

WHAT ECONOMISTS CAN LEARN FROM EVOLUTIONARY THEORISTS

(A talk given to the European Association for Evolutionary Political Economy)

Paul Krugman

Nov. 1996

Good morning. I am both honored and a bit nervous to be speaking to a group devoted to the idea of evolutionary political economy. As you probably know, I am not exactly an evolutionary economist. I like to think that I am more open-minded about alternative approaches to economics than most, but I am basically a maximization-and-equilibrium kind of guy. Indeed, I am quite fanatical about defending the relevance of standard economic models in many situations.

Why, then, am I here? Well, partly because my research work has taken me to some of the edges of the neoclassical paradigm. When you are concerned, as I have been, with situations in which increasing returns are crucial, you must drop the assumption of perfect competition; you are also forced to abandon the belief that market outcomes are necessarily optimal, or indeed that the market can be said to maximize anything. You can still believe in maximizing individuals and some kind of equilibrium, but the complexity of the situations in which your imaginary agents find themselves often obliges you - and presumably them - to represent their behavior by some kind of ad hoc rule rather than as the outcome of a carefully specified maximum problem. And you are often driven by sheer force of modeling necessity to think of the economy as having at least vaguely "evolutionary" dynamics, in which initial conditions and accidents along the way may determine where you end up. Some of you may have read my work on economic geography; I only found out after I had worked on the models for some time that I was using "replicator dynamics" to discuss the problem of economic change.

But there is another reason I am here. I am an economist, but I am also what we might call an evolution groupie. That is, I spend a great deal of time reading what evolutionary biologists write - not only the more popular volumes but the textbooks and, most recently, some of the professional articles. I have even tried to talk to some of the biologists, which in this age of narrow specialization is a major effort. My interest in evolution is partly a recreation; but it is also true that I find in evolutionary biology a useful vantage point from which to view my own specialty in a new perspective. In a way, the point is that both the parallels and the differences between economics and evolutionary biology help me at least to understand what I am doing when I do economics - to get, to be pompous about it, a new perspective on the epistemology of the two fields.

I am sure that I am not unique either in my interest in biology or in my feeling that we economists have something to learn from it. Indeed, I am sure that many people in this room know far more about evolutionary theory than I do. But I may have

one special distinction. Most economists who try to apply evolutionary concepts start from some deep dissatisfaction with economics as it is. I won't say that I am entirely happy with the state of economics. But let us be honest: I have done very well within the world of conventional economics. I have pushed the envelope, but not broken it, and have received very widespread acceptance for my ideas. What this means is that I may have more sympathy for standard economics than most of you. My criticisms are those of someone who loves the field and has seen that affection repaid. I don't know if that makes me morally better or worse than someone who criticizes from outside, but anyway it makes me different.

Anyway, enough preliminaries.

Sister fields

If you are familiar with economics and start reading evolutionary biology in earnest - and presumably vice versa - you quickly realize that these are sister fields. They actually have a remarkable amount in common, not only in terms of the kind of questions they ask and the methods they use, but in terms of the way they relate to and are perceived by the rest of the world.

To begin with, there is the similarity in the basic approach. Let me give you my own personal definition of the basic method of economic theory. To me, it seems that what we know as economics is the study of those phenomena that can be understood as emerging from the *interactions among intelligent, self-interested individuals*. Notice that there are really four parts to this definition. Let's read from right to left.

1. Economics is about what *individuals* do: not classes, not "correlations of forces", but individual actors. This is not to deny the relevance of higher levels of analysis, but they must be grounded in individual behavior. Methodological individualism is of the essence.
2. The individuals are *self-interested*. There is nothing in economics that inherently prevents us from allowing people to derive satisfaction from others' consumption, but the predictive power of economic theory comes from the presumption that normally people care about themselves.
3. The individuals are *intelligent*: obvious opportunities for gain are not neglected. Hundred-dollar bills do not lie unattended in the street for very long.
4. We are concerned with the *interaction* of such individuals: Most interesting economic theory, from supply and demand on, is about "invisible hand" processes in which the collective outcome is not what individuals intended.

OK, that's what economics is about. What is evolutionary theory about?

The answer, basically, is that evolutionists share three of the four concerns. Their field is about the interaction of self-interested individuals - who are often thought

of as organisms "trying" to leave as many offspring as possible, but which are in some circumstances best thought of as genes "trying" to propagate as many copies of themselves as possible. The main difference between evolutionary theory and economics is that while economists routinely suppose that the agents in their models are very smart about finding the best strategy - and an economist is always defensive about any model in which agents are assumed to act with less than perfect rationality - evolutionists have no qualms about assuming myopic behavior. Indeed, myopia is of the essence of their view.

I'll talk later about what difference this makes. My point right now is that because the basic methods are similar if not identical, economics and evolutionary theory are surprisingly similar. It is often asserted that economic theory draws its inspiration from physics, and that it should become more like biology. If that's what you think, you should do two things. First, read a text on evolutionary theory, like John Maynard Smith's *Evolutionary Genetics*. You will be startled at how much it looks like a textbook on microeconomics. Second, try to explain a simple economic concept, like supply and demand, to a physicist. You will discover that our whole style of thinking, of building up aggregative stories from individual decisions, is not at all the way they think.

So there is a close affinity in method and indeed of intellectual style between economics and evolution. But there is another interesting parallel: both economics and evolution are model-oriented, algebra-heavy subjects that are the subject of intense interest from people who cannot stand algebra. And as a result in each case it is very important to distinguish between the field as it is perceived by outsiders (and portrayed in popular books) and what it is really like. We all know that economics is a field in which the most famous authors are often people who are regarded, with good reason, as not even worth arguing with by almost everyone in the profession. Do you remember that global best-seller *The Coming Great Depression of 1990* by Ravi Batra? And I guess it is no secret that even John Kenneth Galbraith, still the public's idea of a great economist, looks to most serious economists like an intellectual dilettante who lacks the patience for hard thinking. Well, the same is true in evolution.

I am not sure how well this is known. I have tried, in preparation for this talk, to read some evolutionary economics, and was particularly curious about what biologists people reference. What I encountered were quite a few references to Stephen Jay Gould, hardly any to other evolutionary theorists. Now it is not very hard to find out, if you spend a little while reading in evolution, that Gould is the John Kenneth Galbraith of his subject. That is, he is a wonderful writer who is beloved by literary intellectuals and lionized by the media because he does not use algebra or difficult jargon. Unfortunately, it appears that he avoids these sins not because he has transcended his colleagues but because he does not seem to understand what they have to say; and his own descriptions of what the field is about - not just the answers, but even the questions - are consistently misleading.

His impressive literary and historical erudition makes his work seem profound to most readers, but informed readers eventually conclude that there's no there there. (And yes, there is some resentment of his fame: in the field the unjustly famous theory of "punctuated equilibrium", in which Gould and Niles Eldredge asserted that evolution proceeds not steadily but in short bursts of rapid change, is known as "evolution by jerks").

What is rare in the evolutionary economics literature, at least as far as I can tell, is references to the theorists the practitioners themselves regard as great men - to people like George Williams, William Hamilton, or John Maynard Smith. This is serious, because if you think that Gould's ideas represent the cutting edge of evolutionary theory (as I myself did until about a year and a half ago), you have an almost completely misguided view of where the field is and even of what the issues are.

This is important, because it is at least my impression that what economists who like to use "evolutionary" concepts expect from evolution is often based on what they imagine evolutionary theory to be like, not on what it is actually like. And conversely, you learn a lot about why conventional economics looks the way it does by seeing how evolutionary theorists have been driven to some of the same positions.

To explain these rather cryptic remarks, let me talk briefly about what - it seems to me, but I am happy to be corrected - economists think an evolutionary approach can give us, then about what evolutionists seem to be saying in practice.

What evolutionary economists want

I don't think that there are many economists, even among the unconventionally minded, who would quarrel seriously with my basic definition of economics as concerning the interactions among intelligent, self-interested individuals. I guess a Marxist would have problems with the whole idea of methodological individualism, and a Galbraithian would have problems with the idea that self-interest can be defined without taking into account the ability of advertisers and so forth to shape preferences. But such quarrels apart I would guess that we do not have much difference with the basic statement.

Where the dissatisfaction sets in is with how we implement the first two terms in my four-part program. Yes, of course economics is about interaction, and the agents are intelligent; but exactly how intelligent are they, and what is the nature of the their interaction?

For there is no question that conventional economics has gone beyond the general ideas of intelligence and interaction to a much harder-edged, extreme formulation. At least since Paul Samuelson published *Foundations of Economic Analysis* in

1947, the overwhelming thrust of conventional theory has been to say that agents are not only intelligent, they *maximize* - that is, they chose the best of all feasible alternatives. And when they interact, we assume that what they do is achieve an *equilibrium*, in which each individual is doing the best he can given what all the others are doing.

Now anyone who looks at the world knows that these are extreme and unrealistic assumptions. I just had some work done on my house; it is painfully obvious, looking at the final bill, that I did not maximize - I did not engage in optimal search for a contractor. In trying to find someone to do the remaining work, I have discovered that local wages and prices have not caught up with the economic boom in Massachusetts, so that it is extremely hard to find anyone to do carpentry or plumbing - the market is definitely not in equilibrium. So can't we get away from the maximization-and-equilibrium approach to something more realistic?

Well, as I understand it that is what evolutionary economics is all about. In particular, evolution-minded economists seem to want the following:

1. They want to get away from the idea that individuals maximize. Instead, they want to represent decisions as the result of some process of groping through alternatives, a process in which it may take a long time to get to a maximum - and in which the maximum you find may well be local rather than global.
2. They want to get away from the notion of equilibrium. In particular, they want to have an approach in which things are always in disequilibrium, in which the economy is always evolving. Latterly there have also been some economists who want to merge evolutionary ideas with the Schumpeterian notion that the economy proceeds via waves of "creative destruction".

Now as I understand it evolutionary economists basically believe that an evolutionary approach will satisfy these desires. After all, real organisms often look to the discerning eye like works in progress - they are full of features that fall short of what would adapt them perfectly to their environment, that is, they have not really maximized their fitness. And they also often seem to be stuck on local maxima: dolphins may look like fish, but they still need to surface for air. Meanwhile, what is evolution but a process of continual change, which has taken us from microbes to man? And if you are a reader of Gould and his acolytes, you have the sense that evolution proceeds through spasms of sudden change that seem positively Schumpeterian in their drama.

So the attractiveness of an evolutionary metaphor - especially if you believe that economics has gotten off on the wrong track by basing itself on physics - is understandable. But before we get too carried away with the prospects for an evolutionary revolution, we had better look at what the evolutionists themselves really do.

What evolutionary theory is really like

To read the real thing in evolution - to read, say, John Maynard Smith's *Evolution and the Theory of Games*, or William Hamilton's new book of collected papers, *Narrow Roads in Gene Land*, is a startling experience to someone whose previous idea of evolution comes from magazine articles and popular books. The field does not look at all like the stories. What it does look like, to a remarkable degree, is - dare I say it? - neoclassical economics. And it offers very little comfort to those who want a refuge from the harsh discipline of maximization and equilibrium.

Consider first the question of maximization. Clearly it is a crucial point about evolution that it must proceed by small steps, which means that maxima must be approached only gradually and that you could easily be trapped on a merely local maximum. But do these observations actually play a large role in evolutionary theory? No, not really.

Look, for example, at William Hamilton's deeply influential paper "The genetical basis of social behavior". In the first part of that paper he introduces a model of population dynamics and shows that a gene will tend to spread if it enhances not an organism's individual fitness, but its "inclusive fitness": a weighted sum of the fitness of all the individual's relatives, with the weights proportional to their closeness of relationship. (An alternative way to think of this is to think of the gene spreading if it is good for its own fitness, never mind the organisms it lives in; this is the theme of Richard Dawkins's book *The Selfish Gene*). Now Hamilton's derivation concerns process - it is a dynamic story about which direction the next small step will proceed in. But when it comes to the second part, in which he uses the idea to discuss the real world - why birds expose themselves to predators by warning their neighbors, why insects have such massively organized societies - he simply assumes that what we actually see can be viewed as the culmination of that process, that the creatures we see have already maximized. In short, even though evolution is necessarily a process of small changes, evolutionary theorists normally take the shortcut of assuming that the process gets you to the maximum, and pay surprisingly little attention to the dynamics along the way.

What about the possibility of being trapped on local maxima? Well, this is a big concern for some theorists, like the Santa Fe Institute's Stuart Kauffman - but Kauffman is not a central player in the field. The general attitude of evolutionary theorists seems to be that Nature can often find surprising pathways to places you would have thought unreachable by small steps; that over a few hundred thousand generations a slightly light-sensitive patch of skin can become an eye that appears to be perfectly designed, or a jaw-bone can migrate around and become a piece of acquisitively sensitive sound-detection equipment. This is the theme of Richard Dawkins's new book *Climbing Mount Improbable*. It is also, if I understand it, the point of what philosopher Daniel Dennett calls Leslie Orgel's Second Law:

"Evolution is smarter than you are". (Alternative version, according to Dennett: Evolution is smarter than Leslie Orgel).

In practice, then, evolutionary theorists generally end up assuming that organisms (or genes, when that is the more useful perspective) do maximize; the process, the necessary caveat that they must get wherever they are going by small steps, gets put to one side.

What about equilibrium? To outsiders, it appears that evolutionary theory must be a theory of continuing, progressive change. Indeed, Stephen Jay Gould's latest book is an argument against the supposed orthodoxy that evolution must be a matter of continuing progress toward ever-higher levels of complexity. But who defends that orthodoxy? The really amazing thing I have found when reading evolutionary theory is how little they talk about evolution as an ongoing process. Instead, they tend to try to explain what we see as the result of a finished process, in which each species has adapted fully to its environment - an environment that includes both other members of its own species and members of other species. It is revealing that the title of the classic book by George Williams that is often credited with a seminal role in modern evolutionary theory - a book that essentially established the principle that social behavior should be explained in terms of the self-interest of genes - is *Adaptation and Natural Selection*. "Evolution" isn't in the title, and certainly isn't in the text if it is taken to mean some kind of inexorable drive toward greater perfection. The working assumption of Williams and most other evolutionary theorists, at least as far as I can tell, is that we should model the natural world not as being on the way but as being already there.

The most telling example of this preference is the widespread use of John Maynard Smith's concept of "evolutionarily stable strategies". An ESS is the best strategy for an organism to follow given the strategies that all others are following - the strategy that maximizes fitness given that everyone else is maximizing fitness, with each taking the others' strategies into account. Does this sound familiar? It should: the concept of an ESS is virtually indistinguishable from an economist's concept of equilibrium.

And by the way: Maynard Smith's textbook is explicitly skeptical of claims that evolution is necessarily an ongoing process, let alone that it need have any particular direction. Not only do the models normally settle down to an equilibrium; so do experiments, for example with RNA evolution. And any evolutionist has got to be aware that life appears to have stayed happily single-celled for several billion years before something led to the next big step.

Now you can understand why I say that a textbook in evolution reads so much like a textbook in microeconomics. At a deep level, they share the same method: explain behavior in terms of an equilibrium among maximizing individuals.

But why does evolutionary theory in practice fail to take advantage, if we can call it that, either of the myopia or of the dynamics inherent in any evolutionary story?

Why evolutionists don't do evolution

What I have argued to this point is that even though evolution is a theory of gradual change, of myopic dynamics, in practice most evolutionary theory focusses on the presumed end result of such dynamics: an equilibrium in which individuals maximize their fitness given what other individuals do. Why should the theory have taken this turn?

The answer is surely the ever-present need to simplify, to make models that are comprehensible. The fact is that maximization and equilibrium are astonishingly powerful ways to cut through what might otherwise be forbidding complexity - and evolutionary theorists have, entirely correctly, been willing to adopt the useful fiction that individuals are at their maxima and that the system is in equilibrium.

Let me give you an example. William Hamilton's wonderfully named paper "Geometry for the Selfish Herd" imagines a group of frogs sitting at the edge of a circular pond, from which a snake may emerge - and he supposes that the snake will grab and eat the nearest frog. Where will the frogs sit? To compress his argument, Hamilton points out that if there are two groups of frogs around the pool, each group has an equal chance of being targeted, and so does each frog within each group - which means that the chance of being eaten is less if you are a frog in the larger group. Thus if you are a frog trying to maximize your choice of survival, you will want to be part of the larger group; and the equilibrium must involve clumping of all the frogs as close together as possible.

Notice what is missing from this analysis. Hamilton does not talk about the evolutionary dynamics by which frogs might acquire a sit-with-the-other-frogs instinct; he does not take us through the intermediate steps along the evolutionary path in which frogs had not yet completely "realized" that they should stay with the herd. Why not? Because to do so would involve him in enormous complications that are basically irrelevant to his point, whereas - ahem - leapfrogging straight over these difficulties to look at the equilibrium in which all frogs maximize their chances given what the other frogs do is a very parsimonious, sharp-edged way of gaining insight.

Now some people would say that this kind of creation of useful fictions is a thing of the past, because now we can study complex dynamics using computer simulations. But anyone who has tried that sort of thing - and I have, at great length - eventually comes to realize just what a wonderful tool paper-and-pencil analysis based on maximization and equilibrium really is. By all means let us use simulation to push out the boundaries of our understanding; but just running a lot of simulations and seeing what happens is a frustrating and finally unproductive

exercise unless you can somehow create a "model of the model" that lets you understand what is going on.

I could multiply examples here, but I think the point is clear. Evolutionary theorists, even though they have a framework that fundamentally tells them that you cannot safely assume maximization-and-equilibrium, *make use* of maximization and equilibrium as *modelling devices* - as useful fictions about the world that allow them to cut through the complexities. And evolutionists have found these fictions so useful that they dominate analysis in evolution almost as completely as the same fictions dominate economic theory.

What is neoclassical economics?

I just said that these fictions dominate economics. But the question in economics is whether we understand that they are fictions, rather than deep-seated truths. For there, perhaps, is where economists have something to learn from evolutionists.

In economics we often use the term "neoclassical" either as a way to praise or to damn our opponents. Personally, I consider myself a proud neoclassicist. By this I clearly don't mean that I believe in perfect competition all the way. What I mean is that I prefer, when I can, to make sense of the world using models in which individuals maximize and the interaction of these individuals can be summarized by some concept of equilibrium. The reason I like that kind of model is not that I believe it to be literally true, but that I am intensely aware of the power of maximization-and-equilibrium to organize one's thinking - and I have seen the propensity of those who try to do economics without those organizing devices to produce sheer nonsense when they imagine they are freeing themselves from some confining orthodoxy.

That said, there are indeed economists who regard maximization and equilibrium as more than useful fictions. They regard them either as literal truths - which I find a bit hard to understand given the reality of daily experience - or as principles so central to economics that one dare not bend them even a little, no matter how useful it might seem to do so.

To be fair, there is some justification in the insistence of some economists on pushing very hard on the principles of equilibrium and in particular of maximization. After all, people are smarter than genes. If I offer a model in which people seem to be passing up some opportunity for gain, you may justifiably ask me why they don't just take it. And unlike the case of genes, the argument that the alternative is quite different from what my imagined agent is currently doing is not necessarily a very good one: in the real world people do sometimes respond to opportunities by changing their behavior drastically. In biology purely local change is a sacred principle; in economics it has no comparable justification.

And yet I think that despite the differences, it would be better if economists were more self-aware - if they understood that their use of maximization-and-equilibrium, like that of evolutionary biologists, is a useful fiction rather than a principle to be defended at all costs. If we were more modest about what we think our modelling strategy is doing, we might free ourselves to accommodate more of the world in our analysis.

And so let me conclude this talk by giving two examples of how a more relaxed, "evolution"-style approach to economics might help us out.

Two economic examples

As you know, one of my areas of research has been the study of economic geography. Perhaps the most basic insight in these models has been the possibility of a cumulative process of agglomeration. Suppose that there are two regions, and one region starts with a slightly larger concentration of industry. This concentration of industry will provide larger markets and better sources of supply for producers than in the other region, perhaps inducing more producers to locate in that region, further reinforcing its advantage, and so on. It's a good story, and I am quite sure that in some sense it is correct. Yet when I and my students try to present this work, we often run into a surprising difficulty: theorists get very upset about the dynamics. Why, they ask, don't individuals correctly anticipate the future location of industry? How can you have such a model without forward-looking agents and rational expectations?

Now the fact is that when you try to do rational expectations in such models they become vastly more difficult, and the basic point becomes obscured. In short, here is a situation in which going all the way to full maximizing behavior - and trying to avoid the disequilibrium, evolutionary dynamics I assume - makes life harder, not easier. It seems to me, at least, that this is a situation where economists would do a better job if they understood that maximization is a metaphor to be used only to the extent that it helps understanding.

And when I run into this sort of critique I am envious of evolutionary theorists who do models like, say, the Fisher theory of runaway sexual selection, and can use myopic, disequilibrium dynamics without apology. (If you don't know that model, it works like this: suppose that there is one gene that makes peahens - that's the female of peacock - like males with big tails, and another that causes males to have big tails. If there is a preponderance of females that carries this gene, then males with big tails will have more offspring even if they have less chance of surviving because of their visibility to predators. But because a male with a big tail is likely to be the son of a female who likes big tails, this success will also tend to spread the gene for big-tail preference The resemblance to agglomeration is obvious- isn't it?)

Another issue: consider the question of whether and how monetary policy has real effects. In the end this comes down to whether prices are sticky in nominal terms. In my view there is overwhelming evidence that they are. But many economists reject such evidence on principle: a rational price-setter ought not to have money illusion, therefore it is bad economics to assume that they do. If neo-Keynesians like me suggest that a bit of bounded rationality would do the trick, the answer is that bounded rationality is too open-ended a concept, and can be used to rationalize too many different behaviors.

And yet in evolution the idea that there are limits to the precision of maximization is adopted cheerfully. When a bird sees a predator, it issues a warning cry that puts itself at risk but may save its neighbors; the reason this behavior "works", we believe, is that many of those neighbors are likely to be relatives, and thus the bird may enhance its "inclusive fitness". But why doesn't the bird issue a warning only its relatives can hear? Well, we just suppose that isn't possible.

In short, I believe that economics would be a more productive field if we learned something important from evolutionists: that models are metaphors, and that we should use them, not the other way around.