction.

The only thing you have to gain from multiple theories about the same phenomenon is that after the fact you can always claim to be able to explain it. You just sort through the various theories to see which one predicts what has already happened. If you have a friend who has a theory that predicts just the opposite of yours in every case, then you can team up, because together you have all the possibilities covered.

I have this theory that under the right conditions, objects should fall upward. While so far there has been little applicability for this theory, if you ever hear of such a thing happening, be sure to tell everyone that I knew it all along.

> Science tries to go up a level above this urge and build "objectivity", i.e. it invents concepts that we can agree about. Sometimes with success (the "simple" objects of physics), sometimes not (the "complex" subjects of psychology).

If psychology deals with such complex subjects and physics such simple ones, why are the theories of psychology so utterly simple-minded, and those of physics so complex?

> One weakness of PCT lies in its proposal of The Test as the tool to decide as to which variable is controlled. The Test is impractical, since it would need to test an infinity of possibilities, some plausible, others less so. A task never truly finished...

Bumblebees can't fly, either. Have you tried the Coin Game?

> That is the source of another weakness of The Test in PCT: in general, we simultaneously control a great many interacting variables. There is in PCT no prescription as to how to create the right experiments that tease those interactions apart so that just one controlled variable can be discovered.

You ought to try some experiments. It's not as hard as it looks.

> Remember that physics is about inanimate matter. As soon as learning, habituation, adaptation or what have you play a role, accurate modelling may become impossible because of the immensity of the task to establish, for a particular individual, that _all possible past and present variables_ have been taken into account.

These armchair pronouncements don't impress me. Everything looks harder when you're only imagining it. In fact, many of our experiments work just the way we thought they would work when we conceived them. All those other unknowable circumstances turn out to make no perceptible difference. Control systems are very good at operating properly in ill-defined situations.

> Physics depends on the reproducibility of experiments, on objects that do not change over time, and on classes of identical objects.

No, it depends on understanding which variables are important and which are not. If a system changes over time, you try to understand the principles of change and take them into account. If the objects are not identical, you try to see how they differ, and learn how to handle the differences. If you give up before you start, you will never understand anything.

>> So we should go to kindergarten and learn how to make simple predictions that are always borne out by observation. There is no short-cut.

> Philosophers have always seen it as their task to come up with a number of statements or explanations that _everyone_ could agree with. This has proved to be an impossible task, even if only truthful people "in their right mind" (these value judgments form a weakness, of course) are allowed to participate in the debate: different people have different opinions. Even Bishop "if I don't perceive you, you don't exist"

Berkeley, whose position is quite unpopular (probably for personal reasons, because people don't like the idea that their existence depends upon someone else's perceptions), cannot be rejected because his point of view is logically unassailable.

That's why logic is insufficient as a way of understanding nature. Nature has its own logic, which immediately becomes apparent when you do experiments and try to predict what will happen. Nature doesn't care that things HAVE to work a certain way, or that certain things CAN'T happen -- according to human logic. What happens happens, quite independently of what we think or hope will happen. People who don't do experiments to test their infallible conclusions get an overinflated idea of the power of formal thinking.

I wasn't talking about philosophical agreements. I was talking about making correct predictions of simple phenomena by actually doing experiments, not just by talking.

> Such a breakdown (reductionism) fails as soon as gambling plays a role. Your first ticket in the sweepstakes doesn't pay off, nor your second or third or hundredth. Your "complex prediction" would logically be that no ticket will ever pay off. Philosophers-logicians will tell you that induction is not an infallible tool. By insisting on certainty you focus your attention, I think, on such a small segment of reality as to become uninteresting.

Oh, balderdash. If organisms behaved as randomly as you seem to believe they do, life would be utter chaos. It is not. Organisms behave in very regular ways most of the time. The only reason this regularity has not been found is that the wrong models of behavior have been used.

> You gambled when you chose a job: would that be the most satisfying one? How could you know? You gambled when you chose a partner, decided for children, on a place to live, on what not. Or were you really driven by an unfailing sense or certainty in all those cases?

What you forget is that if things don't turn out exactly as you wished, you take action to make them become more like what you wanted. You never just roll the dice and live forever with the consequences. If things don't naturally go the way you want, you MAKE them go the way you want. That's the PCT model, and that's how people work. Your world view doesn't take intention into account, or the capacity to control. What good does it to do compute the odds, when you are almost always capable of altering the outcome?

I was not driven by an unfailing sense of certainty, but by the knowledge that whatever happened, I could probably deal with it.

> Hooke's Law is a _tool_, a useful approximation that works under certain conditions. It is not A Truth.

Yes, but it's one of the sharper tools. Are you recommending that we consider all tools to be equally useful just because we call them all tools?

>> The gross mistake made by conventional behavioral scientists was to think that any piece of knowledge, however uncertain, is better than no knowledge at all. It is not; it is far worse than total ignorance.

> This is where gamblers of all type will disagree with you. Any tiny piece of knowledge that someone else does not have improves your odds. What is your favorite game? Chess more likely than bridge or even poker...

I am not impressed by gamblers, in general. They are mostly suffering delusions that they are immune to the laws of chance, to the enrichment of the casino operators, who take no chances. Chance plays very little part in most of human behavior. The wise person does not pretend that adequate control can be maintained under conditions of great uncertainty. The proper thing to do when there is a lot of noise is to change the environment until the noise disappears. You don't go on driving a car with a loose steering linkage. You fix it.

> Wisdom is, I think, having a good set of tools. If you have only a hammer, the whole world looks like a nail. Even with good tools some problems can- not be solved well. But maybe just good enough...

I would not say that many of our solutions to the big problems are good enough. I think that most of them are lousy and getting worse.

What good does it do to have a large selection of tools if you have no way of knowing when or where to use them? Or if you can't distinguish between a sharp new saw and an old rusty dull one?

Best, Bill P.

Date: Tue May 10, 1994 8:38 am PST
Subject: to control or not to control (was: re: harping)

[Hans Blom, 940510] (Bill Powers (940509.0815 MDT))

So much to say, so little time. Just a few things, then.

> You're quite right, of course, in terms of practical considerations in the world as it is now. We always have to do the best we can with what we know, and this always means we settle for less than full understanding.

This may be a misreading of what I mean. I intended to say that full understanding, both in theory and practice, _includes_ an awareness of and a reckoning with our limitations. If not, the understanding is not full.

> But should we then decide that since this is the best we can do, we should focus on defending the way we do things instead of trying to do better?

Two roads are open, as I tried to describe. The first is to conduct experiments in order to be able to control better at a later time -- but not now. The second is to act now -- and realize that afterwards you will always be able to say that you could have done better if only you had known better.

Since we live in a real-time world, acting "optimally" is to combine the two options in the best possible way: do not hesitate for too long, and do not act too stupidly. You and me, we are lucky in that the two aspects coincide: doing experiments _is_ our favorite way of acting!

> When you feel you must act on insufficient knowledge, then you must act. This, however, doesn't mean that acting would necessarily be better than not acting, unless you can count on having many chances to succeed. If you don't bear the costs of your own failures, it's easier to act under the policy that some action is better than no action.

Here speaks a perfectionist. No action, I assure you, _is_ a form of action. Let me try this on you, as an educator:

"Failures do not exist. You are learning. Learning means that you are finding out how to do things that you do cannot do yet. Learning is necessarily a "trial and error" process. Others may say that you make mistakes, but I tell you that those "mistakes" are required in order to truly learn. If you were an expert already, you would not make mistakes. But neither would you learn. Experts have learned already, they have made their "mistakes" already. They were not afraid to make those errors; otherwise they would not be experts now. Do not say "I made a mistake"; say "the outcome was different from what I intend- ed". See the other side of the coin: making a mistake _is_ learning. Be happy to make mistakes; it shows that you learn."

Why is it that learning receives so little emphasis in PCT? You always stress the steady state, _being in control_. I am much more interested in the dynamics, learning, becoming an expert. In our discussions, too. You focus on how well we do, I on how much there is still to learn.

> It's easy to label a scientist who wants to improve knowledge as a "perfectionist," but that can too easily turn into a simple defense of the status quo.

See how you emphasize the steady state situation?

> In PCT we don't believe that it's necessary to perceive "the path to the goal." The goal is not something that exists outside the organism or in the future. It's a reference signal that specifies, in present time, the desired end-state of perceptions, not the means of achieving that state. Since the current difference between what is being perceived and the state defined by the reference signal is always calculable, it is possible to base present actions on that error so the actions tend to make the error smaller.

Yes, that is what I meant to say. But there are situations where a "current difference" is not available, as when you are lost in the woods and don't know where you are. In such cases it is _not_ possible to base present actions on a computable "error" so that subsequent actions make that error smaller. The newspapers just reported about two boys who escaped from an institution and hid in a cave. When their dead bodies were finally found, it turned out that the cave was only some 200 by 600 feet in size.

> As to the relation between gambling and statistics, I fully agree. Acting on the basis of statistical perceptions is always a gamble, and as all gamblers do, one tends to remember the wins more vividly than the losses.

Control theories that explicitly model an error term consider _any_ action to be something of a gamble, because _some_ information is always lacking. The model is, after all, an approximation. As to your remark that humans tend to be optimists: I think that this is due to a deep-seated recognition that learning takes place, and that things that we cannot do today can be done tomorrow. But then, this feeling can be exploited as well...

> You may justify controlling on the basis of statistical perceptions by saying that this is all we have available, but that does not make the quality of control any better.

Not better, _best_. Optimal, given the existing limitations.

> I really don't see why these "other theories" don't fall under PCT. Are they not concerned with making perceptions match some desired state?

Yes and no. Most "other theories" recognize that you generally have many simultaneous and possibly "conflicting" goals and that some "optimal" compromise is required as to how to fulfill which goals by how much. You want to stay home and read that book; you want to visit that friend; you want to see that new movie; you want to make love with your lady; you want to ... What are you going to do _now_?

>> I maintain that _every_ theory is a description only. An 'explanation' is a description in terms of already known concepts. I see a theory as just a concise summary of observations or experimental data.

> Yes, I've heard that idea before. I think that idea is strongly influenced by the statistical approach to nature. It doesn't however, account for theories that propose underlying mechanisms that are not directly observable. "Already known concepts" have to come from somewhere; if one took your definition literally, there could be no descriptions because there would be a grossly insufficient supply of already known concepts and no way of generating more. Where did the first "already known concept" come from?

Our perceptual apparatus appears to abstract raw perceptions into concepts like lines, corners, pitches and what have you. The concept "object" is there, as Piaget has demonstrated. So are the concepts "surprise", "fear", "love", "hate", amongst a great many other things. The basis is our physiology.

> Concepts have to be _invented_, and a very fruitful way of inventing them is to propose models containing imaginary underlying mechanisms.

More complex concepts are built upon the elementary ones. Some of these invention processes have been described in the literature. Kekule is said to have come upon the structure of the benzene ring thanks to a daydream of a snake biting its tail. The results of this creative process of dreaming, fantasizing, imagination seem to be rather haphazard.

> One counterexample can indeed destroy a theory, or at least demand that it be revised. But too much is made of the other logical conclusion, that no number of positive examples can prove the truth of the theory. This is often used as a reason for accepting confidence levels that are ridiculously low. Why should we accept $p < 0.05$, when we might achieve confidence levels of $p < 0.000001$? There is a qualitative difference between such confidence levels, the difference between having an isolated fact and a fact that can be used in extended reasoning with many other equally reliable facts. It is the difference between superstition and science.

I see one bird, and it can fly. The same with a second, a third, a twentieth. Can I now conclude that _all_ birds can fly? The same with a hundredth, a tenthousandth, a millionth. Can I now? When can I? Much seems to depend on where I do my observations and where I will need those conclusions at a later time. When can I be certain about a generalization? I do not see how a _qualitative_ difference arises; it's just a matter of confidence -- or rather reliance.

> There must be _some_ way to challenge a theory through experimental testing, to make it meet the requirement of correctly predicting outcomes. Without that, no theory should be let loose into the world. Laboratory tests give us the opportunity to make the model handle difficult cases and apparent exceptions, if it can.

That does not help. I gave the example of relativity theory versus quantum theory. The lab proves both accurate to full experimental precision. They were let loose into the world. Yet they are in conflict, i.e. they cannot both be true; most likely neither is. The deep discrepancy between "truth" and "useful tool" remains.

> Why does this sound like a cop-out to me? If you have a collection of phenomena, and a collection of different theories about them, then no matter what happens you will be able to find one theory that would have predicted the observation. Now, however, the problem of prediction becomes the problem of knowing which theory to use. If you can predict in advance that theory A should be used, and not B, C, or D, then you have a theory that encompasses A, B, C, and D, and should use it instead. If you have no such encompassing theory, then you might as well flip a coin. Not knowing in advance which theory to use, you remain completely unable to make a prediction.

The solution is called "expertise". I remember my math classes, where I learned the tricks of differentiation thirty years ago. There was, much to my surprise, no "theory" that described how, given a certain formula, to derive its differential. There was just a bag of tricks, each one applicable under certain conditions. Initially I had a very bad match between formula form and applicable method. But through practice I got better. Not always the first method worked, but then the second or third did. You see the same thing in many other circumstances. You have undoubtedly seen the average housewife (excuse my political incorrectness) operate a too big screwdriver on a too small screw. You, as an expert on screwdriving, immediately, without thinking, pick the right screwdriver.

For the moment I see no possibility to translate this kind of expertise into something that you might call "laws".

> The only thing you have to gain from multiple theories about the same phenomenon is that after the fact you can always claim to be able to explain it. You just sort through the various theories to see which one

predicts what has already happened. If you have a friend who has a theory that predicts just the opposite of yours in every case, then you can team up, because together you have all the possibilities covered.

Stories _after_ the fact are pretty useless, isn't it? After the fact, one can defend anything. What do you really mean? That you'd rather talk about your theories than use them?

> I have this theory that under the right conditions, objects should fall upward. While so far there has been little applicability for this theory, if you ever hear of such a thing happening, be sure to tell everyone that I knew it all along.

Let me play the educator again. Can you think of conditions where objects _do_ fall upward? It isn't difficult :-)

>> Physics depends on the reproducibility of experiments, on objects that do not change over time, and on classes of identical objects.

> No, it depends on understanding which variables are important and which are not.

Is "important" an objective notion? Or do you mean the same thing as I do: that the important things should reproduce (no two experiments ever reproduce in all aspects), that the important things should remain constant over time (we probably need not keep track of the identities of the atoms that constitute a body cell), and that identity is taken as that of the important similarities (no two objects can be composed of the same individual atoms)?

> Nature has its own logic, which immediately becomes apparent when you do experiments and try to predict what will happen.

Take care of what you say here! We had just established that people (philosophers, the ones who have thought hard about these things) do not agree. So when you say that nature's logic immediately becomes apparent, it is apparent TO YOU. Someone else might see things differently. What makes you so special?

> What happens happens, quite independently of what we think or hope will happen.

It's all perception. What happens may happen, but what I perceive of what happens has a lot to do with what I pay attention to and what I expect. I am sure that I will see different things -- or things differently.

> I wasn't talking about philosophical agreements. I was talking about making correct predictions of simple phenomena by actually doing experiments, not just by talking.

I'm afraid that you do not take philosophers very seriously. They are, after all, the ones who have paid the most attention to some of the things most central to PCT: how come we do what we do, how come we know what we know, what _can_ we know, etc. They're all modelers, in a way. And they have been very aware of contradictions between what they think/say and what they see -- helped by a centuries spanning debate that, although it has not led to any definite truths, has debunked a great many capricious theories.

> If organisms behaved as randomly as you seem to believe they do, life would be utter chaos. It is not. Organisms behave in very regular ways most of the time.

That is what you perceive. You concentrate on control. I see a lot of utter chaos, a lot of things that we would like to but cannot control. Illness, death, unhappiness, loneliness. Very organic themes all. Congratulations that you have everything under control. I must admit that I do not.

> If things don't naturally go the way you want, you MAKE them go the way you want. That's the PCT model, and that's how people work.

Oh, how I wish that were true!

> I was not driven by an unfailing sense of certainty, but by the knowledge that whatever happened, I could probably deal with it.

That is much more in the spirit of what I intended to say all along. It is what I call being able to live with uncertainty.

> The wise person does not pretend that adequate control can be maintained under conditions of great uncertainty. The proper thing to do when there is a lot of noise is to change the environment until the noise disappears. You don't go on driving a car with a loose steering linkage. You fix it.

Ban war? Ban poverty? Ban discrimination? Illness? Death? I'm afraid that we will have to live with our human limitations, now and forever more.

> I would not say that many of our solutions to the big problems are good enough. I think that most of them are lousy and getting worse.

You just said: "The proper thing to do when there is a lot of noise is to change the environment until the noise disappears. You don't go on driving a car with a loose steering linkage. You fix it." Please do. I would be very, very grateful.

Greetings, Hans

Date: Tue May 10, 1994 3:07 pm PST
Subject: Re: to control or not control

[From Bill Powers (940510.1400 MDT)] Hans Blom (940510)

> If you were an expert already, you would not make mistakes. But neither would you learn. Experts have learned already. They have made their "mistakes" already.

Sounds like something written by an expert.

> Yes, that is what I meant to say.

Wouldn't it have been better to say it, then?

> But there are situations where a"current difference" is not available, as when you are lost in the woods and don't know where you are.

I think you're externalizing "error." Or are you saying that when you perceive yourself to be lost in the woods and don't know where you are, you don't wish to know where you are?

> The newspapers just reported about two boys who escaped from an institution and hid in a cave. When their dead bodies were finally found, it turned out that the cave was only some 200 by 600 feet in size.

Sorry, I don't get the point. Would they have been alive if the cave had been 2000 by 6000 feet in size? Are you saying they didn't perceive themselves to be in a cave and wish to be out of it?

> Control theories that explicitly model an error term consider _any_ action to be something of a gamble, because _some_ information is always lacking.

Nothing to prevent adding an error term to any PCT model, if it improves the match of the model to real behavior. The catch is that the "error term" may not be random in the real behavior: it might reflect a systematic aspect of the control process (perhaps at higher levels) which the current model leaves out. I prefer to leave the "random" hypothesis out until as much of behavior as possible is covered by a systematic model.

> As to your remark that humans tend to be optimists: I think that this is due to a deep-seated recognition that learning takes place, and that things that we cannot do today can be done tomorrow. But then, this feeling can be exploited as well...

I didn't say they were optimists, I said they were poor at estimating actual probabilities. An error on the high side is still an error.

> The results of this creative process of dreaming, fantasizing, imagination seem to be rather haphazard.

Yes, in fact I believe it is random. Insights that are impracticable or impossible of realization are weeded out when the attempt is made, in reality or imagination, to use them. That's the experimental method.

> I see one bird, and it can fly. The same with a second, a third, a twentieth. Can I now conclude that _all_ birds can fly? The same with a hundredth, a tenthousandth, a millionth. Can I now? When can I?

"All" birds is a logic-level concept. Arriving at an abstract conclusion concerning "all" of anything is a logician's hobby, not a real problem. After you have tried walking up and grabbing your 20th bird, and it has flown away, you will decide to perceive bird-like objects as capable of flight, and start sneaking up on them, even if some of them can't fly.

>> Laboratory tests give us the opportunity to make the model handle difficult cases and apparent exceptions, if it can.

> That does not help. I gave the example of relativity theory versus quantum theory.

Help what? It helps to find out if a model fails when confronted with real phenomena, doesn't it? Are you recommending that we NOT test models just because sometimes the experimental results don't let us decide between alternatives? In fact, the only problem between quantum theory and relativity theory (perhaps you meant "continuous wave" theory) is that they seem to apply in different experimental situations. So neither theory is complete.

> The lab proves both accurate to full experimental precision. They were let loose into the world. Yet they are in conflict, i.e. they cannot both be true; most likely neither is.

Most likely. But they do not conflict. Quantum theory does not predict that the velocity of light DOES depend on the choice of inertial frame. Relativity theory does not predict that radiation does NOT occur only in quantized energy levels. The two theories are concerned with different kinds of observations of nature, which so far seem unconnected to each other. That is not a conflict.

> The deep discrepancy between "truth" and "useful tool" remains.

I'm not concerned with "truth." Only with useful tools. But some tools are A LOT more useful than others.

>> If you have no such encompassing theory, then you might as well flip a coin. Not knowing in advance which theory to use, you remain completely unable to make a prediction.

> The solution is called "expertise". I remember my math classes, where I learned the tricks of differentiation thirty years ago. There was, much to my surprise, no "theory" that described how, given a certain formula, to derive its differential. There was just a bag of tricks, each one applicable under certain conditions.

I fail to see any connection between my statement and your statement. Are you saying that you "just know" which theory to apply, when you have a choice of theories that will predict anything you like?

> You have undoubtedly seen the average housewife (excuse my political incorrectness) operate a too big screwdriver on a too small screw. You,

as an expert on screwdriving, immediately, without thinking, pick the right screwdriver.

That's not being politically incorrect, it's being an average European chauvinist pig. I think your faith in your intuitive expert rightness could probably be analyzed into something rather simple, like comparing the sizes of the slot and the screwdriver bit, and controlling for sameness. It's not my problem if you don't understand how you do things.

> For the moment I see no possibility to translate this kind of expertise into something that you might call "laws".

Right, so why not leave such things in the hands of expert theoreticians like me?

> Stories _after_ the fact are pretty useless, isn't it? After the fact, one can defend anything. What do you really mean? That you'd rather talk about your theories than use them?

I guess you really did miss my point, which was that by relying on a "took kit" of incompatible theories, you assure yourself of being able to explain anything, however trivially.

> Let me play the educator again. Can you think of conditions where objects _do_ fall upward? It isn't difficult :-)

Oh, shucks. I thought I had invented that theory.

>> No, it depends on understanding which variables are important and which are not.

> Is "important" an objective notion? Or do you mean the same thing as I do: that the important things should reproduce (no two experiments ever reproduce in all aspects), that the important things should remain constant over time (we probably need not keep track of the identities of the atoms that constitute a body cell), and that identity is taken as that of the important similarities (no two objects can be composed of the same individual atoms)?

No, that isn't what I meant. I meant that an important variable is one which, when varied, has a discernible effect on an outcome you're trying to explain.

>> Nature has its own logic, which immediately becomes apparent when you do experiments and try to predict what will happen.

> Take care of what you say here! We had just established that people (philosophers, the ones who have thought hard about these things) do not agree. So when you say that nature's logic immediately becomes apparent, it is apparent TO YOU. Someone else might see things differently. What makes you so special?

I just meant that nature has an inimitable way of telling you when your logic is wrong, or rather inapplicable. You expect one thing to happen, you stake your reputation and sanity on it's happening, and it still doesn't happen. The way I said this, the referent of "which immediately becomes apparent" was ambiguous. What immediately becomes apparent is that nature has its own logic, but nature does not tell you what that logic actually is.

>> What happens happens, quite independently of what we think or hope will happen.

> It's all perception. What happens may happen, but what I perceive of what happens has a lot to do with what I pay attention to and what I expect. I am sure that I will see different things -- or things differently.

However you perceive, when you expect one perception and fail to get it, or get a different one, you know that your expectation has been incorrect. When you do an experiment expecting to perceive a reading of 5, and read instead -

5, there is no way to pretend that you were right. Point of view has nothing to do with it.

>> I wasn't talking about philosophical agreements. I was talking about making correct predictions of simple phenomena by actually doing experiments, not just by talking.

> I'm afraid that you do not take philosophers very seriously.

Don't be afraid. Philosophers may go into a snit if you don't take them seriously, but they aren't usually dangerous.

> They are, after all, the ones who have paid the most attention to some of the things most central to PCT: how come we do what we do, how come we know what we know, what _can_ we know, etc.

And how many angels can dance on the head of a pin. Just what is there about philosophy that we ought to take "seriously?"

> They're all modelers, in a way.

Oh, come on now. Modelers? What philosopher can you mention who actually constructed a working model of some process and then tested its predictions against observations? From my reading in philosophy, I have concluded that most philosophers are very careful to avoid saying anything that could be tested experimentally.

>> If organisms behaved as randomly as you seem to believe they do, life would be utter chaos. It is not. Organisms behave in very regular ways most of the time.

> That is what you perceive. You concentrate on control. I see a lot of utter chaos, a lot of things that we would like to but cannot control. Illness, death, unhappiness, loneliness. Very organic themes all. Congratulations that you have everything under control. I must admit that I do not.

It's very hard to argue with someone who keeps changing the subject between paragraphs. I have never said or implied that we are capable of controlling everything. What I said was that behavior is very regular. You are completely surrounded by man-made artifacts, every one of which required highly regular and reliable control processes to be produced. When you turn your computer on it works. When you type messages, they show up as you intended them to show up with a very high degree of reliability. Practically everything in your life that is concerned with human behavior is evidence of incredible repeatability and regularity. While you can do little about being dead, there is a lot you can do about being ill, unhappy, or lonely -- and you very reliably will do what you can, usually with success.

The problem is that you are so used to having everything you want to happen come about with an almost imperceptible amount of effort that you have come to take this magical process for granted, and you complain about the few deviations from regularity that do occur. You seem to think that irregularity and chance predominate. In fact, it is hard to find an aspect of your life and your interactions with the environment that is not under continuous and precise control. Don't look at the ripples on the surface: look at the ocean!

> If things don't naturally go the way you want, you MAKE them go the way you want. That's the PCT model, and that's how people work.

> Oh, how I wish that were true!

But it is true! You simply think of it as "nothing happening." You had no difficulty typing "Oh, how I wish that were true!" You had no difficult getting to work (or wherever you are) this morning. You got your breakfast or lunch into your mouth without any serious uncertainties. Practically every single thing you attempt to bring about in your perceptions, today and every day, comes about so nearly like what you want that you ignore this amazing fact: you call this "doing things." How many times in the last month have you fallen down while you were walking? How many times did you intend to utter one sentence and hear something entirely different come out of your mouth? How many times did you attempt to sit in a chair, and fail?

The world you experience is exactly the world you intend to experience, 99.9% of the time or more. That is what control theory explains: the 99.9%, not the 0.1%.

>> I was not driven by an unfailing sense of certainty, but by the knowledge that whatever happened, I could probably deal with it.

> That is much more in the spirit of what I intended to say all along. It is what I call being able to live with uncertainty.

Control systems don't live with uncertainty; they remove uncertainty. They impose order on chaos. They don't require the world to be uniform and repeatable and reliable. In a nonuniform and nonrepeatable and unreliable world, they can produce uniform, repeatable, and reliable consequences minute by minute, day by day, and year by year. They make randomly-growing forests fall down, split into uniform sizes of timber, and line up in rows and layers to make houses. They dig rocks out of the ground and turn them into automobiles. They manipulate entropy and squirt men to the Moon, and back. They turn sand into computer networks like the one we're using.

> Ban war? Ban poverty? Ban discrimination? Illness? Death? I'm afraid that we will have to live with our human limitations, now and forever more.

I have no doubt that we will ban anything we can agree to ban (I'm not sure that wisdom would be on the side of banning death, unless reproduction were also banned -- and evolution). It's up to us, because we are purposive systems. I deplore the idea of simply accepting what _is_ as the necessary natural order of human existence. That is a great way to perpetuate whatever is wrong with human life.

>> I would not say that many of our solutions to the big problems are good enough. I think that most of them are lousy and getting worse.

> You just said: "The proper thing to do when there is a lot of noise is to change the environment until the noise disappears. You don't go on driving a car with a loose steering linkage. You fix it." Please do. I would be very, very grateful.

Well, don't just sit there arguing and waiting for someone else to do it for you. I don't want your gratitude, I want your help.

Best, Bill P.

Date: Tue May 10, 1994 4:30 pm PST
Subject: Re: HARPING

[Dan Miller (940510)] Thomas Baines (940509):

In sociology the theoretical perspective called structural functionalism was developed in the 30s - with a striking parallel to Keynesian economics. The underlying assumption (model?) was that societies exist in equilibrium. Any long term deviation from this equilibrium is pathological, and the society will grow increasingly unstable. Short term deviations that are controlled through social control mechanisms (the military, police, mental health professionals, etc.) can be functional for the system by enforcing rules and conformity to those rules. For this type of stable society to exist there must be agreement on values, rules, and procedures. This consensus and the equilibrium model it supposedly reinforces formed a mostly unquestioned approach to sociology for several decades. This model did not work. It did not pass muster, but it did survive because of political maneuverings in academia.

More recently, rational choice models have taken over. Here the unquestioned assumption is that individuals operate using a rational calculus that translates (roughly) into a cost/benefit analytic procedure. People behave as they do because of self-interest. Even altruistic behavior is reduced to this form of rational calculation. Again, this approach is compatible with your economic equilibrium model. Do these people trade notes, manuscripts, library cards?

The underlying assumptions of each go unquestioned, or more critically, untested. The opposition in sociology can be found among symbolic interactionists (several read the posts on this NET), post-modern critics, and the various branches of neoMarxism.

The opponents to these stifling paradigms lack, I fear, the necessary organizational and political savvy. Like Tom, the central character in Michael Crichton's DISCLOSURE, we are intelligent, but naive. Given Crichton's description of the vicious social acts in that world, I am not sure I would want to play the "necessary games."

Always, Dan Miller

Date: Tue May 10, 1994 4:42 pm PST
Re: HARPING

[Avery.Andrews 940511.1024] (Dan Miller 940510)

If we really want to get the better of the sociologists and economists, the most satisfying way to do it would be to develop better ways of predicting financial market behavior. Like, why waste your time fighting with idiots if you can just take their money away?

Avery

Date: Wed May 11, 1994 12:32 am PST
Subject: Re: social theories

[From Bill Powers (940511.0100 MDT)] Dan Miller (940510)

> More recently, rational choice models have taken over. Here the unquestioned assumption is that individuals operate using a rational calculus that translates (roughly) into a cost/benefit analytic procedure.

This bears some resemblance to a PCT approach, except that the models of this kind that I have seen assume that each person calculates the _actual_ costs and _actual_ benefits, in objective terms. This assumption has the handy feature that an analyst can calculate how different people will behave just by examining the objective facts. This is probably why this theory doesn't work.

My impression is that people work according to a subjective cost-benefit analysis (which translates directly into error signals), with a limited comprehension of consequences and a short-range concept of goal-achievement. An anecdote that illustrates the general principles under which most of my Chicago-bred coworkers seemed to operate back when I worked for a newspaper:

Some years back, the Chicago area had torrential rains that caused much flooding. In the suburbs, the main problem was dealing with runoff. Leaves and junk would accumulate over some drains and cause local overloading of the storm sewers elsewhere which would back up into people's basements. In a lunchtime discussion, I innocently said that the obvious thing to do would be for everyone to take charge of the drains near their houses and keep them clear, so the runoff would be evenly distributed and the storm sewers could work properly. This proposal met with instant derision among my co-workers, all of whom had grown up in Chicago but some of whom were now able to live in the suburbs.

They explained to me that if you went out and unclogged the drains near your house, the water in the street resulting from clogged drains near you would all go into your drain, and flood your basement. So the obvious thing to do was to make sure your own drain stayed clogged until the water went down.

Despite my arguments as to what would be the obvious result, my buddies insisted that this was the smart way to do it, and the only way to avoid being a sucker (like me). The Categorical Imperative offered by a minority of one didn't have a chance against this rational assault.

This is clearly a rational-choice philosophy, but based on a very limited and short-range concept of social interactions (not to mention hydraulic ones). People do not operate according to facts and rational choices, but according to what they _believe_ to be facts and rational choices. My friends considered that they were being highly logical. They were functioning perfectly well at the logic level, and obviously considered behaving according to logic to be a desirable goal. But their view of the world in which they applied their logic extended about 50 yards in every direction, and was based on rather approximate facts.

None of these people were neighbors, but I asked what they would do if they were. The answer was that they would agree not to unclog their drains. Let other people farther up and down the street be the first ones. So the fact that these people worked together did extend their social horizons a little, enough to include two adjacent lots. But including a whole block was apparently unimaginable.

This is a lot like the way the people in the CROWD demo operate. Each individual completely ignores all the others except the ones involved in an immediate interaction that causes error signals. Each individual acts to reduce its own error signals, leaving others to cope with their own.

What my friends at work were doing was very similar, although of course they did have a small circle of others that they took into consideration, their families and relatives first, friends next, and mere acquaintances way down the list. They were a lot more complex than the individuals in the CROWD program. But the circles were small in comparison with the size of the surrounding community.

Also, since some of these people had transplanted from Chicago to the suburbs, they found themselves in larger communities where different values prevailed. There was clearly fertile ground there for conflict: all the good suburbanites going out to clear their drains -- except my friends from Chicago, who would firmly ward off any attempts to "help."

When you think of limited social horizons, the question is no longer what principles operate in a whole society, but how the differing social principles within these relatively confined circles of concern interact with each other. The appearance that there is one large society is deceptive: there are really hundreds of thousands of interacting tribes, each with its own beliefs about what is true and what is logical. And in many contexts, the tribe is simply one family, or one individual.

Perhaps the circles are of different sizes depending on what specific areas of concern are involved. That idea rapidly gets over my head.

Governments bring up an interesting idea, because governments are run by people who specifically become concerned with large bunches of other people - their horizons are broadened by their choice of occupation, or perhaps people with broader horizons self-select for this occupation. If you judge a society by what its governments do, you get quite a different picture of the uniformity of the society. In some cases, governmental functions can temporarily bring people together to resist attempts to create uniformity: witness the Internal Revenue Service. In other cases, people can be brought together temporarily to support a common effort, as in the case of an external threat that disturbs many individuals, even though for different reasons in each tribe. The history I have read seems to be mainly about what politicians and governments do, which creates a picture of concerted action that is probably entirely misleading. The natural unit of organization is a lot smaller than a government.

Trying to find any truly emergent social principles in a country the size of the U.S., it seems to me, would be a very tough job. Our society is composed of every sort of tribe from the Branch Davidians to Wall Street conservatives.

What sort of generalization could possibly apply to such different groups? There must be some, but they're not going to be obvious. They might turn out to be something like the principles of control theory.

Best, Bill P.

Date: Wed May 11, 1994 9:44 am PST
Subject: RESTRING THE HARP

[Tom Baines] Dan Miller (940510):

>> Do these people trade notes, manuscripts, library cards?

> Jaques Elul and Thomas Kuhn would probably say they either do that or trade dreams.

OK. No one has a good clue as to how a socioeconomic system maintains stability. Ted Gurr, M. Skopol, and some others have offered pretty high level "models" of what that kind of stability looks like, and how to recognize that it has been destroyed, but there doesn't seem to be much help if you want to know how relative stability is reached in the first place. Stephen Walt does a pretty good job of outlining why alliances form, and Terrel Arnold has insight into why people support terroristic rebellion. A number of writers (Seabury, Codevilla, DuPuy, Van Crevald) have gone beyond the dated premises of Clausewitz (even as "modernized by Harry Summers) with some sound discussion of what must happen to restore balance when instability turns into conflict. Nobody seems to have a better way of defining and evaluating social instability, however, which would seem to be the necessary starting place for better answers.

So what exactly does PCT tell us about how two or more organisms move toward relative stability in relation to one another? What does Powers' premise that social systems really don't exist - that "Society is [merely] a perception" - really mean in terms of dealing with the needful behavior that leads to riots and war. Does the oriental view of socioeconomic activity as an analog of war inform us in any way?

The "reorganization" discussions all work on reorganization of the INDIVIDUAL'S core references. The only discussion in depth that I've seen about how that gets dealt with by pairs, triads, groups, etc. is that between Martin Taylor & Bill P. about layered protocols.

This is a great time for PCT to offer some new insights into how to analyze public policies with regard to social instability. Some one of you can win a Nobel prize if PCT can give us a better handle on this issue of how to even define social instability in more productive ways.

Tom Baines