

NBER WORKING PAPER SERIES

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

Daron Acemoglu

Working Paper 20060

<http://www.nber.org/papers/w20060>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2014

I am grateful to Anthony Atkinson and Andrea Galeotti for useful suggestions. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Daron Acemoglu. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Localized and Biased Technologies: Atkinson and Stiglitz's New View, Induced Innovations,
and Directed Technological Change

Daron Acemoglu

NBER Working Paper No. 20060

April 2014

JEL No. E25,J31,O30,O31,O33

ABSTRACT

This paper 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.

Daron Acemoglu

Department of Economics, E18-269D

MIT

77 Massachusetts Avenue

Cambridge, MA 02139

and CIFAR

and also NBER

daron@mit.edu

1 Introduction

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 labor). 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 (see Acemoglu, 2009). Atkinson and Stiglitz, instead, noted that it would be much more plausible to assume that technological progress is *localized* and improves the productivity of the techniques (or “activities”) currently being used and perhaps some similar techniques with neighboring capital-labor 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 stylized version of the orthodoxy, while Figure 2 shows the improvement at the technique currently in use—corresponding to the current capital-labor ratio, k_0 —together with an improvement in “neighboring” techniques (as well as a further improvement at k_1 which we will discuss later).

Atkinson and Stiglitz emphasized 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 (p. 575):

“But where technical progress is localized 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 once it has been carried out, the resulting knowledge is specific to one particular process.”

Though they did not emphasize it, another important implication of this new view—




Figure 1: A technological improvement increases productivity by similar amounts at capital-labor ratios.

again using a term that has since then 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 labor. Moreover, as the above quotation indicates, Atkinson and Stiglitz did anticipate that technological change has to be modeled as *directed*—towards specific capital-labor 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 who first discussed these issues in *The Theory of Wages* (1932), in particular when he wrote (p. 124):




Figure 2: A technological improvement increases productivity at the current and neighboring capital-labor ratios, with no effect on productivity at farther capital-labor ratios.

“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), Samuelson (1965), Drandakis and Phelps (1965), and Ahmed (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 criticizing 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 localized, 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 3 I contrast them to the more recent direct technological change literature. In Section 4 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.

2 Modeling the “New View”

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

2.1 Modeling Localized 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-labor ratio—then that specific technique will get better over time. This type of technological progress was emphasized and discussed by, among others, Stewart (1977), and modeled by Basu and Weil (1998). In their model, the aggregate economy operates at a capital-labor ratio determined by the stock of capital (given by past saving decisions) and the inelastically supplied labor in the economy. Then in the next period, productivity improves at the capital-labor ratio currently in use and at neighboring capital-labor ratios—but not at all capital-labor 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-labor ratio k_0 , the economy accumulates further capital, increasing its capital-labor ratio to k_1 . Then with a similar logic, there will be a localized improvement at k_1 . We can now see that capital-labor ratios above k_1 will benefit indirectly from both improvements, but the extent of this will depend on how localized these spillovers are. I will return to the implications of this type of localized technological progress in the context of a specific application in Section 4 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.

2.2 Modeling 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 labor, which is assumed constant (inelastically supplied) throughout the paper, Z is another factor of production, for example capital or another type of labor, 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-maximizing amount of factors, but they choose their technologies to maximize “*the current rate of cost reduction*” for given factor proportions (see, Kennedy, 1964, p. 543, Drandakis and Phelps, 1965, p. 824). This is equivalent to maximizing the rate of output growth, R , taking Z and L as given. This rate of output growth, holding Z and L constant, is

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

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

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

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 labor-augmenting technological change must come at the expense of a lower rate of technological change favoring 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 will be 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 modeling sensibilities, the answer must be no. The objective function posited for the firms, maximizing (2), is difficult to motivate (as already noted by Samuelson, 1965, and Nordhaus, 1973). The natural thing would have been to simply assume that each firm, taking all prices as given, maximizes the discounted value of its profits. This dynamic maximization problem can be written as

$$\max_{Z, L, N_Z, N_L} \int_0^T \exp(-rt) (F(N_L L, N_Z Z) - w_L L - w_Z Z) dt \quad (4)$$

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 optimization problem and does not have an interior

¹To 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 labor is equal to the wage rate and the constant returns to scale assumption on F .




Figure 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.

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 modeling 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 modeling 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—i.e., because much of innovation results from purposeful, profit-maximizing decisions.

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

3 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 modeled. 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 will then show how the problems with which the induced innovation literature concerned itself can be analyzed (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 have access to the production function for producing a unique final good,

$$Y = F(Y_L, Y_Z), \quad (5)$$

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

$$Y_L = \left(\int_0^{N_L} x_L(j)^{1-\beta} dj \right) L^\beta \text{ and } Y_Z = \left(\int_0^{N_Z} x_Z(j)^{1-\beta} dj \right) Z^\beta, \quad (6)$$

where $\beta \in (0, 1)$. The labor-intensive intermediate good is produced from labor and a range of labor-complementary machines, $x_L(j)$ denotes the amount of the j -th labor-

²This 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 $Y_Z = N_Z Z$.

complementary (labor-augmenting) machine used in production. The range of machines that can be used with labor 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 labor, 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-maximizing “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 maximize 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

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

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 (see 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-maximization 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) characterizes dynamic equilibria (and balance growth paths) in this model, emphasizing in particular how the direction (and bias) of technological change depends on relative abundance of different types of labor.

An important generalization of this framework studied in Acemoglu (2007) is worth mentioning. In particular, this paper generalizes the framework mentioned above such

that the overall (aggregate) production function can be written as

$$Y = F(L, Z, \theta),$$

where θ is a vector of technologies (thus generalizing the functional form in (6)). The vector θ 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 localized technological change as in Atkinson and Stiglitz's vision. Acemoglu (2007) characterizes 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.

4 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.

4.1 Appropriate Technology

One of the important applications of Atkinson and Stiglitz's idea that technological progress is localized—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 are absent may be limited. Basu and Weil (1998) modeled 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 does also illustrate some of the major shortcomings of Atkinson and Stiglitz's modeling 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-labor 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 utilize a framework similar to the one presented in Section 3, with $Z = H$ corresponding to skilled labor. 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 labor (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 labor 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 labor, 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 emphasizing. 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 localized. In particular, technological change is localized 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 labor-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 be also 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, 2003b, as well as Diwan and Rodrik, 1991). This suggests that a modeling approach in which new technologies are determined by forward-looking research behavior is not only better able to make contact with stylized patterns and data on innovation and patents, but also generates richer comparative statics and new insights.

4.2 Skill-Biased Technological Change

A major application of emphasized by the induced innovation literature, Atkinson and Stiglitz (1969) 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, Katz and Krueger, 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 labor. Though this was emphasized 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 formalized by the induced innovation literature, can be read in one of two ways. First, “economizing 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 favor 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 stylized 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 H , i.e., 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 formalization of their baseline scenario for the origins of localized 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 localized technological change implies that learning-by-doing will spread to techniques employed in the same sector—i.e., 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 localized 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 (see 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 maximizing the rate of growth of output technology choices while maximizing profits with factor demands.

The importance of microeconomic modeling 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 factorprocess 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.

4.3 Tasks and Technologies

Another modern application of the ideas in Atkinson and Stiglitz (1969) is to models of labor markets with task-specific technologies. In particular, Acemoglu and Autor (2011), building on Autor, Levy and Murnane (2003), Zeira (1998) and Acemoglu and Zilibotti (2001), consider the following model of the labor 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, normalized to lie on the unit interval. For example, it may be given by $\ln Y = \int_0^1 \ln y(i) di$, where $y(i)$ denotes the output of task i .

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

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

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 why a model along these lines provides a much better framework for changes in wage and employment patterns in the United States and other advanced labor markets over the last three decades, in particular allowing for technologies replacing certain sets of tasks and new technologies, while increasing productivity, still 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 localized 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 emphasized 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 localized 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.

4.4 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-labor ratios in order to benefit from different types of localized technological changes. Figure 4 illustrates this possibility by showing that if there are two different types of firms with (very) different capital-labor ratios (the figure k_0 and k_1), then the aggregate production function could become endogenously non-convex, encouraging firms in the future to choose different capital-labor 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-labor ratios, then despite the relatively high price of capital, some firms in developing economies may end up choosing to operate at these high capital-labor ratios, leaving even lower capital-labor ratios for the rest of the economy.




Figure 4: After productivity improvements at the capital-labor ratios used by the two types of firms (k_0 and k_1), the new aggregate production function is no longer concave.

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 (2003b) showed how trade can have a critical impact on the direction of technological change (see also Thoenig and Verdier, 2003, Epifani and Gancia, 2008, and Acemoglu, Gancia and Zilibotti, 2013). The same forces would be present in Atkinson and Stiglitz’s modeling 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 specialization, it will also cause a divergence between the technologies of the trading countries (divergence

in terms of which localized techniques are witnessing improvements). This is, in fact, sharper than in the directed technological change models, because in these models, as emphasized in Acemoglu (1998, 2003b) 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, localized 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, 2003b). Atkinson and Stiglitz’s approach might have different implications. In particular, if the United States 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 localized 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 emphasize the inability of an important fraction of the US (and European) labor force to find jobs or at the very least the stagnation or even the decline in real wages partly caused by these new technologies. Localized 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 labor, 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 em-

pirical 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 microeconomic modeling and the greater availability of micro datasets 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.

5 Conclusion

In many ways, Atkinson and Stiglitz's short paper in the 1969 volume of the *Economic Journal* was ahead of its time in emphasizing localized and biased new technologies and challenging the orthodoxy in the modeling 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 paper 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 Atkinson and Stiglitz's seminal paper 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 localized nature of technological change, is still widespread in much of macroeconomics, several important literatures, perhaps most importantly the literature on changes in the wage structure, have recognized the importance of biased technological change and even modeled 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.

References

Acemoglu, Daron, “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality” *Quarterly Journal of Economics*, CXIII (1998), 1055-1090.

Acemoglu, Daron, “Directed Technical Change” *Review of Economic Studies*, LXIX (2002), 781-810.

Acemoglu, Daron, “Factor Prices and Technical Change: from Induced Innovations to Recent Debates” in Philippe Aghion, Roman Frydman, Joseph Stiglitz and Michael Woodford (editors) *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton University Press, 2003a

Acemoglu, Daron, “Patterns of Skill Premia”, *Review of Economic Studies*, LXX (2003b), 199-230.

Acemoglu, Daron, “Equilibrium Bias of Technology” *Econometrica*, LXXV, (2007), 1371-1410.

Acemoglu, Daron, *Introduction to Modern Economic Growth*, Princeton University Press, Princeton, 2009.

Acemoglu, Daron, “When Does Labor Scarcity Encourage Innovation?” *Journal of Political Economy*, CXIIX, 2010, 1037-1078.

Acemoglu, Daron and David Autor, “Skills, Tasks and Technologies: Implications for Employment Earnings.” In: Ashenfelter, O., Card, D. (Eds.), *The Handbook of Labor Economics*, vol. 4b. Elsevier, Amsterdam, 1043-1171, 2011.

Acemoglu, Daron, Gino Gancia and Fabrizio Zilibotti, “Offshoring and Directed Technical Change,” mimeo, 2013.

Acemoglu, Daron and Fabrizio Zilibotti, “Productivity Differences” (2001), *Quarterly Journal of Economics*, 116, pp.563-606.

Aghion, Philippe and Peter Howitt, “A Model of Growth Through Creative Destruction” *Econometrica*, 110 (1992), 323-351.

Ahmad, Syed, "On The Theory of Induced Invention," *Economic Journal* LXXVI, (1966), 344-357.

Atkinson, Anthony B. and Joseph E. Stiglitz, "A New View of Technological Change," *Economic Journal*, LXXIX (1969), 573-78.

Autor, David, Alan Krueger and Lawrence Katz, "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics*, CXIII (1998), 1169-1214.

Autor, David, Frank Levy and Richard Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* CXIIX, (2003), 1279-1334.

Basu, Susanto and David N. Weil, "Appropriate Technology and Growth," *Quarterly Journal of Economics*, CXIII, (1998), 1025-1054.

Diwan, I. and Dani Rodrik, "Patents, Appropriate Technology and North-South Trade" *Journal of International Economics*, XXX, (1991), 27-47.

Drandakis, E. and Edmund Phelps, "A Model of Induced Invention, Growth and Distribution" *Economic Journal*, 76 (1965), 823-840.

Epifani, Paolo and Gino Gancia, "The Skill Bias of World Trade," *Economic Journal* CXIIX, (2008), 927-960.

Grossman, Gene and Elhanan Helpman, *Innovation and Growth in the Global Economy*, Cambridge, MA, MIT Press, 1991.

Hall, Bronwyn H., "Innovation and Productivity" NBER Working Paper No. 17178.

Hicks, John *The Theory of Wages*, Macmillan, London, 1932.

Kennedy, Charles, "Induced Bias in Innovation and the Theory of Distribution" *Economic Journal*, LXXIV (1964), 541-547.

Kiley, Michael, "The Supply of Skilled Labor and Skill-Biased Technological Progress" *Economic Journal*, CIX (1999), 708-724.

Lucas, Robert E., "On the Mechanics of Economic Development," *Journal of Monetary Economics* XXII, (1988), 3-42.

Nordhaus, William; "Some Skeptical Thoughts on the Theory of Induced Innovation" *Quarterly Journal of Economics*, LXXXVII, (1973), 208-219.

Romer, Paul M., “Increasing Returns and Long-Run Growth” *Journal of Political Economy* 94 (1986), 1002-1037.

Romer, Paul M., “Endogenous Technological Change” *Journal of Political Economy*, IIC (1990), S71-S102.

Samuelson, Paul, “A Theory of Induced Innovations Along Kennedy-Weisacker Lines” *Review of Economics and Statistics*, XLVII (1965), 444-464.

Stewart, Frances, *Technology and Underdevelopment*, The Macmillan Press, Ltd., London, England, 1977.

Thoenig, Mathias and Thierry Verdier, “A theory of defensive skill-biased innovation and international trade,” *American Economic Review* XCIII, (2003), 709-728.

Zeira, Joseph, “Workers, Machines and Economic Growth” *Quarterly Journal of Economics*, CXIII, (1998), 1091-1113.