KEYNOTE ADDRESS
Challenges in the Analysis of the Role of Institutions in Economic Development

Joseph E. Stiglitz

When this year's World Development Report was first discussed, there was great enthusiasm for the topic of institutional economics. But the enthusiasm was in part at least based on a range of interpretations of the topic. While there was widespread recognition of the importance of institutions, there was no consensus on what a WDR on institutions should focus. Focusing on the institutional foundations of a market economy narrows the subject, yet leaves ample room for the WDR to embrace very different views. In reflecting on some of these, I will first talk briefly about the history of institutional economics. Then I will discuss the definition of institutions. Finally, I will discuss two questions related to economic theory and history: How is it possible for bad institutions to persist? And what policy interventions might lead to changes conducive to good institutions?

The History of Institutional Economics

There have been wide swings in the emphasis and importance accorded institutions in economics. Today, not only is their widespread recognition of the key role that institutions play, but institutions have moved "front and center," both in trying to understand how economies function and in trying to improve economic performance. While the economics profession began to pay increasing attention to institutions some quarter a century ago, there is clearly a true surge in interests in the subject today.

A strong German historical school focused on institutions and had some influence on American economics in the late 19th century and early 20th centuries before it fell out of favor. That school described institutions, institutional arrangements, and the organization of societies to make economic decisions but failed to explain why those institutions occurred. Critics said that the view was largely atheoretical and did not go far enough. But I think the real reason the view fell out of favor was quite different—the rise of neoclassical economics.

Neoclassical Economics

Neoclassical economics said that institutions did not matter in a fundamental sense. The equilibrium in the economy, the important variables, prices, outputs—everything economists wanted to know—could be explained without reference to institutions, which were superficial dressing. Banks operate in different kinds of buildings, but the height or size of the buildings was not relevant to what banking was about. In the same way, societies might have different kinds of institutions, but the laws of demand and supply were the relevant issues. No one needed to spend time worrying about how the president of a company was elected or how other decisions were made. The laws of economics, demand and supply, competitive market processes, and general equilibrium theory would explain the determination of economic outcomes.

The neoclassical view prevailed until 30–40 years ago, when people became convinced that the laws of supply and demand did not explain everything about economic equilibria. In a sense the quest of the past 30–40 years has been to discover what was left out of the neoclassical model and what difference those omissions make. Once it was clear that something was missing, people returned to the problem that the old historical perspective did not explain institutions but simply described them historically. The breakthrough came when people began to recognize that economic theory ought to be able to explain the reason for the institutions in a society, the functions they serve and the forms they take.

Limitations to the General Equilibrium Model

One of the early institutional arrangements to come under scrutiny was sharecropping. The relationship between landlords and workers in developing countries was clearly not the same in all countries; but in many, workers were paid shares, not wages, in what was described as a risk-sharing mechanism. Thus the institutional arrangement of sharecropping arose, in this interpretation, as a result of a market failure. Sharecropping involved an institutional arrangement for risk sharing in an environment of incomplete risk-sharing contracts and incomplete or nonexistent markets for risk.

The early work on sharecropping continued, in some critical ways, with the neoclassical perspective: The lack of risk markets might explain sharecropping; but still, sharecropping had no real consequences. In a sense, sharecropping was an institutional adaptation which allowed the standard theory to be maintained; for without some institutional arrangements for risk sharing, clearly there would be a large

market inefficiency. Sharecropping was thus viewed as an effective substitute for the complete set of markets that were part of the standard competitive paradigm. Thus, sharecropping, rather than replacing standard laws and insights of economic theory, only served to reinforce them; the laws of supply and demand, for instance, still applied. The institutions provided insights into why people organized themselves in a particular way, but the only requirement for understanding what determined outputs and inputs and the allocation of resources was the familiar competitive general equilibrium theory. In his famous work at the University of Chicago Cheung (1969) said that sharecropping was an interesting institution but had no effect on behavior other than as a substitute for the absence of markets for incomplete risk. Cheung's basic argument assumed that there was perfect information.

My work on sharecropping emphasized imperfect information—particularly information about inputs—which meant that the landlord looked at the output as an indicator of input. Basing compensation on output rather than input was a fundamental change in the nature of the economic relationship and had important consequences for behavior, and hence outcomes. The institutional arrangement did affect the outcomes. Further analysis yielded insights into many other issues, such as the implications for (and of) cost sharing of inputs, restrictions on technologies and processes, and even the evolution of technology. Economic theory could now not only explain why sharecropping existed, but many of the details of the forms it took; it could explain, for instance, variations in share cropping rates and cost sharing arrangement. The theory provided testable propositions—such as that productivity under sharecropping was lower than under rental arrangements. One could thus test which was closer to the mark: Cheung's hypothesis that sharecropping arrangements had no effect on outcomes, or my hypothesis that they did. Overwhelming, the evidence supported the conclusions that institutions matter.

Scrutinizing institutional arrangements such as sharecropping provided greater insight into the nature of economic processes than the neoclassical model had. But while the new theories did a better job than the old, there remained many puzzles. Some of my early papers on sharecropping were rejected by some journals on the grounds that there was no theoretical reason why sharecropping contracts should be linear—that is, why farmers with a landlord (tenants) should give one-third or one-half of the produce to the landlord. The fact that this happened all over the world did not impress the theoretical journals. Some argued that because it could be explained how the giraffe pumps blood to its head through such a long neck, it must be the case that the giraffe has a short neck; because it could not be proved that contracts should be linear, there must be some hidden non-linearity. Others took a different view: they were amenable to publishing papers which assumed special utility functions (such as constant absolute risk aversion utility functions) for which contracts could be shown to be linear, at least under some circumstances. As much as this seemed to please the journal editors, I was never pleased: there were simply too many other implications of that model that could be shown to be inconsistent with the facts. Moreover, the model, even with that unreasonable utility function, could not explain the rigidities in the shares. The theory of demand and supply says that a change in the relative supply of labor ought to result in a change in relative prices and, in a sharecropping contract, a change in the shares. But the remarkable thing about sharecropping around the world is that tenants give shares of one-third, one-half, two-thirds, or whatever, but the new equilibrium share does not change to, say, 0.378.

Many questions persist about the nature of institutions, and many of them are disturbing. But economics has moved on well beyond the neoclassical paradigm. Today, it is widely recognized that institutions matter. And there is some sense of why they matter.

North (1971) and others hypothesized that institutions fill gaps in the market—they address market failures. For example, imperfect information in a variety of contexts is associated with a moral hazard problem. Moral hazard leads to incomplete insurance markets, and institutions arise to fill that gap. The question is whether they do this well. Is there a welfare theorem for institutions similar to that for markets? Is it the case that because institutions arise to solve market failures, they are therefore efficiency enhancing? And what about the evolutionary process? Is there a dynamic version, that says that as institutions evolve in response to perceived market failures, welfare unambiguously increases? This was, of course, the view that Cheung originally argued for in the context of sharecropping. It might look as if the fact that workers received only 50% of their output would attenuate incentives; but not so. A closer look at the institution of sharecropping (in his view) would show that there were institutional adaptations (e.g. effort requirements) that ensured that the efficient outcomes—the outcomes that the neoclassical model would have predicted—prevail.

The view that institutions arise to fill gaps in the market, and thereby increase economic efficiency, sometimes called the early North view, became very strongly held for a short while in North America. North has now rejected it, but many of his early disciples, including some in the international financial institutions, still believe in it. And they believe in spite of both the theory and evidence which shows that institutions do matter, and institutional equilibrium are, in general, not necessarily efficient. Indeed, in joint work with Richard Arnott, we showed that non-market institutions which arise to fill a market gap may, nonetheless, be dysfunctional: the resulting equilibrium may actually be worse than without the institutions. For instance, because of moral hazard, insurance firms only provide partial insurance; but

"family insurance" which helps fill the gap may actually crowd out market insurance, and in the new equilibrium, individuals actually bear more risk.

There is, however, a far more fundamental attack on this perspective, which holds that a real function of institutions is not to increase economic efficiency but to help preserve some power relationship, for example, inequalities of wealth.

The Definition of Institutions

When institutional economics became a subject for debate 10–15 years ago, one of the most difficult issues was its definition. Since institutional economics is the subject of the World Development Report, it is important that we try to understand more precisely what it means. But I am not sanguine that we will be able to do so, at least to everyone's satisfaction.

One way of approaching the question is to say that people know an institution when they see one. An institution is a primitive concept, and the wisest course is not to define it. Sharecropping, the IMF, and Harvard University are institutions. So are Proctor and Gamble, United Way, and Eléctricité de France. Some institutions can be described as good or bad. The words institutional arrangement are used to refer to the U.S. or Japanese financial systems or systems of corporate governance.

The term institution is also used in other contexts fundamental to the functioning of economies. Jim Crow is an example. Discrimination was an institutional arrangement that persisted for a century in the South of the United States. It would be natural to describe the set of rules and behaviors that make up discrimination in general as an institution. The same is true of family structures, with assignments of who does what, and who is expected to do what. Many people believe those institutions are historically designed to maintain inequalities between genders in our society. Clearly, institutions serve various functions.

Past World Development Reports have included brief discussions of institutions; they have borrowed from the work of North and others, to use words like rules, norms, policies, and conventions to help define institutions. But these definitions are not embracive enough: organizations are clearly institutions. We speak of "good" organizations as good institutions. These different perspectives may not conflict, because you can say that what makes a good organization is good norms, people, policies, and rules. Stanford is a good university because it has good faculty members, and that reflects how personnel are chosen and how they interact with one another—selection and management policies. It may make sense to go back to norms, policies, and rules, including selection policies, and say that those define the organizations and their components.

Assessing a Good Institution

A decade ago when it was said that developing countries ought to have good policies, the meaning of good policies was taken for granted. Now the World Bank, the IMF, and other international institutions have taken a major step forward in saying that developing countries ought to have good institutions. One of the purposes of this forum is to decide what is meant by good institutions.

If I had to forecast the message of the next World Development Report, I would say it will emphasize that successful development requires strong, effective institutions. The challenge, however, will be to go beyond this almost platitude, to describe what are good institutions and to prescribe how one goes about creating such institutions. This will be no mean task.

Fulfillment of a Designed Function

What is meant by good institutions? A first pass would be institutions that fulfill the function for which they are designed, without large negative ancillary effects

But this functional approach is, at best, problematic, because it is an engineering framework—as if someone had sat down to design an institution with specific characteristics and roles. Institutions often evolve without a designer, as was clear in the historical perspective with which I began. The United States designed a constitution and a government, but they evolved afterward. In some cases, it is clear what the function of the institution is: Jim Crow and other discriminatory institutions are designed to maintain a certain set of power relationships. But in many cases, it is difficult to determine what particular institutions were designed for. Good institutions may serve functions beyond those for which they were designed. Against which function should their performance be assessed?

An example of an institution with which all of us are familiar may help illustrate what is at issue. Is the IMF a "good" institution? It could be said the IMF had an architect—John Maynard Keynes. He had a clear set of functions in mind: to promote global economic stability, to help and encourage countries

maintain their economy at full employment. There is near unanimity that the IMF is not only not doing what it was designed to do, but, on the contrary, it has push countries into recessions and encouraged capital market liberalization which has contributed to global economic instability. Surely, from the functional perspective—viewed from the perspective of what it was designed for—it surely fails the criteria of a "good" institution. Note that this functional approach simply values the institution in terms of performance; it does not look at the quality of the people that it hires, the fairness of the rules by which it operates, and so forth. These are relevant only so far as they affect the function.

While viewed from the perspective of the functions for which it was originally designed, the IMF must surely be judged a poor institution, it might be more appropriate to judge it from the perspective of the functions that it actually is performing, or that it has come to assume. But even here, matters are not so simple. Again, consider the IMF: it has taken on the role of promoting development and facilitating the transition of the former Communist countries to a market economy. Neither task is going well. If an organization is designed to perform a function that cannot be performed—for example, the impossible task of ensuring the stability of financial global markets—exchange rates will be volatile, and volatility in exchange rates will cause volatility in economic markets—should it be castigated as a failed institution. Or should we recast its function: it is designed not to ensure global stability, but to reduce instability, and so long as it reduces instability as much as can be reasonably expected, it is a "good" institution.

But what are we to think of an institution that takes on a task that are beyond its capabilities? That has sought out a task, and in persuading others to assign it that task, and claimed more than it can deliver? Surely, a quality of a good institution is to know its own limits and abilities, and one should, accordingly, be less forgiving. There is now, again, a widespread judgment that in terms of the functions—promoting development and facilitating transition—the IMF is again a "poor" institution, and that these functions should be removed from the portfolio of its responsibilities (see, e.g. the unanimous recommendations of the U.S. Congressionally appointed Commission on the international economic institutions, commonly referred to as the Meltzer Commission, the report of the Overseas Development Council, and the report of the Council for Foreign Relations.)

Even an institution which is, from several different perspectives, dysfunctional, can be highly effective from other perspectives. There may be hidden agenda. It may be performing socially important functions that are not fully recognized or, in any case, well articulated. Again, to return to our example of the IMF: From the point of view of financial interests in the G7 countries, it may be a good institution because it serves their interests, it helps their banks get repaid in times of crises, even if it does not serve global economic interests. So the IMF may be a very good institution if we take its function as enhancing the likelihood that creditors be repaid.

To be sure, we can (and should) ask, how did the institution evolve? How did it come to change its function? It may be the governing structure that is responsible for this changed orientation: a change that is not surprising given the central bankers and finance ministers accountable to the financial markets who are running the institution instead of people who care about the welfare of workers and small businesses.

There are further problems in trying to assess an institution, and how well it is performing. Today, we judge institutions not only by how well they perform the functions for which they are designed, that is, in terms of outcomes, but also in terms of process. This is especially true for public institutions. There are some easy elements to an assessment: Public institutions which are corrupt, which give special favors to those who bribe the appropriate officials, are not "good" institutions. Public institutions which operate in a non-transparent manner, which do not give voice to all the affected parties, are not "good" institutions. To be sure, the distinction between process and outcome for public institutions may be somewhat artificial. After all, part of the function of public institutions in a democracy is to promote democratic participation, and institutions which operate secretly might be thought of as thus failing in their democratic function. But what if there are trade-offs, real or perceived, between democratic processes and outcomes. Central Banks often believe that secrecy enhances their effectiveness, though there is little evidence to that effect; as recent experience in the U.K. has shown, there can be marked increases in openness without any evidence of adverse effects. Central Banks seldom are inclusive in their deliberations, i.e. seldom are the voices of workers, who may be strongly affected by their decisions, heard. Central Banks typically argue for greater independence from direct democratic accountability, though there is little evidence that such independence leads to better performance from the perspective of the real side of the economy, i.e. output, growth, employment, or productivity. Similar concerns have been raised concerning the IMF: lack of transparency in process, lack of inclusiveness in deliberations, lack of democratic accountability. From the perspective of public institutions in democratic societies, should these institutions be viewed as "poor" institutions, regardless of their effectiveness in outcome measures? How are we to assess the trade-offs, if there are trade-offs? (To be sure, in the case of the IMF, one of the allegations is that the failures in process have contributed to the failures in performance.)

Distribution

Thus, the challenges in assessing whether a particular institution is a "good" institution are formidable: But as if matters are not difficult enough as I have presented them so far, there is one more set of concerns I want to raise. How can distributional concerns be brought into the assessment? Are institutions that are designed to preserve the interests of the elite good institutions as long as they work? The institution of Jim Crow worked for a hundred years to preserve discrimination—it had a well-defined function and served that function well. No one would deny that it had strong distributional consequences. From the perspective of the people who put it in place, it was a good institution, while people with other ethical views on discrimination would say it was a bad institution. I have not mentioned the economic theory of how that kind of institution can be sustained, but there is no a large body of interesting theoretical work which shows that such institutions—designed not to promote efficiency but to preserve "power" relationships-- can be sustained as a Nash equilibrium. There is no general presumption that the equilibria that emerge and are sustainable have any welfare properties in the conventional sense.

Endogenous Preferences

Standard economic theory takes preferences as given. Given those preferences, economic systems—including institutional arrangements—are evaluated in terms of how well off they make individuals. But many institutions and institutional arrangements can be viewed as having been designed to change preferences. How are we to assess such institutions? Standard theory gives us little guidance. For instance, educational institutions pretend to (and in many cases actually do) improve the efficiency with which people process information; they enhance human capital. But in our societies, they have a perhaps even more important function: they really perform the missionary function of trying to change beliefs and perspectives. This is true too of the international financial institutions in their interactions with developing countries. Corporations, too, try to change the beliefs of those who work for them. Simon (1991) emphasized that trying to change what they call an individual's sense of identification (in effect, changing preferences so that workers view their interests as coinciding with that of the firm) is a far more effective control device in organizational design than trying to change incentive structures.

Measurement Problems

I have discussed some of the theoretical problems with trying to define what is a good institution, and assess whether a particular institution is a good one. One of the hallmarks of the World Development Reports is that they go beyond theory, to look at the data. But here, the problems you will be confronting are even greater. If institutions are assessed only on the basis of outcomes, statements such as "good institutions lead to better performance" become little more than tautologies. That is how we define good institutions. The task ahead is to try to identify characteristics of organizations which are systematically associated with better performance. Many of the relevant characteristics will be hard to quantify. In empirical economics, there is a strong proclivity to measure things and, if they cannot be measured, to assume that they must not be real. But many of the most important characteristics of institutions may be hard to quantify.

Having said this, there are some circumstances in which we can assess quantitatively the value of an institution. The value of most corporations exceeds by a considerable amount the value of the physical assets within it. This is referred to as the good will of the corporation. It is a quantitative measure of the value of the institution. It is simply the difference between the value of the corporation and the value of the underlying assets. Within a broader societal contexts, we have come to think of social capital in an analogous way. This conference will discuss the relationship between social capital and good institutions. Is social capital simply the quantitative measure of the value added by the rules and norms of an institution and the willingness of those in (or that interact with) the institution to adhere to those rules and norms? But what about rules and norms which serve to enhance the welfare of some individuals at the expense of others? Or which are dysfunctional, so that the resulting equilibrium is Pareto inefficient? Is corruption (thought of as a set of rules and norms that guide how individuals relate to each other) an institutional feature of certain institutions that can be thought of as negative social capital?

The General Equilibrium Context

I have raised the issues of how to think about institutions, what is meant by institutions in general and good institutions in particular, and what quandaries, puzzles, and conundrums are involved in institutional analysis. Institutions cannot really be assessed—and good institutions cannot really be defined—outside a general equilibrium context. This makes the task particularly difficult, because it is already hard enough to figure out what is meant by good institutions when they are examined in isolation. But that is not enough: institutions should be viewed in a broader context of general equilibrium models.

In assessing the institutional arrangements of a society, one has to evaluate not only the institutions that exist, but those that are absent. That will not be known without a more systemic point of view, a look at

the whole of society.

The second reason to look at institutions in the context of general equilibrium models is that only that framework can reveal the institutional externalities. Consider two examples. One reason institutions such as Harvard or Stanford are considered good institutions is that they attract bright people, who, within those institutions, prove good theorems, and do good work. They bring together bright people. But brightness is a characteristic of the individual, and those people would be bright no matter where they worked. What these universities have done is to "redistribute" the bright people, to congregate together individuals who might have been teaching at a multitude of other universities. That reallocation has not necessarily increased the productivity of society—it has simply located it in one place. The gains at Harvard and Stanford are, or at least may be, losses to another institution. From a general equilibrium point of view, these institutions may not have really contributed anything. They can only claim to be good institutions if there are complementarities among the good people and the whole is greater than the sum of the parts.

Consider another institution, the family, from the narrow point of view of risk sharing. In most families that I know, when one spouse gets sick, the other continues to transfer income so that the spouse will not starve. This form of mutual insurance is one aspect of the family, though certainly not the only one. If the family is a risk-sharing mechanism, is it welfare enhancing? It may be, because with imperfect information comes the moral hazard problem, and with the moral hazard problem comes imperfect insurance. With perfect insurance markets the family would not be necessary, at least for this risk-sharing function. The husband and wife would have separate insurance policies and be paid by the insurance company in case of illness. But moral hazard is real, and people have partial insurance. A naive economist such as the early North might say that the family was created to solve that problem and is welfare enhancing.

Viewing the family from this narrow risk-sharing economic perspective, we can whether the family improves social welfare. Because the market insurance company knows that the husband and wife insure each other, the incentives to prevent the insured against risks are reduced, and the likelihood of an accident is increased. The insurance company responds by reducing the amount of insurance. Nonmarket insurance, namely the family, crowds out the market insurance. But the family is less efficient than the market at risk sharing, because the market can pool the risk among a large number of people, not just a few, as in a family. Thus it can be shown that, were it possible, forbidding people to share risk within the family could improve social welfare. This is a highly stylized example, but it helps illustrate a general point, which I noted earlier: the fact that an institution seems to be performing a function—that is seems to be filling a gap in the market—does not mean that the institution is welfare enhancing. That can only be ascertained within a general equilibrium model. The result that the institutional equilibrium may not be Pareto efficient should not be surprising to most economists not indoctrinated in to the ideology that free markets must be efficient. It is now recognized that the Nash equilibrium in market economies is almost never Pareto efficient, so long as information is imperfect or markets are incomplete (and information is always imperfect and there are always missing markets.) (See Greenwald and Stiglitz, 1986). Adam Smith's invisible hand theorem holds only under highly restrictive conditions, never satisfied in developing countries. That an institution that arises to address a market failure need not be welfare enhancing should suffice to show that the early North view of institutional evolution was wrong. There was enough history to show that it was wrong, even for a historian. Clearly, evolutionary processes cannot be relied on to ensure that institutions are welfare enhancing, because there can be inefficient equilibria in institutions.

Before turning to a closer look at these evolutionary processes, there is one other attribute of the "family" risk sharing model just described that deserves comment. The problem in the situation just depicted is that the insurance provided by the family led to more risk taking. There was no "shared utility," no sense of "identification," no meaningful "trust." The family was nothing more than a risk sharing institution with no further bonds among the members. But the reason that there is risk sharing within the family, rather than between two randomly chosen members of society, is that there are bonds, so that each internalizes the externality that his risk-taking behavior imposes on the other. Institutional analysis thus break beyond the bounds of individualistic behavior into understanding social relations—a task with which many economists are uncomfortable. But the results of the analysis are highly dependent on such a social analysis. To see this, consider a slight variation of the previous model. To keep our analysis within the time-honored bounds of economics, assume that good behavior (that is, not undertaking large risks simply because of the insurance, which transfers risks to others). In that case, the Nash equilibrium which emerges is Pareto superior to that without the family. That behavior increases social welfare because the market free-rides off their monitoring activity.

Inefficient Evolution as an Explanation for the Persistence of Bad Institutions

Some of the early "institutionalists" saw the new institutional analysis as extending the reach of market

fundamentalism. The assumptions underlying the standard competitive equilibrium model, which ensured the efficiency of the economy, were clearly not satisfied in developing countries. Surely, no one could claim that there was a complete set of risk markets. Those advocating a stronger role for government suggested that this implied not only that the market equilibrium was Pareto inefficient, but that government actions (such as planning, if not production) were required. The early institutionalists said that in spite of the absence of certain key markets, the market equilibrium may still be efficient; government had at most a limited role; and that was because private institutions naturally evolved to fill in the holes. Market fundamentalism seemed to triumph again.

There was a certain intellectual incoherency in such views emanating from policy advisers: Good institutions should have evolved already, and advice should, accordingly, not be necessary. (To be fair, they may have thought of their role as simply trying to speed up the pace of institutional innovation).

The fact of the matter, however, is that it does appear possible for bad institutions to persist, for very long time periods. The discussion of the preceding section should make this no surprise: there may even be equilibrium in which bad (dysfunctional) institutions exist.

But beyond looking at equilibria, we can ask, do institutions evolve in a welfare-enhancing way? In general there are no presumptions about the efficiency of these evolutionary processes. As soon as institutions come into the picture, it is necessary to apply a more historical perspective, an evolutionary perspective.

There is no presumption, no theory concerning the efficiency of evolutionary processes, and indeed there are many examples of inefficient evolutionary processes. One is the adverse affect on biological evolutionary processes of the absence of perfect "capital" markets. I have a reputation for seeing imperfect capital markets everywhere, but it turns out that those imperfections do play a very important part in explaining problems in evolutionary biology. A species might over the long run be very efficient in converting food into biomass, but it cannot borrow against its future productivity to survive ten thousand years of the Ice Age. The result is inefficient specialization with too little diversity and too little broad-range adaptability. Such a species can only survive by borrowing against its future "income," and with imperfect capital markets, it cannot do that. The implication for the market economy should be obvious.

A multitude of examples of a possibly inefficient evolutionary processes comes from the work of Hoff and Stiglitz (2000), who emphasize the perspective that other participants in the market determine the performance, or success, of any behavior, a point that Darwin made forcefully in his discuss of the Galapagos islands. Contrast that view with the real business cycle, neoclassical perspective that says that the outcomes of the economy can be predicted by knowing the so-called deep parameters of technology and preferences. Darwin's observed the huge range of biological diversity in the Galapagos, despite the same deep fundamentals (though Darwin did not use those words) on the islands; on each island, the most important aspects of the environment were defined by the other species in the area. The evolutionary answer to the question why palm trees in tropical climates grow so tall is that they often grow in small oases where each competes for sunlight by growing taller than the others. They might get together and say, "Look, this is inefficient. Let's keep our height down, and then we can devote more of our resources to greenery and more fruit instead of using all our energy to try to outgrow each other." But as in an arms race, that kind of cooperative equilibrium is biologically efficient but unlikely to emerge, or emerge naturally or quickly. Many of the problems in the evolution of species arise precisely from that kind of problem. Agriculture research, the development of hybrids and of genetically modified plants, is based on the premise that direct intervention may not only speed up the process of evolution, but actually produce strands that would never have evolved on their own. The breakthrough in the miracle rices, for example, came from combining a high yielding variety with a dwarf stock that would not tip over.

There is not only not a presumption of efficiency, but strong reasons for inefficiency, at least in many different contexts. From an economist's point of view, change may involve problems of collective action and asymmetric information. One example is contracts. Many contracts do not provide the ideal incentives. In many cases even simple nonlinear contracts would improve on the standard contract, but changing a contract is very difficult. An interesting reason for this is a legal presumption, for which there is good reason based on equity and information asymmetry, that new contracts are always interpreted against the interests of those who write the new clauses. I might want to modify a standard contract to make it more efficient, but language is always ambiguous and subject to misinterpretation. In cases of dispute, where courts are will have to adjudicate the issue, the question is how the courts interpret the language. Because the courts' interpretation will be weighted against the person who has written the clause, there is enormous inertia for not changing, a kind of built in conservatism in the evolutionary process of the legal foundations underlying market economies.

Incentives for Institutional Change

From the perspective of policy makers, perhaps the most important issues are how to encourage the creation of good institutions and how to design institutions that can change as the environment changes. Both questions demand careful thought about organizational design and the sequencing of reforms. How can adaptable learning organizations be designed? What are the incentives to create those organizations? The world changes, and organizations have to change. Otherwise they become dysfunctional. The World Bank is more a learning organization than the IMF, for interesting reasons. One is that learning requires challenging existing practices, which may result in high costs of control. Other reasons have to do with the importance of competition and of decentralization as part of the organizational design. Competition is an important impetus for change. The IMF's dislike of the idea of competition in the marketplace of ideas and of a discussion of alternative perspectives is precisely why it has found it difficult to adapt. It would like to have—and in some areas does have—a monopoly in some of the spheres in which it operates.

Coalition formation and the incentives for institutional reform are interesting in this respect. A monopoly creates a vested interest against competition policies. John D. Rockefeller was not the major impetus for the adoption of antitrust policies in the United States at the turn of the nineteenth century, and Bill Gates has not been a major force for strengthening those laws at the turn of the twentieth. And turning over large amounts of Russia's wealth to 10 people has not created a strong constituency either for the rule of law or for strong competition policies in that country, despite Schleifer's (1998) argument to the contrary. Indeed, there is a growing body of historical evidence that the rule of law is in the interest of the middle class more than of a small oligarchy. Across the world, it has been economies with a strong middle-class, with a broader distribution of wealth, that have been in the forefront in the kinds of social change that have lead to a rule of law.

The dynamics of reform are also interesting and of critical importance—what forces for democratization and reform in the United Kingdom in the early 19th century led to the spread of education in the 19th and 20th centuries? Elites have a clear incentive to maintain their power, including through institutions and institutional arrangements designed not necessarily for promoting efficiency, but rather to maintain existing power and distributional relationships. The question is, what changes can be made today, which might be acceptable to those power elites, or which can be forced upon them—in spite of their power—which will, over time, lead to dynamic changes in society which will erode that power?

The question is: What drives the changes? How can they be initiated? The answer to these questions may in fact be central in the question for enhancing our understanding of how to help move the developing economies to create more efficient and more egalitarian societies. Sermons by themselves are not likely to succeed—they have to be based on a deeper understanding of the interventions that are likely to promote social change.

_____

References

Arrow, Kenneth. 1972. "Gifts and Exchanges". Philosophy and Public Affairs 1(4): 343-362.

Cheung, S. 1969. The Theory of Share Tenancy: With Special Application to Asian Agriculture and the First Phase of Taiwan Land Reform. Chicago: University of Chicago Press.

Darwin, Charles. 1859. The Origin of the Species by Mans of Natural Selection. Modern Library Edition, 1993, New York: Random House.

Hoff, K. and Joseph E. Stiglitz. 2001. "Modern Economic Theory and Development." In G. Meier and J. E. Stiglitz, eds., Frontiers of Development Economics. Oxford University Press.

North, Douglass C. 1971. "Institutional Change and Economic Growth," Journal of Economic History 31(1): 118-25.

———. 1990. Institutions, Institutional Change, and Economic Performance. Cambridge, U.K.: Cambridge University Press.

———. 1994. "Economic Performance Through Time," American Economic Review 84(3): 359-68.

North, Douglass C. and Robert Paul Thomas. 1973. The Rise of the Western World. Cambridge, U.K.: Cambridge University Press.

Shleifer, Andrei and Robert W. Vishny. 1998. The Grabbing Hand: Government Pathologies and their Cures.

Cambridge, Mass.: Harvard University Press.
Simon, Herbert A. 1991. “Organizations and Markets,” Journal of Economic Perspectives 5(2): 25-44.