How to do Economics

But first a relevant digression …

Prediction vs Explanation

Someone does ten experiments and comes up with a theory consistent with the results. The theory predicts the outcome of ten more experiments that have not been done yet. The second ten experiments are done and his predictions are correct. Someone else looks at the results of the experiments and creates a theory consistent with all twenty. We now do one more experiment, for which the two theories give different predictions. Which prediction should we trust?

One might think that, since all we know about either is that it is consistent with the first twenty experiments, they would be equally likely to be correct. You might imagine each possible theory written on a piece of paper and the pieces of paper sorted into barrels according to the results they predict for the first twenty experiments. The first theorist restricted himself to the barrels containing theories consistent with the first ten experiments, drew a theory from one of those barrels, and it happened to be from the barrel containing theories also consistent with the next ten experiments. The second experimenter went straight to that barrel and drew a theory from it. Each has randomly selected a theory from the same barrel.

What is wrong with this model is the assumption that experimenters are drawing theories at random. Suppose we assume instead, as I think much more plausible, that some people are better at coming up with correct theories than others. Only a small fraction of the first group of barrels contained theories consistent with the second ten experiments, so it would be very unlikely for the first experimenter to have chosen one of those barrels by chance. It is much more likely if he is someone unusually good at coming up with correct theories. We have no similar evidence for the second person, since he looked at the results of all twenty experiments before choosing a barrel.

Since we have more reason to believe that the first theorist is good at creating correct theories we have more reason to trust his prediction for the next experiment.

Statisticians may recognize the argument as a version of spurious contagion. Picking the right barrel does not make the theorist any better but the fact that he picked the right barrel increases the probability that he was (even before picking it) a good theorist, hence that the theory he came up with is correct.

Economic Methodology

Some time ago, I came across a talk by Roderick Long, a libertarian philosopher and blogger sympathetic with the Austrian approach to economics. In it he criticizes the argument in my father's essay "The Methodology of Positive Economics," which defends the use in economics of unrealistic models such as perfect competition, on the grounds that the test of a model is not its descriptive accuracy but its ability to make correct predictions.

Roderick, if I understand him correctly, imagines that all that is going on in the Chicago approach is blind curve fitting, looking for patterns in the observed data and assuming that those patterns will continue. A body of data can be fitted with an infinite number of different curves, a set of facts with an infinite number of theories, making the probability of picking the right one by chance low, so the theorist does not pick his theory at random. The test of whether he has done a good job of figuring out what simplified model includes the important factors and excludes the unimportant ones is the ability of the model to make correct predictions.

Crucial to this view of the process is the distinction between explaining facts you already know and predicting facts you do not know, a point emphasized in my father's essay but, I think, missed in Roderick's lecture. Explanation of known facts can be blind curve fitting. But if that is all it is, predictions of facts that did not go into constructing the model are unlikely to be correct.

Roderick offered an elaborate philosophical explanation[1] of why my father rejects what Roderick views as the correct approach to doing economics, the *a priori* approach associated with Ludwig Von Mises and some of his followers. There is a much simpler explanation. The problem with that approach, at least in its extreme version, is that pure *a priori* argument is unable to predict anything of economic interest. If one is completely agnostic about the facts, including both utility functions and production technologies, any physically possible pattern of human behavior is consistent with the theory. As I put it long ago in my Price Theory, explaining why the assumption of rationality is empty unless combined with some knowledge of what humans value:

> Why did I stand on my head on the table while holding a burning $1,000 bill between my toes? I wanted to stand on my head on the table while holding a burning $1,000 bill between my toes.

I conclude that the correct way of doing economics combines *a priori* theory with evidence. You form plausible conjectures on the basis of theory and evidence, where part of forming them is deciding what simplifications, what unrealistic features of the model, assume away inessential complications while retaining the essential features of what you are trying to understand. You find out how good a job you have done by using the conjectures to make predictions and testing them. An added benefit of that process, as I discovered in the course of revising what became my first published journal article in economics in response to an initial rejection, is that finding real world predictions of your model may force you to think through the model itself more clearly.

That is the Chicago School methodology as I understand and practice it.

I have sometimes challenged supporters of the extreme *a priori* approach to describe any real world prediction that can be derived from economic theory with no additional information. One response I sometimes get is that if you raise the minimum wage, all else held constant, employment of workers previously receiving the minimum wage must decline.

My rebuttal used to be to posit a situation where consumers do not like to buy goods produced by very low paid workers, have no easy way of knowing what wage each employer pays, and so are reluctant to buy goods that they know were produced by unskilled labor. After the minimum wage is raised consumers know that all goods are produced by workers making what they consider a tolerable wage. Their demand for goods made by unskilled labor increases, increasing employment for the workers who make such goods. Economic theory does not tell us what tastes consumers will have so cannot eliminate that possibility, however implausible it may seem.

That made the point that the conclusion depends on factual assumptions as well as theory, but I now have a more interesting example of circumstances where raising the minimum wage increases

[1] He explains it in more detail in a webbed text .

instead of decreasing employment for low skill workers.[2] Suppose the employers of unskilled labor are monopsonists, each the only employer of such labor in the relevant market. It pays each to hold down the number of workers he hires in order to hold wages down, just as a monopolist produces less in order to hold price up. Increasing the minimum wage to what it would be in a competitive labor market results in the monopsonist hiring more workers, since doing so no longer increases the wage he must pay.

My previous example showed that it was logically possible for the prediction to be false. This one shows that there are plausible real world circumstances — perhaps not in America today but in some past market societies — where it would be false. That was a fact that had not occurred to the people who offered that response to my challenge — or to me.

Nine to Five Economists

A very long time, I had a colleague and friend whom I concluded was only an economist in working hours. Having come across another example more recently, I thought it would be worth explaining the concept.

Suppose there is some issue on which economics implies a straightforward conclusion, but an implication short of a proof. Minimum wage laws provide a simple example. The conclusion that raising the minimum wage will reduce the employment of those currently receiving a minimum wage is a straightforward implication of the assumption that demand curves, in this case for a particular sort of labor, slope down. It is however possible, as I have just demonstrated, to construct a model of the labor market which yields the opposite conclusion.

The response of a real economist will either be "raising the minimum wage almost certainly reduces employment for low skilled workers" or, less likely but not impossible, "one of those odd models might possibly be right, I will look for a natural experiment by which I can test it."

The response of someone who is only an economist in working hours is to believe whatever he would believe if he was not an economist. If that requires him to believe that raising the minimum wage will not reduce employment for low skilled workers and someone points out the inconsistency with standard economic analysis, he will justify himself on the grounds that the economic argument is not actually a rigorous proof, so its conclusion could be false. For an example on the other side of the political fence, consider a conservative economist who wants to support trade restrictions.

In either of these cases, there may be a way for a real economist to support the policy conclusion he wants. If the supporter of a higher minimum wage is unwilling to claim, implausibly, that employers of low-skill labor are mostly monopsonists, he can argue that the loss of employment to some current minimum wage workers is more than balanced by higher wages to others, so that the total earnings of low skilled workers go up instead of down. If that is his argument he will want to look for evidence on the relevant elasticities. He may support his point by observing that losing a ten dollar an hour job is not really a net cost of ten dollars, since it increases leisure by an hour

[2] I got this point from a article by David Card and Alan Krueger, later discovered it had been made much earlier by George Stigler. The Card and Krueger article is mostly empirical — the theoretically interesting part starts on page 791.

for every ten dollars lost, but will be bothered by the possibility that losing a ten dollar an hour job removes the first step leading to a fifteen or twenty dollar an hour job.

In the trade case, the argument might be that although free trade produces net benefits for Americans for familiar economic reasons, the gains go to richer people than the losses, so a gain in value measured in dollars leads to a loss in value measured in utility. A real economist will want to make some effort to find out if the claim is true, to estimate the income distribution of gains and losses and their relative size. If he cannot find economic support for his position, he may agree in private that the effects of trade restrictions are negative but support them in public as a way of getting the rust belt votes needed to elect politicians who will do other things that, in his view, more than compensate for the loss.[3]

The simple test is whether an economist's views tend to diverge from those of his ideological allies when the ideology clashes with the economics. If they do he is a real economist. If they do not, he is only an economist in working hours.

I came across a different example of the same pattern accompanying my daughter on her college search. Wandering around the economics department of one of the colleges to get a feel for the place I spoke with three faculty members, none of whom appeared to be an economist in my sense of the term, someone for whom economics was part of how he viewed the world. My daughter, having audited an econ class, told me that a student had made a comment which an economist should have responded to with some version of "that sounds plausible but is wrong because." The professor simply let the comment go.

A long time ago I commented to a graduate student at Chicago that it seemed to me that there were a lot of economists who did not really believe in economics; it was what they did in working hours, not how they thought about the world. His response was that some of his fellow graduate students had noticed that when visiting other schools. I do not know enough about economics department to say which ones currently are in which category. Chicago, so far as I can tell, is still a place where some of the economists believe in economics.[4] One simple test would be to have lunch with members of the department, perhaps also with their graduate students, and see what they talk about. A few years ago I did the experiment, repeatedly, at George Mason. It had, at the time, three different campuses with parts of the economics department in them. The one I was visiting easily passed.

My point is not that people should believe in economics because it is true, although I think it largely is, only that economists should believe in economics. I am not inclined to take theology very seriously but, if I did take a course in it, I would expect to learn more and have a more interesting time if the professor was a believer than if he were an atheist.

Getting Disagreement Wrong

An article on my father by Steven Pearlstein starts with a puzzle. Most of the articles published on Milton Friedman after his death agree that he was a great economist; many compare him to John Maynard Keynes, another great economist. But Friedman and Keynes held different, indeed inconsistent, views; an important part of Friedman's accomplishment was to undo Keynes'

---

[3]A commenter on my blog pointed me at an example of behavior along those lines, two economists engaging in public demagoguery, one of them citing his professional qualifications while doing so.

[4] My point is made, specifically for Chicago, in a piece Austan Goolsbee published on my father's death.

accomplishment. If Keynes was right, how can Friedman be a great economist? If Friedman was right, how can Keynes be?

It is an interesting question but the author gets the answer wrong. He concludes that both were right. Keynes' version of economics was correct for the forties and fifties, Friedman's for the seventies and eighties when the Keynesian model "had played itself out."

That is a claim that neither Keynes nor Friedman would have taken seriously. Keynes titled his magnum opus "The General Theory of Employment, Interest and Money" not "The Theory of How Employment, Interest and Money Will Work from 1930 to 1960." Part of the work that earned Friedman his Nobel was *A Monetary History of the United States* (coauthored with Anna Schwartz), in which he demonstrated that the Keynesian analysis of the Great Depression, a centerpiece of the Keynesian view of economics, was based on a historically mistaken account of what actually happened. It is an odd view of science in which the historical facts about the 1930's changed between 1940 and 1970.

Pearlstein claims that economics is less of a science than physics, hence its truths more temporary. Yet the history of physics offers the same puzzle. Newton was a great physicist. Einstein was a great physicist. Part of Einstein's accomplishment was to show that Newtonian physics was, in certain fundamental ways, wrong.

Newton was wrong, wrong not only now but then, but Newtonian physics provided the foundation of ideas on which later generations of physicists, including Einstein, built. Keynes was wrong, but his attempt to make sense of what he believed happened during the Great Depression provided a theoretical foundation on which later theorists, including Friedman, could build. Hence Friedman's comment on Keynes: "In one sense, we are all Keynesians now; in another, no one is a Keynesian any longer"—misquoted by *Time Magazine* as "We are all Keynesians now."

Friedman and Keynes are not the only example in economics of such a pattern. David Ricardo was one of the most important figures in the history of economics, arguably its first great theorist. He starts his book by saying some admiring things about Adam Smith and then demonstrating that Smith's view of how prices are determined could not be true.

The Intensive Margin: Math vs Econ

I was recently told, by an undergraduate at a top school who had been considering majoring in economics and decided not to, that the required courses had turned out to contain a great deal more mathematics than economics. That report was confirmed by a senior faculty member at the same school with whom I raised the issue. He agreed with me that the situation was an unfortunate one.

The content of such courses presumably reflects what professors believe their students must learn in order to go to graduate school, publish articles in leading journals and have a successful career as academic economists. That fits my not very expert impression of the current state of academic economics, that it is heavy with what Gordon Tullock used to refer to as "ornamental mathematics," advanced tools used to demonstrate the author's mathematical sophistication but contributing little to the substance of the analysis.

I have not been much involved with the world of journal submissions for a long time — I prefer to write books and blog posts, although I occasionally get asked to referee articles — so am in a poor position to make blanket judgments. But some years back, reading an interesting article by

Akerlof and Yellin on why changes that should have reduced the number of children born to unmarried mothers had been accompanied instead by a sharp increase, I was struck by the fact that they had used game theory to make an argument that could have been presented equally well, perhaps more clearly, with supply and demand curves. Their analysis was simply an application of the theory of joint products — sexual pleasure and babies in a world without reliable contraception or readily available abortion. Add in those technologies, making the products no longer joint, and the outcome changes, making some women who want babies unable to find husbands to help support them. The argument uses nothing Marshall did not know, arguably nothing Smith, who discusses wool and mutton as joint products, did not know.

Assume, for the moment, that I am right that both economics in the journals and economics in the classroom emphasize mathematics well past the point where it no longer contributes much to the economics. Why?

The answer, I suspect, takes us back to Ricardo's distinction between the intensive and extensive margins of cultivation. Expanding production on the intensive margin means getting more grain out of land already cultivated, expanding it on the extensive margin means getting more grain by bringing new land into cultivation.

In economics, the intensive margin means writing new articles — new enough, at least, to get published — on subjects that smart people have been writing articles about for most of the past century. An example would be the question of why involuntary unemployment sometimes exists and what can be done about it. That is an important question, sufficiently important so that many non-economists seem to consider dealing with it the chief business of economists. But it is also a question which quite a lot of very good economists have been working on for a long time, which makes it difficult to say anything both new and interesting about it.

Another example would be game theory. I like to defend my disinterest in being more than an observer of that field by explaining that, when looking for problems to work on, problems that stumped John Von Neumann, one of the most brilliant thinkers of the Twentieth Century, go at the bottom of my pile.[5]

One consequence of the difficulty of such questions is that anything new is likely to be either uninteresting or wrong. That is an implication of what I referred to in another chapter as the rising marginal cost of originality, a principle I usually illustrate with examples from city planning and architecture. My favorite example of the latter is, for those familiar with it, the Coombs building at Australian National University, a truly inspired piece of bad design in which I once spent part of a summer.

One solution to the problem, assuming you don't have any new and interesting economic ideas on the subject, is to apply a new mathematical tool to an old problem. It has not been done before, that tool not having existed before, so with luck you can get published, whether or not the new tool adds anything useful to analysis of the problem.

The extensive margin, in contrast, is the application of the existing tools of economics, including mathematics where needed, to new subjects. Examples include public choice theory, law and economics, and behavioral economics. A recent example I am fond of is the work Peter Leeson

[5] Eugene Wigner, himself a brilliant man, is reputed to have said "There are two kinds of people in the world: Johnny Von Neumann and everybody else."

has done on applying economics to making sense of 18th century piracy; curious readers will find it in his book *The Invisible Hook*. Because nobody, so far as I know, had thought of doing it before, Leeson was able to produce interesting results by applying conventional economic analysis to a subject that specialist historians had researched but economists knew very little about, to the benefit of both fields.

I have considerable disagreements with Robert Frank, some exposed in exchanges between us in Chapter XXX. But when, in *Choosing the Right Pond*, he showed how the fact that relative as well as absolute outcomes matter to people could be incorporated into conventional price theory he was working new ground and, in the process, teaching the rest of us something interesting.

My conclusion is that, if you want to do interesting economics, your best bet is probably to work on the extensive margin — better yet, if sufficiently clever and lucky, to extend it.

On the Use of Mathematics in Economics: A Letter

Balliol Croft, Cambridge
27. ii. 06

My dear Bowley,

I have not been able to lay my hands on any notes as to Mathematico-economics that would be of any use to you: and I have very indistinct memories of what I used to think on the subject. I never read mathematics now: in fact I have forgotten even how to integrate a good many things.

But I know I had a growing feeling in the later years of my work at the subject that a good mathematical theorem dealing with economic hypotheses was very unlikely to be good economics: and I went more and more on the rules—(1) Use mathematics as a shorthand language, rather than as an engine of inquiry. (2) Keep to them till you have done. (3) Translate into English. (4) Then illustrate by examples that are important in real life. (5) Burn the mathematics. (6) If you can't succeed in 4, burn 3. This last I did often.

I believe in Newton's Principia Methods, because they carry so much of the ordinary mind with them. Mathematics used in a Fellowship thesis by a man who is not a mathematician by nature—and I have come across a good deal of that—seems to be an unmixed evil. And I think you should do all you can to prevent people from using Mathematics in cases in which the English Language is as short as the Mathematical.....

Alfred Marshall

Which leaves me wondering how much of the economics of the next century went into Marshall's fireplace.