

HOW I WORK

My formal charge in this essay is to talk about my "life philosophy". Let me make it clear at the outset that I have no intention of following instructions, since I don't know anything special about life in general. I believe it was Schumpeter who claimed to be not only the best economist, but also the best horseman and the best lover in his native Austria. I don't ride horses, and have few illusions on other scores. (I am, however, a pretty good cook).

What I want to talk about in this essay is something more restricted: some thoughts about thinking, and particularly how to go about doing interesting economics. I think that among economists of my generation I can claim to have a fairly distinctive intellectual style -- not necessarily a better style than my colleagues, for there are many ways to be a good economist, but one that has served me well. The essence of that style is a general research strategy that can be summarized in a few rules; I also view my more policy-oriented writing and speaking as ultimately grounded in the same principles. I'll get to my rules for research later in this essay. I think I can best introduce those rules, however, by describing how (it seems to me) I stumbled into the way I work.

ORIGINS

Most young economists today enter the field from the technical end. Originally intending a career in hard science or engineering, they slip down the scale into the most rigorous of the social sciences. The advantages of entering economics from that direction are obvious: one arrives already well trained in mathematics, one finds the concept of formal modeling natural. It is not, however, where I come from. My first love was history; I studied little math, picking up what I needed as I went along.

Nonetheless, I got deeply involved in economics early, working as a research assistant (on world energy markets) to William Nordhaus while still only a junior at Yale. Graduate school followed naturally, and I wrote my first really successful paper -- a theoretical analysis of balance of payments crises -- while still at MIT. I discovered that I was facile with small mathematical models, with a knack for finding simplifying assumptions that made them tractable. Still, when I left graduate school I was, in my own mind at least, somewhat directionless. I was not sure what to work on; I was not even sure whether I really liked research.

I found my intellectual feet quite suddenly, in January 1978. Feeling somewhat lost, I paid a visit to my old advisor Rudi Dornbusch. I described several ideas to him, including a vague notion that the monopolistic competition models I had studied in a short course offered by Bob Solow -- especially the lovely little model of Dixit and Stiglitz -- might have something to do with international trade. Rudi flagged that idea as potentially very interesting indeed; I went home to work on it seriously; and within a few days I realized that I had hold of something that would form the core of my professional life.

What had I found? The point of my trade models was not particularly startling once one thought about it: economies of scale could be an independent cause of international trade, even in the absence of comparative advantage. This was a new insight to me, but had (as I soon discovered) been pointed out many times before by critics of conventional trade theory. The models I worked out left some loose ends hanging; in particular, they typically had many equilibria. Even so, to make the models tractable I had to make obviously unrealistic assumptions. And once I had made those assumptions, the models were trivially simple; writing them up left me no opportunity to display any high-powered technique. So one might have concluded that I was doing nothing very interesting (and that was what some of my colleagues were to tell me over the next few years). Yet what I saw -- and for some reason saw almost immediately -- was that all of these features were virtues, not vices, that they added up to a program that could lead to years of productive research.

I was, of course, only saying something that critics of conventional theory had been saying for decades. Yet my point was not part of the mainstream of international economics. Why? Because it had never been expressed in nice models. The new monopolistic competition models gave me a tool to open cleanly what had previously been regarded as a can of worms. More important, however, I suddenly realized the remarkable extent to which the methodology of economics creates blind spots. We just don't see what we can't formalize. And the biggest blind spot of all has involved increasing returns. So there, right at hand, was my mission: to look at things from a slightly different angle, and in so doing to reveal the obvious, things that had been right under our noses all the time.

The models I wrote down that winter and spring were incomplete, if one demanded of them that they specify exactly who produced what. And yet they told meaningful stories. It took me a long time to express clearly what I was doing, but eventually I realized that one way to deal with a difficult problem is to change the question -- in particular by shifting levels. A detailed analysis may be extremely nasty, yet an aggregative or systemic description that is far easier may tell you all you need to know.

To get this system or aggregate level description required, of course, accepting the basically silly assumptions of symmetry that underlay the Dixit-Stiglitz and related models. Yet these silly assumptions seemed to let me tell stories that were persuasive, and that could not be told using the hallowed assumptions of the standard competitive model. What I began to realize was that in economics we are always making silly assumptions; it's just that some of them have been made so often that they come to seem natural. And so one should not reject a model as silly until one sees where its assumptions lead.

Finally, the simplicity of the models may have frustrated my lingering urge to show off the technical skills I had so laboriously acquired in graduate school, but was, I soon realized, central to the enterprise. Trade theorists had failed to address the role of increasing returns, not out of empirical conviction, but because they thought it was too hard to model. How much more effective, then, to show that it could be almost childishly simple?

And so, before my 25th birthday, I basically knew what I was going to do with my professional life. I don't know what would have happened if my grand project had met with rejection from other economists -- perhaps I would have turned cranky, perhaps I would have lost faith and abandoned the effort. But in fact all went astonishingly well. In my own mind, the curve of my core research since that January of 1978 has followed a remarkably consistent path. Within a few months, I had written up a basic monopolistic competition trade model -- as it turned out, simultaneously and independently with similar models by Avinash Dixit and Victor Norman, on one side, and Kelvin Lancaster, on the other. I had some trouble getting that paper published -- receiving the dismissive rejection by a flagship journal (the QJE) that seems to be the fate of every innovation in economics -- but pressed on. From 1978 to roughly the end of 1984 I focussed virtually all my research energies on the role of increasing returns and imperfect competition in international trade. (I took one year off to work in the US government; but more about that below). What had been a personal quest turned into a movement, as others followed the same path. Above all, Elhanan Helpman -- a deep thinker whose integrity and self-discipline were useful counterparts to my own flakiness and disorganization -- first made crucial contributions himself, then talked me into collaborative work. Our magnum opus, *Market Structure and Foreign Trade*, served the purpose of making our ideas not only respectable but almost standard: iconoclasm to orthodoxy in seven years.

For whatever reason, I allowed my grand project on increasing returns to lie fallow for a few years in the 1980s, and turned my attention to international finance. My work in this area consisted primarily of small models inspired by current policy issues; although these models lacked the integrating theme of my trade models, I think that my finance work is to some extent unified by its intellectual style, which is very similar to that of my work on trade.

In 1990 I returned to the economics of increasing returns from a new direction. I suddenly realized that the techniques that had allowed us to legitimize the role of increasing returns in trade could also be used to reclaim a whole outcast field: that of economic geography, the location of activity in space. Here, perhaps even more than in trade, was a field full of empirical insights, good stories, and obvious practical importance, lying neglected right under our noses because nobody had seen a good way to formalize it. For me, it was like reliving the best moments of my intellectual childhood. Doing geography is hard work; it requires a lot of hard thinking to make the models look trivial, and I am increasingly finding that I need the computer as an aid not just to data analysis but even to theorizing. Yet it is immensely rewarding. For me, the biggest thrill in theory is the moment when your model tells you something that should have been obvious all along, something that you can immediately relate to what you know about the world, and yet which you didn't really appreciate. Geography still has that thrill.

My work on geography seems, at the time of writing, to be leading me even further afield. In particular, there are obvious affinities between the concepts that arise naturally in geographic models and the language of traditional development economics -- the "high development theory" that flourished in the 1940s and 50s, then collapsed. So I expect that my basic research project will continue to widen in scope.

RULES FOR RESEARCH

In the course of describing my formative moment in 1978, I have already implicitly given my four basic rules for research. Let me now state them explicitly, then explain. Here are the rules:

1. Listen to the Gentiles
2. Question the question
3. Dare to be silly
4. Simplify, simplify

Listen to the Gentiles

What I mean by this rule is "Pay attention to what intelligent people are saying, even if they do not have your customs or speak your analytical language." The point may perhaps best be explained by example. When I began my rethinking of international trade, there was already a sizeable literature criticizing conventional trade theory. Empiricists pointed out that trade took place largely between countries with seemingly similar factor endowments, and that much of this trade involved intra-industry exchanges of seemingly similar products. Acute observers pointed to the importance of economies of scale and imperfect competition in actual international markets. Yet all of this intelligent commentary was ignored by mainstream trade theorists -- after all, their critics often seemed to have an imperfect understanding of comparative advantage, and had no coherent models of their own to offer; so why pay attention to them? The result was that the profession overlooked evidence and stories that were right under its nose.

The same story is repeated in geography. Geographers and regional scientists have amassed a great deal of evidence on the nature and importance of localized external economies, and organized that evidence intelligently if not rigorously. Yet economists have ignored what they had to say, because it comes from people speaking the wrong language.

I do not mean to say that formal economic analysis is worthless, and that anybody's opinion on economic matters is as good as anyone else's. On the contrary! I am a strong believer in the importance of models, which are to our minds what spear-throwers were to stone age arms: they greatly extend the power and range of our insight. In particular, I have no sympathy for those people who criticize the unrealistic simplifications of model-builders, and imagine that they achieve greater sophistication by avoiding stating their assumptions clearly. The point is to realize that economic models are metaphors, not truth. By all means express your thoughts in models, as pretty as possible (more on that below). But always remember that you may have gotten the metaphor wrong, and that someone else with a different metaphor may be seeing something that you are missing.

Question the question

There was a limited literature on external economies and international trade before 1978. It was never, however, very influential, because it seemed terminally messy; even the simplest models became bogged down in a taxonomy of possible outcomes. What has since become clear is that this messiness arose in large part because the modelers were asking their models to do what traditional trade models do, which is to predict a precise pattern of specialization and trade. Yet why ask that particular question? Even in the Heckscher-Ohlin model, the point you want to make is something like "A country tends to export goods whose production is intensive in the factors in which that country is abundant"; if your specific model tells you that capital-abundant country Home exports capital-intensive good X, this is valuable because it sharpens your understanding of that insight, not because you really care about these particular details of a patently oversimplified model.

It turns out that if you don't ask for the kind of detail that you get in the two-sector, two-good classical model, an external economy model needn't be at all messy. As long as you ask "system" questions like how welfare and world income are distributed, it is possible to make very simple and neat models. And it's really these system questions that we are interested in. The focus on excessive detail was, to put it bluntly, a matter of

carrying over ingrained prejudices from an overworked model into a domain where they only made life harder.

The same is true in a number of areas in which I have worked. In general, if people in a field have bogged down on questions that seem very hard, it is a good idea to ask whether they are really working on the right questions. Often some other question is not only easier to answer but actually more interesting! (One drawback of this trick is that it often gets people angry. An academic who has spent years on a hard problem is rarely grateful when you suggest that his field can be revived by bypassing it).

Dare to be silly

If you want to publish a paper in economic theory, there is a safe approach: make a conceptually minor but mathematically difficult extension to some familiar model. Because the basic assumptions of the model are already familiar, people will not regard them as strange; because you have done something technically difficult, you will be respected for your demonstration of firepower. Unfortunately, you will not have added much to human knowledge.

What I found myself doing in the new trade theory was pretty much the opposite. I found myself using assumptions that were unfamiliar, and doing very simple things with them. Doing this requires a lot of self-confidence, because initially people (especially referees) are almost certain not simply to criticize your work but to ridicule it. After all, your assumptions will surely look peculiar: a continuum of goods all with identical production functions, entering symmetrically into utility? Countries of identical economic size, with mirror-image factor endowments? Why, people will ask, should they be interested in a model with such silly assumptions -- especially when there are evidently much smarter young people who demonstrate their quality by solving hard problems?

What seems terribly hard for many economists to accept is that all our models involve silly assumptions. Given what we know about cognitive psychology, utility maximization is a ludicrous concept; equilibrium pretty foolish outside of financial markets; perfect competition a howler for most industries. The reason for making these assumptions is not that they are reasonable but that they seem to help us produce models that are helpful metaphors for things that we think happen in the real world.

Consider the example which some economists seem to think is not simply a useful model but revealed divine truth: the Arrow-Debreu model of perfect competition with utility maximization and complete markets. This is indeed a wonderful model -- not because its assumptions are remotely plausible but because it helps us think more clearly about both the nature of economic efficiency and the prospects for achieving efficiency under a market system. It is actually a piece of inspired, marvellous silliness.

What I believe is that the age of creative silliness is not past. Virtue, as an economic theorist, does not consist in squeezing the last drop of blood out of assumptions that have come to seem natural because they have been used in a few hundred earlier papers. If a new set of assumptions seems to yield a valuable set of insights, then never mind if they seem strange.

Simplify, simplify

The injunction to dare to be silly is not a license to be undisciplined. In fact, doing really innovative theory requires much more intellectual discipline than working in a well-established literature. What is really hard is to stay on course: since the terrain is unfamiliar, it is all too easy to find yourself going around in circles. Somewhere or other Keynes wrote that "it is astonishing what foolish things a man thinking alone can come temporarily to believe". And it is also crucial to express your ideas in a way that other people, who have not spent the last few years wrestling with your problems and are not eager to spend the next few years wrestling with your answers, can understand without too much effort.

Fortunately, there is a strategy that does double duty: it both helps you keep control of your own insights, and makes those insights accessible to others. The strategy is: always try to express your ideas in the simplest possible model. The act of stripping down to this minimalist model will force you to get to the essence of what you are trying to say (and will also make obvious to you those situations in which you actually have nothing to say). And this minimalist model will then be easy to explain to other economists as well.

I have used the "minimum necessary model" approach over and over again: using a one-factor, one-industry model to explain the basic role of monopolistic competition in trade; assuming sector-specific labor rather than full Heckscher-Ohlin factor substitution to explain the effects of intraindustry trade; working with symmetric countries to assess the role of reciprocal dumping; and so on. In each case the effect has been to allow me to tackle a subject widely viewed as formidably difficult with what appears, at first sight, to be ridiculous simplicity.

The downside of this strategy is, of course, that many of your colleagues will tend to assume that an insight that can be expressed in a cute little model must be trivial and obvious -- it takes some sophistication to realize that simplicity may be the result of years of hard thinking. I have heard the story that when Joseph Stiglitz was being considered for tenure at Yale, one of his senior colleagues belittled his work, saying that it consisted mostly of little models rather than deep theorems. Another colleague then asked, "But couldn't you say the same about Paul Samuelson"? "Yes, I could", replied Joe's opponent. I have heard the same reaction to my own work. Luckily, there are enough sophisticated economists around that in the end intellectual justice is usually served. And there is a special delight in managing not only to boldly go where no economist has gone before, but to do so in a way that seems after the fact to be almost child's play.

I have now described my basic rules for research. I have illustrated them with my experience in developing the "new trade theory" and with my more recent extension of that work to economic geography, because these are the core of my work. But I have also done quite a lot of other stuff, which (it seems to me) is also in some sense part of the same enterprise. So in the remainder of this essay I want to talk about this other work, and in particular about how the policy economist and the analytical economist can coexist in the same person.

POLICY-RELEVANT WORK

Most economic theorists keep their hands off current policy issues -- or if they do get involved in policy debates, do so only after the midpoint of their career, as something that follows creative theorizing rather than coexists with it. There seems to be a consensus that the clarity and singleness of purpose required to do good theory are incompatible with the tolerance for messy issues required to be active in policy discussion. For me, however, it has never worked that way. I have interspersed my academic career with a number of consulting ventures for various governments and public agencies, as well as a full year in the US government. I have also written a book, *The Age of Diminished Expectations*, aimed at a non-technical audience. And I have written a pretty steady stream of papers that are motivated not by the inner logic of my research but by the attempt to make sense of some currently topical policy debate -- e.g., Third World debt relief, target zones for exchange rates, the rise of regional trading blocs. All of this hasn't seemed to hurt my research, and indeed some of my favorite papers have grown out of this policy-oriented work.

Why doesn't policy-relevant work seem to conflict with my "real" research? I think that it's because I have been able to approach policy issues using almost exactly the same method that I use in my more basic work. Paying attention to newspaper reports or the concerns of central bankers and finance ministers is just another form of listening to the Gentiles. Trying to find a useful way of defining their problems is pretty much the same as questioning the question in theory. Confronting supposedly knowledgeable people with an unorthodox view of an issue certainly requires the courage to be silly. And of course, ruthless simplification is worth even more in policy discussion than in theory for its own sake.

So doing policy-relevant economics does not, for me, mean a drastic change in intellectual style. And it has its own payoffs. Let's be honest and admit that these include invitations to fancier conferences and speaking engagements at much higher fees than an academic purist is likely to get. Let's also admit that one of the joys of policy research is the opportunity to shock the bourgeoisie, to point out the hollowness or silliness of official positions. For example, I know that I was not the only international economist to have some fun pointing out the absurdities of the Maastricht Treaty, and was not above some wicked pleasure when the ERM crisis I and others had long predicted actually came to pass in the fall of 1992. The main payoff to policy work, though, is intellectual stimulation. Not all real-world questions are interesting -- I find that almost anything having to do with taxation is better than a sleeping pill -- but every couple of years, if not more often, the international economy throws up a question that gives rise to exciting research. I have been stimulated to write theory papers by the Plaza and the Louvre, by the Brady Plan, NAFTA, and EMU. All of them are papers that I think could stand on their own, even without the policy context.

There is, of course, always a risk that an economist who gets onto the policy circuit will no longer have enough time for real research. I certainly write an awfully large number of conference papers; I am a very fast writer, but perhaps it is a gift I overuse. Still, I think that the big danger of doing policy research is not so much the drain on your time as the threat to your values. It is easy to be seduced into the belief that direct influence on policy is more important than just writing papers -- I've seen it happen to many colleagues. Once you start down that road, once you begin to think that David Mulford matters more than Bob Solow, or to prefer hobnobbing with the Ruritanian finance minister to talking theory with Avinash Dixit, you are probably lost to research. Pretty soon you'll probably start using "impact" as a verb.

Fortunately, while I love playing around with policy issues, I have never been able to take policy *makers* very seriously. This lack of seriousness gets me into occasional trouble -- like the time that a gentle parenthetical joke about the French in a conference paper led to an extended diatribe from the French official attending the conference -- and may exclude me from ever holding any important policy position. But that's OK: in the end, I would rather write a few more good papers than hold a position of real power. (Note to the policy world: this doesn't mean that I would necessarily turn down such a position if it were offered!)

REGRETS

There are a lot of things about my life and personality that I regret -- if things have gone astonishingly well for me professionally, they have been by no means as easy or happy elsewhere. But in this essay I only want to talk about professional regrets.

A minor regret is that I have never engaged in really serious empirical work. It's not that I dislike facts or real numbers. Indeed, I find light empirical work in the form of tables, charts, and perhaps a few regressions quite congenial. But the serious business of building and thoroughly analyzing a data set is something I never seem to get around to. I think that this is partly because many of my ideas do not easily lend themselves to standard econometric testing. Mostly, though, it is because I lack the patience and organizational ability. Every year I promise to try to do some real empirical work. Next year I really will!

A more important regret is that while the MIT course evaluations rate me as a pretty good lecturer, I have not yet succeeded in generating a string of really fine students, the kind who reflect glory on their teacher. I can make excuses for this failing -- students often prefer advisers who are more methodical and less intuitive, and I all too often scare students off by demanding that they use less math and more economics. It's also true that I probably seem busy and distracted, and perhaps I am just not imposing enough in person to be inspiring (if I were only a few inches taller ...). Whatever the reasons, I wish I could do better, and intend to try.

All in all, though, I've been very lucky. A lot of that luck has to do with the accidents that led me to stumble onto an intellectual style that has served me extremely well. I've tried, in this essay, to define and explain that style. Is this a life philosophy? Of course not. I'm not even sure that it is an economic research philosophy, since what works for one economist may not work for another. But it's how I do research, and it works for me.