KNOWLEDGE AS A GLOBAL PUBLIC GOOD

Joseph E. Stiglitz
Senior Vice President and Chief Economist
The World Bank[1]

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 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.

The purpose of this paper is to review the concept of global public goods, to explain the sense in which knowledge is a public good, and to explore the implications for *international* public policy that derive from the fact that knowledge is a global public good. In particular, I shall emphasize the role of knowledge for development, articulated forcefully in this year's World Development Report,[2] and the consequences that follow.

I. Basic Concepts

This paper combines two concepts developed over the past quarter of century: the concept of global public goods and the notion of knowledge as a global public good.[3]

A public good has two critical properties, non-rivalrous consumption--the consumption of one individual does not detract from that of another--and non-excludability--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.

Non-rivlarousness

[1] The views presented here are solely those of the author and not those of any institution with which he is or has been affiliated.
[2] World Bank (1998a).
[3] See Stiglitz (1995), Economic Report of the President (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 that of Stiglitz (1977) and Romer (1986). For an early textbook discussion, see Stiglitz (1986).

The fact that knowledge is non-rivalrous--there is a 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 are no marginal cost to sharing its benefits. If information is to be efficiently utilized, it cannot be privately provided as efficiency implies charging a price of zero—the marginal cost of another individual enjoying the knowledge. However, at zero price, only knowledge that could be produced at zero cost would 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 can remain free.

Excludability

While its non-rivalrous 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), non-excludability implies that no one *can* be excluded. This too has important implications: it means that knowledge cannot be provided privately. For assume someone produced, say, a theorem. Assume 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, the firms run a risk: a competitor, observing the new alloy, could analyze 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, yogurt, others cannot be excluded from using that knowledge to put on the market their own yogurts.

Patents provide the exclusive right to the inventor to enjoy the fruits of his innovative activity over a limited period of time (17 years), but in return, the inventor must disclose the details of his invention. 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 enormous information: it demonstrated the feasibility of a synthetic fiber - knowledge which itself was of enormous commercial value and which enhanced incentives for others to look for other synthetic fibers. 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 prefer the seemingly less protective route of trade secrecy.

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

Global public goods

Shortly after Samuelson 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.[4] Of course, the public goods earlier theory had focused on --such as national defense-- were also limited geographically to a particular country. At the same time there are several public goods which are not so limited--the benefits of which accrue to everyone in the world. In my earlier paper[5], I identified 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, there are some kinds of knowledge which are of value only or mostly to those living the country, e.g. knowledge particular to a country's institutions, weather, or even geography. But scientific truths--including many of the propositions of the social sciences-- are universal in nature. The problems with which economics deals, such as scarcity, are ubiquitous, and the "laws" of economics are accordingly 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 do so: in deciding how much to invest, they will only look 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 an enormous, extending well beyond benefits accruing to those who made or financed these innovations and discoveries.

Governments have pursued two different strategies in addressing these concerns. The first is to increase the degree of appropriability of the return 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 (non-rivalrous consumption). Inventors obtain a return on their innovative activity, either through 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, as a result of the greater innovative activity, is intended to balance out the losses from *static inefficiency*, from the under-utilization of the knowledge, or 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

[4] See Tiebout (1956) and Stiglitz (1977, 1983).

very long life to a patent 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, since the inventor would never be subjected to competitive pressure. Typically, patents last for 17 years, and as it turns out, in many cases, by the time the patent has expired, its value is limited, as 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 there are other aspects of the patent system that 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 obviously have profound implications.

Initial knowledge is a key input into the production of further knowledge, and the design of the patent system can thus affect dramatically the overall pace of innovation. An excessively broad patent system (e.g. 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.[6] Worries about the adverse effects of excessively strong intellectual property protection have been brought home by the recent anti-trust suit against Microsoft, which (it is alleged) has attempted to leverage the power associated with its control of the dominant operating system (itself a consequence of important network externalities which result in huge advantages associated with the establishment of an industry standard)[7] to a broader dominance in application software. Many industry experts believe that in doing so, the overall pace of innovation in the industry may have been retarded.[8]

These concerns are of particular importance to developing countries. New innovations (research and development expenditures) are even more concentrated in the advanced industrial countries than are incomes,[9] and many of the advances in less developed countries consist of adapting the technologies of the more advanced countries to the circumstances of the developing world.

---

[5] See Stiglitz (1995).

[6] In theory, if the original innovator were a perfectly discriminating monopolist, the adverse effects might be limited, since, 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 which might result in the loss of some of his own monopoly rents.

[7] Katz and Shapiro (1985).

[8] Aaron Edlin of the University of California at Berkeley (and a former staff economist at the Council of Economic Advisers) 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 the period of 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 continue to innovate at a fast pace, in order 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.

[9] See Figure 1.

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 results 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 R & D; but the tax is not an "optimal tax".[10] 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 misjudgment, not by the public at large. The system thus provides strong incentives for individuals to engage in due diligence in assessing the merits of alternative research proposals. It is because of these strong incentive/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 true in the case of 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, in spite of 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 judgment calls and trade-offs, with different people and different countries all affected differently by alternative decisions. There are conflicts of interests between the developed and the less developed countries. But unfortunately, many of the key issues cannot even by summarized by a set of simply stated principles; in practice, decisions are made on a case-by-case approach.

The stance sometimes taken by the 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 issues are far more implicated: stronger, in the sense of "tighter" protection, could not only have large distributive consequences (between say developed and less developed countries), but also large efficiency consequences, with the pace of innovation actually impeded and living standards in less developed countries diminished.

Within some of the advanced industrialized countries there are effective competition policies, which work to mitigate the risks that results from the abuse of monopoly power associated with a patent. But most countries do not have comparably effective anti-trust policies. For instance, drug companies can, and have an incentive to, act like discriminating monopolists, charging higher prices where the consumer surplus is higher, and/or where they can extract

[10] According to the standards of optimal tax theory, which seeks to minimizing dead weight 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.

more of the consumer surplus. Some European countries have policies that offset these monopolistic powers: given the large role of government in the health care sector, they can effectively exercise their monopsonistic powers. It is thus conceivable (and there are anecdotes supporting this possibility) that consumers in less developed countries be charged higher prices for drugs than consumers in far richer countries. (In doing so, it is the consumers in the less developed countries that, in effect, are paying the fixed cost of research; consumers in the 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 which protects the poor country. Well-designed (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 favored 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.[11]

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 pure public good, the argument for charging a fee is largely based upon 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 are 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 bio-diversity, 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 within the plants which give the desired effects.

The contrast 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 could not be more stark. In the former case, all of the return is credited to the "discoverer," with none to the pre-existing knowledge; in the latter, the patent holder is allowed to act as a perfectly discriminating monopolist, regardless of the extent to which his own innovation built on pre-existing knowledge.

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.

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

The effective use of knowledge developed in the advanced industrialized 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.

II. Combining Local and Global Knowledge

As I have just mentioned, 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 the less developed countries. But there are many other aspects of the "knowledge infrastructure" within less developed countries that 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. Korea and other successful newly industrialized countries that have closed the knowledge gap between themselves and the more advanced industrialized 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. What is at issue is not the importance of higher education. The criticism is *what* is taught, *the quality* of the education, and *how it is financed.* Science and technology are vital. It must be taught at international standards--otherwise it does little good in closing the knowledge gap and it would 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, through repaying student loans.

Governments in the newly industrialized countries often played another set of important roles in facilitating the transfer of knowledge. They did this by establishing standards laboratories that helped countries 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 actively recruited those forms of FDI most likely to have knowledge spill-overs and designed employment and other programs to enhance the likelihood of such spill-overs. Licensing policies also played a role in the transfer of knowledge.[12]

As essential as the adaptation and creation of new knowledge within a country is the dissemination of knowledge throughout the country. There are huge differences between average and best practices. The movement of ideas within a country is affected by the effectiveness of its communication system. Recent advances in telecommunication have brought the costs of communication down tremendously and have made possible the development of communications networks in parts of the world where it would have been decades, at best, before such

[12] 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 importation of essential inputs into more advanced technological processes.

systems would have been developed with the older technologies. These new technologies mean that there is no longer a natural monopoly, and by using competitive, market forces, access can be enhanced and prices lowered.[13]

This communication 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 most rapid in the United States, and not surprisingly, slowest in the less developed countries. The enhanced ability to share and acquire knowledge in the advanced industrialized countries may increase the knowledge gap as the less developed countries may become even more disadvantaged.

At the same time, they can tap into a larger knowledge pool than they ever had access to before. Today, a child anywhere in the world who has access to the internet has access to more knowledge than a child in the best schools of the industrial countries did a quarter century ago. He 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 the 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",[14] that is, creating the capacity to close the knowledge gap, an essential part of a successful development strategy.

[13] Competition remains, however, far from perfect, so there is still an important role for an effective regulator. Chapter 2 of the 1998 World Development Report documents the success of the countries that have used market basic competition with regulation. 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 access to the internet was raised to the point that the University could not afford maintaining connectivity. The "reform" thus reduced the ability of those in the country to avail themselves of global knowledge.

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

III. 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 organized societies, how to live healthier lives in ways which support the environment. It involves knowledge that affects fertility and knowledge about the design of economic policies that promote economic growth.

We in the development institutions acquire much of this knowledge as a by-product of our general development activities. It is a form of learning-by-doing.[15] 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. Research is thus 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. The international institutions, including the World Bank and the UNDP, have an especial role in the production and dissemination of this knowledge. We at the World Bank are increasingly conceiving of ourselves as a Knowledge Bank,[16] and are organizing ourselves in ways which 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*[17] shows that aid has a substantial impact on economic growth in countries which put into place good policies, while it has a negligible effect in countries which do not. Knowing what those policies are and adapting them to the circumstances of the individual country is thus a central part of a successful lending program.[18] More broadly, knowledge, aid, and private capital all work together in a successful development program; they are all complementary.[19] [20]

---

[15] See Arrow (1962).

[16] The concept of the Knowledge Bank was introduced in Wolfensohn's address to the Annual meetings in 1996. Wolfensohn (1996).

[17] World Bank (1998b).

[18] *Assessing Aid* points out that foreign aid money is only significantly correlated with positive impacts in those developing countries with sound economic policies and institutions. In particular, in countries with sound overall economic management policies, 1 percent of GDP assistance leads to a sustained increase in growth of 0.5 percent and reduces poverty by 1 percent. 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).

[19] Thus, in countries which pursue good economic policies, aid "crowds in" private capital; $1 of aid helps bring in $2 of private capital. This helps explain its strong role in promoting economic growth.

[20] 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 provides the stimulation for higher levels of investment, and the new investment embodies new technology. Without improvements in knowledge, presumably the countries of East Asia would have quickly experienced diminishing returns; as it was, they could maintain high rates of investment for an extended period of time 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

But there is more: We have increasingly realized that isolated projects will have only limited effects in the transformation of societies which 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 which can constitute the basis of economy-wide transformations.[21]

IV. Concluding Remarks

The concept of global 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.[22] 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 which 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, a balancing of equity and efficiency concerns among all 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. Efficiency requires public support. And this public support must be at the global level.

I have argued that knowledge is one of the critical 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.

The current arrangements can be made to work effectively, but if they are to succeed, we must all 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 requires collective actions. The challenge facing the international community is whether they can make our current system of voluntary, cooperative governance work in the collective interests of all.[23]

---

entirely by investments, including investments in people—are so misleading. It was a miracle that they were able to maintain high returns with the levels of savings and investment—few if any other countries in the world had 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 I think, 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 (1998a).

[21] These ideas are developed more fully in Wolfensohn's speech at the World Bank/International Monetary Fund Annual Meetings in 1998. See Wolfensohn (1998) and in Stiglitz (1998).

[22] See Stiglitz (1989).

[23] We can and should be more precise: since there are likely to be trade-offs, with some arrangements serving to advantage some groups relative to others, the two key questions are standard efficiency and equity issues. Can the 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 which comport with basic notions of equity?

References

Arrow, K. 1962. "The Implications of Learning by Doing". *Review of Economic Studies.* 29: 155-173.

Economic Report of the President. 1997. United States Government Printing Office. Washington.

Katz, M. L. and C. Shapiro. 1985. "Network Externalities, Competition, and Compatibility." American Economic Review 75:424-40 June.

Klenow, P.J. and A. Rodríguez-Clare. 1997. "Increasing Returns and Long-Run Growth" in *Journal of Political Economy*. 40:597-617.

Romer, P.M. 1986. "Increasing Returns and Long-Run Growth". *Journal of Political Economy.* 94:5. 1002-1037.

Samuelson, P. 1954. "The Pure Theory of Public Expenditure," *Review of Economics and Statistics*, 36, 387-89.

Stiglitz. J.E. 1977. "Theory of Local Public Goods," In *The Economics of Public Services*, M.S. Feldstein and R.P. Inman (eds.), MacMillan Publishing Company, 1977, pp. 274-333. (Paper presented to IEA Conference, Turin, 1974.)

———. 1983. "Public Goods in Open Economies with Heterogeneous Individuals," In *Locational Analysis of Public Facilities*, J.F. Thisse and H.G. Zoller (eds.), North-Holland Publishing Company. 55-78.

———. 1983b. "The Theory of Local Public Goods Twenty-Five Years After Tiebout: A Perspective," In *Local Provision of Public Services: The Tiebout Model After Twenty-Five Years*, G.R. Zodrow (ed.), Academic Press. 17-53.

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

———. 1987. "On the Microeconomic of Technical Progress". *Technology Generation in Latin American Manufacturing Industries*, Jorje M. Katz (ed.) The Macmillan Press Ltd. 1987, 56-77. (Presented to IDB-Cepal Meetings, Beunos Aites, November 1978)

———. 1987b. "Learning to Learn, Localized Learning and Technological Progress," Economic Policy and Technological Performance, P. Dasgupta and Stoneman (eds.), Cambridge University Press. 125-153.

———. 1989. "On the Economic Role of the State". *Economic Role of the State.* A. Heertje (ed.). Basil Blackwell, Oxford, 9-85.

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

———. 1998. *Towards a New Paradigm for Development: Strategies, Policies, and Processes.* Given as Raul Prebisch Lecture at United Nations Conference on Trade and Development (UNCTAD). Geneva. October 19.

Tiebout, C.M. 1956. "A Pure Theory of Local Expenditures". *Journal of Political Economy*. LXIV. 416-24. October.

Wolfensohn. J.1996. Annual Meetings Address. www.worldbank.org/html/extdr/extme/jdwams96.htm

———. 1998. Annual Meetings Address. www.worldbank.org/html/extdr/extme/jdwams98.htm

World Bank. 1998a. *Knowledge for Development*. World Development Report. Oxford University Press.

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

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

Figure 1: GDP and R&D by Region

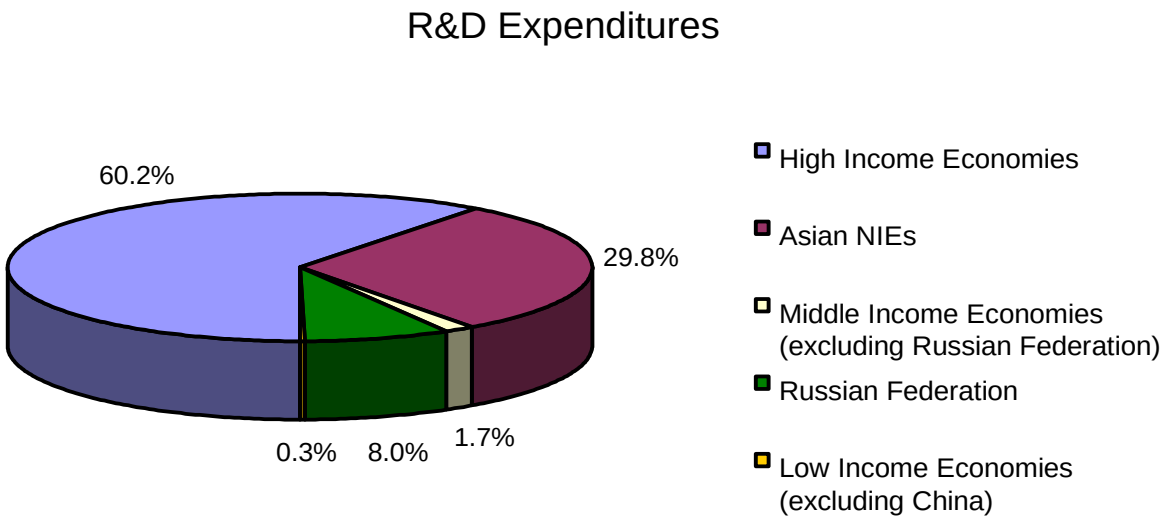

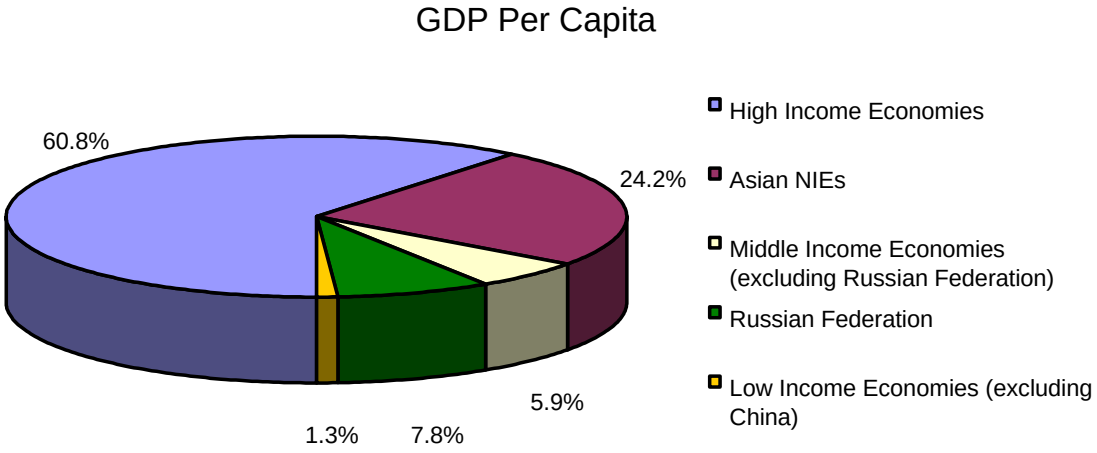

Source: World Development Report 1998.