LOCALISED AND BIASED TECHNOLOGIES: ATKINSON AND STIGLITZ'S NEW VIEW, INDUCED INNOVATIONS, AND DIRECTED TECHNOLOGICAL CHANGE*

Daron Acemoglu

This study revisits the important ideas proposed by Atkinson and Stiglitz's seminal 1969 paper on technological change. After linking these ideas to the induced innovation literature of the 1960s and the more recent directed technological change literature, it explains how these three complementary but different approaches are useful in the study of a range of current research areas – though they may also yield different answers to important questions. It concludes by highlighting several important areas where these ideas can be fruitfully applied in future work.

Atkinson and Stiglitz's seminal paper in the ECONOMIC JOURNAL 1969, 'A New View of Technological Change', took an important departure from the orthodoxy of its time, which assumed that technological improvements could be viewed as increasing productivity at all factor proportions (in particular, at all combinations of capital and labour). In the language of modern growth theory, technological progress was neutral – in the simplest form, Hicks neutral, creating the same proportional gain in output regardless of factor proportions, though the Harrod neutral version that became more central to growth theory is also broadly similar (Acemoglu, 2009). Atkinson and Stiglitz, instead, noted that it would be much more plausible to assume that technological progress is localised and improves the productivity of the techniques (or 'activities') currently being used and perhaps some similar techniques with neighbouring capital–labour ratios, rather than all techniques regardless of whether or not they are being used and how far they are from the current practice. This situation is illustrated in Figures 1 and 2, adapted from Atkinson and Stiglitz (1969). Figure 1 represents a stylised version of the orthodoxy, while Figure 2 shows the improvement at the technique currently in use – corresponding to the current capital–labour ratio, k_0 – together with an improvement in 'neighbouring' techniques (as well as a further improvement at k_1 which we will discuss later).

Atkinson and Stiglitz (1969, p. 575) emphasised how their new view would be natural in the context of 'learning-by-doing' but did not neglect the case where technological progress results from research and development activity, writing:

> But where technical progress is localised to one technique, there is a second important question [in addition to the optimal amount of research and development] that we must answer – which technique should we improve? Research activity can be directed towards the improvement of any process but
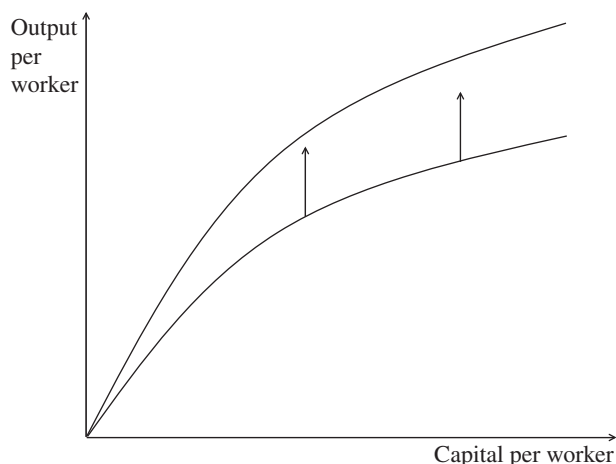
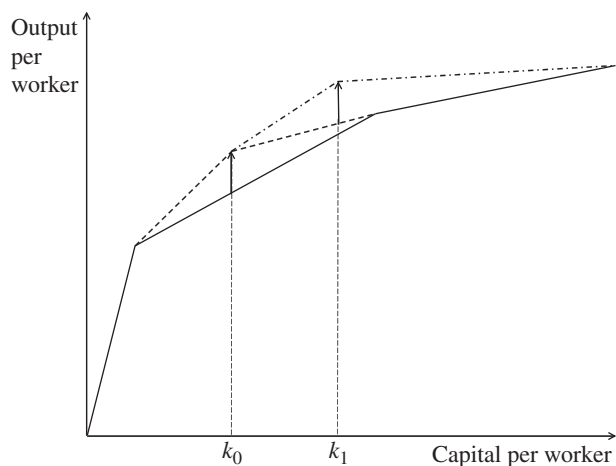Fig. 1. *A Technological Improvement Increases Productivity by Similar Amounts at Capital–Labour Ratios*

Fig. 2. *A Technological Improvement Increases Productivity at the Current and Neighbouring Capital–Labour Ratios, with No Effect on Productivity at Farther Capital–Labour Ratios*

once it has been carried out, the resulting knowledge is specific to one particular process.

Though they did not emphasise it, another important implication of this new view – again using a term that has since become more widely used – is that technological change is biased. This can readily be observed from Figure 2, where the slope of the production function, corresponding to the marginal product of capital in this case, changes very differently at different ratios of capital to labour. Moreover, as the above quotation indicates, Atkinson and Stiglitz did anticipate that technological change has to be modelled as directed – towards specific capital–labour ratios or the specific techniques.

Atkinson and Stiglitz were not the first to make these observations. As they note in passing, an earlier literature on 'induced innovations', which can be more directly viewed as the harbinger of the endogenous growth literature of the 1980s and 1990s,

also concerned itself with the same questions. Just like Atkinson and Stiglitz's new view, the induced innovation literature was ahead of endogenous growth in one sense. It went beyond studying the process of growth at the aggregate and also strove to understand what type of innovations the economy would generate and what the implications of these innovations were for factor prices and the factor distribution of income (but differently from Atkinson and Stiglitz's approach, it did not attempt to unpack the production function by working with technological progress at the level of techniques or 'activities'). It seems to have been Hicks (1932, p. 124) who first discussed these issues in *The Theory of Wages*, in particular when he wrote:

> A change in the relative prices of the factors of production is itself a spur to invention, and to invention of a particular kind – directed to economizing the use of a factor which has become relatively expensive

This was followed up with the significant progress by, *inter alia*, Kennedy (1964), Drandakis and Phelps (1965), Samuelson (1965) and Ahmad (1966) – interestingly, all except Samuelson's article were also published in the ECONOMIC JOURNAL in the 1960s.

Despite these important contributions, the orthodoxy that Atkinson and Stiglitz were criticising is still fairly influential. But important advances that are closer in spirit to Atkinson and Stiglitz's vision have also been made, in large part because several central questions, including the role of appropriate and inappropriate technology in economic development, skill-biased technological change and wage inequality, the relationship between new technologies and tasks and the impact of trade on technological change, necessitate a clear departure from the conception of neutral technological change towards localised, biased and directed technological change.

In the rest of this short article, I first explain the logic of both the induced innovation literature and Atkinson and Stiglitz's seminal paper. In Section 2, I contrast them to the more recent directed technological change literature. In Section 3, I discuss how Atkinson and Stiglitz's insights have been part of certain modern analyses – even though in many cases these analyses themselves did not directly build on their work.

1. Modelling the 'New View'

In this Section, I briefly discuss the modelling approaches that Atkinson and Stiglitz (1969) and the induced innovation literature adopted.

1.1. *Modelling Localised Technological Progress*

Atkinson and Stiglitz formulated how their new view would correspond to a different type of shift of a production function but did not present a detailed model of where this type of technological progress comes from. There are two obvious ways of approaching this question. The first, which Atkinson and Stiglitz stress, is through learning-by-doing. If the firm (or the economy) uses a specific technique – corresponding to a specific capital–labour ratio – then that specific technique will

get better over time. This type of technological progress was emphasised and discussed by, among others, Stewart (1977), and modelled by Basu and Weil (1998). In their model, the aggregate economy operates at a capital–labour ratio determined by the stock of capital (given by past saving decisions) and the inelastically supplied labour in the economy. Then in the next period, productivity improves at the capital–labour ratio currently in use and at neighbouring capital–labour ratios – but not at all capital–labour ratios.

We can see how this type of technological progress works with the help of Figure 2 above. Suppose that after the technological improvement starting at the capital–labour ratio k_0, the economy accumulates further capital, increasing its capital–labour ratio to k_1. Then with a similar logic, there will be a localised improvement at k_1. We can now see that capital–labour ratios above k_1 will benefit indirectly from both improvements but the extent of this will depend on how localised these spillovers are. I return to the implications of this type of localised technological progress in the context of a specific application in Section 3 below.

Though learning-by-doing is undoubtedly important for the improvement of certain technologies, it seems unsatisfactory to assume that this is the only, or the primary, form of technological progress. A great number of innovations in the world today are a result of purposeful activity, as evidenced by the fact that companies (and governments) spend a significant fraction of their resources for research and development, patent this research and use and defend the resulting patents (see Hall, 2011, for a recent survey). But it is far from straightforward to have purposeful research and development activity be the engine of economic growth together with this new view – or for that matter together with the induced innovation literature.

While Atkinson and Stiglitz never attempted to formulate a model in which biased technological change arises endogenously from research and development activities, this was the main focus of the induced innovation literature. The difficulties that they faced illustrate a number of important issues, as I discuss below.

1.2. Modelling Induced Innovations

I now consider a prototypical model of induced innovation. Suppose that the economy is populated by many firms, each with a constant returns to scale production function,

$$Y = F(N_L L, N_Z Z), \tag{1}$$

where L is labour, which is assumed constant (inelastically supplied) throughout the article, Z is another factor of production, for example capital or another type of labour, and N_L and N_Z are factor-augmenting technology terms, which are controlled by each individual firm. The induced innovation literature assumes that firms hire the profit-maximising amount of factors but they choose their technologies to maximise 'the current rate of cost reduction' for given factor proportions (Kennedy, 1964, p. 543; Drandakis and Phelps, 1965, p. 824). This is equivalent to maximising the rate of output growth, R, taking Z and L as given. This rate of output growth, holding Z and L constant, is
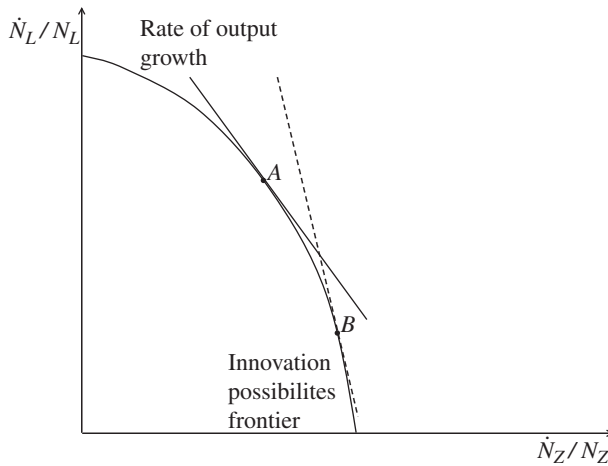
Fig. 3. *An Increase in the Share of Factor Z in GDP, s, Makes the Counters Representing the Rate of Output Growth Steeper. This Changes the Point of Tangency with the Innovation Possibilities Frontier from A to B, Inducing Further Z-augmenting Technological Change*

$$R = (1 - s)\frac{\dot{N}_L}{N_L} + s\frac{\dot{N}_Z}{N_Z}, \tag{2}$$

where $s = w_Z Z/Y$ is the share of factor Z in GDP.[1] The constraint facing firms, and thus the maximisation problem is given by the 'innovation possibilities frontier', a concept first introduced by Kennedy (1964), which I write as

$$\frac{\dot{N}_L}{N_L} = \Gamma\left(\frac{\dot{N}_Z}{N_Z}\right), \tag{3}$$

where Γ is a strictly decreasing, differentiable and concave function. This innovation possibilities frontier captures the trade-off that, at the frontier, a higher rate of labour-augmenting technological change must come at the expense of a lower rate of technological change favouring the other factor, Z. Once formulated in this way, the solution to the problem of determining endogenous technology is straightforward and some of its implications are discussed in the next Section. For now, it is also useful to note that the determination of endogenous technology in this model can be represented as a point of tangency between the contours of (2) and the innovation facilities frontier given by (3) as shown in Figure 3, which also shows how this point of tangency changes when the share of factor Z in GDP, s, increases as I discuss further below.

But does this formulation make sense? Given our modern modelling sensibilities, the answer must be no. The objective function posited for the firms, maximising (2), is difficult to motivate (as already noted by Samuelson (1965) and Nordhaus (1973). The

---

[1] To obtain this expression, differentiate both sides of (1) with respect to time, holding Z and L constant, and use the fact that the marginal product of labour is equal to the wage rate and the constant returns to scale assumption on F.

natural thing would have been simply to assume that each firm, taking all prices as given, maximises the discounted value of its profits. This dynamic maximisation problem can be written as

$$\max_{Z,L,N_Z,N_L} \int_0^T \exp(-rt)[F(N_L L, N_Z Z) - w_L L - w_Z Z]\mathrm{d}t \tag{4}$$

subject to (3) and for given factor prices (relative to the price of the final good), w_L and w_Z (where I set the planning horizon arbitrarily to some T, which could be infinity). But a moment's reflection will show why induced innovation literature shied away from it. In general, (4) is not a convex optimisation problem and does not have an interior solution because the production function exhibits increasing returns to scale in factors of production and technology terms combined.

The induced innovation literature came before the advances in the modelling of monopolistic competition, which formed the basis of Romer's (1990) model of endogenous growth, as well as that of Grossman and Helpman (1991) and Aghion and Howitt (1992). Without modelling monopolistic competition (or monopoly or oligopolistic competition), there was no obvious way of overcoming the non-convexity of (4), unless one adopted the assumption of Romer (1986) and Lucas (1988) that technological progress was generated as a byproduct externality of investment in physical capital or education, which is ultimately unsatisfactory for the same reason as noted above – that is, because much of innovation results from purposeful, profit-maximising decisions.

The same difficulty that dogged the induced innovation literature would have prevented progress in modelling research and development activity within the context of Atkinson and Stiglitz's new view.

2. Directed Technological Change

A more recent literature, one with which I have been heavily involved, shows how induced innovations and Atkinson and Stiglitz's new view of technological change can be microeconomically modelled. The approach, first developed in Acemoglu (1998) and Kiley (1999), is similar to that used in the first-generation endogenous technological change models such as Romer (1990), Grossman and Helpman (1991) and Aghion and Howitt (1992), but with a richer menu of technologies towards which innovations can be directed – hence the designation directed technological change. Here, I present a stripped-down version similar to Acemoglu (2002). As in the first-generation endogenous technological change models, new technologies create monopoly power for their inventor, thus leading to monopolistic competition. But differently from these models, the implications are much richer as in the induced innovation literature and Atkinson and Stiglitz's vision. In this Section, I briefly overview this approach. In the context of the applications presented in the next Section, I then show how the problems with which the induced innovation literature concerned itself can be analysed (but with some notably different conclusions) when new innovations can be directed to factor-augmenting technologies, and how Atkinson and Stiglitz's new view can be obtained when new innovations are directed to techniques or tasks.

Suppose that a continuum of firms has access to the production function for producing a unique final good,

$$Y = F(Y_L, Y_Z), \tag{5}$$

where Y_L and Y_Z are two intermediate goods that are themselves produced competitively with the following production functions.[2]

$$Y_L = \left[\int_0^{N_L} x_L(j)^{1-\beta} \mathrm{d}j \right] L^\beta \text{ and } Y_Z = \left[\int_0^{N_Z} x_Z(j)^{1-\beta} \mathrm{d}j \right] Z^\beta, \tag{6}$$

where $\beta \in (0,1)$. The labour-intensive intermediate good is produced from labour and a range of labour-complementary machines, $x_L(j)$ denotes the amount of the j-th labour-complementary (labour-augmenting) machine used in production. The range of machines that can be used with labour is denoted by N_L. The production function for the other intermediate differs only because it uses Z-complementary machines. That these two sets of machines are different is important because it implies that some technologies will be augmenting labour, while others increase the productivity of factor Z. In fact, the functional form in (6) is immaterial and is only adopted for transparency here as I discuss below.

Suppose that machines in both sectors are supplied by profit-maximising 'technology monopolists'. Each monopolist will set a rental price $\chi_L(j)$ or $\chi_Z(j)$ for the machine it supplies to the market in order to maximise its profits. For example, we can assume that technology monopolists have access to production functions for new technologies – or to the innovation possibility frontier in Kennedy's terminology – given by

$$\dot{N}_L = \eta_L N_L^{(1+\delta)/2} N_Z^{(1-\delta)/2} S_L \quad \text{and} \quad \dot{N}_Z = \eta_Z N_L^{(1-\delta)/2} N_Z^{(1+\delta)/2} S_Z, \tag{7}$$

where $\delta \in (0,1)$ is a parameter that determines the relative degree to which advances in one technology depend on the current state of knowledge in that technology relative to the other one (Acemoglu, 2002).

Observe that for given N_L and N_Z, the production functions in (6) exhibit constant returns to scale. However, when N_L and N_Z are also treated as choice variables, there will be increasing returns to scale in the aggregate as in (4) above. Despite this non-convexity, the reason why the profit-maximisation problems of firms in this economy are well defined is because technology choices, N_L and N_Z, will be made by a different set of agents – the technology monopolists – than the producers of Y_L and Y_Z (see the discussion in Acemoglu (2007), on the importance of this assumption).

Acemoglu (2002) characterises dynamic equilibria (and balanced growth paths) in this model, emphasising in particular how the direction (and bias) of technological change depends on relative abundance of different types of labour.

An important generalisation of this framework studied in Acemoglu (2007) is worth mentioning. In particular, this article generalises the framework mentioned above such that the overall (aggregate) production function can be written as

$$Y = F(L, Z, \boldsymbol{\theta}),$$

---

[2] This formulation is identical to (1). To see this, rewrite that problem as in here. In particular, suppose that the final good is produced from two intermediates as in (5) with these two intermediates being produced competitively as $Y_L = N_L L$ and $Y_Z = N_Z Z$.

where $\boldsymbol{\theta}$ is a vector of technologies (thus generalising the functional form in (6)). The vector $\boldsymbol{\theta}$ could include factor-augmenting technologies, like N_L and N_Z above, but also potentially includes other types of technologies that affect, say, the elasticity of substitution between factors, or create localised technological change as in Atkinson and Stiglitz's vision. Acemoglu (2007) characterises the nature of technological change in this economy and shows that the relationship between factor supplies and equilibrium bias of technology derived in Acemoglu (2002) extends to this more general environment.

3. Applications

In this Section, I provide several applications of the ideas presented so far, stressing how the micro-founded directed technological change model gives different or more nuanced implications than the induced innovation literature and Atkinson and Stiglitz's approach, but also highlighting certain important parallels and interesting new directions based on Atkinson and Stiglitz's ideas.

3.1. *Appropriate Technology*

One of the important applications of Atkinson and Stiglitz's idea that technological progress is localised – rather than being neutral – is one that the authors themselves mentioned: the potential inappropriateness of frontier technologies to less developed economies (Stewart, 1977). Consider the case of agricultural technology for example. Because frontier technologies are developed in rich, capital-intensive countries, the resulting agricultural technologies will be capital-intensive and tend to improve the effectiveness of such capital-intensive techniques (such as tractors and mechanical harvesters), and even use computer and GPS technology to improve productivity. But the contribution of these technologies to a capital-scarce agricultural economy where such machinery is absent may be limited. Basu and Weil (1998) modelled this idea in a framework where all technological progress results from learning-by-doing as in Atkinson and Stiglitz's baseline approach.

Though this idea has obvious appeal, it also illustrates some of the major shortcomings of Atkinson and Stiglitz's modelling approach. The reason why advanced countries are so important for frontier technologies is that they undertake research and development to create and develop new technologies. But this cannot be captured in a model in which all technological progress results from learning-by-doing.

The induced innovation and directed technological change literatures may, therefore, offer a more appealing framework for studying these effects. Acemoglu and Zilibotti (2001) were the first, to the best of my knowledge, to have done so using this complementary approach. Acemoglu and Zilibotti focused on differences in human capital rather than capital–labour ratio. This is motivated both by the increasing importance of human capital for modern technology (and the sharp differences in human capital endowments across countries) and by the link that this creates between ideas of appropriate technology and skill-biased technological change (which I discuss next).

I now present the Acemoglu–Zilibotti model briefly to illustrate how, for this application, these different approaches lead to similar and complementary conclusions. Acemoglu and Zilibotti utilise a framework similar to the one presented in Section 2, with $Z = H$ corresponding to skilled labour. The world economy consists of two sets of countries, the North and the South, the first corresponding to richer countries with a higher relative endowment of skilled to unskilled labour (i.e. higher H/L). All research and development takes place in the North using an innovation possibilities frontier similar to (7) above. But the South can costlessly copy frontier technologies from the North. Because intellectual property rights are not enforced in the South, Northern technology monopolists cannot make any profits from the Southern market. This implies that all technological change will be directed to the relative skill abundance of the Northern market. However, because skilled labour is scarcer in the South, this implies that technologies developed in the North are partly inappropriate to the South. In particular, Acemoglu and Zilibotti show that measured productivity will be endogenously higher in the North (controlling for differences in capital intensity). Relatedly, in a hypothetical world in which new technologies were directed to the world ratio of skilled to unskilled labour, the South would have been richer relative to the North than it is in this equilibrium.

The similarities between this story and Atkinson and Stiglitz's model are worth emphasising. In both models, technology can be inappropriate to the needs of poorer economies (because of differences in capital intensity or scarcity of skill). In both cases, this is because technological change is biased and localised. In particular, technological change is localised in Acemoglu and Zilibotti's model as well because new technologies are created with an innovation possibilities frontier, just like (7), where research directed to skill-intensive products generates limited spillovers to labour-intensive products and *vice versa*. However, there are also important differences between these approaches. For example, Acemoglu and Zilibotti's analysis shows why these results depend on the international enforcement of intellectual property rights. If these were enforced, then research would also be directed more towards Southern needs and the extent of this effect would be more limited. In addition, international trade can exacerbate the inappropriateness of technology by creating a price effect (see also Acemoglu, 1998, 2003*b*, as well as Diwan and Rodrik, 1991). This suggests that a modelling approach in which new technologies are determined by forward-looking research behaviour is not only better able to make contact with stylised patterns and data on innovation and patents but also generates richer comparative statics and new insights.

3.2. *Skill-biased Technological Change*

A major application of the new view emphasised by Atkinson and Stiglitz (1969) as well as of the induced innovation literature and the directed technological change literature is to skill-biased technological change. There is now a huge empirical and theoretical literature on this topic (see Autor *et al.*, 1998, for the empirical relationship between technology and wages, and Acemoglu and Autor, 2011, for a recent review). Much of this literature, however, treats technology as exogenous. But an idea already clear from Hicks's quote given above is that the type of technologies, and hence skill bias of technology, should respond to various factors in the economy, primary among

them the abundance or skill of different types of labour. Though this was emphasised by Hicks, neither the induced innovation literature nor Atkinson and Stiglitz dwelt much on the endogenous skill bias and the inequality implications of their approach.

Hicks's statement above, which was partly formalised by the induced innovation literature, can be read in one of two ways. First, 'economising on a factor' can be interpreted as 'biased against that factor', in which case Hicks's statement can be read as suggesting that technology should become endogenously biased in favour of a factor that becomes more 'expensive' or scarce. Second, it can be interpreted as 'augmenting that factor', in which case Hicks would appear to argue that technology will naturally tend to further augment a factor that becomes more 'expensive'.

As shown in Acemoglu (2003a), a variant of the second interpretation can be derived from the above stylised model of induced innovations. In particular, from Figure 3 introduced above, we can see that the relative rate of growth of N_Z (relative to that of N_L) and thus N_Z/N_L will be increasing in the factor share of Z (i.e. in $s = w_Z Z/Y$): a higher s leaves the innovation possibilities frontier unchanged but makes the contours of (2) steeper as shown in the Figure. This then increases N_Z/N_L, which is also shown in the Figure. Setting $Z = H$ again, this implies that technology will become relatively more augmenting to H, that is N_H/N_L will increase, when $s = w_H H/Y$ increases.

However, this result implies that things are a little more nuanced than a simple reading of Hicks's statement because what is important is not the scarcity of a factor but its factor share, s, which is a non-monotonic function of its relative abundance (or conversely, scarcity) H/L. The factor share s is increasing in H/L if the elasticity of substitution between the two factors is greater than 1, thus making N_H/N_L increasing in H/L, and decreasing in H/L when this elasticity is less than 1 (Acemoglu, 2003a), hence making N_H/N_L decreasing in H/L in this case. In summary, whether relative scarcity of H will encourage further advances in N_H or N_H/N_L depends on the elasticity of substitution between the two factors. The answer is yes when elasticity of substitution is greater than 1 and no when it is less than 1.

What about Atkinson and Stiglitz's approach? Though these authors do not provide a framework that can readily speak to these issues, the following would be a natural formalisation of their baseline scenario for the origins of localised technological progress which, as already mentioned above, relies on learning-by-doing. Suppose, in a similar spirit to the discussion in the previous subsection, that localised technological change implies that learning-by-doing will spread to techniques employed in the same sector – that is, greater H will imply more learning-by-doing in the H-sector and thus higher N_H. This of course implies that N_H/N_L will be increasing in H/L, so that relative scarcity of H will discourage further advances in N_H or N_H/N_L, which is quite different from the induced innovation literature's conclusion. Which perspective is correct?

It turns out that the correct answer is closest, but not identical, to the induced innovation literature – and thus somewhat more subtle than what Hicks anticipated or the localised learning-by-doing *à la* Atkinson and Stiglitz would imply.

In Acemoglu (2007), I show that, under very weak regularity conditions, whenever a factor becomes more scarce, technology becomes endogenously biased against that factor. This conclusion is independent of whether the menu of technological changes is assumed to be factor augmenting. It clearly goes against the first reading of Hicks's statement.

If, in addition, we assume that there are only two factors and only factor-augmenting technologies are allowed, then whenever a factor becomes more scarce, technology will tend to further augment that factor if the elasticity of substitution between the two factors is less than 1 and the opposite conclusion (technology endogenously augmenting the factor becoming more abundant) holds when the elasticity of substitution is greater than 1 (Acemoglu, 2002, 2007). This is related closely to the result from the version of the induced innovation literature presented above. But crucially, it is derived from a coherent, micro-founded model of innovation incentives rather than the induced innovation literature's *ad hoc* rule of maximising the rate of growth of output technology choices, while maximising profits with factor demands.

The importance of microeconomic modelling and the fact that a micro-founded model does give different and richer results can be highlighted by considering another set of results derived in Acemoglu (2002, 2007), which have no equivalents in the induced innovation literature or in Atkinson and Stiglitz's approach. If, in addition, the elasticity of substitution between factor process is sufficiently large (in particular, larger than a threshold that is equal to $2 - \delta$ in the model highlighted above, where δ is the parameter in the innovation possibilities frontier, (7), above), the relative demand curves for factors with endogenous technology is upward sloping (and under a related condition, demand curves for factors are also upward sloping). In particular, this implies that, in this case, technology responds to changes in skill supply so much that when H/L increases, w_H/w_L also increases – rather than decreasing as in basic producer theory – in the long-run equilibrium after technology has fully adjusted to the change in supplies and prices.

3.3. *Tasks and Technologies*

Another modern application of the ideas in Atkinson and Stiglitz (1969) is to models of labour markets with task-specific technologies. In particular, Acemoglu and Autor (2011), building on Zeira (1998), Acemoglu and Zilibotti (2001) and Autor *et al.* (2003), consider the following model of the labour market. Suppose that there are four factors of production, high, medium and low-skilled workers and capital, with respective supplies given by L, M, H and K. Aggregate output is produced by combining the output of a continuum of tasks, normalised to lie on the unit interval. For example, it may be given by $\ln Y = \int_0^1 \ln y(i)\mathrm{d}i$, where $y(i)$ denotes the output of task i.

Let us assume that each task can be produced according to the production function

$$y(i) = A_L\alpha_L(i)l(i) + A_M\alpha_M(i)m(i) + A_H\alpha_H(i)h(i) + A_K\alpha_K(i)k(i), \qquad (8)$$

where A terms represent factor-augmenting technology, and $\alpha_L(i)$, $\alpha_M(i)$, $\alpha_H(i)$ and $\alpha_K(i)$ are the task productivity schedules, designating respectively, the productivity of low, medium and high-skill workers and capital in different tasks. This specification implies, for instance, that $\alpha_L(i)$ is the productivity of low-skill workers in task i, and $l(i)$ is the number of low-skill workers allocated to task i. Acemoglu and Autor (2011) argue that a model along these lines provides a much better framework for changes in wage and employment patterns in the US and other advanced labour markets over

the last three decades, in particular because this allows for technologies replacing certain sets of tasks while at the same time reducing the wages of certain groups of workers.

This model shares with Atkinson and Stiglitz (1969) the feature that new technologies exist at the task/activity level. It is potentially richer than Atkinson and Stiglitz's formulation, however, because some new technologies can increase productivity of a task when performed by one type of worker without affecting the productivity of other factors at this task (e.g. $\alpha_H(i)$ may increase while the other α terms remain constant).

As we have already noted, Atkinson Stiglitz did not provide a detailed analysis of how localised technological change could take place. In this framework, however, where task-based technologies come from is as important as their implications for wages and employment. Acemoglu and Autor (2011) use results from Zeira (1998) and Acemoglu (2007, 2010) to show how this can be done. First, as emphasised by Zeira (1998) and Acemoglu (2010), high wages could be an impetus for more rapid technological change. But more importantly, and in line with Acemoglu (2010), these changes will be directed towards certain tasks, and hence will be localised and biased. Though these results take some steps towards a general framework for endogenous task-based technological change, much more remains to be done here, and Atkinson and Stiglitz's approach will remain relevant for these new investigations.

3.4. *Other Applications*

Several other applications of ideas stemming from the induced innovation literature and from Atkinson and Stiglitz are important, even if most of them have been even less well developed than those already discussed. Here, I mention a few.

First, Atkinson and Stiglitz are not clear on whether what matters is capital per worker at the economy level or at the firm level. If the latter, then equilibria could be asymmetric. In particular, firms could choose different capital–labour ratios in order to benefit from different types of localised technological changes. Figure 4 illustrates this possibility by showing that if there are two different types of firms with (very) different capital–labour ratios (k_0 and k_1), then the aggregate production function could become endogenously non-convex, encouraging firms in the future to choose different capital–labour ratios, perpetuating this non-convexity. To the best of my knowledge, this topic has not been investigated systematically in any of the literatures discussed so far. This type of non-convexity might be particularly important for less-developed economies, creating another reason why dual economies may emerge in such societies. In particular, if technologies imported from the world technology frontier have undergone much improvement only in high capital–labour ratios, then despite the relatively high price of capital, some firms in developing economies may end up choosing to operate at these high capital–labour ratios, leaving even lower capital–labour ratios for the rest of the economy.

Second, another important set of applications of these ideas is to the relationship between trade and technology. Grossman and Helpman (1991) investigated the relationship between trade and the overall rate of technological progress. Acemoglu (2003*b*) showed how trade can have a critical impact on the direction of technological
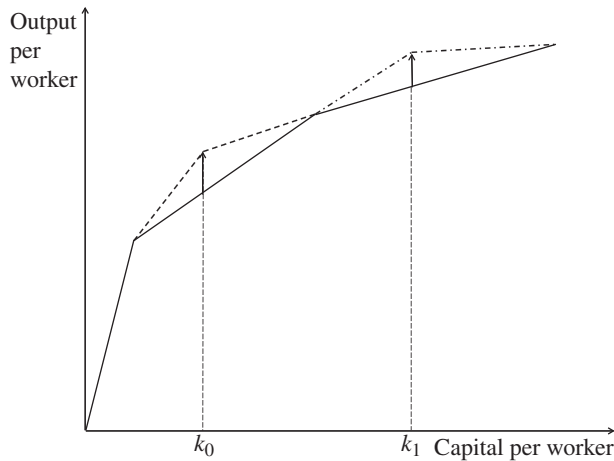
Fig. 4. *After Productivity Improvements at the Capital–Labour Ratios Used by the Two Types of Firms (k_0 and k_1), the New Aggregate Production Function is No Longer Concave*

change (Thoenig and Verdier, 2003; Epifani and Gancia, 2008; Acemoglu *et al.*, 2013). The same forces would be present in Atkinson and Stiglitz's modelling approach based on learning-by-doing on the basis of the current factor proportions but with some notable differences. The type of learning-by-doing envisaged by Atkinson and Stiglitz would imply that as international trade leads to greater specialisation, it will also cause a divergence between the technologies of the trading countries (divergence in terms of which localised techniques are witnessing improvements). This is, in fact, sharper than in the directed technological change models, because in these models, as emphasised in Acemoglu (1998, 2003*b*) and Acemoglu and Zilibotti (2001), the effect of trade on technology depends on how well intellectual property rights are enforced internationally. As already noted above, when such property rights are enforced, trade need not lead to a divergence of technologies because technology monopolists would find it profitable to produce machines useful for both countries. In contrast, localised technological change following from the techniques currently in use will necessarily lead to a divergence regardless of whether intellectual property rights are enforced internationally.

Third, and relatedly, directed technological change models do not generally lead to a major technological divergence between advanced economies (though see the last Section in Acemoglu, 2003*b*). Atkinson and Stiglitz's approach might have different implications. In particular, if the US and European economies operate at different factor proportions because of differences in relative supplies or institutional differences, this could shape the path of learning-by-doing and create different types of localised technological improvements.

Fourth, a major current debate concerns whether new technologies are creating 'technological unemployment' whereby many workers are displaced by new technologies and find it difficult to become employed again. Critics of this view point to the fact that employment has increased steadily in the face of major (and often disruptive) new technologies throughout the last 200 years, while its proponents emphasise the inability of an important fraction of the US (and European) labour force to find jobs

or at the very least the stagnation or even the decline in real wages partly caused by these new technologies. Localised technological change might put important new ideas on the table here. If technological change has historically taken place at techniques that increase the wages and demand for labour but because of institutional, technological and supply reasons it now takes place at techniques that do not do so, then the presence of technology-driven unemployment today may not contradict the historical record.

Finally, the potential impact of these ideas – in any application – will ultimately depend on their ability to make contact with data and receive support from careful empirical work. This has been the case in the study of changes in wage structures, which has focused on the theoretical and empirical investigation of implications of skill-biased technological changes (and sometimes also of the origins of such technological change) as reviewed in Acemoglu and Autor (2011). With advances in microeconometric modelling and the greater availability of micro data sets with rich information on firm, worker and technology characteristics, much more fruitful work combining new theoretical ideas with creative empirical strategies is likely to be forthcoming.

4. Conclusion

In many ways, Atkinson and Stiglitz's short article in the 1969 volume of the ECONOMIC JOURNAL was ahead of its time in emphasising localised and biased new technologies and challenging the orthodoxy in the modelling of technological change. At the same time, it also had many similarities with the prior literature on induced innovations which was also flourishing in the 1960s (and mostly in the pages of the ECONOMIC JOURNAL). Both of these literatures are also closely related to the more recent directed technological change literature, developing microeconomic models of endogenous bias and direction of new technologies and their macroeconomic implications.

After overviewing and linking these three literatures, this article has shown how they generate complementary but sometimes also quite different answers in the context of issues related to appropriate technology, the response of (skill) biased technological change to the abundance of different types of factors, and task-based technologies, and technological change. It has also highlighted several important areas where the ideas in Atkinson and Stiglitz's seminal article and the induced innovation and direct technological change literatures can be fruitfully applied in future work. Though the orthodoxy that Atkinson and Stiglitz challenged, which ignores the biased and localised nature of technological change, is still widespread in much of macro-economics, several important literatures, perhaps most importantly the literature on changes in the wage structure, have recognised the importance of biased technological change and even modelled the endogenous determination of this bias. On this trend, Atkinson and Stiglitz's ideas may apply even more centrally in the next 50 years than the last 50 years.

Massachusetts Institute of Technology

References

Acemoglu, D. (1998). 'Why do new technologies complement skills? Directed technical change and wage inequality', *Quarterly Journal of Economics*, vol. 113, pp. 1055–90.

Acemoglu, D. (2002). 'Directed technical change', *Review of Economic Studies*, vol. 69, pp. 781–810.

Acemoglu, D. (2003*a*). 'Factor prices and technical change: from induced innovations to recent debates', in (P. Aghion, R. Frydman, J. Stiglitz and M. Woodford, eds.), *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, pp. 464–91, Princeton NJ: Princeton University Press.

Acemoglu, D. (2003*b*). 'Patterns of skill premia', *Review of Economic Studies*, vol. 70, pp. 199–230.

Acemoglu, D. (2007). 'Equilibrium bias of technology', *Econometrica*, vol. 75, pp. 1371–410.

Acemoglu, D. (2009). *Introduction to Modern Economic Growth*, Princeton, NJ: Princeton University Press.

Acemoglu, D. (2010). 'When does labor scarcity encourage innovation?', *Journal of Political Economy*, vol. 118, pp. 1037–78.

Acemoglu, D. and Autor, D. (2011). 'Skills, tasks and technologies: implications for employment earnings', in (O. Ashenfelter and D. Card, eds.), *The Handbook of Labor Economics*, vol. 4b, pp. 1043–1171, Amsterdam: Elsevier.

Acemoglu, D., Gancia, G. and Zilibotti, F. (2013). 'Offshoring and directed technical change', CEPR Discussion Paper 9247.

Acemoglu, D. and Zilibotti, F. (2001). 'Productivity differences', *Quarterly Journal of Economics*, vol. 116, pp. 563–606.

Aghion, P. and Howitt, P. (1992). 'A model of growth through creative destruction', *Econometrica*, vol. 110, pp. 323–51.

Ahmad, S. (1966). 'On the theory of induced invention', ECONOMIC JOURNAL, vol. 76, pp. 344–57.

Atkinson, A.B. and Stiglitz, J.E. (1969). 'A new view of technological change', ECONOMIC JOURNAL, vol. 79, pp. 573–78.

Autor, D., Krueger, A. and Katz, L. (1998). 'Computing inequality: have computers changed the labor market?', *Quarterly Journal of Economics*, vol. 113, pp. 1169–214.

Autor, D., Levy, F. and Murnane, R. (2003). 'The skill content of recent technological change: an empirical exploration', *Quarterly Journal of Economics*, vol. 118, pp. 1279–334.

Basu, S. and Weil, D.N. (1998). 'Appropriate technology and growth', *Quarterly Journal of Economics*, vol. 113, pp. 1025–54.

Diwan, I. and Rodrik, D. (1991). 'Patents, appropriate technology and north-south trade', *Journal of International Economics*, vol. 30, pp. 27–47.

Drandakis, E. and Phelps, E. (1965). 'A model of induced invention, growth and distribution', ECONOMIC JOURNAL, vol. 76, pp. 823–40.

Epifani, P. and Gancia, G. (2008). 'The skill bias of world trade', ECONOMIC JOURNAL, vol. 118, pp. 927–60.

Grossman, G. and Helpman, E. (1991). *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press.

Hall, B.H. (2011). 'Innovation and productivity', Working Paper No. 17178, NBER.

Hicks, J. (1932). *The Theory of Wages*, London: Macmillan.

Kennedy, C. (1964). 'Induced bias in innovation and the theory of distribution', ECONOMIC JOURNAL, vol. 74, pp. 541–7.

Kiley, M. (1999). 'The supply of skilled labor and skill-biased technological progress', ECONOMIC JOURNAL, vol. 109, pp. 708–24.

Lucas, R.E. (1988). 'On the mechanics of economic development', *Journal of Monetary Economics*, vol. 22, pp. 3–42.

Nordhaus, W. (1973). 'Some skeptical thoughts on the theory of induced innovation', *Quarterly Journal of Economics*, vol. 87, pp. 208–19.

Romer, P.M. (1986). 'Increasing returns and long-run growth', *Journal of Political Economy*, vol. 94, pp. 1002–37.

Romer, P.M. (1990). 'Endogenous technological change', *Journal of Political Economy*, vol. 2C, pp. S71–102.

Samuelson, P. (1965). 'A theory of induced innovations along Kennedy-Weisacker lines', *Review of Economics and Statistics*, vol. 47, pp. 444–64.

Stewart, F. (1977). *Technology and Underdevelopment*, London: The Macmillan Press Ltd.

Thoenig, M. and Verdier, T. (2003). 'A theory of defensive skill-biased innovation and international trade', *American Economic Review*, vol. 93, pp. 709–28.

Zeira, J. (1998). 'Workers, machines and economic growth', *Quarterly Journal of Economics*, vol. 113, pp. 1091–113.

Appendix A. Atkinson, A.B. and Stiglitz, J.E. (1969). 'A new view of technological change', ECONOMIC JOURNAL, vol. 79(315), pp. 573–8.

A NEW VIEW OF TECHNOLOGICAL CHANGE[1]

THE recent literature on technological progress has almost entirely been based on the assumption that its effect can be represented as shifting the production function outwards—as illustrated in Fig. 1. Technical advance is assumed to raise output per head for all possible techniques. The advocates of this approach seem, however, to have forgotten the origins of the neo-classical production function: as the number of production processes increases (in an activity analysis model), the production possibilities can be more and more closely approximated by a smooth, differentiable curve. But the different points on the curve still represent different processes of production, and associated with each of these processes there will be certain technical knowledge specific to that technique. Indeed, both supporters and critics of the neoclassical theory seem to have missed one of the most important points of the activity analysis (Mrs. Robinson's blue-print) approach: that if one brings about a technological improvement in one of the blue-prints this may have little or no effect on the other blue-prints. If the effect of technological advance is to improve one technique of production but not other techniques of producing the same product, then the resulting change in the production function is represented by an outward movement at one point and not a general shift—see Fig. 2. This figure
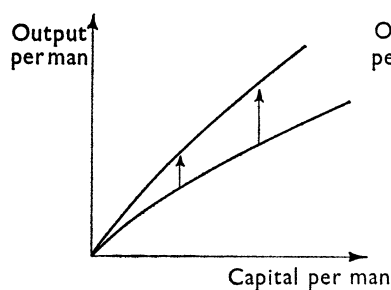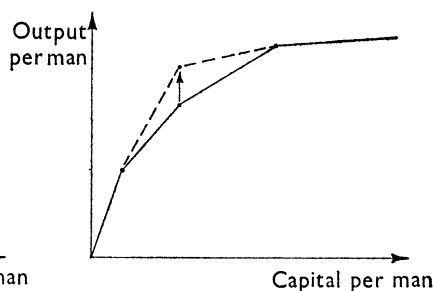

FIG. 1.   FIG. 2.

shows the extreme case where technical progress is completely " localised " to one technique: there are no spillover improvements in other techniques. It reality we should expect that a given technical advance would give rise to some spillovers and that several techniques would be affected. However, we would reach the traditional position only if there were spillovers to *every* technique. This means that a technical advance would have to be

such as to raise productivity on, say, every type of textile loom from the fully automated to the crudest hand loom.[1] In this note we shall examine some of the implications of the " localisation " of technical progress and contrast this view with the traditional assumption that technical progress leads to a general shift in the production function.

<div align="center">I</div>

For the most part, increases in technical knowledge involve either experience in production (" learning by doing ") or research activity. These factors have both received considerable attention in the literature, but where technical progress is localised there are several important implications which have not yet been brought out. We consider first the case of learning by doing. If the knowledge acquired through learning is localised, then the shift in the production function will be located at the point where the firm (or economy) is now operating. This in turn means that when a firm is deciding which technique to use it must take account of the effect of its current choice of technique on future production possibilities. Consider the case of a firm for which learning is effectively internalised and which is choosing between two techniques, one more capital-intensive than the other. If it adopts the more labour-intensive technique, then this means that the productivity of this technique will be increasing through learning, while that of the capital-intensive technique will be unchanged. It cannot, therefore, base its choice of technique solely on current factor prices, but must take account of the value of the increase in knowledge associated with each technique—in other words, it cannot behave completely myopically. The cost-minimising condition is

$$C_1 - p_1 = C_2 - p_2$$

where C_i is the marginal cost on the ith technique; p_i is the value of the gain in knowledge from producing one more unit using the ith technique. If the firm places a sufficiently high value on increases in knowledge about the second technique, then it may use this, even though it has a higher cost at current factor prices. (Such an argument could be used to justify, for example, the adoption of atomic power plants: even though the cost per kilowatt hour is higher, the knowledge gained in their construction offsets the additional cost.)   Suppose that the firm expects wages to rise and that at some point in time it will switch from the labour-intensive technique (1) to the more capital-intensive technique (2) (and will never switch back). When should it make the switch?   At the point of switching, p_1 will be zero—knowledge about the first technique has no value, since it will never be used again.   On the other hand, p_2 will in general be positive, so that

---

[1] Indeed, the usual assumption of Harrod neutrality assumes that the percentage reduction of labour requirements per unit of output be the same for all techniques.

the firm will switch when the marginal cost on the capital-intensive technique is higher than that on the labour-intensive technique. This would not, of course, happen where technical progress was not localised.[1]

In the above discussion it was assumed that the benefits from learning by doing were internal to the firm. As Arrow has pointed out,[2] where the benefits are external, there is a case for government intervention to raise investment or output to the socially optimal level. But where technical progress is localised, the Government should be concerned not merely with the level of investment or output, but must make sure that firms are directed towards the " right " technique on long-run considerations. This is particularly relevant to underdeveloped countries trying to encourage infant industries: where technical progress is localised, the Government should be concerned with subsidising infant *techniques* rather than infant industries. In this case a tariff may not merely be inefficient but may be totally ineffective: the Government will want, for example, to encourage an " intermediate " manufacturing technique, but neither traditional nor highly automated methods of producing the same product.

Turning to the case where technical progress is the result of research and innovational activity, there has been a considerable literature on the optimal amount to spend on adding to technical knowledge. But where technical progress is localised to one technique, there is a second important question that we must answer—which technique should we improve? Research activity can be directed towards the improvement of any process, but once it has been carried out, the resulting knowledge is specific to one particular process. In terms of the recent growth-theory literature, technical knowledge has the same characteristics as " putty-clay " capital. This means that when the firm is choosing which technique to develop it must take account of future as well as present factor prices. A firm will not necessarily allocate research expenditure to the technique in current use when it expects that rising wages will lead it to develop and use a more capital-intensive technique in the future. Again, the firm cannot behave myopically.

II

The concept of localised technical progress also throws light on the question as to whether underdeveloped countries should devote resources to developing new techniques of production.

It is sometimes argued that the problem of the allocation of resources to

---

[1] Whether technical change is localised or not, if there is learning by doing, the firm will always produce at a level where marginal revenue is less than short-run marginal cost; equilibrium will require marginal revenue to equal short-run marginal cost minus the present discounted value of the marginal reduction in future costs from the learning.

[2] K. Arrow, " The Economic Implications of Learning by Doing," *Review of Economic Studies*, Vol. 29, June 1962, pp. 155–73.

research is not relevant for a present-day underdeveloped country, since it will benefit from technical progress in the advanced countries, and any independent research would simply be a duplication of effort. But if, as we have suggested, technical knowledge is highly specific to particular production processes this will not be the case. Where technical progress is " localised," technical progress in the advanced countries, whether from research or learning by doing, will leave relatively unaffected the less-capital-intensive techniques that the underdeveloped country would choose in the light of its factor endowment. Indeed, in some industries the effect of localisation has been so strong that the advanced techniques dominate the less capital-intensive ones, requiring both less labour and less capital.[1] Where technical progress is localised to one technique, and there are positive rates of labour and capital augmentation, it is clear how this can happen— see Fig. 3. It is also clear from Fig. 3 that localised technical progress is likely to lead to reductions in the short-run elasticity of substitution.
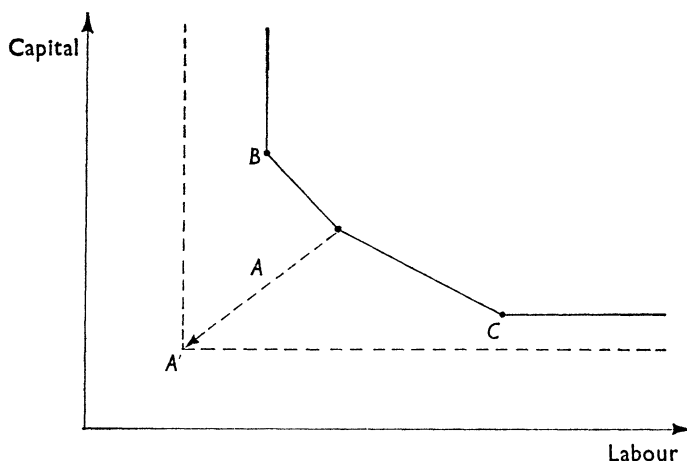

FIG. 3.

The dominance of these advanced capital-intensive techniques has led to a widespread feeling that under developed countries should adopt these rather than more labour-intensive processes. But this takes no account of the possibility of undertaking research: even though techniques with a lower capital–labour ratio may at present be inefficient, it could pay to devote resources to improving them. This is relevant to the recent debate on " intermediate technology," in which it has been argued that under-developed countries should use techniques which require capital of the order of £100 per man rather than £1,000 per man. We are suggesting that even where these techniques do not exist, it may pay these countries to develop

[1] For examples of dominance, see A. K. Sen, *The Choice of Techniques* (Oxford: Basil Blackwell, 1962), and R. S. Eckaus, " The Factor Proportions Problem in Underdeveloped Areas," *American Economic Review, Supplement*, Vol. L, May 1960, pp. 642–8.

them.   Whether or not it will in fact do so depends on the extent to which the technique is dominated by the advanced technique, on the resources required to improve it, on the time horizon of the planners and so on.

III

Our approach may be contrasted with some recent versions of the theory of induced innovation, in which there is an innovation possibility schedule: the firm is assumed to be faced with the choice between different degrees of factor-augmenting technical progress.   In this case present decisions will affect future production possibilities, just as where technical progress is localised, but at least two unreasonable assumptions are employed: (a) In the Drandakis-Phelps [1] and Samuelson [2] models the choice of the kind of factor-augmenting technical progress is made on the basis of current factor prices only.[3] (b) More fundamentally, the firm's choice is restricted to purely factor-augmenting technical progress.[4] Would it really want to raise productivity on handcarts as well as forklift trucks? [5]

The view of technical progress presented in this note contrasts sharply with the dominant mathematical theories of growth, which are essentially ahistorical in character.   Where technical progress is localised, history is very important.   Suppose, for example, that an economy is in long-run equilibrium using a relatively labour-intensive technique when suddenly a plague wipes out a large proportion of the labour force, so that wages rise and a more capital-intensive technique is adopted.   Technical progress is now localised to this technique, and it is possible that the economy will continue to use it rather than return to the more labour-intensive one.   In this case the history of the economy is qualitatively and quantitatively different from what it would have been had there been no plague.   For instance, the capital–output ratio in the new long-run equilibrium may be

---

[1] E. M. Drandakis and E. S. Phelps, " A Model of Induced Invention, Growth, and Distribution," ECONOMIC JOURNAL, December 1966, pp. 823–40.

[2] P. A. Samuelson, "A Theory of Induced Innovation on Kennedy–von Weizsäcker Lines," *Review of Economics and Statistics*, 1965.

[3] The early literature on induced innovation made much more sense from this point of view. See, *e.g.*, W. Fellner, " Two Propositions in the Theory of Induced Innovation," ECONOMIC JOURNAL, June 1961, pp. 305–8.   W. Nordhaus, " The Optimal Rate and Direction of Technical Change," *Essays on the Theory of Optimal Economic Growth*, K. Shell, ed. (Cambridge: M.I.T. Press, 1967), has made some progress in eliminating this unrealistic assumption.

[4] If there is only one technique available at any point of time the whole question of localisation is, of course, irrelevant.   See C. Kennedy, " Induced Bias in Innovation and the Theory of Distribution," ECONOMIC JOURNAL, September 1964, pp. 541–7.

[5] The recent article by Drandakis and Phelps (*op. cit.*) illustrates some of the other difficulties which this approach runs into.   As they recognise, there is no reason to restrict the innovation possibility schedule to the positive orthant; if, however, there is negative augmentation of one factor the " new " isoquant will intersect the " old " one.   They then worry whether the new isoquant is really the envelope of the two, or whether one must tell some story about forgetting techniques.   The whole difficulty arises because they assume that technical progress applies to *all* processes.

larger than in the old; if there is more " learning by doing " associated with more capital-intensive techniques the rate of technological change may be increased.   But if history is important, so then is planning present activities with a view to their long-run consequences.   In this note we have discussed some of the implications for the firm and for a planned developing country.

ANTHONY B. ATKINSON

St. John's College, Cambridge.

JOSEPH E. STIGLITZ

*Cowles Foundation for Research in
    Economics, Yale University.*