8

Evolutionary Economics: Cultural Evolution in the Marketplace

Cultural evolutionary concepts and methods can also be used to explain various aspects of economic change. Unlike some other branches of the social sciences, economics as a field has no problem with quantitative, mathematical models, and economists routinely construct complicated mathematical models of economic processes. These models concern both microeconomic phenomena, such as the interactions of individuals or firms within a market in response to different levels of supply and demand, and macroeconomic phenomena, such as changes in the GDP or unemployment rates of entire countries. Indeed, many of the methods that economists have developed have been so useful that biologists have used them to explain biological phenomena. As noted in chapter 2, Thomas Malthus's economic model of exponential population growth served as the basis of Darwin's key insight that animals face a constant and competitive "struggle for existence," while in the 1960s biologists such as John Maynard Smith used game theory, originally developed to model the strategic interactions of individuals within economic markets, to model the strategic interactions between individual organisms competing over mates or resources.

Yet there has been much less transference of ideas and methods the other way, from biology to economics. In recent years this has been slowly changing. Two emerging subfields that have used cultural evolution theory to inform our understanding of economic processes are evolutionary economics and behavioral economics. Evolutionary economics is concerned with modeling economic change as a Darwinian process in which firms are selected in the marketplace and behavioral routines are transmitted from individual to individual within firms. This contrasts with the traditional economic assumption that economies are in a constant state of equilibrium, and that individuals can always calculate the optimal long-term strategy to follow. Behavioral economics, meanwhile, has amassed growing evidence that human economic behavior is far more altruistic and far less self-regarding than is assumed by traditional economic models. This altruistic behavior, several researchers argue, is the result of a process of cultural group selection, where altruistic groups outcompete less altruistic groups. Although both subfields are generally identified with different researchers, there are strong connections between them, and both are united by the core principle that cultural change, in this case economic change, constitutes a Darwinian evolutionary process.

Evolutionary Economics: Challenging the Myth of Perfect Foresight

The history of evolutionary thinking in economics somewhat parallels that of other social sciences.[1] In the late nineteenth and early twentieth century, several economists sought to use Darwinian principles to explain economic phenomena, most notably Thorstein Veblen and Joseph Schumpeter. Following the Second World War, however, interest in evolutionary theory among mainstream economists virtually disappeared. It was not until 1982 that a formal evolutionary theory of economic change appeared, in the form of Richard Nelson and Sidney Winter's book *An Evolutionary Theory of Economic Change*.[2] In this book, Nelson and Winter outlined what they believed were serious shortcomings of mainstream economic theory and models before presenting an evolutionary theory of economic change that addresses these shortcomings. In particular, they argued that mainstream economics is unrealistically focused on static equilibria. A typical economic model specifies a set of decisions that a firm can carry out in response to a set of external conditions, such as supply and demand in the market, and internal conditions, such as stock levels, assuming that firms are attempting to maximize profits. Rigorous mathematical modeling techniques can be used to determine the static equilibrium, or stable state, of the economic system at which all economic forces balance out, such as when supply matches demand. While such equilibria describe real economic conditions reasonably well at a single point in time, Nelson and Winter argued that they are not so good at explaining changes in economic systems over time. Economic growth that is driven by technological or scientific change is particularly poorly described by static equilibria, given that computing or pharmaceutical technology, for example, exhibits rapid and unpredictable growth and change. As famously expressed in Moore's Law, computer processing power doubles approximately every two years, and firms that rely on computer technology must adapt to this technological change. Change this rapid does not easily translate into models featuring static equilibria.

A second problem with mainstream economic theory pointed out by Nelson and Winter is its assumption that people are perfectly rational. That is, people are assumed to be able to accurately evaluate every possible behavioral option available to them (e.g., what stocks to buy, what product to launch onto the market) and choose the option with the highest long-term payoff (e.g., in terms of sales or profit). This assumes quite a lot about people's cognitive abilities: that they are aware of every possible behavioral choice, that they can tally up all of the potential costs and benefits of each choice, and that they possess almost omniscient powers of foresight, able to predict long-term changes in markets. As Herbert Simon pointed out long ago, however, people are not this clever.[3] In Simon's terminology they possess *bounded rationality*, which means that they are reasonably rational but operate within certain constraints imposed by the limitations of human cognition and the sheer complexity of many economic choices (think about how even the most complicated mathematical models and computer simulations failed to predict the 2007 global recession). Furthermore, we have encountered much evidence in the previous chapters, including models, experiments, and field studies, that people often don't engage in individual learning at all, whether it is bounded or unbounded. Instead, they simply copy the behavior of others, which, after all, is often a cheap shortcut to acquiring adaptive behavior.[4]

Routines not Rationality. So Nelson and Winter outlined an evolutionary theory of economic change that, they argued, can more accurately capture the observed dynamics of real economic systems compared to orthodox static equilibrium models that assume unbounded rationality. The first element of their theory was what they called "routines," defined as well-learned and automatically executed sequences of behaviors carried out by workers or managers in a firm. As examples, Nelson and Winter list "well-specified technical routines for producing things, through procedures for hiring and firing, ordering new inventory, or stepping up production of items in high demand, to policies regarding investment, research and development (R&D), or advertising, and business strategies about product diversification and overseas investment."[5] Routines are transmitted to new workers when they arrive at a firm and are transmitted from managers to workers when new strategies are implemented. As such, routines constitute the inheritance mechanism in Nelson and Winter's theory (as encapsulated in the slogan "Routines as genes"), which as we saw in chapter 2 is an essential prerequisite of Darwinian evolution. They also embody a more accurate assumption regarding human behavior than the extreme rationality of traditional economic models. Rather than each individual economic actor independently calculating the optimal action to perform in a given situation, they assumed that people simply engage in tried-and-tested routines that they acquired or were taught by other members of the firm. Indeed, recall from chapter 6 Andrew Schotter and Barry Sopher's experimental finding that participants are strongly influenced in their behavior in an economic game by advice given to them by a previous participant. This advice can be seen as an example of a culturally transmitted routine. And whereas this advice allowed players to successfully coordinate their choices in the form of stable conventions, providing participants with the full behavioral history of every previous player of the game had virtually no effect. If people were rational maximizers, then they should have used this behavioral history to calculate the optimal choice. Instead they went with the (possibly biased and inaccurate) advice.

An economic theory built on culturally transmitted routines rather than individual rational calculation can explain hitherto puzzling aspects of economic change. An example is the failure of Polaroid to shift into the digital camera market in the late 1990s. By interviewing key Polaroid employees and analyzing financial records and internal research reports, Mary Tripsas and Giovanni Gavetti showed that this failure to adapt was due to managers applying existing routines to a novel situation where they were no longer appropriate.[6] Polaroid's problem was not technological: from the early 1980s the company had invested huge sums of money into researching digital camera technology, and by the early 1990s when the digital camera market was taking off they had a working prototype of a high-resolution megapixel camera that was superior to any existing camera on the market. Yet by 1998 Polaroid had a limited product range and a small share of the digital camera market. To explain this failure to capitalize on their initial technological advantage, Tripsas and Gavetti point to the Polaroid management's insistence on the "razorblade" business model. Traditionally, Polaroid had made money not from their cameras, which were sold at a loss, but from photographic film, which customers had to keep buying in order to use the camera (a business model also adopted by, as the name suggests, razorblade companies, who make their money on disposable blades rather than the razor itself). But digital cameras do not easily fit into the razorblade business model, because images can be stored digitally and require no film. While the people at Polaroid were trying to develop digital imaging technology that could work with the razorblade business model, other companies overtook them in the digital camera market. As late as 1997, the CEO of Polaroid stated in an interview that "[i]n the digital world we believe that hard copy is required . . . Unless there is a consumable component, the business model falls apart."[7] This cultural inertia reflected an unwillingness to abandon previously successful routines in response to a novel technology.

Firm Competition in the Marketplace. Routines provide the inheritance in Nelson and Winter's evolutionary theory of economic change, but, as outlined in chapter 2, Darwinian evolution also needs variation and competition. Variation comes in the form of new technological innovations that result from firms' research and development (R&D) efforts. Larger firms are assumed to be able to devote more resources to R&D, and so generate more novel variation, than smaller firms. Competition might occur at the level of routines, with more-effective routines replacing less-effective routines. More commonly, however, competition is modeled at the level of the firm, with firms containing inefficient or inappropriate routines making fewer profits than firms containing more efficient routines. This ultimately leads to less profitable firms going bust. In this sense Nelson and Winter's theory is an example of cultural group selection, where selection acts on groups of individuals (i.e., firms) causing entire groups to either increase in size or go extinct.

Putting these ideas concerning variation and competition together, economist Steven Klepper has put forward an evolutionary theory concerning how and why the number of firms in an industry changes over time.[8] Klepper's model starts with the appearance of some new technological or scientific innovation, which is potentially exploitable by firms. Yet because this innovation is unfamiliar and novel, firms are not aware of how best to exploit it (given their limited foresight). Many different firms appear, each one exploiting the innovation in a different way. Eventually one or more firms hit upon a product design that is particularly effective, or at least one that consumers become familiar with. This design comes to dominate the industry and firms that do not adopt this design go bust. Once product design has converged on this single dominant form, surviving firms can invest in R&D to develop the dominant design further. As a result, new companies that try to enter the market will be at a disadvantage compared to the incumbent firms, which form an oligopoly.

So Klepper's theory predicts a distinct industry life cycle: initially the number of firms rapidly increases until the dominant design is discovered, then the industry becomes dominated by a small number of established firms. Klepper and colleagues have shown that this life cycle can be observed in many real-life industries, from the automobile industry to television manufacturing.[9] The tire industry, for example, showed a steady increase in the number of firms in the initial 25 years since the first tires were produced for the first cars in 1896, peaking at 274 in 1922. At around this time the tire industry fixed on a particular "cord and balloon" design in which the tire rubber is supported by cords rather than fabric as it had been prior to 1922, and an internal balloon is used to maintain the tire's shape. Over the next fifteen years, the number of tire firms declined more than 80 percent, with just a handful of large firms remaining. By 1950, just four companies—Goodrich, Goodyear, the United States Rubber Company, and Firestone—accounted for almost 80 percent of the tire market, a situation that has not changed much to this day. And as predicted by Klepper's model, the older and larger a firm was, and the more it spent specifically on R&D, the more likely it was to survive this culling process. The aforementioned four largest firms in 1950, for example, had all entered the market in its first ten years of existence, by 1906.

Two further points are worth highlighting here. First, although the appearance of technological innovations in Klepper's model is treated exogenously (i.e., as external events not accounted for by processes within the model), such as the initial invention of the automobile tire or the development of the new cord and balloon design, we should not forget that such technological innovations are also the result of an evolutionary process. This was argued in chapter 7, where we encountered David Hull's theory that scientific change constitutes an evolutionary process in which concepts are transmitted from scientist to scientist and selection occurs as conceptual groups of scientists compete. Given that technological developments are partly driven by scientific progress, such as the chemical process of vulcanization leading to the development of more durable tires, we would expect technological change to also take the form of an evolutionary process. There is as well a large literature on technological evolution more specifically, as mentioned in the context of the accumulation of modifications in chapter 2.[10] So a full account of industry evolution would be one of coevolution between, on the one hand, science and technology as specified by Hull and others, and, on the other, firms within industries as specified by Nelson, Winter, Klepper, and others. Such a coevolutionary process has not, to my knowledge, been formally modeled or empirically studied but might provide valuable insights into both evolutionary processes.

Second, the pattern of industry evolution predicted and documented by Klepper—an increase in the number of firms in response to a novel technology followed by a rapid reduction in the number of firms—has an intriguing parallel in biological evolution, specifically the phenomenon of adaptive radiation. When a new environmental niche opens up, such as when a mass extinction event wipes out existing species in a region, or when a geological event such as a volcanic eruption creates a new and uninhabited island, there follows a rapid increase in the number and diversity of species that fill the newly created niches. But these adaptive radiations do not continue forever. At some point all available niches have been filled, and the number of species stabilizes. Exactly this process has been observed in the lab. Biologist Michael Brockhurst and colleagues have found that when bacteria are experimentally introduced into a new environment, adaptive radiation is more likely to occur if there are no resident bacteria already occupying the available niches.[11] When resident bacteria are already present, then the newly introduced bacteria are unlikely to diversify. This resembles Klepper's finding that companies are more likely to diversify and increase in number when entering a vacant market (e.g., the tire market at the start of the automobile industry) than when entering a filled market (e.g., the tire industry after 1923 following the convergence on the cord and balloon design).

Although still somewhat at the fringes of economics, Nelson and Winter's evolutionary theory of economic change is becoming increasingly influential (more so in Europe than the United States, interestingly). Empirical case studies illustrate the value of their evolutionary theory, such as cultural inertia within Polaroid caused by maladaptive, out-of-date routines, and the diversification-then-stabilization in the tire industry in response to a new technological niche. Traditional economic theory predicts that the Polaroid management should have exhibited rather more accurate foresight and switched from the old, inappropriate business model to a new, more appropriate business model. Under an evolutionary theory of the firm, the observed cultural inertia is unsurprising, given people's reliance on culturally transmitted routines. Similarly, if the managers of tire firms possessed perfect foresight then there should not have been so many failed firms in the automobile industry. According to an evolutionary theory, strong selection during an initial period when the environment has yet to be exploited is to be expected.

Behavioral Economics: Challenging the Myth of Pure Self-Interest

Another problem with the traditional economic model of human behavior is its assumption that people are entirely self-interested: that is, human behavior is guided solely by one's own economic payoff. This assumption has been challenged recently by a wealth of experimental and field studies that show that people are not ultimately self-interested. Most people seem to possess a strong sense of fairness, which in many cases results in unselfish behavior. Take the ultimatum game, for example, which was discussed in chapter 1. Recall that in the ultimatum game a proposer divides up a sum of money between themselves and a responder. This might be an entirely fair split of 50/50, or a selfish split, such as a 70/30 split in favor of the proposer. The responder can then choose to either accept the split, in which case both players get the specified amounts of money, or reject the offer, in which case neither player gets anything. A purely self-interested responder should accept any non-zero offer, given that any offer is better than nothing, which is what they get if they reject the offer. Proposers, knowing this, should therefore offer the smallest possible amount. Yet when Western undergraduates play the ultimatum game, the most common offer is a fair 50/50 split. This is because responders typically reject any offer less than 30 percent, and only reliably accept offers greater than 40 percent. Responders therefore appear to have an irrational sense of fairness: they are willing to forego a significant payoff (30 percent of an often quite large sum of actual money) if they think that the proposer is being unfair. As detailed in chapter 1, there is a substantial amount of cross-cultural variation in this finding, with some societies such as the Machiguenga of Peru exhibiting weaker fairness norms than Western undergraduates.[12] Yet despite this cross-cultural variation, no society ever studied exhibits purely self-interested behavior, as standard economic theory predicts they should.

This finding is not specific to the ultimatum game, however. A variety of other experimental setups, as well as actual labor markets, also shows that fairness concerns are a significant motivator of people's economic decision making.[13] For example, in the gift exchange game, an "employer" can offer a contract to a "worker" in which the employer agrees to pay the worker a fixed wage in return for a specified amount of effort in an experimental task. This contract is nonbinding, and workers receive the full preagreed wage no matter how much effort they put in. Employers receive a payoff proportional to the worker's effort minus their agreed wage, while workers receive a payoff proportional to the agreed wage minus their effort. A purely self-interested worker would take any nonzero wage and put zero effort in, giving them the maximum possible payoff (assuming that the game is played only once and participants are anonymous, which are the conditions of most actual gift exchange experiments). Employers, knowing this, should offer the minimum possible wage in order to minimize their losses. Yet real people don't behave in this way. The higher the offered wages are, the more effort the worker puts in. Again, a sense of fairness is at play here, motivating workers to honestly repay the wages offered to them in terms of effort even when this is unnecessary, and indeed reduces their payoff.

Such experimental findings are reinforced by field studies and real-life case studies. When workers feel that they are being treated unfairly, they protest by decreasing their work effort. One study found that tires produced at a plant where workers had been threatened with wage cuts and temporary contracts were of significantly lower quality than tires produced during the same period at the same company's other plants, which were unaffected by employment problems.[14] Another study found significantly more flight delays for an airline that had cut its pilots' wages simply in order to increase already-healthy company profits, compared to other airlines that cut pilot pay by the same amount but in order to avoid bankruptcy.[15] Finally, the introduction of a minimum wage has been shown to significantly, and irrationally, alter people's perceptions of a fair wage.[16] After it is introduced, a minimum wage is perceived by workers to be unfairly low, even if they were happy with the same wage level before the minimum wage was introduced. And when the minimum wage is removed, perceptions of a fair wage do not return to pre–minimum wage levels. They remain high, as if people's standards of fairness have been shifted. None of these findings are consistent with the assumption that people are purely self-interested. People are more cooperative than they should be, more concerned with what other people are receiving, and more motivated to punish selfish free riders than they should be if they were solely maximizing their own economic payoff. How can such a mismatch be explained?

One potential explanation centers on the theory of cultural group selection, as developed by anthropologists and economists including Robert Boyd, Ernst Fehr, Herbert Gintis, Joseph Henrich, and Peter Richerson.[17] These researchers argue that the tendencies to cooperate and punish free riders arose in our evolutionary past as a result of gene-culture coevolution. Specifically, cultural groups in which people cooperate with one another and punish selfish free riders would have, during human evolutionary history, outcompeted cultural groups that were less internally cooperative and allowed free riders to exploit collective rewards. This intergroup competition might be through direct conquest, because internally cooperative groups are more effective in intergroup warfare (e.g., members of cooperative groups are more likely to sacrifice themselves for the rest of the group and punish deserters or cowards). Or it might be through more indirect means, such as when people are more likely to migrate to groups that exhibit prosocial norms (e.g., providing welfare for the poor or sick). Whether direct or indirect, this cultural group selection may then have favored a set of genetically specified psychological dispositions such as those exhibited in the experiments and field studies discussed above, such as a tendency to cooperate with other members of one's group, and a sense of fairness that motivates people to punish selfish free riders. These genetically specified psychological dispositions would have then facilitated the cultural evolution of various large-scale cooperative institutions. At first, these large-scale institutions took the form of egalitarian hunter-gatherer societies. Then larger social groups emerged, such as the empires that were discussed in chapter 5. Indeed, Peter Turchin's theory of empire expansion and competition is an example of the general process of cultural group selection, where empires high in internal cohesiveness expand at the expense of empires that are less internally cohesive.[18]

Business firms may be a modern-day example of organizations that are made possible by culturally group selected cooperative tendencies.[19] Traditional economic theory (e.g., transaction cost theory) views business firms as a set of explicit rules that counteract people's intrinsic self-interest. Employment contracts that specify the minimal amount of required effort from an employee, for example, ensure that employees do not free ride. Yet we have already seen plenty of experimental and real-world evidence, such as the gift exchange game or the effect of minimum wages, that people are motivated not just by explicit contracts, but also by their sense of fairness and their unselfish, cooperative tendencies. Without the prosocial motives that have emerged as a result of cultural group selection, firms could not exist.

Economist Christian Cordes and colleagues have recently modeled the growth and contraction of firms using predictions derived from the cultural group selection hypothesis.[20] In their mathematical model, a firm exists that is composed of a number of workers. Workers are either cooperative, in which case they add profit to the firm by doing work, or selfish, in which case they contribute nothing to the firm's profits. All workers cost the firm wages. Two processes change the relative frequency of cooperative and selfish workers in the firm: genetically evolved psychological predispositions and prestige-biased cultural transmission. Regarding the former process, selfish psychological dispositions cause cooperative workers to become selfish: this represents the rational temptation to free ride. Cooperative psychological dispositions, on the other hand, cause selfish workers to become cooperative: this is the result of cultural group selection. Regarding the latter process, workers are assumed to copy the behavior of the entrepreneur who founded and leads the company (who is, by definition, "prestigious" within the organization). Entrepreneurs are assumed to be cooperative in their behavior. However, their influence depends on the size of the firm. It is more difficult for entrepreneurs to influence a large firm of hundreds or thousands of employees than a small firm with tens of employees. Every cultural generation, these two sets of processes act jointly on the frequency of cooperative and selfish employees in the firm. Firms grow when adding new employees increases profit and shrink when they are making a loss.

Cordes et al.'s model showed, first, that firms grow in size as prestige bias effectively transmits cooperative behavior from the entrepreneur to the employees. Once the firm reaches a certain size, however, the firm stops growing as the entrepreneur's influence cannot reach enough employees to make it cost-effective to grow further. Second, the higher the costs of selfish behavior relative to altruistic behavior, the smaller the firm size that can evolve. High costs of selfish behavior produce small firms of mostly cooperative workers motivated by their culturally group-selected altruistic dispositions. Low costs of selfish behavior allow firms to be larger but with a higher proportion of free riders. So Cordes et al.'s model suggests that firms initially grow in size due to cooperative behavior on the part of their members, before they reach a maximum size due to the temptation to free ride and the inability of leaders to invoke cooperation via direct social influence.

One implication of this model is that business managers can influence the cooperativeness of their firms by using prestige bias to invoke their employees' innate psychological dispositions to act cooperatively. Ultimately, this should result in higher profits, as employees voluntarily put more effort into their work. Indeed, a recent meta-analysis found a significant correlation between a firm's financial performance and its "social responsibility," such as the degree to which it fosters cooperation (rather than competition) between its employees.[21] Over time, a further process of cultural group selection may then act, as more internally cooperative firms outcompete less-cooperative firms in the marketplace.

Conclusions: Cultural Evolution Explains Economic Phenomena Better Than Standard Economic Theory

In this chapter we have seen how cultural evolutionary theory has been used to explain certain aspects of economic change that traditional economic theory cannot explain well. Cultural inertia prevents firms such as Polaroid from adopting new and better routines, while industries such as tire manufacturing exhibit a wasteful culling of firms that have not adequately exploited a new technology. Neither of these makes much sense under the standard economic assumption that people (and firms) can effectively and independently determine the most adaptive long-term strategy. Employees within firms also work harder than they should do if they were, as standard economic theory predicts, solely maximizing their own economic payoff. Instead, they are motivated by a sense of fairness and a tendency to cooperate even when there is no payoff for doing so. Such altruistic tendencies can be explained by a gene-culture coevolutionary theory based on the process of cultural group selection.