12 Technological Learning, Social Learning and Technological Change

Raaj Kumar Sah
YALE UNIVERSITY
and
Joseph E. Stiglitz
PRINCETON UNIVERSITY

A central puzzle facing development economists is why it is that the growth rates and income levels of various countries have not converged faster than they have. Indeed, there is some evidence that there has been divergence for many less developed countries (LDCs), rather than convergence. Traditional neoclassical growth theory (Solow, 1956) predicts that, in the long run, the growth rates in all countries should be related only to the rate of technological progress and of population growth; growth rates in per capita incomes should be related only to the rate of labour augmenting technological progress; and differences in levels of per capita consumption should be related to differences in savings rates. Even if the LDCs adopt the best practices of the developed countries with a lag, the rates of technological progress will be the same, and differences in levels of per capita income will then be related also to the length of the lag in the diffusion of technology.

This paper presents two different perspectives providing alternative explanations of the non-convergence. One is based on certain characteristics of technology. The other is based on socioeconomic considerations. The two perspectives have quite different policy implications.

1. THE TECHNOLOGICAL PERSPECTIVE

This perspective is based on three aspects of technology which have received insufficient attention in the literature:

Reprinted from:
The Balance Between Industry and Agriculture, Proceedings of 8th World Congress of the IEA, Vol. 3. Ed: S. Chakravarty. Macmillan, London, 1989, Ch. 12, pp: 285-98.

(i) Much of learning is a by-product of production. Though Arrow (1962) called attention to this phenomenon of learning-by-doing, it has played a minor role in the subsequent literature on growth theory.

(ii) Much of learning is localised, that is, accretions in knowledge that are relevant to one technology may have little bearing on other technologies. Spill-overs are far from uniform.[1]

(iii) The process of learning is, itself, learned. Thus, just as Adam Smith emphasised the importance of specialisation in production, what has become increasingly apparent is the importance of specialisation in learning. Just as the worker who specialises in producing pins becomes more proficient in pin production, unless he suffers from boredom, so too the individual who specialises in research may become more proficient in doing it.[2]

These three simple observations concerning technological learning have important implications for economic theory and policy. We summarise the major implications below.

1.1 Optimality of Non-myopic Policies

The optimal policies (both at the firm level, and for the economy) are non-myopic; that is, the technology that a firm should employ, as well as its level of production (and, indeed, whether it should produce or not) cannot be decided solely by looking at current factor and commodity prices. If the firm believes that, at some future date, it will pay to switch to a more capital-intensive technology, then it pays to switch at some date prior to the date at which, at current factor prices, the more capital-intensive technology has lower costs.

If there are significant gains in learning-to-learn, or learning-by-doing, then it pays to produce more than the level at which price (or marginal revenue) equals short-run marginal costs. A corollary of this is that it may pay to enter into an industry, even when at current market prices the firm makes a loss, because of the learning or the increase in learning capacity which results.

These results follow from the localised nature of technological progress, and from the fact that there is learning-by-doing. The presence of learning-to-learn simply strengthens the arguments. It implies that it may pay to produce some commodity for which the country not only does not currently have a comparative advantage, but

for which it may never have a comparative advantage. But what it learns from producing that commodity may result in the country eventually having a comparative advantage in some new product. For instance, it paid Japan to enter the production of computer chips at a time when it had a comparative disadvantage, because what it learned in the process enabled it to attain a comparative advantage in the next set of products to be developed.

1.2 Non-convexities and Imperfect Competition

As Arrow recognised in his early essay, learning-by-doing gives rise to non-convexities. These non-convexities are exacerbated by the presence of localised learning and learning-to-learn effects.

Non-convexities are of importance for at least two reasons. First, they imply that specialisation is often advantageous. While in the absence of non-convexities, the country might pursue a gradual process of transition from, say, labour-intensive to capital-intensive technologies, the non-convexities may make this gradual approach non-optimal.[3]

Secondly, market equilibria are not likely to be perfectly competitive. Dasgupta and Stiglitz (forthcoming), for instance, have observed that, in the presence of learning-by-doing effects, the market equilibrium with free entry may be characterised by a single dominant firm earning monopoly profits; previous decisions concerning output and technology represent effectively sunk costs. Further, as is by now known (Dasgupta and Stiglitz, 1985), with even small sunk costs, potential competition is not sufficient to ensure that profits are driven to zero.

1.3 Historical Nature of the Growth Process

A feature of neoclassical models which makes them useful for analytical purposes, but suspect for other purposes, is their ahistorical character: history does not matter. In the long run the economy converges to where it would have converged in the absence of all those events on which historians focus, such as wars and plagues. This is not true of the model of technology which we have formulated. The fact that some technique was employed at some time changes for ever the shape of the production possibilities schedule. If the Black Plague induced firms to use more capital-intensive techniques than they

otherwise would have produced, then it is those techniques which got developed. If the Black Plague had not occurred, some other technique would have been developed.

Again,. learning-to-learn effects strengthen these conclusions. Previous lack of experiences may limit the capacity of individuals in some economies to take full advantage of the capacities for productivity improvement associated with certain technological changes.

1.4 Multiplicity of Equilibria

It is possible to show that one consequence of our technological hypotheses is that there may be multiple long run equilibria. Assume, for instance, that more capital-intensive technologies have a greater capacity for learning. That is, the rate of increase in productivity associated with any increment in production is greater in such technologies. Then, some economy may be trapped in a low-level equilibrium, with a low capital labour ratio and a low rate of technological progress, even though there exists another equilibrium with a high capital labour ratio, and a high rate of technological progress. Thus, our model provides an explanation of the non-convergence of rates of growth as well as levels of income.

This explanation is, however, not completely convincing. In the context of an international economy in which knowledge can move across national borders, less capital intensive economies have potential access to the information available in more developed countries. At the same time, the theory of localised technological change provides a part of the explanation for why such knowledge, developed for more developed economies, may have limited relevance for LDCs. The latter face the direct costs as well as other problems of adapting these technologies to local conditions, such as a higher variance in the quality of inputs. Those with the less capital-intensive technologies may thus be deterred from using the technologies developed in the more developed countries.

1.5 Some Policy Conclusions

Our analysis has some important potential implications for the appropriate technology policy and for industrial policy. We draw these conclusions with caution, since their validity depends, not only on the

existence of the theoretical possibilities, but also on the establishment of the quantitative importance of the effects which we have noted.

If the learning potential is lower for the labour-intensive technologies (say because they have been around for a much longer period, so that all the 'obvious' insights have already been exhausted), then it may be desirable to use a more capital-intensive technique. Thus, if firms are myopic, their myopia (leading to excessively labour-intensive technologies, from a social viewpoint) may partly offset the bias towards excessive capital-intensive technologies resulting from excessively high real wages.[4]

On the other hand, it may be desirable to employ labour-intensive technologies, even though such technologies are currently dominated by more capital-intensive technologies. This will be the case if, for instance, there remains large unexploited technological possibilities with these techniques. Though these techniques may have been around for a long time, scientific attempts to improve these technologies may have been limited.[5]

The arguments presented so far for the use of non-myopic rules apply equally well to the firm as to society. Unless one is convinced that government bureaucrats are more foresighted than business entrepreneurs, there is no convincing case for government intervention. There are, however, four related potential sources of market failure, suggesting the possible desirability of government intervention:

(a) *Externalities*

The firm making or adapting an innovation seldom reaps all of the benefit; there are important spill-overs to other firms in the same industry[6] or to other industries. In agriculture, these externalities may be particularly important, where the success of some varieties of seeds may quickly lead to the adoption of the innovation by other farms, driving down the price and lowering the total expected returns for the farmer who risked trying the new seed. The fact that technical change is localised may reduce the significance of these externalities. Indeed, appropriability problems may result in excessively localised technical progress.

(b) *Co-ordination problems*

The adoption of new technologies often requires the existence of a 'technological infrastructure' to facilitate the adoption of new tech-

niques. The absence of organisations facilitating such changes is both a cause and an effect of the lack of technological progress. Thus, there are multiple equilibria to the economy: there is one equilibrium in which there is little technological change, and another one with more rapid technological change.

The structure of the argument is a familiar one. In the absence of a complete set of markets, there frequently exist multiple Nash equilibria for the economy. If individuals only consume coffee with sugar, and only consume sugar with coffee, then in the absence of coffee, there will be no sugar produced, and in the absence of sugar, no coffee produced. There can exist a market equilibrium with neither produced, which is inferior to one in which both are produced.

The necessity for co-ordination has often been viewed as a reason for government intervention. For the kinds of example just presented, such arguments are not entirely persuasive, because if the interactions are sufficiently obvious, then the market can do the necessary co-ordination just as well as any government bureaucracy. When US Steel decided to construct a steel mill on the southern shore of Lake Michigan, they internalised all of these externalities, by simultaneously constructing the railroad, the necessary mines, as well as providing housing and public goods for their employees.

But the kinds of externalities with which we are concerned in this paper are diffuse. The technological infrastructure that facilitates change and adaptation, reflecting the specialisation in learning (one of the three aspects of technical change which we have emphasised) may serve a large number of enterprises, provided it is not too localised. Of course, there may be instances of large enterprises of sufficient scale that much of the 'technological change externalities' may be internalised; that is, even in the absence of the availability in the market of specialised information gathering services, it pays the firm to establish its own specialised internal unit.[7,8]

(c) *Imperfect competition*

The non-convexities with which we have been concerned give rise to natural monopolies, and more generally to imperfect competition. Because the gains which accrue to consumers from lower prices are not appropriated by the firm, firms' decisions with respect to the level of production and the choice of technique will not always be socially optimal.

(d) *Imperfect capital markets*

Socially optimal levels of production and choices of techniques may entail losses. Credit and equity constrained firms may thus be induced to produce at lower levels and with inappropriate techniques. Though it is now recognised that these capital market constraints do not just represent the capriciousness of the market-place, but rather may be a consequence of imperfections of information and the real costs of obtaining information (Stiglitz and Weiss, 1983; Greenwald, Stiglitz and Weiss, 1984). Still it is also clear that the resulting market equilibrium is not constrained Pareto-optimal (Greenwald and Stiglitz, 1986). Thus, there exist in principle, government interventions which are Pareto-improving.

It has sometimes been argued that the appropriate remedy for this market failure is for the government to intervene in the capital market, not in the production decisions of the firms directly. However, the history of government interventions in credit markets does not make us sanguine about its ability to remedy the consequences of imperfect capital markets directly. In particular, it should be remembered that the government faces the same or worse imperfections of information and costs of acquiring it that the private sector does. In addition, government credit programmes provide opportunities for hidden subsidies, whenever interest rates do not accurately reflect actuarial probabilities of repayment, with all the political consequences which follow.

2 SOCIAL LEARNING

The previous discussion focused on certain critical aspects of technology, which might result in their being multiple equilibria. In this part of the paper we present an alternative perspective, based on a theory of social learning due to which some societies may be characterised by high levels of innovation, and others by low levels. This view should be contrasted with the kind of determinism so popular in the late eighteenth and nineteenth centuries.

The main feature of our argument is that a central determinant of the survival value of certain types of characteristics in the population is the nature of the environment. At the same time the economic environments are themselves endogenous, and an important determinant of the environment is the population mix of characteristics. These are the arguments to be derived from more basic economic logic.

Sah (1985, 1986, 1987, 1988) has developed a new class of models to explain a number of observed patterns of phenomena such as corruption, crime and dishonesty which hitherto have not been adequately explained. In these models, individuals face concrete economic reasons why they must learn, at least to some degree, from the economic environment of the past periods. The past thus exerts stochastic (but systematic) influences on the current economic choices of heterogeneous individuals. The precise form of influence is, of course, derived from individual microeconomic considerations, and it differs from context to context depending on, for instance, who the relevant economic actors are, what their choices are, and what the nature of their interactions is. When the individuals' choices in each period are aggregated to the economy-wide level,[9] then one obtains dynamic relationships relating the current period's economic environment (and associated variables such as the levels of different types of economic activities) to the past variables of the economy.

We wish to argue here that analogous dynamic externalities might be important in understanding certain aspects of the lack or the presence of technological change in LDCs. We have developed this view more fully in Sah and Stiglitz (1988). In the rest of this section, we briefly present the qualitative arguments.

We argue that the demand for 'innovativeness' is an increasing function of the fraction of the population which is innovative; and the current supply of innovativeness is an increasing function of the fraction of individuals who were innovative in previous periods. As a result, there may exist multiple equilibria, in one of which there is a relatively small fraction of individuals who are innovative, in another the fraction may be relatively large. Such positive feedback effects among innovators, and correspondingly negative feedback effects between innovators and non-innovators (say, 'bureaucrats') may arise from several sources.

First, as we emphasised earlier, the process of learning itself is learned. The frame of mind which is associated with asking, 'How can this task be performed better?' is different from the frame of mind which is associated with asking, 'How am I supposed to perform this task?' And there are better and worse ways of going about learning about how to perform a task better. Moreover, much of learning, and learning-how-to-learn, occurs in a social context. If there are more individuals with high learning capacities, the chances of learning-how-to-learn are improved. Also, not only the capacity to learn is acquired, but also tastes for innovation or routinisation. Thus an innovative

person placed in a bureaucratic environment may switch to being a routine-lover. The probability of such a switch is a function of the fractions of innovators and bureaucrats in the economy.

Secondly, most changes in methods of production within organisations require the acquiescence of many individuals. A worker who discovers a better way of making widgets must get the approval of his boss. An owner who finds a better way of producing a widget must get his workers to go along. Most of production occurs within organisations, and organisations can facilitate or inhibit the introduction of innovations.

The nature of organisations, in turn, depends on the demands imposed on the organisation for adaptation, which itself may be a function of the state of innovativeness of the society. The supply of innovation gives rise, in a sense, to its own demand. In unchanging environments, there is little demand for individuals and institutions who know how to cope with changes or who are specialised in the collection, processing and dissemination of information.

Thirdly, the kinds of regulations and reporting requirements often imposed by bureaucracies may serve to inhibit innovativeness, creating negative feedback effects.[10] Thus, innovators and bureaucrats each may make the life of the other more difficult: a bureaucrat may have the power to suppress an innovation, or at least make it more difficult for innovations to occur; innovative individuals can create situations where bureaucratic routines do not apply. Thus, in bureaucratic environments, bureaucratic individuals and organisations thrive, and innovators may wither away, and conversely in innovative environments.

Fourthly, for a number of reasons including those stated above, the optimal behaviour of each individual depends on his beliefs about the nature of those with whom he interacts. These expectations are, in turn, a function of the past mix of individuals in the population. If there are more innovators, it becomes more likely that the person with whom one interacts is an innovator. Also, there may be strategic considerations. If it pays to behave in an 'innovative' way when dealing with a particular innovator, it becomes more profitable for the individual to behave in that way. Note, however, that strategic considerations are not essential for an individual's choice to depend on his beliefs. Also, such considerations may not be central to an understanding of those aspects of social environment in which the change in the environment due to the actions of any *one* individual is small.

A reduced-form description of the above ideas is as follows. Con-

sider an economy in which individuals are characterised by some variable, say, the degree of innovativeness. For simplicity, we shall assume that the variable can take only two values; that is, an individual is either an innovator or a bureaucrat. The proportion of innovators in the population at date t is denoted $x(t)$.

In the reduced-form, the economy will be characterised by dynamic relationships linking $x(t)$ to the variables of different past periods, and the nature of dynamic relationships will be determined by the micro-structure outlined above. In some simple cases it will be possible to depict the economy using transition rules under which $x(t)$ depends only on $x(t-1)$. In such cases, the transition rules themselves will be a function of the population mix at time t. Thus we write

$$x(t+1) = A(x(t))x(t).$$

If A were independent of x, then under standard conditions there would exist a unique steady-state x^* satisfying

$$x^* = Ax^*.$$

However, since the environment determines the nature of the transitions, A in general depends on x. Accordingly, there may be multiple solutions to the equation

$$x^* = A(x^*)x^*.$$

Whenever there are multiple equilibria, one of which is more 'socially desirable' than the other, there is scope for government interaction, to shift the economy from one equilibrium to another. Whether, and at what cost, the government can do this depends, of course, on the instruments it has at its disposal. An investigation of these questions would take us beyond the scope of this brief paper.

Those brought up in the Spencerian tradition of Social Darwinism may wonder how it is possible that these two alternative forms of social organisation can both survive. Surely, one must be better than the other, and the better one will survive? But that analysis missed the central point that we have emphasised: the survival of organisations depends on the environment, which, to a large extent, is endogenous. Bureaucrats create an environment in which bureaucratic modes of behaviour have survival value; innovators create an environment in which innovative modes of behaviour have survival value. One cannot

rely on any natural selection argument for the evolution of efficient social organisations.

3 CONCLUDING REMARKS

There is an old saying that some of the best things in life are free. Only a fraction of our interactions are completely regulated by the price system. Any parent recognizes that his child picks up attitudes not only from his parents but from peers, and from a wide variety of environmental influences. More than that, one picks up modes of thought and behaviour. In all cultures information is constantly exchanged among individuals. But the nature of the information exchanged may differ from place to place. At some places like the Silicon Valley, it is about the most recent developments in computers, at other places it may be the latest piece of juicy gossip. Whether economically motivated or not, these social interactions may have a profound effect on economic behaviour. There are, to use the economist's traditional jargon, important externalities arising out of these social interactions. This is not to say that social interactions are not affected by economic returns. Yet the central point that we have emphasised, the possibility of multiple outcomes, remains.

A conventional economist might object by saying that if information about computers has economic value, firms will enter to provide that information; indeed, that it is more efficient to have such information provided by the market mechanism just like any other good. This objection is significantly inaccurate for the issues discussed in the present paper. As has been repeatedly emphasised elsewhere, information is not like an ordinary good. In many ways, it is like a public good; non-convexities are important; and there are intrinsic problems associated with the sale of information. It may be more efficient to disseminate information as a by-product of other forms of interaction. Given that there is a low level of innovation, there is no incentive for information-producing and disseminating firms to enter; and given the absence of such firms, innovation remains at a low level.

The two perspectives agree on two fundamental points. First, there has been a long-standing view that there is a well-defined process of development; that the less developed countries must go through the same process of industrialisation and urbanisation that the currently developed countries went through in the past. This view is wrong because it ignores the historical nature of technology development,

which we stressed in the first part of this chapter. The technological opportunities facing the LDCs today are different from those facing the more developed countries at comparable stages of their development. For instance, rapid improvements in agriculture technology have occurred during the past quarter of a century. Also, since today's LDCs are in a position of being imitators, their optimal strategy is markedly different from those that are on the leading edge of the development of new technologies, for whom the appropriability problem is central.

Secondly, the two perspectives we have discussed also agree that the development strategy should not be based simply on static comparative advantage, and that dynamic comparative advantages may differ markedly from static comparative advantages.

There is one respect in which these two perspectives may differ in a fundamental way. If the argument of the second part of this paper is correct, then central to the process of technological change is the transformation of a society into an innovative and adaptive culture. What is required for that is more than the shipment of capital, or the construction of oil refineries and tyre factories. Indeed, the central role that the government has played in selecting out the more able individuals, educating them, and then hiring them to work as bureaucrats may have served, in the long run, to suppress the development of such a culture. If this argument is correct, it has profound implications for how we should think of development strategies.

Notes

1. Again, although the notion of localised learning was discussed almost two decades ago by Atkinson and Stiglitz (1969), its full consequences have yet to be explored.
2. Note, however, that there are dangers from excessive specialisation. Specialised individuals may be less able to adapt to changes that are sufficiently far removed from their specialisation, though their greater specialised knowledge may make them more able to adapt to changes that are within their sphere of specialisation. The trade-offs between specialisation and adaptability have not been adequately studied.
3. Risk aversion provides a further rationale for gradual transitions. However, the non-convexities with which we are concerned here can be sufficiently great to overcome the effects of risk aversion.
4. Stiglitz (1976) has argued that there may be good economic reasons for high real urban wages; these wages may increase the productivity of workers in the urban sector. See Sah and Stiglitz (1985) for a treatment of the consequences of these and other considerations on the relationship between the shadow wage and market wage.

5. In this case, learning is not just a function of accumulated experience, but how this experience is processed. The environmental factors which are crucial in determining how experiences are processed are the concern of the second part of this chapter.
6. In small LDCs, where there may be no other firm in the same industry, this may not be so important.
7. The importance of a technological change infrastructure has been particularly apparent in the rapid development of small-scale office computerisation, where a whole industry has developed specialising in adapting software and hardware to the particular needs of users. This provides an example where the co-ordination problem was solved without government intervention; but whether it will always be solved in this way is not so obvious.
8. It should also be noted that the co-ordination problem we have focused on is quite different from that addressed in much of the planning literature, which has focused on the material balances between different industries.
9. If the full dynamic economic consequences are to be captured, then the appropriate model to work with is the overlapping generations model. This can be seen in the papers cited earlier.
10. As usual, there are exceptions. There have been instances, over the short run, of successfully-run bureaucratic enterprises, though these have generally had well-focused objectives. Also, it should be clear that, throughout this section, we have in mind a stylised model of bureaucrats and innovators. In practice, lines are never so neatly drawn.

References

Arrow, K. J. (1962) 'The Economic Implications of Learning by Doing', *Review of Economic Studies*, vol. 29, pp. 155–93.

Atkinson, A. B. and Stiglitz, J. E. (1969) 'A New View of Technological Change', *Economic Journal*, vol. 79, pp. 573–8.

Dasgupta, P. and Stiglitz, J. E. (1985) 'Sunk Costs, Competition and Welfare', mimeo, Princeton University.

Dasgupta, P. and Stiglitz, J. E. (forthcoming) 'Learning-by-Doing, Market Structure, and Industrial and Trade Policies', *Oxford Economic Papers*.

Greenwald, B. and Stiglitz, J. E. (1986) 'Externalities in Economics with Imperfect Information and Incomplete Markets', *Quarterly Journal of Economics*, vol. 100, pp. 229–64.

Greenwald, B., Stiglitz, J. E. and Weiss, A. (1984) 'Informational Imperfections and Macro-economic Fluctuations', American Economic Review, vol. 74, pp. 194–9.

Sah, R. K. (1985) 'What Affects the Level of Honesty in an Economy?' Working paper, Series D, School of Organization and Management, Yale University.

Sah, R. K. (1986) 'Rapid Economic Changes versus Stagnation: The Dynamic Roles of Individual and Social Beliefs', Working paper, Yale University.

Sah, R. K. (1987) 'Persistence and Pervasiveness of Corruption: New Perspectives', paper presented at the Conference on Political Economy, Washington, DC (June).

Sah, R. K. (1988) 'Social Osmosis and Patterns of Crime: A Dynamic Economic Analysis', paper presented at Columbia, Maryland, Toronto and Yale (Spring).

Sah, R. K. and Stiglitz, J. E. (1985) 'The Social Cost of Labour, and Project Evaluation: A General Approach', *Journal of Public Economics*, pp. 135–63.

Sah, R. K. and Stiglitz, J. E. (1988), 'Technological Progress and Socioeconomic Structure', forthcoming in the conference volume in the memory of Carlos Diaz-Alejandro.

Solow, R. M. (1956) 'A Contribution to the Theory of Economic Growth', *Quarterly Journal of Economics*, vol. 70, pp. 65–94.

Stiglitz, J. E. (1976) 'The Efficiency Wage Hypothesis, Surplus Labour and the Distribution of Income in LDCs', *Oxford Economic Papers*, pp. 185–207.

Stiglitz, J. E. and Weiss, A. (1983) 'Incentive Effects of Terminations: Applications to the Credit and Labour Markets', *American Economic Review*, vol. 72, pp. 912–27.