

PRELIMINARY AND INCOMPLETE

Evolutionary Theory and the Current Economic Crisis

Joseph E. Stiglitz

The global financial crisis has been hard on a number of economic theories. Notions that markets are always efficient and that there are no such things as bubbles have taken a beating. In this paper, I want to argue that the crisis has also been hard on another major strand in economics, the simplistic evolutionary theories that held that unguided and unfettered evolutionary processes lead to economic efficiency. Natural selection would ensure that only the best firms would survive. With new creative firms continually originating, the economy gets populated with increasingly productive firms.¹

The paper is divided into two parts. In the first, I review some theoretical perspectives on the welfare economics of evolutionary processes. In the second, I show how the recent crisis illustrates some of these perspectives.

Some general theoretical perspectives

Evolutionary Theory and the Theory of Screening

¹ In providing this characterization of the evolutionary model, I do a great deal of injustice to the many more subtle analyses in modern evolutionary theories. I apologize in advance for any offence. My point is to draw out some common implicit assumptions in at least the “folk” characterizations of evolutionary economics.

Some 35 years ago (Stiglitz 1975), I described some of the insights that screening theory (which I was just then developing) held for evolutionary theory. Screening attempts to discriminate among “good” and “bad” firms (or workers or buyers of insurance, etc.) Those that are “good” might get allocated more resources (or pay a low interest rate, or their products might fetch a higher price.) So too evolutionary processes seek to discriminate between good (“efficient”) firms and bad firms. Those that are good survive, those that are not die, and in the process, the average efficiency of the economy is enhanced. In a sense, evolution can be thought of as a screening process; and economic systems that “screen” better in some sense will perform better.

Imperfect Screening and its consequences

Screening itself is inevitably imperfect, and I analyzed the consequences of this in a series of papers with Raaj Sah. There are two types of errors—identifying a good firm as bad or a bad firm as good.² Better screening eliminated fewer good firms and more bad firms. The costs of each type of mistake can, however, differ. The (economic) environment may itself affect these costs. For instance, there may be a shortage of good firms (projects, ideas), in which case there is a high (opportunity) cost from eliminating a good firm (project, idea) relative to that of not eliminating a bad one. Different economic structures (polyarchy, associated with a decentralized market economy, versus hierarchy, associated, say, with central planning) make mistakes of one type more frequently than those of the other type, and accordingly might perform *systemically*

² This formulation involved only two types. Obviously, in a more general formulation where there are a continuum of types, better screening simply means a more accurate labeling of the types.

better.

There is a second way in which one economic system may be better at screening: one system may do it faster. There is a third, related characteristic: the consequences. A firm that has been identified as “bad” can be instantly killed, or it can be given a “second chance.” If, of course, screening were infallible, there would be little reason to give it a second chance. But if it is fallible, then giving it a second chance is like requiring “failure” on two exams before execution. Under certain parameters, such a “humane” screening system may actually prove to be more productive in the long run.

In evolutionary processes, bad firms die, good firms prosper, and resources get reallocated to and by the better firms that survive. Firms that go bankrupt are presumed to be “worse,” in some sense than firms that do not.

Screening, Natural Selection and the Business Cycle

Advocates of this kind of Schumpeterian creative destruction even celebrated recessions; for in such downturns, the economy “cleanses” itself of bad firms.

But a closer look at what happens in economic downturns provides a less sanguine view, suggesting even that death (evolutionary selection) mechanisms may not filter for the right characteristics. This alternative perspective was illustrated by a study of Korean firms facing bankruptcy during that country’s crisis at the end of the last decade. These firms did *not* have lower overall profitability or productivity in the period prior to the

bankruptcy. On average, they made one key mistake: they took on too much debt. When interest rates skyrocketed (in a way that could not have been anticipated), they were not able to service the debt. If someone should have been fired, it was the corporation's chief financial officer. Operations—other than finance—were good. Indeed, new firms—the source of new job creation and much of the new innovation in an economy—tend to have a lower cash cushion, so a recession or a sudden increase in interest rates discriminates against such firms *even when their future potential is strong*.

One might argue, of course, that if capital becomes scarcer—interest rates increase—then the design of the screening system *should* change. The observed pattern of death (screening) is “efficient.” It might be expected that if the social rate of discount were to increase, then optimal screening should be “tighter.” But the higher private cost of capital in a recession is, in general, not a result of an increase in the social rate of discount; it is, rather, a reflection of an (increased) divergence between the private and social rates of return in the presence of imperfect information. Business cycles themselves are a reflection of a series of market failures (many of which themselves can be related to imperfections of information). There is no reason to presume that there has suddenly been a change in the optimal set of filters. Or more generally, if there is a shift, it is not obvious that it should be in the direction observed. A firm that becomes unprofitable in a business downturn may be ill suited to that state of nature. But that may be a relatively rare state of nature. If so, one would want to assess its likely performance in future goods states.

If capital markets were perfect—and information were perfect—a firm whose only flaw was a mistaken financial policy should be able to borrow against its future prosperity. To some extent, it may be able to do so. Investors may have confidence in his business model, giving him the wherewithal to survive the temporary bad news. But capital markets at best are likely still to be more pessimistic than the entrepreneur, and as a result, the entrepreneur will (on average) be unable to get the requisite capital.

More generally imperfections in capital markets and in the inability to screen perfectly between firms which just have a problem in financial structure (or whose profits are more cyclically sensitive, with strong opportunities in good states) and those which have deeper problems mean that a large fraction of the firms that close in a recession are “good firms,” consistent with the observed results for Korea. The screening process is highly imperfect; the evolutionary “natural selection” would appear to be excessively tight in a recession.³

Ecology and Evolutionary equilibrium

³ The social consequences of the fact that a good project with a bad timing simply can't get finance may be less than the private consequences to the particular entrepreneur whose fortunes have been diminished. Failures get a second (and sometimes a third and a fourth chance.) If there are relatively good ideas around the demise of which was due to timing, then some entrepreneur will recognize this, and the new idea will be tried at a more opportune time. The first entrepreneur, whose timing was unfortunate, suffered a capital loss. But the second entrepreneur, who now has less research to do to undertake the project, gets a windfall gain. But if cyclical fluctuations are unpredictable, and if innovators risk being able to appropriate the returns to their ideas if their timing is unlucky, the supply of innovations will be less than optimal.

There are three other aspects of evolutionary processes that this discussion overlooks.

The first is one that Darwin himself noted: the most important aspect of the environment facing any species was created by nature itself, the ecology. When one says, “survival of the fittest,” it is a *relative* statement. But the point is broader: whether there is a niche for a species depends on the niches taken up by others. Even with the same external environment, there can be different ecological equilibrium. As Charles Darwin put it near the end of *The Origin of Species*, reflecting on the Galapagos Islands:

[The plants and animals of the Galapagos differ radically among islands that have] the same geological nature, the same height, climate, etc.... This long appear to me a difficulty, but it arises in chief part from the deeply seated error of considering the physical conditions of a country as the most important for its inhabitants; whereas it cannot, I think, be disputed that the nature of the other inhabitants, with which each has to compete, is at least as important, and generally a far more important element of success. (Darwin [1859] 1883L 540)

As Hoff and I [2000] argued, “The economy is like an ecosystem, and Darwin was implicitly recognizing that ecosystems have multiple equilibria.”

(Formally, we can write that the growth of firm i , whose size we denote by x^i depends on that of other firms and state S :

$$X_{t+1}^i = A(X_t^i, S_T) X_t^i$$

If the transition matrix did not depend on the state or the “endogenous environment”,

then there would be a unique equilibrium given by

$$X^* = AX^*.$$

But with A being a function of X, with S fixed, there can be multiple solutions to

$$X^* = A(X^*, S)X^*$$

The existence of multiple evolutionary equilibrium raises fundamental questions about the welfare economics of evolution. It is hard to claim simultaneously that evolution leads to “superior” outcomes, but that it can lead (randomly) to different outcomes. *The fact that something emerges as an evolutionary equilibrium may have no welfare significance*, other than that such an equilibrium may be superior to some other “outcomes” which might not have been able to have been sustained.

Adaptability and myopia

The second observation is that there are different environments, and what is most “fit” for one environment may not be the most fit for another. This is obvious; but what may not be so obvious is that environments change (in ways which may or may not be predictable), and that means that *adaptability* may be a critical characteristic. But when the pace of selection does not match well with the pace of variability, natural selection may yield results that may not seem so good. There may be insufficient appreciation for the value of variability.

Consider two species, one that can eat a variety of foods but is less efficient in processing any of them, and a highly specialized species, that eats only one food, but is very efficient

in doing so. Assume the “environment” in which they compete changes randomly, and the force of natural selection works rapidly, but growth rates starting with low numbers are low. The particular food on which the “well adapted” species depends periodically disappears. Assume our criteria for a good evolutionary system is the time-weighted biomass. Then it should be clear that it is possible that “natural” selection will not work well. The “efficient” species drives out the “diverse” one, until the food on which it depends suddenly disappears. Then, assuming that a few of the Type B survive, that group enjoys a *slow* resurgence, until nature reverts to the former situation. Type B might do better *on average*, but during the period in which he has a competitive disadvantage, he can’t “borrow” from his future strength. And because recovery takes time, but death is rapid, much of the time that he should be the dominant species, his numbers are low relative to his well-suited competitor.

Irrationality and “inefficient” screening/evolution

More generally, there is no reason that any such myopic selection process would maximize any intertemporal welfare measure. Indeed, matters may be even worse. Even without a changing environment, a “species” which is maladapted for long run survival can drive out a good species, before the maladapted species actually succumbs itself. A firm that has enough resources with a bad product, but convinced that its product is superior, can undercut a less well-funded competitor, extinguishing it before its own limitations get exposed. The real costs are associated with hysteresis: One cannot simply “unbankrupt” a firm that has gone out of existence; organizational capital has been destroyed. It is expensive to recreate it.

(Of course, in a world with perfect information, the bad firm's owner would realize that it is bad; lenders would recognize that the good firm is good. The bad firm's owners would not squander their capital in the hope that, in time, the firm would show its merits; the good firm would be able to borrow funds to keep it alive, as investors would realize that eventually, its competitor would not survive. But information is imperfect. In general, firms have to form prior beliefs about how much better than they are relative to existing products. At a zero interest rate, it would always pay those that believe that they have good product to push their products, until they get enough evidence to the contrary, in which case they abandon the market penetration exercise. Existing firms would receive enough financial support during the battle to keep them going. There would be no premature deaths. But capital markets are imperfect, and firms are not always rational. Some entrepreneurs are overly enthusiastic about their products—just as most of my students believe that they are in the top half of the class. Thus, irrationally optimistic firms can drive out better firms, and while such entrepreneurs and those who provide them capital bear the brunt of the costs, those who are driven out of business and their workers bear some of the costs.)

Other intertemporal inefficiencies in selection

These are only some of several ways in which natural selection processes may not be intertemporally efficient. I just provided an example where the implicit interest rate in natural selection is high. But there are other instances in which it is too low. The result is that it might yield outcomes that are in the long run efficient, but at some short run

costs—greater than can be justified by any discounting. Consider a two-armed bandit being played by a Bayesian. It is well known that he can get trapped playing the poorer arm as a result of a bad set of outcomes on the “good” arm. But in an evolutionary model, the good arm might nonetheless continue to be tried over and over again—and eventually it would win out. (In a sense, if the genetic mutation process remains unchanged, the set of new candidates to challenge a dominant species remains the same, so that a superior mutation will eventually become dominant.) In evaluating the evolutionary process, we do not focus on the cost of the failed attempts to challenge the dominant species—these are all part of the rubbish of history gone by. We look only at the “quality” of the resulting dominant species. And because the external environment is always changing (and it is even possible that the dominant species or institutions themselves might change, there is also a quest to see if some other species is better for the new environment. There is no place for—no opportunity for-- institutional or individual ossification in the story we have told.

Market failures

Evolution screens out firms on the basis of profits, but private profits often do not correspond to social returns, and if that is the case, firms that do well—and survive—are not necessarily maximizing social efficiency. This is the case, for instance, whenever there are externalities.

Again, it is a case of Gresham’s law: bad firms drive out good firms. Firms that treat the environment appropriately will not survive in competition with firms that do not, so long

as there are not appropriate charges imposed for environmental degradation.

A key aspect of the standard evolutionary story has to do with innovation (about which I will have more to say later.) But innovation is rife with externalities—in the case of important innovations, the innovator seldom appropriates more than a fraction of the social returns.⁴ In other cases, private returns may exceed the marginal social returns. The rules of the game (e.g. relating to intellectual property) determine the extent of appropriability. But if social returns typically deviate from private returns in this key area, how can there be a strong presumption that evolutionary processes will lead to long run dynamic efficiency?

Another important source of market failures arises when there are imperfections of competition. Monopoly profits do not accord with social returns. Moreover, in such situations firm profitability may be increased by raising rivals' costs (more than by reducing one's own costs.)

Public bailouts

A particularly relevant set of market failures arises from government bailouts.

Individuals do not bear the full consequences of the risks which they have undertaken.

That means that private profitability (ex ante) may differ markedly from social returns.

⁴ For a simple exposition, see J. E. Stiglitz, 2006 (chapter 5) and 2008. Greenwald and Stiglitz [2006] present an argument for government intervention in the structure of the economy to promote technological progress.

Firms and enterprises that expand may do so not on the basis of greater abilities or better understanding of risk, but on the basis of a greater willingness to be bailed out by the government. Again, bad firms may drive out firms that are less willing (or able) to impose costs on others.

Distortions in the supply of mutations

There are, in addition, two fundamental ways in which the analysis of natural selection in economics differs from that in “nature.” The first is in the generation of “mutations,” the innovations among which natural selection is supposed to choose the “best.” The innovations are random, and while they might be affected by certain physical conditions (e.g. exposure to radiation), those conditions themselves are not directly related to the environmental conditions which determine which mutations are “best.” This is not so in economics, where innovations are endogenous, and very much affected by the same conditions which affect natural selection. New firms (innovations) do not arise randomly, but in response to a perceived deficiency in the market. In some cases, that’s a good thing. It means that there is a “favorable” bias in the creation of new enterprises, so long as there are not market distortions. If wages are high, so that technologies that are labor saving “survive” better than other technologies, then innovators will have an incentive to design new technologies that are labor saving.

But if there are certain market failures, with private rewards not well aligned with social returns, then incentives for innovation will also be distorted. The firms that arise to exploit “gaps” in the markets may actually be dysfunctional. They may have adverse

effects on other firms, just as a weed that might develop to fill a “hole” in an ecological space might then overrun other species, with long run adverse effects. Not only with technologies which are socially the most profitable not be privately most profitable—so that such technologies may not survive—but there will be a corresponding bias in innovation. (The poster child of such distortions in earlier decades may have been the cigarette companies, who exploited new knowledge about how to increase addiction, and chose to ignore the new knowledge about how cigarettes cause cancer and other diseases. The companies that were best at developing and marketing addictive products, regardless of the adverse effects on health and of the costs imposed on the rest of society, prospered. As we note below, the financial sector has become the poster child of such distortions in the twenty-first century.)

Teleology, welfare, and public policy

Finally, the *objective* of natural selection in biology is generally viewed to be simply survival (e.g. of the gene). Natural selection does not, however, in general, maximize other objectives, such as the average biomass which can be supported by a given (fluctuating) environment. In economics, those who argue for the virtues of a competitive marketplace on the basis of evolutionary “natural selection” have presumably some more welfare-oriented measure. Natural selection, it is argued, leads to efficiency, as inefficient technologies cannot survive. But the examples given above suggest that the view of efficiency is short sighted: it may select for a firm (technology) best adapted to the current environment, not one which is most flexible, and can best adapt to the changing environment. The result is that *average* efficiency may be lower. More

generally, in the absence of some variant of a perfect capital market, the outcomes of natural selection are unlikely to be consistent with longer term social objectives; the problems get heightened as the discrepancy between social rates of discount and market rates of interest increase, as often seems to be the case in recessions.

Every society determines certain rules of the economic game, e.g. those relating to bankruptcy; there may be, in addition, corporate safety nets. These rules of the game determine how the process of natural selection plays out. To put it another way, in economics, there is no such thing as *natural* selection: all selection is a consequence of a set of government policies. The earlier discussion highlighted certain biases in the selection process, which can either be ameliorated or exacerbated through the design of the rules of the game.

The current crisis and the theory of natural selection

The recent crisis illustrates several of these limitations in evolutionary theory. We have noted that Schumpeter may have trumpeted the creative destruction of the business cycle, but there is little reason to believe that many of the firms that have died or will die before recovery sets in are the least productive.

Market failures and financial markets

It is perhaps not surprising that evolutionary processes have not worked well in the lead up to the crisis: we argued that they typically don't work well when there are important

market failures, and financial markets are rife with market failures. The failure of a financial of a financial institution has externalities on others; the failure of the financial system has large externalities on the functioning of the rest of the economic system. A core function of the sector is gathering, analyzing, and disseminating information, and markets in which information is at the center are especially prone to market failure. Moreover, in certain areas, there is considerable evidence of imperfections of competition.

Financial markets did not perform their core functions of managing risk and allocating capital. They are supposed to manage risk and allocate capital well—and do so at low transactions costs. Ours failed in their social functions, but imposed high transactions costs.

In the paragraphs below, I comment on the insights provided by the death and survival of certain institutions and practices may shed light on evolutionary processes.

Survival of the fittest or of the fattest?

Short run profits during a recession may be an imperfect indicator of long run performance, and it is the older, more mature firms—the cash cows—that have the easiest time in survival. It is the newer, more innovative firms that are most likely to die, simply because they have less of a cash cushion.⁵

⁵ Again, in the presence of perfect information, none of this would matter.

Gresham's Law: Bad banks drive out good banks

Moreover, the above analysis argued that natural selection does not necessarily choose those firms that are best *in the long run*. Natural selection can be myopic—just as many of America's financial firms were myopic. Some of those in the failed investment banks claim that they realized that their high leverage exposed them to high risk—even a small change in asset prices could bankrupt them. But, they claim, they had no choice. If they had not undertaken such leverage, the return on equity would be low, and analysts and investors alike would have subjected them to unrelenting criticism. Their officers (or even the enterprise as an independent corporation) would not have survived. Their firms would have been taken over and they themselves would have been fired. Rather than “survival of the fittest” one had a version of Gresham's law: bad firms (undertaking excessive risk) drove out good firms. But now that the recklessness of the banks like Citibank has been exposed, the more prudent firms that could not have survived cannot be brought back from the dead.

Moreover, when there are market distortions, private and social profitability differ, and this too has been evident in the crisis. Much of the profits represented predatory lending and taking advantage of imperfections in competition. Usurious and non-transparent interest rates and fees, with products targeted at the poor and less well educated, marked the most successful firms.

Assessing profits and profitability

But in finance, in particular, there is a further problem—measured profitability (in the short run) may have little to do with true long run profitability, because of creative accounting. This is most evident in insurance companies, which take in premiums today, in return for payments to the insured if the insured against event happens—sometime in the future. In the short run, one can't tell whether the firm is profitable unless one knows whether the firm has set aside enough money to pay the promised benefits, and one doesn't know that unless when knows the probability of the insured against events occurring. When they are events like death, there may be a statistical basis for making those probabilistic inferences; but in other cases there may not be. This is where AIG got into trouble—and got the country into trouble. But the problem continues in the credit default swap markets.

With imperfect information, firms that prosper (and expand) are not necessarily those that are doing well, but those that *appear* to be doing so, and in finance, there are large opportunities for manipulating appearances through creative accounting, evidenced both in the Enron/WorldCom scandals earlier in the decade and in this crisis. Financial institutions showed their mastery as this creating accounting—simultaneously claiming low returns to the tax authorities and high returns to potential investors—by moving certain risks and liabilities off balance sheet. Such create accounting has “survival value,” i.e. if one does it successfully, the corporate coffers are richer on both accounts.

Typically, information that's hid in the process of creating accounting eventually does

come out. (Those undertaking these acts of creativity hope that this will occur after they have sold their shares. That is what happened in the case of America's banks, whose officers managed to still walk away with millions.) The social consequences of these deceptive accounting practices can be great, as we saw in the earlier episode and even more forcefully in this crisis: credit markets froze as each bank knew that it could not assess the credit worthiness of the other banks to whom it might lend. But banks that did not engage in these deceptive practices may not have been able to have survived, as their returns would have *appeared* to be lower year after year.

Exploiting economic ideologies and investor ignorance

There are other examples in which investor ignorance—or mistaken ideas about what makes for a good firm—can be exploited by firms; firms that act in a way which is consonant with prevalent ideologies will do well. And it is in the interests of these firms and their officers to promote these ideologies. In short, if there are investor irrationalities, “survival of the fittest” may entail playing into the irrationalities; the firms that prosperous are not necessarily the firms that are more efficient.

Consider, for instance, the mantra surrounding executive compensation. It became a fashion that executives should be paid on the basis of performance. It seemed reasonable that if you simply paid an executive \$5 million a year, he would only devote a quarter of his attention to the job. The CEO, evidently, was in no position to ascertain the lackadaisical effort. So too for the board in judging the performance of the CEO. He could only judge by results. And so stock options became a prevalent form of

compensation, a natural evolution towards a more efficient compensation scheme. Firms that were more aggressive in adopting these “better” incentive driven compensation schemes were viewed as more dynamic. They prospered, at the expense of those that did not. But did any of this make sense?

Most of the increase in the stock price was unrelated to the firm’s performance, but to events outside the firm’s control (lower interest rates help drive up stock prices, lower oil prices help drive up the prices of products that use oil as inputs.) In terms of standard theory—devising “high power” incentives with low “noise” (i.e. with returns unrelated to the performance of the manager) there are far better systems (e.g. relative stock performance using an appropriately defined comparator group) (Nalebuff and Stiglitz, 1983a). Relatively few firms use these better techniques. Indeed, most firms use a variant of the executive compensation scheme just described, for it leaves open the question of what happens when the stock market crashes. The company is willing, at that point, to admit that some variations in market price have little to do with the efforts of the management. They plausibly argue that they should have to bear these costs. A more accurate description is that senior executives get paid highly in good times through stock bonuses, and in bad times through other mechanisms, like retention bonuses, or special pay to recognize the difficulties of managing one’s ways through these difficult times. Looking more carefully at actual pay as it relates to, say, market performance—taking into account the variety of escape clauses in practice—it appears that pay is little related to performance. (See Stiglitz, 2003, and the literature referred to there.) This discrepancy between the claims of incentive pay and actual practice were exposed by the crisis.

But the compensation practices, including the ideology of pay for performance, has allowed senior executives of many large firms to extract far more rents from the corporations which they head. If what matters is the survival of the CEO, not in his position, but in the maintenance of his wealth, then the current system may have a kind of evolutionary robustness. Even those that lost their jobs walked off with hundreds of millions of dollars. The newly established firms as well as the older firms that survived seem quickly to be gravitating towards similar compensation schemes. While overall, a financial institution with a less rapacious executive compensation scheme might perform better over the long run, it has two strikes against it in the short run: if investors believe that such schemes are necessary in the financial market, they will provide it less capital (or demand a higher cost for their capital.) Moreover, the higher returns paid to the CEO may result in their ability to higher better CEOs—able to generate higher turns, maybe not enough to compensate for the higher risk and the extra pay, but good enough to deceive the undiscerning who have come to believe in the necessity of “incentive pay.” They filter the information about returns in ways that confirms their prior bias.⁶

While those in the industry, supposedly experts in economics and risk, touted the benefits of incentive pay did they discuss alternative more efficient systems. Clearly, it was not beyond their intellectual ability, at least of those who were supposedly managing large amounts of risk. But even worse, they did not discuss the distortions and risks associated with the commonly employed schemes. These included on average, rewards increased

⁶ This is called confirmatory bias. See Rabin and Schrag, 1999.

with increasing beta, i.e. exposing their firms to greater systematic risk and in some cases creating greater systemic risk, rather than on increasing alpha, improving performance at any given beta. There is no skill to the former, a lot of skill in the latter; and yet no differentiation was made. In the end, not surprisingly, most of the executives chose the easy route, increasing beta—e.g. through an increase in leverage. Did the banks, supposedly experts in risk management, not realize this obvious perversity in the design of their compensation scheme? Or was it simply that greed prevailed: they discovered that these were a form of words that gave comfort to their shareholders as they robbed the corporate rents. In any case, from an evolutionary perspective, given the myopia of investors and their lack of understanding of these markets these firms “worked well”. For a number of years they delivered high returns, with stock prices correspondingly high. They are now trying to tell a story to their investors of a once in a hundred year flood, a framework that basically worked well, and that there is little reason to change the basic business model.

And indeed, the officers of the firms, even of those that have since died have, for the most part, done well for themselves—though, to be sure, they are not as well off as they would have been had their strategies paid off.

As firms realize that investors have “bought into” certain aspects of business ideology, there are ample opportunities to exploit investor irrationalities. And while the *eventual* poor performance of these firms compared to those following another ideology (rule) in the long run suggests that these models may not have long run survival value, there is

nothing to suggest that they are replaced by firms that are truly efficient over the long run, but rather by other firms playing to other ideologies, delivering, in many cases, only more returns in the short run.

Organizational versus individual incentives/selection: Too big to fail

The private rewards and social returns to the actions of corporate officers are not well aligned, but the real problems in evolutionary processes arise from the disparity between the private returns of the “organization” and the social returns. If it were not for this disparity, then firms that evolved incentive structures that did a better job at maximizing organizational profits would survive.

Thus any firm that grew to the size that it was too large to fail knew that (at least with a high probability) it would be bailed out. There was uncertainty about the extent and form of bailout: would shareholders get bailed out, bondholders, those holding counterparty risks? It was unconceivable that none of these will be bailed out, so that a rational actor would undertake excess risk. The more prudent firm, not wanting to face the risk of a public bailout, would perform more poorly, and would not survive.

Thus, there were organizational incentives to become too big to fail or too intertwined to fail. It is not the best firms that have survived, but those that were best in realizing this.

Financial Innovation and evolution

The financial markets touted their innovation. But in fact, apart from the important

institutional innovation of venture capital firms, the *social* returns of much of the innovation remain elusive. Some of the innovations lowered transactions cost; but they also increased problems of information asymmetries and complexity, contributing to the uncertainties in firms' balance sheets, which in turn contributed to the freezing of credit markets. If there were significant social benefits, they were almost surely far outweighed both by the cost to the taxpayers of the bailouts and the costs to the economy in the resultant loss output resulting from the recessions.

Indeed, the financial markets have been criticized for not only the failure to innovate, but for even resisting key innovations, e.g. that would have lowered transactions costs (through the creation of an efficient electronics payment mechanism) and provided better mortgage products (e.g. Danish mortgages and mortgage bonds) (Stiglitz, 2010).

Indeed, some of the “newer” mortgages were a step backwards—they made it more difficult for poor individuals to manage the risk of homeownership. Bad mortgages drove out good; bad mortgage dealers and brokers drove out good. The saving grace may be that we have learned our lesson, and that going forward, good mortgage products will dominate in the market. And yet, both the political and the economic provide little comfort. Good products, like Danish mortgage bonds, are hardly being discussed. Attempts to restrict the worst mortgage products have been beaten back. The financial institutions may have lost much of their capital, but they have retained a considerable amount of political capital, and have used it in ways that limited the scale and scope of reform. The institutional evolution is less than one would have hoped for.

Regressive institutional innovations

Changes—“innovations”—in financial markets may have made the problems posed by myopia worse. Today, much of financial capital is managed by “agents,” i.e. by individuals on behalf of others (e.g. the huge pension funds.) These are short term investors, and are judged by their quarterly returns; so in return, they focus on the short term returns of the firms in which they invest. For short term investors, a firm that appears to be doing well—so its stock price rises—*is* doing well. They can cash in on this short term performance. Hence, markets seem to have evolved in ways which encourage short termism, distorting further the selection process.

How Financial Engineering may make matters even worse

Earlier, we noted that if investors are not fully rational, then markets have an incentive to exploit their foibles. Advances in science have increased the ability to exploit these foibles. But, again, this increase in profitability is not necessarily a sign of increased societal well-being.

But even if investors are highly rational, financial institutions have an incentive to devise products that are, at best, hard to evaluate—and financial engineering has given them the tools to do so. Consider a financially engineered product that yields an above normal return almost all the time—but once in 20 years, there is a 90 percent chance of a large loss in value. In buying such a product, no one knows whether his expected return is higher or lower than the expected return on, say, a safe asset. He will only know whether

he should have bought the product at the end of the day, when he sees the magnitude of the loss. Of course, the purveyor may claim that the loss is limited, so that on average, the return is far higher than on the safe asset. But why should he be believed? But the complexity of the product is such that few not in the industry can assess the likely return. An economist ingrained in the notion that there is no such thing as a free lunch would be skeptical. But the ordinary layman, seeing a “safe” asset beating the market year after year might be tempted to purchase the product.

I raise this specter because it illustrates the difficulties facing evolutionary processes in making good *long run* selections—difficulties that only will increase with time.

Government regulation can help—Ponzi schemes and quasi-ponzi schemes can be detected. But most of these schemes are not Ponzi schemes in the conventional sense—though by promising more than a “normal” return they are not too different. In short, natural selection may favor those who continue to devise such products—rather than those who are making the real investments that really matters.

Is the failure of Lehman Brothers, Bear Stearns, and Merrill Lynch Vindication of Evolutionary Theory?

The fact that some of the worst culprits have died, and some are on government life support programs, suggests to some that evolutionary processes work—slowly, but eventually. But the fact that some of the worst practitioners of these deceptive practices have survived raises deeper questions. If what matters is relative performance, all that one cares about is that one be the last one standing. Our financial system was intertwined

in a complex way, such that how one began to unravel the mess could make a big difference. If the government rescued some of the weaker institutions in ways that preserved some of the stronger banks, then the latter would actually emerge stronger. If the government had let AIG go under, whether Goldman would have survived may be a matter of dispute, that that it would not have gotten gratis \$13 billion in settlement of an AIG contract is not. It is also apparent that Goldman has benefitted not only the AIG bailout, but also from the demise of some of its key rivals. What happened may bear more testimony to political doctrines about the survival of the politically connected than to standard evolutionary notions.

Concluding Comments: Natural selection and economic evolution

Nature sets its own rules; but in economics, we can set the rules that affect the force of selectivity. As we have noted, laws and customs determine, for instance, the nature of competition and what happens when firms “fail.” These laws and customs evolve over time. And as they evolve, they affect the evolutionary process.

Unfortunately, just as natural selection does not, at least in the relevant time frame, necessarily choose the enterprises that are “socially efficient,” so too it does not select the institutions, rules, and customs that are best for the long run, including the “best” evolutionary process. One of the main criticisms of financial markets is that they have become increasingly shortsighted. Some of the institutional changes (such as investors’ focus on quarterly returns) have made it more difficult for firms to take longer-run

perspectives. If markets become increasingly shortsighted, with increasing focus on quarterly returns, the “quality” of natural selection may deteriorate. The fact that selection is based on quarterly observed returns affects firm behavior. Short sighted behavior (and accounting gimmicks) are effectively encouraged. Those that are most successful in adapting to this new “selection process” will survive best. This is especially so in a market economy characterized by a series of agency problems, in which many of the “principals” are themselves shortsighted.

By contrast, “family firms” can take a long horizon. They pay little attention to day-to-day market values. But there are limits in the size to which such firms can typically grow. Historically, family firms had to raise outside capital. Many of the widely traded firms are slaves to what happens to their market capitalization, and have become increasingly so. Their executives are paid with stock options, which provides them with incentives to design policies and provide information that maximizes current stock market value. This suggests the evolution of an economic ecology in which there are different kinds of firms—each with its limits. The family firm may do a better job at intertemporal wealth maximizing, but it is constrained in its access to capital. As it becomes publicly listed, it becomes subjected to the short sightedness that has been so obvious in the current crisis.

Firms and their managers inevitably will exploit investor ignorance and help develop ideologies that justify their rent extraction activities and take actions that enhance their scope to do so (Edlin and Stiglitz, 1995). Innovations, as we have seen, do not

necessarily make the market work better. Some innovations enable markets to exploit better investor and consumer ignorance. Incentives to innovate are themselves distorted, not just by the broader misalignment of private rewards and social returns, but by a more specific distortion: the difficulty in appropriating the marginal social return to innovation. This is especially the case for social innovations.

Adam Smith held out the promise that the market, would lead the economy, as if by an invisible hand, to efficient (welfare maximizing) outcomes. A major insight of the past quarter century is to show that that is not the case. Unfettered markets are not, in general efficient. And while there are government interventions which could make everyone better off, governments too fail, and often their interventions are directed at helping the strong and powerful at the expense of the weak and poor.

Those who lost faith in the mechanistic processes just described—or in the assumptions of rationality on which they were based-- held out the hope that “good” outcomes could be achieved in a more “groping” decentralized way, through evolutionary processes. By trial and experiment, we do things better, and as we do that, society gets better off. If individuals were living in isolation—exploring how to improve the productivity of their farm—the evolutionary model might provide important insights into improvements in standards of living over time. But we are interested in the evolution of societies, and it should come as no surprise that when private and social returns differ markedly, evolutionary processes cannot be relied upon. The firms that survive may not even be the fittest—in any relevant sense. And over time, the ability of evolutionary processes to

produce better outcomes may be attenuated. (In a sense, this is a pessimism which even Schumpeter shared: he was not confident that capitalist institutions, which he thought as superior at promoting innovation and the process of creative destruction, would survive.)

In this paper, I have outlined some of the limitations of evolutionary processes and suggested how these may have played out in the current financial crisis. The market did not select for the firms that held out the promise of the highest long term economic growth.

By the same token, there is little reason to believe that, going forward, evolutionary responses to the crisis—without strong government intervention—will suffice to protect us against a recurrence. Nor is there reason to believe that, at least in the relevant time frame, natural selection will necessarily work against those countries in which governments that fail to respond in the appropriate way.

More generally, this crisis has not provided comfort to those who believe in the efficacy of evolutionary processes in promoting economic efficiency. Bad mortgages drove out good; bad mortgage dealers and brokers drove out good. Banks and financial institutions that learned how to exploit better the unwary poor prospered—at least for a time—at the expense of more ethical institutions. Knowledge (in this case about consumer weaknesses) was used to increase profits, but not societal well-being; and just as the financial sector had demonstrated its ability earlier in the decade to manipulate information (“creative accounting”) in ways that may have well done for its officers, but

not necessarily for society as a whole, new forms of “creativity,” often based on deception and lack of transparency, so too did they do so once again. The saving grace may be that we have learned our lesson, and that going forward, good mortgage products will dominate in the market. So too, one might hope that banks that are more competitive, do a better job at risk management, and provide better the key functions that banks should provide will dominate. And yet, both the political and the economic provide little comfort. Good products, like Danish mortgage bonds, are hardly being discussed. Attempts to restrict the worst mortgage products have been beaten back. The surviving banks are not making money by lending to enterprises, but by speculative activity. They continue to resist the creation of an efficient electronics payments system. The financial institutions may have lost much of their capital, but they have retained a considerable amount of political capital, and have used it in ways that limited the scale and scope of reform.

REFERENCES

- Greenwald, Bruce and J.E. Stiglitz, "Helping Infant Economies Grow: Foundations of Trade Policies for Developing Countries," *American Economic Review: AEA Papers and Proceedings*, Vol. 96, No. 2, May 2006, pp. 141-146.
- Hoff, Karla and J. E. Stiglitz, "Modern Economic Theory and Development," *Frontiers of Development Economics: The Future in Perspective*, G. Meier and J. Stiglitz (eds.), Oxford: Oxford University Press, 2001, pp. 389-459.
- Sah, Raaj and J. E. Stiglitz, "Human Fallibility and Economic Organization," *American Economic Review*, 75(2), May 1985, pp. 292-296.
- _____, "The Architecture of Economic Systems: Hierarchies and Polyarchies," *American Economic Review*, 76(4), September 1986, pp. 716-727. Also a shortened version in *AEA Papers and Proceedings*, also published in *Personnel Economics*, Edward P. Lasear and Robert McNabb (eds.) 2002.
- _____, "Committees, Hierarchies and Polyarchies," *The Economic Journal*, 98(391), June 1988, pp. 451-470.
- _____, "Qualitative Properties of Profit-Maximizing K-out-of-N Systems Subject to Two Kinds of Failure," *IEEE Transactions on Reliability*, 37(5), December 1988, pp. 515-520.
- _____, "The Quality of Managers in Centralized Versus Decentralized Organizations," *Quarterly Journal of Economics*, 106(1), February 1991, pp. 289-25
- Stiglitz, J. E., "Information and Economic Analysis," in *Current Economic Problems*, J.M. Parkin and A.R. Nobay (eds.), Cambridge: Cambridge University Press, pp. 27-52. (Proceedings of the Association of University Teachers of Economics, Manchester, England, April 1974.)
- _____, *Making Globalization Work*, New York: WW Norton, 2006.
- _____, "The Economic Foundations of Intellectual Property," sixth annual Frey Lecture in Intellectual Property, Duke University, February 16, 2007, published in *Duke Law Journal*, Vol. 57, No. 6, April 2008, pp. 1693-1724.