United Nations

Economic Commission
for Latin America

United Nations
Development Program

Inter-American
Development Bank

Research Programme
in Science and Technology

Working Paper N° 32

ON THE MICRO-ECONOMICS OF TECHNICAL PROGRESS

J.E. Stiglitz

Distr.
RESTRICTED
BID/CEPAL/BA/51
April 1979
ORIGINAL: ENGLISH

Joseph Stiglitz is Drummond Professor of Economics, All Souls College, Oxford. This paper was written while the author held the Oskar Morgenstern distinguish Fellowship at Mathematica and was visiting professor at the Institute for Advanced Study, Princeton.

The present paper was originally presented at the "IDB/CEPAL Seminar on Studies on Technology and Development in Latin America" held in Buenos Aires, Argentina on 6-10 November, 1978.

Programa BID/CEPAL/PNUD
Oficina de la CEPAL en Buenos Aires
Callao 67 - Piso 3°
1022, Buenos Aires - Argentina

ON THE MICRO-ECONOMICS OF TECHNICAL PROGRESS

This paper surveys some recent developments in the theory of the economics of technological change.[1]

Although I have written the survey with the questions of how technological change gets transfered from developed to the less developed countries, how it diffuses within the ldc's, and how it gets adapted there to local conditions clearly in mind, I have not surveyed the literature which directly addresses those questions; it is my belief, however, that the insights which can be obtained from this more general discussion of the nature of technological change will have considerable bearing on those questions, as I hope my discussion will make evident.

I have divided this survey into two parts. In the first, I discuss the basic "underpinnings" of the theory of technological change, what I call the micro-micro-economics of technological change: what is technological change, how is it produced, how is its production similar and dissimilar from other economic activities? In the second section, I

[1]Much of the literature which I survey is unpublished, some of it is unwritten. In this draft, because of lack of time, I have been less than thorough in making a complete survey of the literature. This paper should rather be viewed as the conceptual outlines of the Theory of Technical Progress.

discuss the micro-economics of technological change; within the past few years, the traditional questions of the relationship between market structure and technological change has been re-examined. This has, I think, produced a significantly different view of the relationship, a reformulation which has, I think, some fundamental implications for both patent, anti-trust and protectionist policies.

I. ON THE NATURE OF TECHNOLOGICAL PROGRESS

2.1. Alternative Views of R&D

In Solow's famous 1957 paper, over 8% of the growth in the United States over the past 50 years was attributed to technological change.. Clearly, if that were the case, an understanding of the nature of technological progress is required for an understanding of the processes of economic growth. Three views of what technological change was developed in the subsequent years:

2.1.1 The Residual Approach

The first view was an essentially negative one: technological change was what one called a change in the production function which one could not explain in some other way; it was, in other words, a "residual." Thus, to "understand" technological progress meant "to explain it away." The attempt to do so became a favorite past time of both clever econometricians and skilled number crunchers. Part of the residual could be attributed to improved education, part of the residual could be eliminated if one "adjusted" the price of capital goods correctly, part of it was related to scale economies, part of it was due to shorter working hours, etc. So successful was this quest for "explaining technological change by explaining it away" that Griliches and Jorgenson did precisely that: in their formulation there was essentially no residual.

This result did not, I think, convince anyone that technological change had not occurred in the United States over the past one hundred

years. What it did do, however, was to suggest that the residual approach to quantifying technological change was not likely to shed much light on the question of the importance of technological change, let along on the processes by which it occurred.

2.1.2 R&D As An Investment

The second view of technological progress also developed out of the concern for understanding the interaction of technological change and growth: it observed that expenditures on R&D were like an investment, the returns to these expenditures occurring over extended periods into the future. Economies were viewed as having a choice between investing in machines and investing in R&D. In analyzing this choice, capital goods and R&D were modelled in almost the same manner; the only critical difference was that the production function did not exhibit diminishing returns to knowledge in the same way that it did to capital. More formally, the aggregate production function was usually written in the form

(1) $Q = AF(K,L)$

where Q is output

A is the state of technology

K is the capital stock

L is the labor force.

Thus doubling A doubled output for all values of K and L; while doubling K, for a given value of A and L, would less than double output.

This simple observation has some very important implications. For most of traditional economic analysis (e.g. the viability of competitive

equilibrium) requires diminishing returns to all "factors". If knowledge is like a "factor of production" then we have to inquire how can it be produced and supplied? Clearly, traditional models of the supply of factors will be inapplicable. These questions, implicit (or perhaps explicit) in this literature were not (until recently) pursued.

2.1.3 R&D As A Public Good

The third view observed that technological change was like a public good: giving one person a piece of knowledge did not detract from the one amount of knowledge that others had. In this sense, it seemed like a rare case of a pure public good. But it was a public good the returns to which could be appropriated with patents. It was obviously not socially desirable to do so, since to do so would restrict output (since there was no marginal social cost to using the knowledge); but if R&D had to be produced in the private sector, then there was an optimal patent, balancing the incentives which a patent provided for doing research and the returns resulting from the research, on the one hand, against the losses associated with the restriction in output.

2.1.4 R&D As Information

The view which has been developed in the last few years is closely related to the two preceding views: it is that R&D is like "information" or perhaps more precisely, it is a kind of information. Expenditures on information are often investments; the returns to information increase with the scale of production; and information is a public good. Viewing technological change as information does, however, give us considerable

further insight into the nature of technological progress. For it provides us some further insights into the nature of the processes by which technological change is acquired, and the manner and extent to which the return can be appropriated.

2.2 Differences Between R&D and Other Economic Activities

The production of knowledge is different from the production of ordinary goods in several important ways.

a) Knowledge is often produced as a by-product of other economic activities, in a way which is inseparable from those activities. The classic formulation of this is what is referred to as "learning by doing", the process of production leads to improved efficiency in doing whatever is being done. But there are other important examples as well: the fact that a firm succeeds in producing something conveys what in many cases turns out to be an extremely valuable piece of information: that the thing which the firm is producing can be produced. (The knowledge that synthetic fibers can be produced in an economicly viable manner has a significant effect on the expected returns to searching for a synthetic fiber, and thus on the allocation of resources to the discovery of other synthetic fibers.)

b) The analogy to information has, in fact, also proved useful in generating a better understanding of the "micro-micro" structure of research. There has been some interesting work modeling the problem of research as a "search problem". This

allows a detailed analysis of the equilibrium "search strategy" compared to the optimal search strategy. Such comparison underlie some of the analysis of section 2.4.3 where we suggest for instance that the market undertakes excessively risky research strategies.

c) There is a general theorem which shows that the production of information is, in general, subject to increasing returns to scale: regardless of the parameterization of information it never pays a Bayesian to undertake a little bit of research. This means that if it is worth conducting research, it ought to be conducted at a scale no smaller than some strictly positive level. It also means that there will be specialization of research strategies, rather than the comple diversification which would predominate if there were imperfectly correlated risks with no increasing returns.

d) Since every piece of information produced must be different from any other piece of information produced (otherwise it is not new "knowledge"), there is a fundamental sense in which there cannot exist competitive markets for information (knowledge). This can be put another way: since before being told the "piece of information" the buyer cannot know what the seller is selling, and since after being told precisely what it is that the seller has to sell, the buyer has no reason to pay the seller for it (since it has, by then, already received it), markets for information cannot work. (This is obviously an extreme way of putting the matter: in practice, the seller gives some indications to the

buyer of what he has to sell, but does not provide all the details; the point is that markets for information are, inherently, imperfect.)[1]

(e) There is another related characteristic of the production of knowledge which it shares with some other commodities but which is of critical importance here: the activity of research and development is inherently extremely risky, but the risk are uninsurable; it is virtually impossible to ascertain whether the inventor failed to discover the invention because the invention was not feasible, or because he did not work hard enough, or because he was an incompetent investigator. Insurance markets are likely to be non-existent or highly imperfect whenever there is a large potential for asymmetric information (the buyer of the insurance knows more about the likelihood of success than does the seller) and when the actions of the buyer affect the outcomes, and are difficult if not impossible to observe. Both of these conditions are satisfied here. Moreover, unlike other forms of investment, there is no collateral, so that the costs of failure to a lender are likely to be greater than for ordinary investment opportunities. The consequence of these observations is that inventors are likely to require self-financing;

[1]The arguments of the preceding paragraph would lead to the conclusion that there would never exist a market for information. Some information is, however, purchased. The seller establishes a reputation for the quality of information he sells. Salop and Stiglitz have recently analyzed the nature of markets with imperfect information but repeat sales.

many of the earlier discussions of innovation and industrial concentration emphasized this point (although without much detail as to the economic explanation of why the appropriate insurance and capital markets would not exist); but the importance of this for modern industrialized economies is moot; all that is required is that the firm have enough physical capital to serve as collateral; since funds are fungible, firms can borrow for capital investment, but use the proceeds for R&D.[1]

2.3. The Problem of Appropriability

We noted earlier that knowledge (information) is a pure public good, in the sense that giving knowledge to another individual does not detract from that which the others has. But unlike some public goods, in many cases the returns to invention are appropriable. However, the degree of appropriability will differ for different kinds of information (invention), and this will affect the pattern and direction of the allocation of resources to R&D and innovation.

Traditional discussions of appropriability have focused on the use of the patent system. (We mentioned earlier, for instance, the trade off between the incentives for research provided by R&D and the restriction in output which results from a patent system.) But it is now becoming widely recognized that the patent system is relevant for

---

[1]Since the amount of R&D affects the riskiness of all loans outstanding to the firm, the implicit cost of expenditures on R&D are likely to be greater than the cost of capital used for other kinds of investment. The relationship between the private cost of capital and the social cost for R&D expenditures requires further investigation.

only a small fraction of inventions and innovations. One reason for this is that the very process of patenting conveys important information to competitors: the design specifications which are required to be submitted in conjunction with a patent application provide the basis for the competitors to "invent" around the patent. Firms may be able to increase their profits by attempting to keep the design specifications secret.

A second reason is that many important innovations are not really patentable. The improvement in the technology of selling hamburgers provided by McDonald's was evidently of enormous value (measured by the profitability of the enterprise), yet this technological innovation was not patentable.

Indeed, almost all research, whether successful or not successful, contributes to the common pool of knowledge, by providing other researchers with information about where to look and where not to look for discoveries. The returns to this information cannot be appropriated, or can only be partially appropriated.

The fact that private firms do engage in research, most of which is unpatentable, suggests that there are other methods of reaping returns. The basic method of obtaining returns is closely related to the property of increasing returns that we noted earlier: significant discoveries (leading to significant lowering of costs or to a new product) give firms some degree of monopoly power; this is temporary--others may enter the market, and bid away their monopoly profits--but until this occurs, the firm earns some monopoly profits. The knowledge is not sold; it is goods which are sold; but the nature of market in which the goods are sold is affected by the production of knowledge.

The fact that the returns to R&D are related to the size of the firm means that R&D is likely to lead to industrial concentration.[1] But there is a limit on the degree to which firms may exercise the monopoly power which may be associated with any degree of concentration: for in the absence of direct restraints on trade, threat of entry will limit the extent to which, even in a highly concentrated market, the largest firm can act like a monopolist.[2] We shall return to this question in section 3.

There are other methods of reaping returns to innovation. In some cases, the knowledge of a discovery provides the firm (inventor) with "prior" or "inside" knowledge, which allows him to profitably speculate on the change of price of some commodity or firm (Hirschliefer, for instance, suggests that Eli Whitney could have made a killing on the cotton market; but to appropriate even a small fraction of the returns requires that there exists markets extending at least several years into the future).

In many cases, the researcher can charge for the transmission of knowledge, and appropriate rents associated with the transmission.

The management fees associated with introducing new technologies can be interpreted as a method of appropriating returns on R&D.

---

[1]That is, if a process invention which cannot be patented reduced the cost of a firm by $x per unit, the total value of that invention is proportional to the number of units produced by the firm; and this fact means that the allocation of resources to R&D will be greater the larger the firm.

[2]The fact that he cannot act like an unbridled monopolist does not mean that he acts competitively. The threat of entry, as we shall argue below, affects actions other than price determination.

Although there are several methods of appropriating returns, it should be emphasized that the returns received are not directly related to the marginal social value of the research. On the one hand, (a) there are frequently spillovers of the knowledge gleaned in one area to another; and (b) whenever a product invention occurs, there is some consumer surplus associated with it. Unless firms are perfectly discriminating monopolies, they do not capture this surplus in its entirety. Much of the literature has focused on this partial appropriability of the returns, with the implied underinvestment in R&D which results.

But there are several other effects which makes the nature of the bias more ambiguous. Three of these we have already noted: the rents associated with transmission are not necessarily equal to the marginal costs of production and transmission of knowledge (but are obviously less than the value of the piece of knowledge); the capital gains associated with the change in prices, and the returns which can thereby obtained, may be greater or less than the social value of the piece of information; and the monopoly power associated with large changes in the costs of production yields monopoly rents which are not directly related to the value of the invention (although again, these must be less than the total value of the invention, although not necessarily less than the marginal value).

There are two other effects, however, which should be noted. First, the actual contribution of a particular inventor is the present discounted value of having the invention earlier than it otherwise would have occurred. But this is not observable. There is a sense in which all potential inventors draw upon a common pool of knowledge.

Much invention consists of transforming basic ideas into marketable commodities, and some of the return, in a system in which basic ideas are not patentable, but the produced commodities are, is really a return on the basic ideas. Thus, individuals compete to produce marketable patentable commodities, to acquire a share in the common pool of rents. As in any common pool type of situation, this may lead to excessive entry.

Secondly, in economies with monopolistic competition and product differentiation prices in general do not equal marginal costs; the introduction of a new way commodity may have consequences for the viability of other commodities, e.g. the introduction of a new commodity may lead to the withdrawal of others. In this case, the gain in consumer surplus due to the newly invented commodity has to be matched by the loss due to the withdrawal of others. The net effect is ambiguous; and what the research can appropriate could conceivably exceed the next gain to society.

It is because the degree of appropriability will differ so markedly among different types of invention that generalizations such as "there is too much or too little invention" seem inappropriate; what is probably more important is that the differences in the degree of appropriability will affect the direction of R&D; and firms will attempt to take actions which increase the degree of appropriability associated with any particular innovation.

2.4. Biases in the Pattern of R&D and the Characterization of Technological Change

Earlier discussions attempting to characterize technological progress focused on the effect that technological progress had on the demand for different factors; technological change was characterized as being labor saving or capital saving or labor augmenting or capital augmenting. Traditional discussions of the relationship between R&D in developed and less developed countries has emphasized that technological progress in developed countries is biased towards saving labor, a bias which is inappropriate for l.d.c.'s. Without disagreeing with the conclusion of that argument, what I would like to do in this section is to point out (a) there are a number of implicit assumptions in the traditional argument which seriously limit the force of the argument; and (b) there are a number of other characteristics of technological change besides the labor biasedness, biases in which may be more important than the traditional factor choice bias.

2.4.1 Remarks on the Factor Bias of Technological Progress

Older discussions of induced technological progress (e.g. Habakkuk) refer to "labor scarcity" as having induced labor saving invention. These earlier discussions were subjected to criticism, on the grounds that in equilibrium, all factors are scarce. The fact that wages are rising suggests that labor is becoming more scarce, and might lead to an argument that, if technological change is related to factor scarcity, that it might become more labor saving than it had been previously. But the mere fact that wages are higher than they were at some arbitrary date in the past is

not a reason for arguing that firms should attempt to economize on labor; they should attempt to economize on all factors.

To make sense out of an argument for factor bias, one must introduce a "trade-off"; it turns out more convenient to phrase our analysis in terms of labor or capital augmenting technological change, but the conversion to labor or capital saving (in the Hicksian terminology) is straightforward. We thus postulate for that for some industry

$$Q = F(\mu K, \lambda L),$$

output is a function of the capital and labor inputs (we could add additional terms for the inputs of raw materials); and of the efficiency with which those are utilized, which is summarized by μ and λ .

We postulate that, for a given expenditure on R&D, the firm has a choice of various possible values which μ_{t+1} and λ_{t+1} can take next period, as depicted in Figure 1.

We write

$$\mu = \mu(\lambda)$$

We have drawn the function to be concave, but we make no argument for that shape.

For any given input, it wishes to maximize its output next period i.e.

$$\max \quad F(\mu(\lambda) K, \lambda L)$$

so

$$F_1 \mu_\lambda K + F_2 L = 0$$

or

$$\frac{F_2 L \lambda}{F_1 \mu K} = \frac{\alpha}{1 - \alpha} = -\frac{\mu_\lambda \lambda}{\mu}$$

where α = share of labor

In a competitive economy, the elasticity of the invention possibilities schedule must be equal to the relative share of labor divided by the relative share of capital.

More generally, we can write the cost of production facing any firm as

$$C = C(w/\lambda, r/\mu)$$

where w = the cost of labor

r = the cost of capital.

Thus cost minimization, for <u>competitive</u> or <u>non-competitive firms</u>, is given by

$$\frac{C_1 w}{\lambda^2} + \frac{C_2 r}{\mu^2} \mu_\lambda = 0$$

But

$$C_1 = \lambda L, \quad C_2 = \mu K$$

Hence we obtain

$$\frac{\mu_\lambda \lambda}{\mu} = \frac{wL}{rK} = \frac{\alpha}{1 - \alpha}$$

We thus immediately obtain the result that if in two economies, the relative shares are the same (even if factor suppies are markedly different), if they face the same invention possibilities schedule, then they would have the same factor bias in technological progress.

This would suggest that technological change would be somewhat less labor augmenting, were the l.d.c.'s to determine its direction, than it is at present, but perhaps not dramatically so.

A natural retort by those who believe that technological change in l.d.c.'s should be much less labor augmenting is that factor prices in l.d.c.'s do not correctly reflect true factor scarcities; that labor is in surplus, so that the appropriate α to use from a social view is "O". The recognition that market prices and shadow prices may differ has played an important role in cost benefit analysis of projects, and should certainly not be ignored is this context. At the same time, it is important not to confuse the existence of unemployment with a zero shadow price. As Harberger and Stiglitz have shown in models in which there is rural-urban migration, the shadow price on labor is approximately the urban wage, if the rural wage is relatively unaffected by the out-migration.

There is another sense in which shadow prices for labor, particularly for skilled labor, may differ from market prices, which is especially relevant in situations where markets are thin. Assume that a particular machine requires labor with particular characteristics. Each machine requires one unit of labor. Then, when that laborer leaves, the machine will be unemployed until a new laborer is hired. The firm can either keep a stock of reserve capacity of workers available, or it can

let the machine remain unemployed until the worker is replaced. It can attempt to induce more workers to apply (in order to obtain, from the applicant pool, someone with the desired characteristics) so that the length of time that the machine is unemployed is reduced. In any of these cases, the cost of labor (relative to the cost of capital) is greater than a simple calculation using the wage rate of those actually working on the machine would suggest. In particular, if skilled labor markets are thin, as they are in many l.d.c.'s, the fact that there is unemployment of general labor (or even skilled labor in general, but not of labor with the particular skills required by the firm in question) does not mean that the shadow wages are not equal to or greater than market wages.

The traditional factor bias argument needs to be qualified in two other directions:

a) In the traditional analysis, as in our presentation above, firms developing the new techniques took a completely myopic view. On the one hand, if they recognize that the machines and technologies which they are presently developing for the developed countries, will eventually become utilized by the less developed economies, then with suitable appropriability assumptions, they will take that into account in determing the direction of factor bias.

This makes particular sense for those multinational companies who regularly export technology to the l.d.c.'s.
Note that if this argument is correct, then if the speed of diffusion of new technology to the l.d.c.'s is increased, then the l.d.c.'s may benefit, not only as a result of the improved technology, but also as a result of the new technology which is being exported being more appropriate for the l.d.c.'s.

On the other hand, if the differences in factor shares between the developed and less developed countries are expected to increase in the future, and if firms in the developed countries take a non-myopic view in their choice of direction of factor bias in technological change, then the magnitude of the bias in the direction of technological change may be even greater than we have postulated above.

b) In the traditional analysis, no attention is paid to the nature of the returns to scale, either in the production or utilization of R&D.

2.4.2 <u>Other Biases in the Direction of R&D</u>

There are two categories of choices with respect to R&D strategy that have to be separated: those that affect how research is conducted and those which affect the objective of the research program. The preceding discussion concerned one dimension of the latter. Several other dimensions need to be mentioned:

a) R&D and the optimal scale of production. In many industries, technology is best described by a U-shaped cost curve. Changes in technology affect the entire shape of the cost curve, and in particular, they effect the scale of output corresponding to minimum average cost. This, for instance, is particularly true of the chemical industry, where better plants almost always have larger capacities. We can envisage a trade-off between reductions in the fixed costs associated with any plant, reductions in the variable cost, and increases in the maximum feasible output from any plant, as depicted in Figure 2.

Our concern here is how the choice among these potential directions of technical change in the market might differ from the socially optimal direction, and how it might differ between developed countries and l.d.c.'s. If markets are imperfectly competitive, then firms may, by increasing the optimal scale at which production occurs, effectively reduce competition and deter entry. There would seem to be a bias in favor of excessive scale. Since markets in l.d.c.'s for non-exported goods are smaller than corresponding markets in developed economies, the ensueing degree of monopoly within the l.d.c. may be even larger and the distortion in the direction of R&D even more costly.

b) Imitability. Some inventions or innovations can easily be imitated, even with a patent system; others may not be. Firms make profits over the interval of time during which they can

exercise monopoly power in the market; if they can exercise this power for a long time, there will be a greater incentive for R&D.

c) Consumer communicability. For product inventions, consumers have to be informed of the new invention and of its characteristics. For some commodities, this information may be conveyed easily and cheaply, say by distributing samples. But for durable goods, this is not a feasible method of conveying information. In that case, there is a bias in the direction of technical progress towards those commodities with characteristics which can be readily and easily communicated.

d) Present state of competition. In some commodities, there is close competition, while in others, there may be relatively few firms. In areas with heavy competition, the returns to a new invention are likely to be bid away quickly.

There is an important interaction between (b), (c) and (d): Firms are likely to direct their research to commodities in which ther is relatively little competition, and in which they can quickly establish a large market share before imitation occurs. It should be apparent that this particular pattern of innovation may not bear a close correspondence to the socially optimal pattern.

e) Localization of technical progress and spillovers. Many changes in technique have implications for a number of different production processes. For instance, a better conveyor belt affects a vast array of production processes. Other changes in technique are very particular; they have very little spillovers.

This is sometimes depicted, as in Figure 4, in terms of the effect of technical change on the isoquant. Some technical improvements move the entire isoquant down; others shift only a point. The latter type of technical progress is sometimes referred to as localized technical progress (Atkinson and Stiglitz).

There may be some choice about the degree of localization of technical progress; that is, there may be a trade off between looking for an invention that shifts the isoquant from its present state to the new isoquant A or the new isoquant B. This choice is affected by a number of considerations. First, if the firm knows precisely the factor prices it faces, it is clear that A is preferable to B. On the other hand, if it is uncertain about its factor prices, B may be preferable. B has, in a sense, greater flexibility. Secondly, localization may make the invention less imitable, and thus make it easier for the firm to appropriate the returns. In particular, the firm will attempt to localize the technical progress in directions which are related to the peculiar situation which it finds itself in. There is a sense, then, in which firms, more than simply ignoring spillovers, may attempt to minimize them.

Clearly, the degree of localization has an important bearing on the extent to which l.d.c.'s benefit from improvements in technology in the developed economies.

We have focused, in the preceding paragraphs, on localization with respect to factor bias (the choice of technique of production). But there are other choices, e.g. improvements which might affect the quality of a wide range of products or only of a subset of commodities, for which the concept of "localization" is applicable.

In the preceding discussion, we presented a number of characteristics of commodities, which affect the return to their invention and innovation. The return to inventing or developing a particular commodity may differ between different countries or economic environments. For instance, if an electric generating firm is connected with a large network of other generating firms, the return to reliability may be markedly different than if it is the sole producer; and if it is the sole producer, the return to reliability will differ markedly depending on whether it has a reserve generating capacity which it keeps on hand or does not.

2.4.3 Biases in the Pattern of Research

In this section, we discuss briefly biases associated with the manner in which research is conducted. Four potentially significant kinds of biases have been identified:

a) Speed of Research. Since the patent system (as well as the method of appropriating returns through exercising monopoly power) rewards the first entrant disproprotionately ("winner takes all"), there may be a significant bias in encouraging excessively fast research.

b) Riskiness of Research Strategy. The returns to research are, however, extremely stochastic, and what is relevant therefore is the probability that one's research strategy pays off before one's competitors succeed. Undertaking riskier research strategies may increase the probability of success at an earlier date, and may be undertaken even if the expected date at which success occurs is thereby postponed. The market, in other words, undertakes, from a socially point of view, excessively risky research strategies.

c) Size of Projects. Since the degree of monopoly power increases more than in proportion to the reduction in risk (that is, a small invention reduces costs, but gives the inventor little monopoly power; a large invention may make the firm dominant in the industry), there may be a bias for "large" relative to small inventions.

Both of these arguments have to be qualified when account is taken of risk aversion on the part of inventors and the existence of imperfect markets for risks.

d) Independence of research strategies. Different firms can pursue research strategies the outcomes of which differ in their correlation. Arguments have been put forward suggesting that the market will be insufficiently diversified: since what firm A cares about is the probability that B succeeds before A, if A could undertake a perfectly correlated research strategy (success of B implies success of A and conversely) but only

at a slightly faster speed, he would. This question, however, has not yet been fully analyzed.[1]

3.1 Market Structure and Innovation

It has long been recognized that there are important relations between market structure and the technical progress. There has been, for instance, considerable controversy over whether monopoly encourages invention or innovation or discourages it.

Many of the simple models (e.g. that of Arrow) which have been formulated to compare the two are misleading or at least overly simple.

First, they often assume that the market situation prior to the invention or innovation was identical; in fact in different market structures, firms will take different actions prior to the invention or innovation. For instance, in the case of natural resources, the rate of depletion of the natural resource will be different between competitive and monopolistic markets. In commodities, which require durable capital goods for their production, firms will differ in the durability and size of the capital stock.

Secondly, the earlier analysis ignored the fact that there may be competition for research. The number of researchers is, itself, an endogenous variable, so that even if they were a single firm presently producing the commodity, there might be a number of potential researchers.

---

[1]There is an analogy here to monopolistic competition which may prove useful: there we are concerned with the distribution of firms across product space, here we are concerned with the distribution of firms across "research strategy space".

Firms recognize this, and their research and production strategies will be directed to take this into account. A monopolist will attempt both to pre-empt potential competitors and to take actions which will discourage them from undertaking research, i.e. lowering the return to their invention if he fails to pre-empt them. It is easy to show that if both the monopolist and other researchers have the same cost function, then the monopolist will always find it profitable to pre-empt his potential competitors, so long as the joint returns to the two duopolists are less than the profits of the monopolist (which they presumably would be if there were any competition between them, i.e. unless they acted completely collusively). This is true even if research outcomes are stochastic.

Indeed, it may pay them so much to pre-empt that they obtain patents considerably before the optimal time (given their capital stock) for introducing the innovation; i.e. there are sleeping patents.

But for one reason or another, firms often are not able to pre-empt all potential competitors. There are a number of actions which they can take to discourage potential competitors. They can have excess capacity, so that the equilibrium price, were the competitors successful, would be low; in the case of natural resources they threat of entry leads the monopolist to increase his price, conserving on the use of natural resources, so that at the time of invention, the price to which the market falls will be low.

In the case of R&D for the development of a substitute for a natural resource, a detailed comparison between the socially optimal allocation of resources to invention and the timing of invention and

innovation, and that in a number of alternative market structures has been carried out. This has shown that in markets with competitive entry into research, if the stock of natural resources is low, there may be excessive research, but in other market structures there will be too little. Moreover, a monopoly engages in less research than does a market with competitive entry into the research sector; but as we noted earlier, unrestricted entry into the research sector may lead current prices in the resource monopoly to be higher. In certain market structures (e.g. a monopoly of the resource, competition in the research sector, but a patent which gives the successful researcher a monopoly over the substitute) the date of first innovation given that the invention has already occurred, may be too early.

There is no comparable detailed investigation of other kinds of commodities, but the results should be directly applicable to commodities produced by means of durable capital goods.

Thirdly, and perhaps most importantly, the market structure itself needs to be viewed as endogenous, a result of the research undertaken by different firms and their success (itself a stochastic variable).

In this view, it is only the technology of technical change and demand functions, and the legal structure which should be taken as exogenous. This view reverses the traditional causal structure: a particular industry may have a high pace of technical change not because it is highly concentrated; rather, it may be high concentrated because of the high pace of technical change which is occurring there. In fact, as in any simultaneous equation system, neither sentence by itself makes much sense. The degree of concentration, in equilibrium, is

just that required to generate the monopoly profits to pay for the R&D expenditures which, given the degree of concentration, it is optimal for each firm to undertake. In this situation, there is no incentive for entry or exit. (There in fact may be profits associated with cost differentials, implictly a rent on earlier obtained knowledge.)

The policy question is not so much what would be the effect of more competition on the equilibrium price, as in traditional anti-trust discussions, but what would be the effect of a particular policy (patent policy, anti-trust policy, protection policy) on the rate of technical progress, the degree of concentration, and the level of prices in the long run. The short run and long run consequences may differ markedly.

We shall not, however, pursue these questions further here.

4. Concluding Remarks

In this paper, I have attempted to survey a number of recent developments in the economics of R&D. Adam Smith's invisible hand is not only not visible, but there is considerable evidence that it is not even there. But that does not mean that the visible hand of Government Intervention would improve the allocation of resources. The question is far more subtle than that.

What the theories we have presented may do is provide an explanation of the particular patterns of innovation and diffusion of techniques that have been observed, an explanation of why in certain industries there may be rapid innovation, in others less rapid. Moreover, our analysis may provide a rationale for certain types of intervention and an argument against other types of intervention. Let me briefly illustrate.

a) We suggested earlier that knowledge that a certain innovation is economically viable is of considerable economic value; the first innovator may not be able to capture these returns. There are further "advantages to being late." If that is the case, it may be in the private interests of each firm to postpone innovation, waiting for someone else to "try out" this innovation. Such delay may be far from socially optimal. Auctioning off temporary monopoly rights may in this situation, constitute a Pareto Improvement.

b) Models of diffusion of technological change have largely focused on the lags in learning about the new technology. The usual logistics curve is obtained using the standard "epidemic" model based on contacts; in our analysis, the

rate of introduction is affected by the distribution of risk attitudes within the population and by the rate by which uncertainty about the economic viability of the new technology or product is reduced as a result of its introduction by others.

c) Similar considerations are involved in an analysis of the optimal method of adapting a technology which was originally designed for developed countries to the l.d.c.'s. A sequence of adaptations, first to environments which are similar to those in developed countries, and then to environments which are less similar, may reduce both the total cost of adaptation and the risk associated with the success of each stage in the process of adaptation. (The distributional implications of such a "trickle down" approach to knowledge needs to be investigated.)

d) Policies which affect the amount of returns which multinationals can obtain in the introduction of new technologies to l.d.c.'s may not only reduce the flow of technology to l.d.c.'s, but may affect the pattern of technology development in a way which may be disadvantageous to the l.d.c.'s.

These remarks are put forward as suggestions of how the basic theory we have developed, an information theoretic approach to R&D, focusing on the biases introduced by varying degrees of appropriability associated with varying kinds of R&D, and on the critical relationships between R&D and market structure

may enhance our understanding of the processes by which new technologies and products become transfered to l.d.c.'s and adapted there to the conditions within the countries.

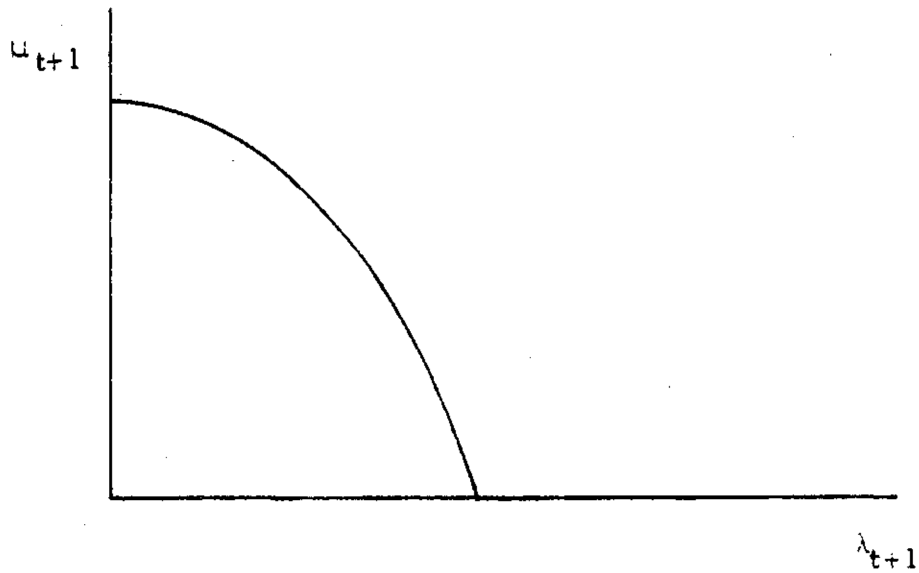

Figure 1. Invention Possibility Schedule: Tradeoff Between Labor and Capital Augmenting Technical Progress

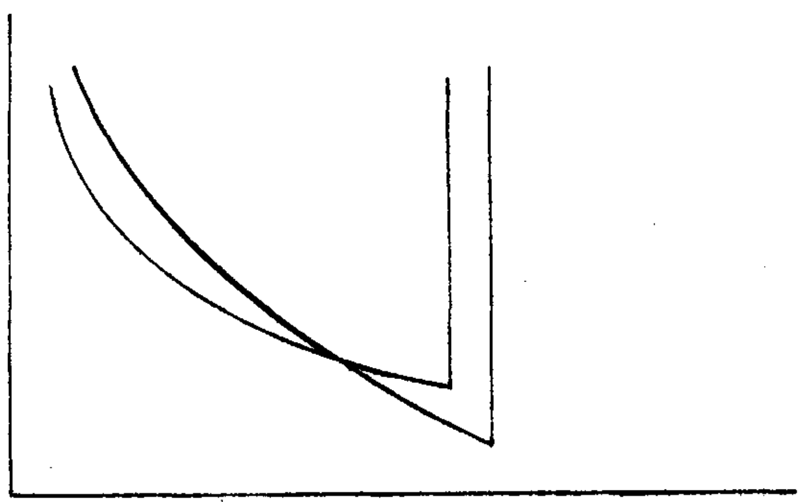

Figure 2. Scale Biased Technical Progress

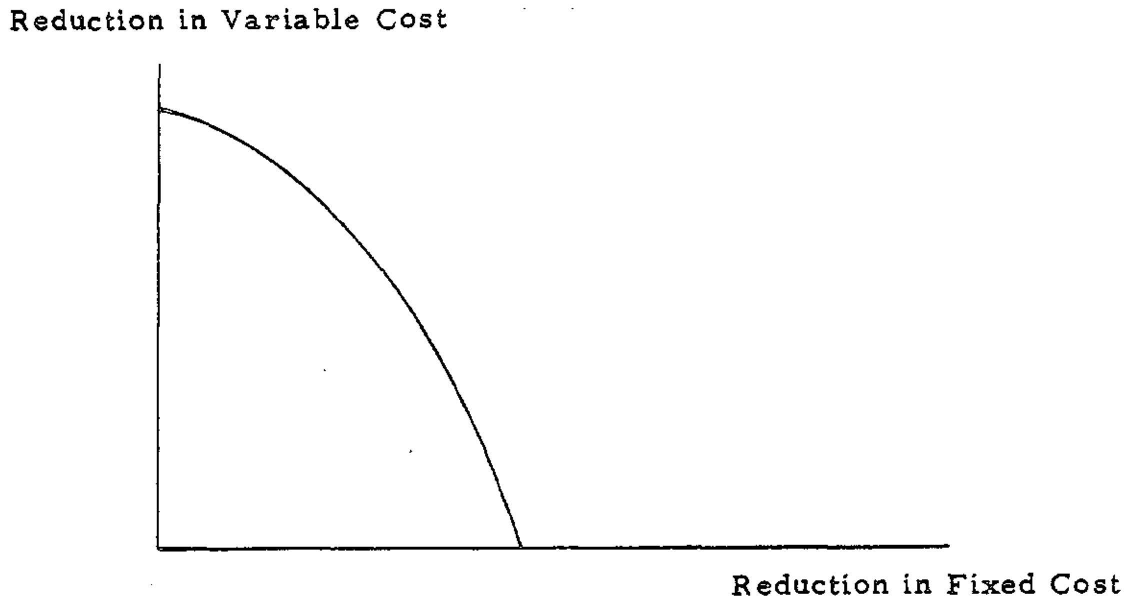

Figure 3. Trade-offs Between Reduction in Fixed and Variable Costs

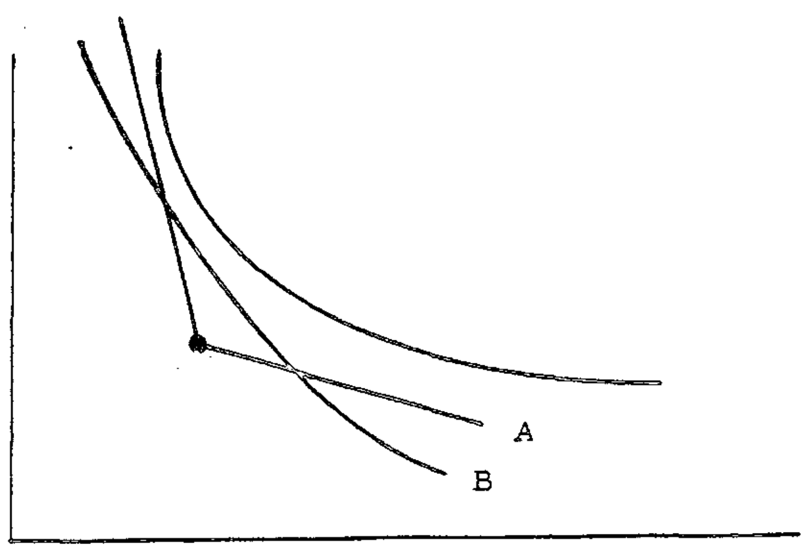

Figure 4. Localized Technical Progress

REFERENCES

Arrow, K. J. (1962), "Economic Welfare and the Allocation of Resources for Invention", in The Rate and Direction of Inventive Activity: Economic and Social Factors, (NBER) Princeton University Press.

Barzel, Y. (1968), "Optimal Timing of Innovation", Review of Economics and Statistics, (August), pp. 348-355.

Dasgupta, P., R. Gilbert and J. E. Stiglitz (1977), "Market Structure, Research and Development and Exhaustible Resource", mimeo, London School of Economics.

Dasgupta, P. and J. E. Stiglitz (1977), "Market Structure and Research and Development", mimeo, London School of Economics.

Dasgupta, P. and J. E. Stiglitz (1978a), "The Design of Incentive Structures for R&D", mimeo, London School of Economics.

Dasgupta, P. and J. E. Stiglitz, "Market Structure and the Nature of Innovative Activity", paper presented to Tokyo Meeting at International Economic Association, August 1978b.

Evenson, R. and Y. Kieslev (1976), "A Model of Technological Research", Journal of Political Economy.

Kamien, M. and N. Schwartz (1972), "Timing of Innovations under Rivalry", Econometrica, 40, No. 1, pp. 43-60.

Kamien, M. and N. Schwartz (1975), "Market Structure and Innovation: A Survey", Journal of Economics Literature, 13, No. 1, pp. 1-37.

Kamien, M. and N. Schwartz (1978), "Optimal Exhaustible Resource Depletion with Endogenous Technical Change", Review of Economic Studies.

Loury, G. (1976), "Market Structure and Innovation", Northwestern University, Center for Mathematical Studies in Economics and Management Science, Discussion paper No. 256.

Nelson, R. R., S. G. Winter and H.L. Schuette (1976), "Technical Change in an Evolutionary Model", Quarterly Journal of Economics, Vol. 90, No. 1, pp. 90-118.

Radner and J. E. Stiglitz (1975), "Fundamental Non-Convexities in the Value of information", Stanford University, mimeo.

Scherer, F. M. (1970), Industrial Market Structure and Economic Performance, Rand McNally & Co., (Chicago).

Publicación Impresa por
Organización Norte S.C.A.
Avda. Pte. R. S. Peña 744
Buenos Aires