# KNOWLEDGE AS A GLOBAL PUBLIC GOOD

JOSEPH E. STIGLITZ

Thomas Jefferson, the third president of the United States, described knowledge in the following way: "he who receives an idea from me, receives instruction himself without lessening mine; as he who lights his taper at mine, receives light without darkening me". In doing so, Jefferson anticipated the modern concept of a public good. Today we recognize that knowledge is not only a public good but also a global or international public good. We have also come to recognize that knowledge is central to successful development. The international community, through institutions like the World Bank, has a collective responsibility for the creation and dissemination of one global public good—knowledge for development.

This chapter reviews the concept of global public goods, explains the sense in which knowledge is a public good and explores the implications for international public policy that derive from the fact that knowledge is a global public good. In particular, I emphasize the role of knowledge for development, articulated forcefully in the *World Development Report 1998/99* (World Bank 1998b), and the consequences thereof.

## BASIC CONCEPTS

This chapter combines two concepts developed over the past quarter century: the concept of global public goods and the notion of knowledge as a global public good.[1]

A public good has two critical properties: nonrivalrous consumption—the consumption of one individual does not detract from that of another—and nonexcludability—it is difficult if not impossible to exclude an individual from enjoying the good. Knowledge of a mathematical theorem clearly satisfies both attributes: if I teach you the theorem, I continue to enjoy the knowledge of the theorem at the same time that you do. By the same token, once I publish the theorem, anyone can enjoy the theorem. No one can be excluded.







They can use the theorem as the basis of their own further research. The "ideas" contained in the theorem may even stimulate others to have an idea with large commercial value.

### *Nonrivalrousness*

The fact that knowledge is nonrivalrous—that there is zero marginal cost from an additional individual enjoying the benefits of the knowledge—has a strong implication. Even if one could exclude someone from enjoying the benefits of knowledge, it would be undesirable to do so because there is no marginal cost to sharing its benefits. If information is to be efficiently utilized, it cannot be privately provided because efficiency implies charging a price of zero—the marginal cost of another individual enjoying the knowledge. However, at zero price only knowledge that can be produced at zero cost will be produced.

To be sure, to acquire and use knowledge, individuals may have to expend resources—just as they might have to expend resources to retrieve water from a public lake. That there may be significant costs associated with transmission of knowledge does not in any way affect the public good nature of knowledge itself: private providers can provide the "transmission" for a charge reflecting the marginal cost of transmission while at the same time the good itself remains free.

### *Nonexcludability*

While its nonrivalrous property says that no one should be excluded from the enjoyment of a public good (since the marginal cost of benefiting from it is zero), nonexcludability implies that no one *can* be excluded. This too has important implications: it means that knowledge cannot be provided privately. Assume that someone produced a theorem. Assume that the theorem is valuable in providing insights into how to solve practical problems. But assume also that the theorem cannot be kept secret and must be immediately available. Then, since anyone can immediately enjoy the theorem, the individual could make no profit from it. Competition would drive its price to zero. At any positive price, it would pay someone to get the information (which by assumption he could do) and undercut the seller.

Some forms of knowledge are (or can be made) excludable. For example, in some industries, such as metallurgy, trade secrets are used. To be sure, firms relying on secrets run a risk: a competitor, observing a new alloy, could analyse its composition and infer the mix of metals (and with modern techniques,

even the relative proportion of the atoms). The firm might have a hard time inferring precisely how the alloy is made, but there is no way that rivals can be excluded from knowledge of the chemical composition and the properties of the alloy. By the same reasoning, when a firm discovers that consumers love, say, yoghurt, others cannot be excluded from using that knowledge to put their own yoghurts on the market.

Patents provide the exclusive right to inventors to enjoy the fruits of their innovative activity over a limited period (17 years in the United States). In return, inventors must disclose the details of their invention. But the fact of the invention, let alone the details provided in the patent application, make an enormous amount of knowledge freely available. The development of rayon provided other researchers with important information: it demonstrated the feasibility of a synthetic fibre—knowledge that was of enormous commercial value and that enhanced incentives for others to look for other synthetic fibres. Indeed, research in chemicals often consists of looking for slight variations of the original chemical. It is precisely because of the high value of the knowledge disclosed through the patent process (and the limited duration of the patent) that some firms decide not to patent their inventions and to rely on trade secrecy—even though this may seem to offer less protection at first sight.

But because the returns to some forms of knowledge can to some extent be appropriated (there is some degree of nonexcludability), knowledge is often thought of as an *impure* public good.

### *Global public goods*

Shortly after Samuelson (1954) articulated the general theory of pure public goods, it became recognized that the benefits of some public goods were limited geographically. These were called local public goods (see Tiebout 1956 and Stiglitz 1977, 1983). Of course, the public goods earlier theory had focused on—such as national defence—were also limited geographically to a particular country. At the same time, there are several public goods that are not so limited—the benefits of which accrue to everyone in the world. In Stiglitz (1995) I identify five such global public goods: international economic stability, international security (political stability), the international environment, international humanitarian assistance and knowledge.

Most knowledge is a global public good: a mathematical theorem is as true in Russia as it is in the United States, in Africa as it is in Australia. To be sure, some kinds of knowledge are of value only or mostly to those living in a






certain country—for example, knowledge particular to a country's institutions, weather or even geography. But scientific truths—from mathematical theorems to the laws of physics and chemistry—are universal in nature. The problems that economics deals with, such as scarcity, are ubiquitous, and accordingly the laws of economics are universally applicable, even if idiosyncratic institutions exist within each country.

### *The role of the state*

The central public policy implication of public goods is that the state must play some role in the provision of such goods; otherwise they will be undersupplied. If firms cannot appropriate the returns to producing knowledge, then they will have limited incentive to produce it: in deciding how much to invest, they will look only at the return that they acquire, not the benefits that accrue to others. The benefits that have accrued from the development of the transistor, the laser or the mathematical algorithms that underlay the modern computer have been enormous, extending well beyond the benefits accruing to those who made or financed these innovations and discoveries.

Governments have pursued two strategies in addressing these concerns. The first is to increase the degree of appropriability of the returns to knowledge, by issuing patents and copyright protection. In doing so, governments are engaged in a careful balancing act: after all, one of the basic properties of knowledge as a public good is that the marginal cost of usage is zero (nonrivalrous consumption). Inventors obtain a return on their innovative activity either by charging through the use of a patent (licensing) or by charging a monopoly price on the product. In either case there is an inefficiency. The gain in *dynamic* efficiency from the greater innovative activity is intended to balance out the losses from *static inefficiency* from the underutilization of the knowledge or from the underproduction of the good protected by the patent.

One part of the balancing act is to limit the duration of the patent. A very short patent life would imply a low level of appropriability—such that the limited returns to innovative activity would imply low levels of innovation. A very long patent life would mean large losses in static efficiency; most of the fruits of the innovation would accrue to the innovator, with little passed on to consumers (say, in the form of lower prices) because the inventor would never be subjected to competitive pressure. Patents typically last for 17 years, and in many cases by the time a patent has expired its value is limited because new products and innovations have superseded it. This is not the case, however, for many drugs (partly because there may be a long testing period before the drug is actually marketed.)

But other aspects of the patent system play an important role in how the dynamic efficiencies are balanced with the static inefficiencies: the breadth and scope of a patent claim (whether a patent for a new genetically altered tomato covers all genetically altered vegetables, all genetically altered tomatoes or only this particular genetically altered variety) can have profound implications.

Initial knowledge is a key input into the production of further knowledge, and thus the design of the patent system can dramatically affect the overall pace of innovation. An excessively broad patent system (for example, with long-lived patents of broad scope) can raise the price of one of the most vital inputs into the innovative process and thus reduce the pace of follow-on innovations, even as it may provide returns to those making the original innovation. As a result the overall pace of technical progress may be slowed.[2] Worries about the adverse effects of excessively strong intellectual property protection have been brought home by the US government's recent antitrust suit against Microsoft, which has allegedly attempted to leverage the power associated with its control of the dominant personal computer operating system (itself a consequence of important network externalities that result in huge advantages associated with the establishment of an industry standard; Katz and Shapiro 1985) to a broader dominance in application software. Many industry experts believe that in doing so, Microsoft may have retarded the pace of innovation in the computer industry.[3]

These concerns are of particular importance to developing countries. Innovations (research and development expenditures) are even more concentrated in advanced industrial countries than are incomes (figure 1), and many of the advances in less developed countries consist of adapting the technologies of more advanced countries to the circumstances of the developing world.

The second strategy for dealing with the appropriability problem entails direct government support. If government could costlessly raise revenues for financing the support and if government were effective in discriminating between good and bad research projects, clearly this strategy would dominate that of enhancing intellectual property rights, for the latter strategy entails static distortions (the monopoly prices associated with patent rights result in prices exceeding marginal costs) and the inefficient utilization of knowledge. The static distortions can be thought of as a tax used to finance the research and development; the tax, however, is not an optimal tax.[4] But the patent system provides an effective self-selection mechanism: those who are convinced







that they have a good idea invest their own money and the money of those whom they can persuade of the attractiveness of their idea. Such selection mechanisms may not only be more effective than, say, government bureaucrats attempting to assess various applications, but the costs of mistakes are borne by those making the misjudgement, not by the public at large. Thus the system provides strong incentives for individuals to engage in due diligence when assessing the merits of alternative research proposals. It is because of

FIGURE 1

**GDP and research and development expenditures by region, 1987**

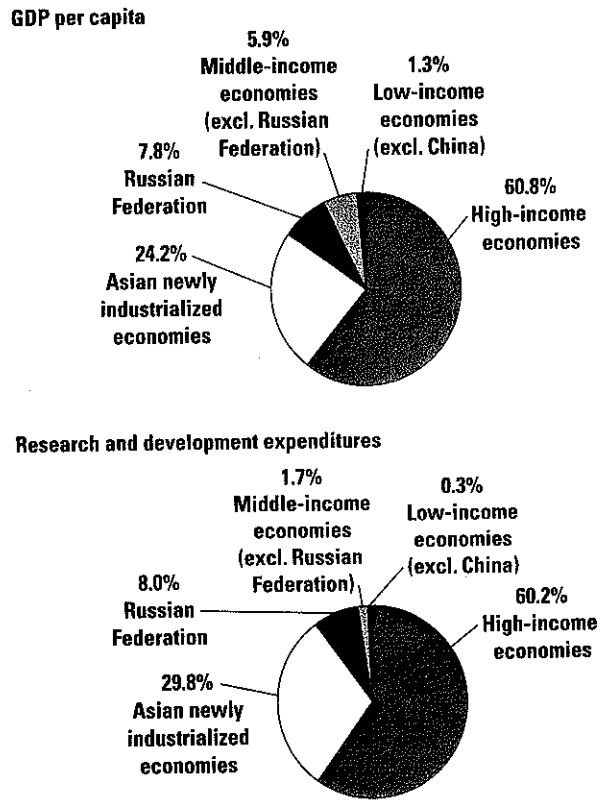


*Source:* World Bank 1998b.

these strong incentive and selection properties that most economists believe that for a wide range of areas, the strategy of enhancing intellectual property rights is preferable to that of government subsidization.

But there are some important situations where the costs of the improved appropriability strategy are high. This is particularly the case for basic research because its benefits are widespread and diffuse, and because attempts to appropriate its returns may significantly slow the overall pace of innovation. Indeed, many advances in basic knowledge—such as mathematical theorems—are not patentable despite their importance and their potential practical applications.

This discussion should have made clear one central point: the concept of intellectual property—the breadth, scope and applicability of patent protection—is not just a technical matter. There are judgement calls and trade-offs, with different people and different countries all affected differently by alternative decisions. There are conflicts of interests between developed countries and less developed countries. But unfortunately, many of the key issues cannot even be summarized by a set of simply stated principles. In practice, decisions are made on a case-by-case approach.

For instance, two key issues in the granting of a patent are scope and novelty. Does the first person to develop a genetically engineered tomato get a patent for that specific tomato, for all genetically engineered tomatoes or for all genetically engineered vegetables? Is the idea of a genetically engineered plant sufficiently obvious that it is simply not patentable, with only specific and nonobvious genetic engineering processes being patentable? The consequences of answering these questions in different ways are enormous, as in the case of the automobile patent. In the early days of the automobile a lawyer-inventor named Selden received a patent for a horseless, self-propelled carriage. He attempted to use the patent not only to extract a royalty but also to enforce an industry cartel. Had Selden succeeded, he would have suppressed innovations such as those of Henry Ford, who subsequently attempted to provide a low-cost automobile. While most of the industry was willing to go along with Selden (because he offered the prospect of an industry cartel, which would raise their profits), Ford challenged the patent and won. Had he lost, there could have been a long delay before cars became a method of mass transportation.

The stance sometimes taken by producers of knowledge, that we need "strong" intellectual property rights, masks this underlying debate. Strong, in this context, becomes equivalent to "good", with the implication that the







"stronger" the better. But I hope this discussion has brought out that the issues are far more complicated. Stronger, in the sense of "tighter" protection, could not only have large distributive consequences (between, say, developed countries and less developed countries), but also large efficiency consequences, with the pace of innovation actually impeded and living standards in less developed countries diminished.

Some industrial countries have effective competition policies that mitigate the risks that result from the abuse of monopoly power associated with a patent. But most countries do not have effective antitrust policies. For instance, drug companies can, and have an incentive to, act like discriminating monopolists, charging higher prices where the consumer surplus is higher or where they can extract more of the consumer surplus. Some European countries have policies that offset these monopolistic powers: given the large role of government in health care, they can effectively exercise their monopsonistic powers. Thus it is conceivable (and there are anecdotes supporting this possibility) that consumers in less developed countries could be charged higher prices for drugs than consumers in far richer countries. (In doing so, the consumers in less developed countries are in effect paying the fixed cost of research; consumers in more developed countries are partial free riders.)

Within the United States such price discrimination (not fully justified by differences in transactions costs) would probably be illegal. But there is no international competition policy that protects the poor country. Well-designed (but not excessively strong) intellectual property regimes can provide some protection. It is not clear the extent to which effective competition policies *within a country* might provide safeguards: presumably a country could pass a "most favoured nation" provision—no firm, enjoying the benefit of intellectual property protection, could charge the consumers of that country a higher price than the price charged for the same good elsewhere in the world.[5] (Today there is concern within the United States that Microsoft is using intellectual property protection in ways that thwart innovation, making it difficult for small, rival software firms to enter the market. The recent federal court case has presented allegations of a variety of anticompetitive practices—practices that stifle new entrants, sometimes with superior products.)

There are other issues in the design of an intellectual property regime. Every innovation makes use of previously accumulated knowledge—it draws on the global commons of pre-existing knowledge. How much of the returns to the innovation should be credited to this use of the global commons?

Current practice says zero—because it is a commons, there is no price. But this is not the way things need be. In many parts of the world there is a recognition that charges can and should be imposed for the use of commons (whether they are forests, grazing lands or fisheries). Such charges can be justified on both efficiency and equity grounds. The international community could similarly claim the right to charge for the use of the global knowledge commons. Because knowledge is a public good, the argument for charging a fee is largely based on an equity rationale. However, by recycling funds to support further research, an efficiency argument could also be developed. There are obvious practical problems in the implementation of such a scheme: what fraction of the returns to the innovation is due to the use of the global commons? But even a rough rule of thumb, in which a certain fraction of the returns to innovations is used to finance a "replenishment" of the global knowledge commons, might be an improvement.

This issue of the use of the global knowledge commons has been brought home forcefully in the context of biodiversity, where private firms have prospected for valuable drugs in natural settings. In many cases local people have long recognized the value of these local drugs, though they have not identified the particular chemicals in the plants that give the desired effects.

The contrast could not be more stark between the way this unpatented knowledge is treated and the way adaptations of innovations in developing countries of patented ideas from developed countries are treated. In the first case all of the return is credited to the "discoverer", with none to the pre-existing knowledge. In the second case the patent holder is allowed to act as a perfectly discriminating monopolist, regardless of the extent to which his or her innovation built on pre-existing knowledge.

The effective use of knowledge developed in industrial countries typically involves substantial elements of adaptation—combining global and local knowledge. Yet the intellectual property regime, as it has been evolving, assigns most of the bargaining power associated with how the fruits of these combinations are shared to the developed country, especially in larger developing countries, where there may be effective competition for the use of the patented idea.

An international intellectual property regime, designed to facilitate the production and use of the global public good—knowledge—in a way that sustains high rates of growth and is consistent with broad notions of equity, must balance a variety of subtle concerns, including dynamic and static efficiency and the use of the global knowledge commons.







## Combining local and global knowledge

As I have just noted, a key part of successful development is combining global knowledge with local knowledge. The intellectual property regime affects how the gains are shared, and in doing so affects the pace of development within less developed countries. But many other aspects of the "knowledge infrastructure" within less developed countries can affect the pace of development and the extent to which developing countries can avail themselves of the fruits of the global public good of knowledge.

Perhaps the most important is education. The Republic of Korea and other newly industrialized countries that have closed the knowledge gap between themselves and more advanced industrial countries invested heavily in secondary and tertiary education, especially in science and technology. Poor developing countries have rightly stressed the importance of primary education, for primary education is the base of the entire education system. Even primary education can have a large impact on the pace at which innovations in agriculture or better fertility and health practices are spread. But a significant closing of the knowledge gap requires more than a strong primary education system.

In the past, some poor countries have been rightly criticized for investing too much in higher education, the benefits of which go to a small elite. But the criticism has been misinterpreted. The issue is not the importance of higher education. The criticism is of what is taught, the quality of the education and how it is financed. Science and technology are vital. They must be taught at international standards—otherwise the instruction does little good in closing the knowledge gap and it would be better to send students to study abroad. And the students should be made to bear as much of the costs as possible, if not now then later, by repaying student loans.[6]

Governments in newly industrialized countries often played other important roles in facilitating the transfer of knowledge. For example, they established standards laboratories to attain the kinds of international standards required for participation in global markets for high-technology commodities. Some countries not only showed an openness to foreign direct investment but also actively recruited those forms most likely to have knowledge spillovers and designed employment and other programmes to enhance the likelihood of such spillovers. Licensing policies also played a role in the transfer of knowledge.[7]

As essential as the adaptation and creation of new knowledge within a country is the dissemination of knowledge throughout a country. The

movement of ideas within a country is affected by the effectiveness of its communications system. Recent advances in telecommunications have brought the costs of communication down tremendously and made possible the development of communications networks in parts of the world where it would have been decades, at best, before such systems would have been developed with older technologies. These new technologies mean that there is no longer a natural monopoly on communications, and by using competitive, market forces, access can be enhanced and prices lowered.[8]

This communications revolution, at the same time that it has made great strides in facilitating communication within countries, has also enhanced the ability of less developed countries to tap into the global knowledge pool. The Internet is proving to be a tool of immense power in sharing knowledge. Today developing countries face both great risks and great opportunities. Internet growth has been fastest in the United States and, not surprisingly, slowest in less developed countries. The enhanced ability to share and acquire knowledge in industrial countries may widen the knowledge gap because less developed countries may become even more disadvantaged.

At the same time, less developed countries can tap into a larger knowledge pool than they ever had access to before. Today a child anywhere in the world who has Internet access has access to more knowledge than a child in the best schools of industrial countries did a quarter century ago. He or she is no longer isolated. It is too soon to see how these contrasting forces will play out—whether the knowledge gap will be widened or narrowed. But it is clear that it is incumbent upon less developed countries to do everything they can to enhance their ability to tap into the reservoir of global knowledge.

Creating the knowledge infrastructure entails learning how to learn[9]—that is, creating the capacity to close the knowledge gap, an essential part of a successful development strategy.

## Knowledge for Development

Much of the knowledge that is required for successful development is not patentable; it is not the knowledge that underlies new products or new processes. Rather, it is equally fundamental knowledge: how to organize firms, how to organize societies, how to live healthier lives in ways that support the environment. It involves knowledge that affects fertility and knowledge about the design of economic policies that promote economic growth.






Those of us working in development institutions acquire much of this knowledge as a by-product of our general development activities. It is a form of learning by doing (see Arrow 1962). But knowledge for development goes beyond the collection of best practices and the accumulation of successful anecdotes and into analysis—why do certain policies and practices work in some circumstances and not others? Thus research is a central element of knowledge for development.

The ideas presented so far make clear that such knowledge is a global public good, and without active public support, there will be underprovision of this good. International institutions, including the World Bank and the United Nations Development Programme (UNDP) play a special role in the production and dissemination of this knowledge. We at the World Bank are increasingly thinking of ourselves as a Knowledge Bank,[10] and are organizing ourselves in ways that enhance our ability both to produce this knowledge and to disseminate it widely.

There is a natural complementarity between these new roles and the more traditional role of the World Bank in providing capital to less developed countries. Knowledge enhances the productivity of capital. Our research department's recent report on *Assessing Aid* shows that aid has a substantial impact on economic growth in countries that put into place good policies, while it has a negligible effect in countries that do not (World Bank 1998a). Knowing whether good policies are in place in developing countries and adapting World Bank lending programmes to reflect these realities is thus an important element of a successful lending programme.[11]

While we already know many elements of what makes for good policies, much needs to be learned. We need, for instance, to be able to better tailor policies to the different conditions and changing circumstances of individual countries. We have gradually come to recognize the adverse consequences of corruption, but we are only beginning to understand how to reduce corruption. While in the past we have focused, for instance, on the efficiency and equity aspects of tax structures; we are only beginning to pay attention to the susceptibility of different tax structures to corruption. Similarly, while there is widespread recognition of the advantages of privatization of certain public enterprises, we have only gradually come to recognize the problems that arise when privatization occurs prior to the establishment of effective regulatory or competition regimes. We have only slowly come to perceive the pervasiveness of corruption in the privatization process and the long-lasting adverse effects of that corruption. And we have all too late recognized that privatization prior

to the establishment of effective market institutions may not necessarily lead to a vibrant market economy—because the incentives provided by privatization may be directed more towards stripping assets than creating wealth. It should be clear that this kind of knowledge is essential to all of the World Bank's lending programmes, in its project lending, in its sectoral lending and, perhaps most important, in its adjustment lending. More broadly, knowledge, aid and private capital work together in a successful development programme; they are complementary.[12]

But there is more. We have increasingly realized that isolated projects will have only limited effects in the transformation of societies that we call development. We have to go beyond projects, and we have to scale up projects. An essential aspect of this strategy is the design of projects from which we can learn, from which we can garner knowledge and that can constitute the basis of economy-wide transformations (see Wolfensohn 1998 and Stiglitz 1998).

## Conclusion

The concept of global public goods is a powerful one. It helps us think through the special responsibilities of the international community. National public goods provide one of the central rationales for national collective action and for the role of government. Efficiency requires public provision and, to avoid the free-rider problem, the provision must be supported by compulsory taxation (see Stiglitz 1989). Similarly, global public goods provide a central rationale for international collective action. But today governance at the international level entails voluntary, cooperative actions. These include agreements to support an international property regime that facilitates the private production of certain kinds of knowledge. (We have raised questions of whether the current regime adequately reflects the broad interests of the international community, balancing equity and efficiency concerns among the affected parties.) But basic research and many other fundamentals forms of knowledge are not, and almost certainly should not be, protected by an intellectual property regime. In these areas efficiency requires public support. And this public support must be at the global level.

I have argued that knowledge is one of the keys to development and that knowledge is complementary to private and public capital. Knowledge is a global public good requiring public support at the global level. Current arrangements can be made to work effectively, but if they are to succeed we must be aware of the dangers and pitfalls. Some countries may try to free ride







on others; they may try to capture more of the returns that are available from the use of the global knowledge commons; they may see their self-interest enhanced more by taking out of the global knowledge commons than contributing to it, in supporting research to design patentable applications rather in supporting basic research.

The efficient production and equitable use of global knowledge require collective action. The challenge facing the international community is whether we can make our current system of voluntary, cooperative governance work in the collective interests of all.[13]

## Notes

The views presented here are solely those of the author and not those of any institution with which he is or has been affiliated.

1. See Stiglitz (1995) and US Council of Economic Advisers (1997). While the public good properties of knowledge had long been noted (Arrow 1962), early articulations of knowledge as a public good (in the sense defined by Samuelson 1954) include Stiglitz (1977) and Romer (1986). For an early textbook discussion, see Stiglitz (1986).

2. In theory, if the original innovator were a perfectly discriminating monopolist, such adverse effects might be limited because, it is alleged, he would never charge a fee for the usage of knowledge that would actually discourage a productive utilization (he would simply extract all of the users' producer surplus). But in practice there is not perfect discrimination, partly because the original innovator simply does not have the information required to be a perfectly discriminating monopolist. Moreover, competition in the product market is imperfect, and the innovator will discourage innovations that might result in the loss of some of his monopoly rents.

3. Aaron Edlin of the University of California at Berkeley (and a former staff economist at the US Council of Economic Advisors) has proposed an ingenious solution to spur innovation and limit the undue exercise of monopoly power: Microsoft would have to release its code, and the duration of its intellectual property protection would be limited to three years. If Microsoft continued to improve its product, the update versions of its software would be protected (for three years). Consumers would have a choice: they could avail themselves of the outdated (three-year-old) software or pay for the more advanced software. Microsoft would thus be forced to innovate at a fast pace to justify its dominant position in the market. Applications using the slightly outdated operating system would compete with those using the newer; and consumers would only be willing to pay for the new operating system if the improvements were worth the price.

4. According to the standards of optimal tax theory, which seeks to minimize deadweight losses. Moreover, the peculiar property of patents—imposing a high tax rate for a short period, followed by a zero tax rate—would (apart from the other considerations discussed in this section) appear to be far from optimal in terms of standard tax considerations. On the other hand, the tax is a "benefit" tax: those who enjoy the good pay the tax, and such taxes can be motivated by equity concerns.

5. This would, in a sense, be the opposite of antidumping laws, which stop firms from selling products at lower prices in international markets than they do domestically. While antidumping laws have the effect of hurting consumers at the same time that they protect producers, these "price gouging" laws would protect consumers.

6. One should note that to the extent that there are externalities associated with this education, there is an argument for public subsidies. The key question is, at the level of investment in education that maximizes a student's net present discounted value of income, is there a marginal externality—that is, is it desirable for government to encourage still further investment? Even without such marginal externalities, capital market imperfections provide a compelling argument for government intervention, but the interventions should be directed at reducing the impact of the imperfection.

7. This list is not meant to be exhaustive. For instance, some governments also created industrial and research parks, facilitating the exchange of ideas. Another important policy was the reduction of tariffs on intermediate goods, which allowed the import of essential inputs into more advanced technological processes.

8. Competition remains, however, far from perfect, so there is still an important role for an effective regulator. Chapter 2 of *World Development Report 1998/99* documents the success of countries that have used basic market competition with regulation (World Bank 1998b). Countries that have privatized without adopting a competitive framework have, at least in some cases, seen prices rise and access restricted: the private producer is more efficient in acting as a monopolist than the government was. In one instance the price of Internet access was raised to the point that a university could not afford to maintain connectivity. Thus the "reform" reduced the ability of those in the country to avail themselves of global knowledge.

9. I developed the concept of "learning to learn" and its implications for economic growth in Stiglitz (1987).

10. The concept of the Knowledge Bank was introduced in World Bank President James E. Wolfensohn's address to the 1996 Annual Meetings of the World Bank and International Monetary Fund (Wolfensohn 1996).

11. *Assessing Aid* points out that foreign aid is only significantly correlated with positive impacts in developing countries with sound economic policies and







institutions. In particular, in countries with sound economic management policies, 1% of GDP in aid leads to a sustained increase in growth of 0.5% and reduces poverty by 1%. In contrast, for those countries with poor economic environments, aid has no significant impact (the coefficient for growth as a result of aid inflow is actually negative, although not statistically different from zero).

12. Thus in countries that pursue good economic policies, aid "crowds in" private capital: $1 of aid helps bring in $2 of private capital. This helps explain aid's strong role in promoting economic growth. Similarly, the strong complementarity between knowledge and capital is one of the reasons that it is so difficult to parse out the extent to which growth is due to capital accumulation and the extent to which it is due to closing the knowledge gap. Improved knowledge stimulates higher investment, and the new investment embodies new technology. Without improvements in knowledge, East Asian countries presumably would have quickly experienced diminishing returns. As it was, they could maintain high rates of investment for an extended period without their incremental output-capital ratio falling. That is (only) one of the reasons that studies such as those of Young (1995), which purport to show that there was no East Asian miracle—that the region's growth can be explained entirely by investments, including investments in people—are so misleading. It was a miracle that these countries were able to maintain high returns with the levels of savings and investment—few if any other countries in the world succeeded in doing so. They did succeed in closing the knowledge gap, though to be sure, some of this knowledge was "purchased", like physical capital. For an alternative and more convincing interpretation (as well as a technical critique showing how sensitive Young's results are to the particular and unconvincing ways in which the variables entering the analysis are measured), see Klenow and Rodríguez-Clare (1997) and World Bank (1998b).

13. We can and should be more precise: since there are likely to be tradeoffs, with some arrangements providing an advantage to some groups relative to others, the two key questions are standard efficiency and equity issues. Can international arrangements lead to a reasonably high level of efficiency (that is, not too large an undersupply of the global public good knowledge and not too high a level of "static inefficiency" from restrictive utilization of knowledge) in ways that comport with basic notions of equity?

## References

Arrow, Kenneth J. 1962. "The Implications of Learning by Doing." *Review of Economic Studies* 29: 155–73.

Katz, Michael L., and Carl Shapiro. 1985. "Network Externalities, Competition, and Compatibility." *American Economic Review* 75: 424–40.

Klenow, Peter J., and Andres Rodríguez-Clare. 1997. "Economic Growth: A Review Essay." *Journal of Monetary Economics* 40: 597–617.

Romer, Paul M. 1986. "Increasing Returns and Long-Run Growth." *Journal of Political Economy* 94(5):1002–37.

Samuelson, Paul. 1954. "The Pure Theory of Public Expenditure." *Review of Economics and Statistics* 36: 387–89.

Stiglitz. Joseph E. 1977. "Theory of Local Public Goods." In Martin S. Feldstein and Robert P. Inman eds., *The Economics of Public Services.* New York: Halsted Press.

———. 1983. "Public Goods in Open Economies with Heterogeneous Individuals." In Jean-François Thisse and Henry G. Zoller, eds., *Locational Analysis of Public Facilities.* Amsterdam and New York: North-Holland.

———. 1986. *Economics of the Public Sector.* New York: W.W. Norton.

———. 1987. "Learning to Learn, Localized Learning and Technological Progress." In Partha Dasgupta and Paul Stoneman, eds., *Economic Policy and Technological Performance.* New York and Sydney: Cambridge University Press.

———. 1989. "On the Economic Role of the State." In Arnold Heertje, ed., *Economic Role of the State.* Oxford: Basil Blackwell.

———. 1995. "The Theory of International Public Goods and the Architecture of International Organizations." *United Nations Background Paper 7.* New York: United Nations, Department for Economic and Social Information and Policy Analysis.

———. 1998. "Towards a New Paradigm for Development: Strategies, Policies, and Processes." Given as Raul Prebisch Lecture at the United Nations Conference on Trade and Development (UNCTAD), 18 October, Geneva.

Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy* LXIV(October): 416–24.

US Council of Economic Advisers. 1997. *Economic Report of the President.* Washington, DC: US Government Printing Office.

Wolfensohn. James E. 1996. "1996 Annual Meetings Address." www.worldbank.org/html/extdr/extme/jdwams96.htm

———. 1998. "The Other Crisis: 1998 Annual Meetings Address." www.worldbank.org/html/extdr/am98/jdw-sp/am98-en.htm

World Bank. 1998a. *Assessing Aid: What Works, What Doesn't and Why.* A Policy Research Report. New York: Oxford University Press.

———. 1998b. *World Development Report 1998/99: Knowledge for Development.* New York: Oxford University Press.







Young, Alwyn. 1995. "The Tyranny of Numbers: Confronting the Statistical Realities of the East Asian Growth Experience." *Quarterly Journal of Economics* 110(3): 643–80.