Wage Inequality Within and Between Groups: A Schumpeterian Perspective.*

by Philippe Aghion[†] and Peter Howitt[‡]

January 12, 2001

1 Introduction

In this paper we examine the increase in wage inequality that has taken place in much of the world during the past three decades, from the point of view of Schumpeterian growth theory,[1] which has been developed to understand the mutual determination of technological change and economic growth. Many economists have come to the conclusion that the increase in inequality is attributable to an increased pace of skill-biased technological change. Here we put forth a somewhat different technology-based argument; namely that the increase in inequality might be caused by technological change that is not skill-biased in the usual sense but that raises the reward to adaptability. The two key elements in this argument are (a) the recognition that skills are not just an input to the production of goods but also a factor in the process of technological change, and (b) the exogenous driving force behind the evolution of the wage structure is the arrival of a new general purpose technology. The theory helps to account for some otherwise puzzling facts concerning wage-inequality both between and within groups, and has policy implications that differ from those that would follow if technological change were skill-biased.

*Preliminary paper prepared for the CSLS-IRPP Conference on Growth and Inequality: The Linkages, Ottawa, January 26-27, 2001. The paper draws heavily on current research being undertaken with Daron Acemoglu and Gianluca Violante.

†Harvard University, CIAR and UCL.

‡Brown University and CIAR

[1] Aghion and Howitt (1992, 1998)

2 Alternative explanations of rising wage-inequality

We start by discussing some of the major facts concerning the recent evolution of wage inequality. Consider first the rise in wage inequality between educational groups. What is perhaps most striking about this rise is that the increase in the wages of college-graduates relative to high-school-graduates that took place in Canada, the US and the UK between the early 1980s and the mid-1990s coincided with an increase in the relative supply of college-graduates. In the US, for example, Autor, Katz and Krueger (1998) show that the ratio of "college-equivalents" (defined as the number of workers with a college degree plus 0.5 of the number of workers with some college education) to "non-college equivalents" (defined as the complementary set of workers) has increased at an average rate of 3.05% between 1970 and 1995, up from an average rate of 2.35% between 1940 and 1970. In the meantime, the ratio between the average weekly wages of college- and high-school graduates has gone up by more than 25 percent during the period 1970-1995 whilst it had fallen by 0.11% a year on average during the previous period.

Consider next the rise in wage inequality *within* educational and age groups. Machin (1996) finds that the *residual* standard deviation in hourly earnings has increased by 23% in the UK and by 14% in the US over the period between 1979 and 1993. Equally intriguing is the fact that the rise in within-group wage inequality started to occur *before* the rise in between-group inequality and accounts for a substantial fraction of the overall increase in income inequality (Katz and Autor, 2000). The final part of this puzzle is that the increase in within-group inequality has mainly affected the *temporary* component of income whereas the increase in between-group inequality has mainly affected the *permanent* component of income (Blundell and Preston, 1999).

While the rise in residual inequality has barely been addressed by the economic literature so far,[2] economists have come up with various explanations for the observed upsurge in educational wage inequality, in particular: trade liberalization, deunionization, and skill-biased technical change.

The trade explanation was fairly straightforward and directly inspired by Heckscher-Ohlin theory: in a nutshell, a globalization boom should drive up the demand for skilled labour in the developed countries where skilled labour is cheap relative to developing countries, and it should drive down the relative demand for unskilled labour which is relatively expensive

[2]See footnote 22 below for references.

in developed countries. Unfortunately, the trade-liberalization explanation is not supported by the evidence. First, as argued by Krugman and others, it is hard to see how trade liberalization could have such a big impact on wage inequality in a country like the US where trade with non-OECD countries represent no more than 2% of GDP. Second, this explanation would imply a fall in prices of less skill-intensive goods relative to prices of more skill-intensive goods in developed countries, but empirical studies find little evidence of this in either the US or Europe during the 1980s. A third implication of the above trade explanation is that labour should be reallocated from low-skill to high-skill industries, or from those sectors in developed countries that are most exposed to international competition to the other sectors. However, Berman, Bound and Griliches (1994) for the US and Machin (1995) for the UK, found that only a minor part (about 20%) of the shift away from manual/blue-collar workers to non-manual/white collars was due to between-industry changes, the remaining 70% or 80% being entirely attributable to within-industry shifts. Fourth, by this argument the rise in wage inequality in OECD countries should have been accompanied by falling inequality in low-wage countries, whereas in fact the rise in inequality seems to be global in nature.

To the extent that unionization is often positively correlated with wage compression,[3] some economists also perceived deunionization[4] as an important source of the observed increase in wage inequality.[5] However, the attempt to attribute the increase in wage inequality to deunionization failed largely on the basis of the following 'timing' considerations: in the UK the rise in wage inequality started in the mid-seventies whilst union density kept increasing until 1980; on the other hand, in the US deunionization began in the 1950s at a time when wage inequality was relatively stable.[6]

Meanwhile, a number of empirical studies, have pointed to a significant impact of skill-biased technical change on the evolution of wage inequality. For example, using R&D expenditures and computer purchases as measures of technical progress, Berman, Bound and

[3] For example, Freeman (1993) showed that the standard deviation of within-firm log wages in the US was 25% lower in unionized firms compared to non-unionized firms.

[4] For example, according to Machin (1997), in the UK union density among male workers fell from 54% in 1980 to 38% in 1990; in the US the percentage of private sector workers that are unionized fell from 24% in 1980 to less than 12% in 1990.

[5] For example, Card (1996) and Fortin and Lemieux (1997).

[6] Although deunionization (organizational change) and trade liberalization do not fully explain the recent evolution in wage inequality, nevertheless we believe that these factors can become more significant when analyzed in relation to skill-biased technical change (see for example Aghion, Acemoglu and Violante, 2000 on deunionization and skill-biased technological change, and Garcia-Penalosa and Koebel, 1999 and Acemoglu, 1999 on trade liberalization and skill-biased technological change).

Griliches (1994) found that these two factors could account for as much as 70% of the move away from production to non-production labour over the period 1979–1987. Murphy and Welch (1993) find that the share of college labour has increased substantially in all sectors since the mid-seventies, which, together with the observed increase in the college premium, provides further evidence of skill-biased technical change. More recently, based on the data reported in Autor, Katz and Krueger (1998) and assuming an elasticity of substitution of 1.4 between skilled and unskilled labour, Acemoglu (2000) estimates that the relative productivity of college graduates has increased from 0.157 in 1980 up to 0.470 in 1990 (whereas this relative productivity had risen at a lower rate prior to the early 1980s).

2.1 Skill-biased technological change: some remaining puzzles

While this technological explanation helps to account for several of the facts that are hard to reconcile with the other explanations, in particular why we observed a sharp increase in the college premium in countries like the US at a time when the relative supply of skilled labour was also increasing, nevertheless it raises some questions to which it seems to provide no easy answer. First, there is the question of why the pace of skill-biased technological change increased sometime after the late 1970s.

One attempt to account for the increased pace of technological change is that of Krusell, Ohanian, Rios-Rull and Violante (1994),[7] who pointed to the increased rate of decline in the relative price of production equipment goods since the mid-1970s,[8] which they interpreted as an increase in the pace of capital-embodied technological progress. They argued that this increase could account for the rise in the college premium because capital and skills are complementary in the aggregate production function. However, their argument begs the question of what caused the increased rate of capital-embodied technological progress. Also, the argument implies that the rise in wage inequality should have been accompanied by a rise in productivity growth, whereas in fact productivity growth was much slower in the two decades following 1975 than during the previous two decades.

Acemoglu (1998, 2000) has provided an alternative explanation for the rise in the pace of skill-biased technological change, according to which the increased relative supply of college-educated workers in the 1970s was responsible for a shift in the direction of technological

[7] See also Stokey (1996) who analyzes the implications of Capital-Skill complementarity for trade, using a similar modelling approach.

[8] See Gordon (1990).

change, which became more skill-biased than before because of a "market-size" effect. That is, suppose that final output is produced by two kinds of intermediate product; one that requires college graduates to operate it (a "skill-intensive product") and one that can be operated by high-school graduates, and technological progress comes from innovations that improve the quality of intermediate products. An R&D firm must direct its efforts towards improving one kind of intermediate product or the other. This choice will be governed by profitability considerations. When the relative supply of college-graduates rises, the relative profitability of improving the skill-intensive product will also increase, provided that the elasticity of substitution between the two kinds of intermediate products in the final goods sector is sufficiently large.

The mechanism works as follows. In the short run, the increased relative supply of college workers will reduce their relative wage (the "skill premium"). This has two opposing effects on the relative profitability of improving the skill-biased product. On the one hand, the lower skill premium will reduce the relative price of the skill-biased product, which reduces the relative cost per unit that can be saved by improving the product. But on the other hand, the increased relative supply of college-graduates implies that in equilibrium there will be an increase in the relative quantity produced of the skill-intensive product, so the per-unit gain from a quality improvement will be earned on a larger volume. The latter effect is the market-size effect, and it will dominate when the elasticity of substitution is large enough. Although the short-run impact, as we have seen, is to reduce the skill premium, the fact that technical progress will be increasingly directed towards the skill-biased product means that the skill-premium will start to grow at a faster rate, because every improvement in a skill-biased product raises the derived demand for skilled labour. Thus the long-run effect will be to increase the skill-premium.

The result is a manifestation of the increasing returns to scale that typically prevail in an economy with endogenous technology, under which an increase in supply of some good or service can lead to an increase in its relative price. It is reinforced by an additional mechanism that is exemplified by the "robot model" of Aghion and Howitt (1998, ch.9). In this model, technological change is presumed always to be skill-biased, taking the form of improved intermediate products that substitute for unskilled labour. Whereas Acemoglu takes as given the overall level of innovation in the economy, Aghion and Howitt takes into account that an increase in the supply of college-graduates will increase the overall innovation rate, given that R&D is a skill-intensive activity. Thus an increase in the relative supply of skilled workers will speed up the pace of skill-biased technological change, thus raising the rate at which the skill-premium rises and leading to a long-run increase in the skill-premium.

These increasing-returns arguments appear to fit the evidence of a wage premium first decreasing (during the early 1970s) and then sharply increasing (starting in the late 1970s), following the increase in relative skilled labour supply in the late 1960s. On the other hand, they raise two difficulties. First, they do not explain why wage inequality failed to increase in other historical episodes involving an increase in the relative supply of educated labour. For example, Goldin and Katz (1999) show that in spite of a substantial increase in relative supply between 1900 and 1920, the relative wage of white collar workers fell continuously during the first half of the century. Moreover, while mentioning a "strong association between changes in the use of purchase in electricity and shifts in employment toward more educated labour," Goldin and Katz report no sharp widening of the wage distribution prior to the 1970s. Obviously, any explanation of the recent patterns in wage inequality needs to integrate the distinguishing features of the past twenty years from previous episodes if it is to be taken as comprehensive. This does not invalidate the importance of market size and labour supply effects, but it does suggest that any explanation that would rely primarily upon these effects may not be fully satisfactory from a historical point of view.

Second, as in with the "capital-skill complementarity" argument, these increasing-returns arguments have difficulty accounting for the productivity slowdown that started in the 1970s. As Jones (1995) points out, this slowdown appears to be at odds with all R&D-based models of growth that predict that the innovation rate should significantly increase when the supply of skilled labour s increases. Howitt (1999) provides a response to Jones, according to which a combination of product proliferation, capital deepening, and diminishing returns to the production of ideas can be invoked to reconcile R&D-based models with the fact that productivity growth has not increased since the 1970s. But something more is needed to reconcile the above arguments with the fact that productivity growth actually appeared to fall.

Moreover, none of the above technology-based arguments can account for the rise in residual wage inequality. The best they can do is to suggest that residual wage inequality is just another manifestation of the skill-premium; skill is a function not just of education and experience but also of innate ability, which is correlated with education, experience and other variables, but not perfectly. The problem with this argument is twofold. First, why did residual inequality rise in the United States even during the period of the 1970s when the education premium was falling? If the different dimensions of skill are substitutable, as is implicitly assumed by such arguments and as seems likely, then the increase in their overall supply should have had the same short-run effect of reducing the premium to unmeasured skills as it did with measured skills. Second, if increased residual inequality reflects an increased premium on innate ability why was it mainly transitory income whose inequality rose? Innate ability is not a transitory characteristic of a worker. Its relative demand should be raised permanently by skill-biased technological change, which should therefore raise the permanent income of those who are abundantly endowed it.

3 Human capital and the process of technological change

All of the arguments involving skill-biased technological change regard skills as a (multi-dimensional) form of human capital, and human capital as a factor of production in the technology for producing goods. Under given technological conditions, and assuming perfectly competitive factor and output markets, an increase in the relative supply of skills cannot coincide with an increase in their relative price, because factor-demand schedules are always downward-sloping. Thus the observed coincidence of an increased relative supply of college-graduates and an increased supply of their relative wage is *prima facie* evidence that the technology for producing goods has changed; the skills of college-graduates must have become more productive on the margin, relative to the services that high-school graduates supply, for any given quantities employed.

However, Schumpeterian growth theory is built on the idea that human capital can be used not just for producing goods but also for producing ideas. Technological progress often comes in the form of innovations, many of which are the result of R&D, an activity that is intensive in the use of human capital. Learning by doing is another skill-intensive source of technological progress. Moreover, the absorption and diffusion of technological change are obviously activities in which human capital plays an important role. More generally, as Nelson and Phelps (1966) emphasized, human capital can be viewed not just as a factor used in producing goods but a factor used in generating and implementing technological change.

In our *Endogenous Growth Theory* (1998, ch.10) we found it useful to distinguish between two basic frameworks for analyzing the relationship between skills and growth. The first approach, initiated by Lucas (1988) and inspired by Becker's (1964) theory of human capital as a factor in the production of goods, portrays growth as primarily driven by the accumulation of human capital, so that differences in growth rates across countries are mainly attributable to differences in the rates at which those countries accumulate human capital over time. The second approach, which goes back to the seminal contribution of Nelson and Phelps and which has been revived by the Schumpeterian growth literature, describes growth as being driven by the *stock* of human capital, which in turn affects a country's ability to innovate or catch up with more advanced countries. Differences in growth rates across countries are then primarily due to differences in human-capital stocks and thereby in those countries' abilities to generate technical progress.

Benhabib and Spiegel (1994) have tried to decompose the contribution of human capital and education to growth in cross-country regressions. They showed that past educational attainment (as a measure of the current stock of human capital) is essentially uncorrelated with growth if one uses an augmented Solow model à la Mankiw, Romer, and Weil (1992), in which human capital is modeled as an ordinary input in the aggregate production function. But the effect becomes significant if one follows Nelson and Phelps in assuming: (i) that growth is positively affected by the rate of technological innovations and also by the rate of diffusion or adoption of existing innovations; (ii) that the stock of human capital affects both of these rates. The Benhabib-Spiegel analysis is interesting, not only because it provides additional support to the Schumpeterian approach to growth theory but also because it suggests that the divergence in growth rates across countries could be due not so much to differences in the rates of accumulation of human capital, as suggested by Lucas (1988), as to differences in the stocks of human capital, which in turn will affect the various countries' ability to innovate and/or catch up with more advanced countries' technologies.

The robot model we described in section 2.1 above incorporates elements of both frameworks. For it supposes that educated workers can be used as a factor of production both in the technology for producing intermediate products and also in the R&D technology for producing innovations that result in improvements in the quality of intermediate products. But it maintains the assumption of the other technology-based arguments discussed above to the effect that the rise in wage inequality reflects skill-biased technological change in the usual sense of twisting the isoquants of the technology for producing goods.

However, a recent literature has grown up using the Nelson-Phelps framework of Schumpeterian theory, which implies that an increase in the skill premium together with an increase in the relative supply of skills does not necessarily imply skill-biased technological progress of the usual sort. That is, it is not necessary for skilled labour to become more productive on the margin than unskilled labour, for given supplies, in the technology for producing goods. Standard trade-theory arguments can be used to produce models in which all that is required is an increase in the value of skilled labour in the technology for producing ideas, given that the production of ideas uses skilled labour more intensively (relative to unskilled labour) than does the production of goods.[9] So, for example, a (Hicks-neutral) increase in the productivity of R&D will raise the skill premium even if R&D generates skill-*un*biased technological progress.

Dinopoulos and Segerstrom (1999) have used this idea to revisit the trade-versus-technology debate. They argue that trade liberalization may be the driving force underlying the recent rise in wage inequality, despite the counterevidence described in section 2 above. They present a non-Hecksher-Ohlin model of specialized trade, in which specialization results from endogenous country-specific innovations. The model is similar to that of Rivera-Batiz and Romer (1991) except that innovations are quality-improving, as in Aghion and Howitt (1992), and the distribution of skills between and within countries is endogenous. They argue that trade liberalization raises the return to R&D because of the familiar scale effect of endogenous growth theory; that is, a successful innovator can now monopolize a larger international market, and hence can earn a greater flow of profits (until displaced by a subsequent innovator) than before liberalization. Thus, by the argument outlined in the preceding paragraph, the skill premium will rise even though there has been no skill-biased technological progress in the usual sense.

If the supply of skills were exogenous then this argument would founder, like all other trade-based explanations of rising wage inequality, on the counterfactual implication that the relative employment of skilled labour should fall in response to the increased premium. But if the supply responds strongly enough over time to the increased wage premium then the long-run effect can be both an increased wage premium and increased relative employment of skilled labour. Indeed, Dinopoulos and Segerstrom show that their argument does not require any change in the price of skill-intensive products; in their model all products are equally skill-intensive, so no such differential appears. More generally, their point is that in the

[9] Also given no factor-intensity reversals.

Nelson-Phelps framework differences in skill-intensity across produced goods are not needed in order for trade to affect the skill premium, just differences between the production of goods and ideas. All sectors will hire more skilled labour, regardless of their skill-intensity, and despite the increased relative wage of skilled labour, to cope with and profit from a faster pace of technological change. Finally, the fact that the rise in skill-premium appears to be global in scope is consistent with the Dinopoulos-Segerstrom argument, since trade liberalization affects the reward to R&D in all countries.

The Dinopoulos-Segerstrom argument offers yet another explanation for why wage inequality started rising when it did; in this case because the period since the 1960s has been a period of drastic trade liberalization. However, it fails on the same grounds as do those arguments based on skill-biased technological change. For it implies, counterfactually, that productivity growth should have risen along with wage inequality because of the increased R&D that is central to the causal mechanism outlined above, and it offers no explanation for the rise in residual wage inequality. Moreover, like the increasing-returns arguments of Acemoglu and Aghion-Howitt, it implies, counterfactually, that wage inequality should also have risen during the first third of the 20th century, when there was also a dramatic expansion in international trade.

These shortcomings in the Dinopoulos-Segerstrom argument have been addressed in another set of papers that have adopted the Nelson-Phelps framework for modelling human capital. These papers all make use, in one form or another, of the idea that technological progress sometimes comes not in smooth incremental fashion but in waves of fundamental innovations associated with a new General Purpose Technology.

4 General purpose technology and wage inequality

A General Purpose Technologies (GPT)[10] is a fundamental technological breakthrough that dramatically affects the entire economic system. Examples of GPTs include the steam engine, the electric dynamo, the laser and more recently the arrival of the new information technologies embodied in information and communication equipment. The diffusion of a new GPT therefore consists of a wave of secondary innovations, each of which creates a new product or process in a particular sector, that improves upon, but yet is closely related to

[10] The term 'General Purpose Technology' was first introduced by Bresnahan-Trajtenberg (1995). See the articles in Helpman (1998) for theoretical and historical analyses of the role of GPTs in the growth process.

recent adaptations of the same GPT in other sectors.

The demand for skills increases when a new GPT diffuses throughout the economy, because experimentation and adoption of a the new technology requires skilled labour to be hired and employed by the relevant sector of the economy.[11] Greenwood and Yorukoglu (1997) point out that this feature of a GPT can help to account not only for the rise in the skill premium but also for the more or less simultaneously observed slowdown in measured aggregate productivity growth that none of the above technology-based arguments have been able to account for. That is, the increased pace of technological change associated with the GPT raises the relative demand for skills as an agent of change, thereby raising the equilibrium skill-premium, even though it is not skill-biased in usual sense. At the same time, the costs of adopting, diffusing, learning and improving the new technology takes resources out of producing measured output and can thus show up as a decreased rate of productivity growth.[12] They also argue that although there was a slowdown in aggregate measured productivity, the evidence is consistent with an increase in the pace of *embodied* technological progress, which is what one would expect to observe as a consequence of the computer revolution.[13]

A similar argument is presented by Galor and Moav (1999) in a model that can be solved analytically and which they use to account for a number of facts that the other technology-based arguments described above cannot account for. In their model, as in the Dinopoulos-Segerstrom model, the increased relative supply of skilled labour is a consequence of the increased skill premium, which comes from an increased pace of skill-*un*biased technological change. In this case the increased pace of technological change springs not from trade

[11] We are thus implicitly focusing on technological breakthroughs, like the information revolution, which themselves generate a whole set of new ideas (or 'secondary innovations') which require skilled labor in order to be finalized and then implemented. This, in turn, distinguishes the recent technological revolutions from previous GPTs, such as the invention of electricity, which also increased the scope for employing unskilled labor in mass production.

[12] Howitt (1998) argues out that induced obsolescence is another channel through which the wave of innovations associated with the introduction of a new GPT can cause a long-lasting, but ultimately transitory, slowdown in observed growth rates.

[13] For example, the annual growth rate in the US has declined by 1.8% on average since the 1970s. The decline has been most pronounced in the service sector, and more generally the productivity slowdown appears to be mainly attributable to a decline in disembodied productivity growth. Indeed, since the early 1970s the rate of embodied technical progress has accelerated (see McHugh and Lane, 1987; Greenwood and Yorukoglu, 1997; Horstein and Krusell, 1997; and Comin, 1999), and the bulk of this acceleration, e.g., as measured by the decline in the quality-adjusted price of equipment goods, appears to be attributable to computers and other information processing goods. This, again, points to the important role played by the new information technologies and their diffusion during the past twenty years.

liberalization but from a new GPT. The non-monotonic cyclical dynamics of their model is invoked to account for the failure of wage-inequality to rise following the "second industrial revolution" in the early 20th century. Galor and Moav account for residual inequality as an increase in the return to innate ability, under the assumption that economic growth at that time had the effect of reducing the capital-market imperfections that resulted in sub-optimal investment in education. Their explanation for the productivity slowdown is similar to that of Greenwood and Yorukoglu.

One puzzle concerning the GPT explanation of rising wage inequality is how the spread of a new GPT that takes place over a period of several decades[14] could have quantitatively significant macroeconomic consequences. It is easy to imagine how the introduction a fundamentally new technology can have a substantially disruptive effect at the level of the firm or industry where the technology is being introduced. But if the diffusion of the technology takes place slowly, and at a uniform rate, one would not expect to see the cost of these disruptions showing up at the macro level. During any year the adjustments would hardly be noticeable from a macro perspective, being limited to a small subset of the economy. The slowness of diffusion might imply a long time before the rise in aggregate productivity was noticeable, but it is hard to understand why aggregate productivity growth should ever show a significant slowdown.

There are several reasons to believe that endogenous innovation could lead to the diffusion of a new GPT that was gradual, as the evidence suggests, but not uniform over time; exhibiting instead the sort of logistic diffusion pattern observed by epidemiologists when studying the spread of a disease. For example, the existence of strategic complementarities (or network externalities) between various sectors of the economy could generate temporary lock-in effects, of a kind similar to the implementation cycles in Shleifer (1986). It may then take some exogenous factor (e.g. a continuous increase in real labour costs,[15] trade liberalisation, an intensification in product market competition, or a sharp increase in skilled labour supply[16]!) before a critical number of sectors will decide to jump on the bandwagon of the new GPT. Another potential source of non-linearity in the diffusion of a new GPT,

[14] For example, David (1990) talks about a pre-paradigm phase of twenty-five years in the case of the electric dynamo.

[15] See Caballero and Hammour (1997).

[16] The 'baby boom'-driven increase in skilled labor supply in the mid-seventies, might indeed be one of the factors that fostered the diffusion of the new information technologies in countries like the US. In this case, the theory developed in the previous section would maintain some relevance provided it is adequately embedded into a GPT framework of the kind outlined in this section.

which we formalize in Aghion and Howitt (1998, ch.8), lies in the phenomenon of social learning, according to which people in one sector learn from those in other sectors how to use the new GPT. As we shall explain in more detail in the next section, even a diffusion process that takes decades to reach near completion can be characterized by a relatively brief peak period of intense experimentation and learning throughout the economy, a period in which the new GPT begins to snowball. This social learning mechanism will be explained in more detail in the following section.

5 Wage inequality as a premium to adaptability

Our argument to this point has led us to consider explanations of rising wage inequality that combine the Nelson-Phelps idea of skills as a factor of technological change together with the idea of general purpose technology. In this section we sketch a pair of models that show in more detail how these factors can help to account for facts concerning rising wage-inequality both between and within educational groups. The key to both of these models is the idea that how quickly a worker can adapt to working with a new technology is partly a matter of education and partly a matter of luck.

In the long run, we assume that anyone can adapt to the new technology. Our reason for making this assumption is that in the long run what technological progress does is to automate skills. Penmanship is made easier by ball-point pens, copying documents is much easier by modern photo-copiers than by lithograph machines, horsemanship is certainly more difficult to master than is operating a motor vehicle, mental addition more difficult than using an electronic calculators. Even the ability to read is being replaced by computer graphics, as in the case of fast-food clerks. Thus in the long run we assume there is no skill bias to technological change. In the short run, however some catch on much faster than others, and thus earn a premium for their adaptability. The introduction of a new GPT can enhance this premium in several ways, as we explain below.

5.1 Explaining the rise in between-group inequality

Our first model is a simplified version of the one involving social learning that we presented in our (1998, ch.9). In this model, the way a firm or sector typically learns to use a new technology is not to discover everything on its own but to learn from the experience of other firms in a similar situation. For it to be worthwhile for a firm to try to use the procedures of these successful firms as a "template", it must be able to learn from other firms. This will occur only when the problems it needs to solve to implement a GPT bear enough resemblance to the problems already solved by other firms. For a long time, improvements in knowledge will take place slowly, because these are independent discoveries with little guidance from other sectors; eventually, a point will be reached when enough other firms are using the new technology to make it possible to use their experience and experimentation will become much more widespread. This results in an acceleration in the demand for skilled labour and therefore the skill premium will start to rise.

There are two technologies that can be used in any sector; an old technology and a new one that has just arrived in the form of a GPT. Aggregate final output is produced by labour according to:

$$y = \left\{ \int_0^1 A(i)^{\alpha} x(i)^{\alpha} di \right\}^{\frac{1}{\alpha}},$$

where $A(i) = 1$ in sectors where the old GPT is still used, and $A(i) = \gamma > 1$ in sectors that have successfully innovated, while $x(i)$ is the flow of intermediate good i currently used in the production of final output. Manufacturing labour produces intermediate goods using a one-for-one technology, so that $x(i)$ also denotes the labour demand flow in sector i. The total labour force L is divided into skilled and unskilled workers. Whereas old sectors are indifferent between skilled and unskilled workers, the experimentation and implementation of the new GPT can only be done by skilled labour.

For simplicity, we assume that the supply of skilled workers is monotonically increasing over time, partly as a result of schooling and/or training investments, and partly as a result of the technology becoming more familiar:

$$L_s(t) = L - (1 - s) \cdot L \cdot e^{-\beta t},$$

where $s < 1$ is the initial fraction of skilled workers and β is a positive number measuring the speed of skill acquisition. Thus in the long run everyone will be skilled in using the new GPT. All that differs across individuals is the speed with which they learn. Those sL workers who are skilled to begin learn as soon as the new GPT arrives. The rest learn randomly with a Poisson arrival rate β.

We now have to analyze the demand side of the labour market, and in particular determine at any point in time how many sectors are still using the old GPT and therefore do not have any specific need for skilled workers, and how many sectors are experimenting with and already using the new GPT.

We assume that in each sector i, moving from the old to the new GPT requires two steps. First, a firm in that sector must acquire a "template" on which to base experimentation; second, the firm must succeed in making the transition to the new GPT. Let n_0 denote the fraction of sectors that have not yet acquired a template, n_1 denote the fraction of sectors that are currently experimenting on the new GPT, and $n_2 = 1 - n_0 - n_1$ the fraction of sectors that have completed the transition to the new GPT.

Let $\lambda(n_2)$ denote the Poisson arrival rate of *templates* for the new GPT in a given sector and suppose that it is increasing in n_2, to reflect the social learning process by which firms acquire their templates. A special case is when:

$$\lambda(n_2) = \begin{cases} \lambda_0 \text{ if } n_2 \leq \overline{n} \\ \lambda_0 + \Delta \text{ if } n_2 \geq \overline{n} \end{cases}$$

Now, suppose that for a templated firm to actually succeed in implementing the new GPT, it must employ at least H units of skilled labour per period. We can think of this labour as being used in formal R&D, informal R&D, or in an experimental start-up firm. In any case it is not producing current output. Instead, it allows the sector to access a Poisson process that will deliver a workable implementation of the new GPT with an arrival rate of λ_1. Thus the flow of new sectors that can implement the new GPT will be the number of experimenting sectors n_1, times the success rate per sector per unit of time λ_1.

The evolution over time of the two variables n_1 and n_2 is then given by the autonomous system of ordinary differential equations:

$$\begin{aligned} n_1 &= \lambda(n_2) \cdot (1 - n_1 - n_2) - \lambda_1 n_1 \\ n_2 &= \lambda_1 n_1 \end{aligned}$$

with initial condition $n_1(0) = 0$, $n_2(0) = 0$. The time path of n_0 is then given automatically by the identity $n_0 \equiv 1 - n_1 - n_2$.

Figure 1 depicts the kind of dynamic pattern followed by n_1 and n_2 when λ_0 is small and Δ is sufficiently large.[17] Not surprisingly, the time-path of n_2 follows a *logistic* curve, accelerating at first and slowing down as n_2 approaches 1, with the maximal growth rate occurring somewhere in the middle. Likewise, the path of n_1 must peak somewhere in the middle of the transition, in as much as it starts and ends at zero.

[17] The figure represents the case in which $\lambda_0 = 0.01, \lambda_1 = 0.3, \Delta = 0.5$ and $\pi = 0.1$.

The transition process from the old to the new GPT can then be divided into two subperiods. First, in the early phase of transition (i.e. when t is low) the number of sectors using the new GPT is too small to absorb the whole skilled labour force, which in turn implies that a positive fraction of skilled workers will have to be employed by the old sectors at the same wage as their unskilled peers. Thus, during the early phase of transition the labour market will remain unsegmented, with the real wage being the same for skilled and unskilled labour and determined by the labour market clearing equation:

$$(1 - n_2) \cdot x_O + n_2 \cdot x_N + n_1 \cdot H = L.$$

where x_O, x_N, and H denote the labour demands respectively by an old manufacturing sector, a sector using the new GPT, and an experimenting sector.[18]

In the later phase of transition, however, where the fraction of new sectors has grown sufficiently large that it can absorb all of the skilled labour force, the labour market will become segmented, with skilled workers being employed exclusively (and at a higher wage) by new sectors while unskilled workers remain in old sectors. Let w_u and w_s denote the real wages respectively paid to unskilled and skilled workers. We now have $w_s > w_u$, since the two real wages are determined by two separate labour market clearing conditions. The skilled wage is determined by the labour-market clearing equation for skilled labour:

$$L_2 = n_1 \cdot H + n_2 \cdot x_N,$$

while w_u is obtained from the market-clearing equation for unskilled labour, namely:

$$L_1 = (1 - n_2) \cdot x_O,$$

where $L_1 = L - L_2$.[19]

[18] For any sector i, profit maximization by the local monopolist in such a sector, gives:

$$x_i = \arg\max_x \{p_i(x)x - wx\},$$

where:

$$p_i(x) = \frac{\partial y}{\partial x_i} = (A(i))^\alpha x^{\alpha-1} y^{1-\alpha}.$$

The first-order condition for this maximization, respectively for $A(i) = 1$ and $A(i) = \gamma$, yields:

$$x_O = (\frac{w}{\alpha})^{\frac{1}{\alpha-1}} y; x_N = (\frac{w}{\gamma^\alpha \alpha})^{\frac{1}{\alpha-1}} y.$$

[19] Substituting for x_O and x_N in these two labor market clearing equations, we get:

$$w_s = \gamma^\alpha \alpha (\frac{n_2 y}{L_2 - n_1 H})^{1-\alpha}; w_u = \alpha (\frac{(1-n_2)y}{L - L_2})^{1-\alpha}.$$

Figure 2 depicts the time-path of real wages, assuming a relatively large cost H of experimentation.[20] The skill premium, here measured by the ratio (w_s/w_u), starts increasing around the same time as the diffusion of the new GPT across sectors accelerates, as a result of the increased demand for skilled labour in production and experimentation. The premium keeps on increasing although more slowly during the remaining part of the transition process. Since everyone ends up earning the same (skilled) wage, standard measures of inequality first rise and then fall, as in the previous section.[21]

This explanation of increased inequality between skill groups is also consistent with the observed dynamic pattern of the wage premium in the US or the UK since 1970, namely a reduction of the wage premium during the early 1970s, followed by a sharp increase in that premium between the late 1970s and the mid-1990s. In particular, a one-time increase in skilled labour supply, occurring during the acceleration phase in the diffusion of new information technologies, would also result in a short-run reduction followed by a medium-term increase in the skill-premium, of the kind experienced in the US respectively during the 1970s and the 1980s.

With regard to the comparison between the recent period and the early 1900s, an increase in the supply of educated labour alongside with adoption of the electric dynamo did not result in a comparable increase in the skill premium as the diffusion process of that earlier GPT was not nearly as skill-biased as that for the new information technologies, even in the transitional sense used in the Nelson-Phelps framework. For example, whilst workers operating steam engines would need to know how to maintain and repair their own engine, the maintenance of new electrical machinery would only require firms to hire a limited number of skilled workers

[20] In addition to the parameter values specified in footnote 17 above, Figure 2 is plotted with $\gamma = 1.5$, $\alpha = .5$, $L = 10$, $H = 6$, $\beta = 0.05$ and $s = 0.25$.

[21] This simple model of GPT diffusion and between-group wage inequality can be extended easily to accomodate the existence of productivity spillovers among sectors that currently adopt the new GPT. For example, in line with other multisector models of endogenous growth (e.g., see Aghion-Howitt, 1998, ch.3) we could assume that the productivity γ of a sector that has just adopted the new GPT depends positively upon the current flow of adoptions, e.g., according to:

$$\gamma = \gamma_0 + \lambda(n_2)\sigma$$

where σ is a positive number that reflects the extent of the cross-sector spillovers. In such an extension of the above GPT model, the speed of technological change as measured by the derivative $d\gamma/dt$ will increase during the acceleration phase in the GPT diffusion; this, in turn, will only magnify the increase in skill-premium during that phase. That the speed of technological change should increase when a new GPT hits the economy is a plausible assumption which we shall also make in the next section when discussing the effects of GPT diffusion on within-group inequality.

specialized in that task. Thus the appearance of bottlenecks that segments the labour market in our story may have been less of a factor in accounting for wage movements in the earlier example.

The clustering of experimentation that results from social learning makes this story is easily reconcile also with the experience of the productivity slowdown. Figure 3 shows the time-path of GDP, relative to its initial value, assuming the same parameter values as in Figures 1 and 2. GDP falls at the beginning of the acceleration phase in the diffusion of the new GPT, as a result of the diversion of labour away from production of goods (which is measured in GDP) and into the production of technological change (which is not measured in GDP), before growing again at a later stage in this diffusion process. In addition, the explanation appears to be also consistent with the observed deceleration in the increase in between-group wage inequality during the past three years: one may indeed interpret this embryonic trend reversal as reflecting the fact that the diffusion of new communication and information technologies is now entering a mature phase. Perhaps we are now beyond the experimentation phase and ready to start reaping the benefits of a technology that is getting easier to use and that we are all getting used to.

5.2 Explaining the rise in within-group inequality

Existing explanations for the rise in residual wage inequality[22] are based on the idea that whatever raises the demand for observed skills also raises the demand for unobserved skills, which are generally treated as one form or another of innate ability. As mentioned above, this explanation is at odds with recent econometric work, e.g by Blundell and Preston (1999), which shows that the within-group component of wage inequality in the US and UK is mainly *transitory* whereas the between-group component accounts for most of the observed increase in the variance of *permanent* income. These explanations also fail to explain why the rise in within-group inequality has been accompanied by a corresponding rise in individual wage instability (see Gottschalk and Moffitt, 1993). In the remaining part of this section, we shall argue, using the Nelson-Phelps framework, that the diffusion of a new technological paradigm can account for the evolution of within-group wage inequality in a way that is consistent with these and other puzzling facts.

[22] See Acemoglu (1998), Heckman et al. (1998), Galor-Moav (1999) and Rubinstein-Tsiddon (1999) for models of within-group inequality based on differences in innate ability.

The basis of this argument is the model of Aghion, Howitt and Violante (2000)[23], which is built on prior work of Violante (1996). The model, like that of the preceding section, involves differential rates of adaptability, which in this case are entirely random. Thus consider a group of *ex ante* identical workers with the same educational background; each period technological change takes place as the new fundamental technology diffuses to a new sector and is embodied in new machines; wage inequality arises in this framework because only a random fraction of workers get the opportunity to adapt at once to the most recent vintage of machines. Moreover, those workers who get the opportunity to adapt faster to the newest vintage several periods in a row, obtain an additional premium as they can more easily transfer to the new leading-edge machines skills that they acquired through learning-by-doing on their previous job: indeed, the technological distance between their previous job and their current job, is smaller than for other workers. The diffusion of a new technological paradigm then raises (within-group) wage inequality for at least two reasons: first, the rise in the speed of embodied technical progress associated with the diffusion of the new GPT increases the market premium to those workers who adapt quickly to the leading-edge technology; second, to the extent that a new GPT generates a wave of secondary innovations that are closely related to one another, its diffusion raises the ability of workers that are adaptable several times in a row to transfer recently acquired knowledge to the newest vintage.

5.2.1 Basic Framework

We construct an infinite-horizon discrete-time model with sequential productivity-improving innovations in a one-good economy; each period an innovation occurs. This allows a new vintage of machine to be produced and used for final good production: indeed the new technology must be embodied in a machine for it to be used. We focus here on the simplest possible case, where, in accordance with Krusell et al. (2000), machines last for only two periods but do not depreciate in the first period. Once the machine is in place, the new (or leading-edge) technology at date t allows the production of final output according to

$$y_t = A_t \cdot x_{0t}^{1-\alpha}. \tag{1}$$

where x_{0t} is the labour input working with technology t.

[23]The model in this subsection is a stripped-down version of the framework developed in Aghion, Howitt and Violante (2000). In that paper, we provide: (i) a more explicit formalization of the notion of skill transferability in relation to the nature (general versus sector specific) of the process of technological progress; (ii) an explicit analysis of the evolution of between- and within- group wage inequality during the transition from a "sector-specific" to a "general" technological change regime.

Since capital goods last for just two periods, there is only one other technology to consider, which produces output according to

$$z_t = A_{t-1} \cdot ((1+\eta)x_{1t})^{1-\alpha}. \tag{2}$$

where η is a constant exogenous rate of learning by doing and x_{1t} is the labour input used on the old technology at date t.

Each new technology is $(1+\gamma)$ times more productive than the previous one, so

$$A_t = (1+\gamma)A_{t-1} \qquad \forall t.$$

Let 0 (respectively 1) denote the leading-edge (respectively the old) technology at any time in steady-state, and x_0 and x_1 the amount of efficiency units of labour employed respectively by new and old firms in steady-state, and let $y_0 = y_t/A_t$ and $y_1 = z_t/A_t$ denote the productivity-adjusted outputs from new and old sectors respectively. Then, we can re-express the above equations (1) and (2) as

$$y_0 = x_0^{1-\alpha}; y_1 = \frac{1}{1+\gamma}((1+\eta)x_1)^{1-\alpha}. \tag{3}$$

Adaptability constraint: Not all workers can move at once to the leading-edge and this is the primary source of within-cohort inequality; more specifically, a randomly selected fraction σ of all workers at most can relocate at once to the leading-edge.

Transferability parameter: An individual who worked in the leading-edge last period and is moving to the new leading-edge this period, can transfer some of the knowledge she acquired on the previous leading-edge. That is, every unit of labour services by this individual on the new leading-edge technology generates $(1+\tau)$ units of labour input on that technology, where $\tau < \eta$ (the rate of learning-by-doing enjoyed by those workers that remain on the same technology for two periods in a row); on the other hand, an individual working on the old technology last period and who moves to the new leading-edge this period, carries less relevant experience with her, and for simplicity we assume that such an individual can only generate one unit of labour input per unit of labour services. We then have

$$\begin{aligned} x_0 &= (1+\tau)n_{00} + n_{10}, \\ x_1 &= n_{01} + n_{11}; \end{aligned}$$

where n_{ij} is the steady-state labour flow from the i-technology $(i = 0, 1)$ last period to the j-technology $(j = 0, 1)$ this period. In steady-state equilibrium these labour flows satisfy:

first, the labour market clearing condition

$$n_{00} + n_{10} + n_{01} + n_{11} = 1;$$

second, the adaptability constraints

$$n_{00} \leq \sigma(n_{00} + n_{10}),$$

$$n_{10} \leq \sigma(n_{01} + n_{11});$$

and third, the stationarity condition

$$n_{10} = n_{01}.$$

Preferences: We assume a logarithmic specification for households' instantaneous preferences, that is:

$$u(c_t) = \ln c_t,$$

and we denote the discount factor by $\beta < 1$. Thus, in steady-state equilibrium, the real rate of interest satisfies the following Euler equation:

$$\frac{1}{1+r} = \beta \frac{c_t}{c_{t+1}} = \beta \frac{1}{1+\gamma}.$$

Labour demand schedule: The plants operating old and new machines will set their demands for labour by equalizing the marginal product of labour to the wage level which they take as given. Hence, if $\omega_i = w_{it}/A_t$ denotes the steady-state productivity-adjusted wage in sector i, the ratio ω_0/ω_1 is equal to the ratio of marginal products in sectors 0 and 1, which in turn leads to the relative labour demand schedule:

$$\frac{\omega_0}{\omega_1} = \frac{1+\gamma}{(1+\eta)^{1-\alpha}} \left(\frac{x_0}{x_1}\right)^{-\alpha} \tag{4}$$

and the productivity-adjusted wage rate of an individual working on the leading-edge for two periods in a row is simply equal to:

$$\omega_{00} = (1+\tau)\omega_0$$

whereas $\omega_{10} = \omega_0$ and $\omega_{01} = \omega_{11} = \omega_1$, where ω_{ij} denotes the normalized wage of a worker who moved from a sector of age i last period to a sector of age j this period.

5.2.2 Labour supply decisions

Every period workers supply one unit of their time inelastically to the market, so their only choice is whether to stay with their existing job (and machine) or to move to another job. In order to make this decision, a worker who last period was on a machine of age i, this period compares the discounted future income from working on the new technology of age 0 with the value of remaining on the same technology, now of age 1. We denote the present discounted value of these two options by v_{i0} and v_{i1} respectively. We then have the following Bellman equations:

$$\begin{aligned} v_{i0} &= \omega_{i0} + \beta \left\{ \sigma \max \left[v_{00}, v_1 \right] + (1-\sigma) v_1 \right\} \\ v_1 &= \omega_1 + \beta \left\{ \sigma \max \left[v_{10}, v_1 \right] + (1-\sigma) v_1 \right\} \end{aligned}$$

where we make use of the Euler equation: $1+\gamma = \beta(1+r)$. There are three possible solutions to the above discrete choice problem:

(a) $v_1 = v_{10}$, which in turn is equivalent to: $\omega_0/\omega_1 = \Omega = 1/(1+\beta\sigma\tau)$. In this case, the individuals who worked in sector 1 last period are indifferent between the two sectors this period. But then, because they get an additional transferability premium if they move to the new leading-edge, those individuals who worked in sector 0 last period will strictly prefer to move to the new leading-edge this period. Therefore, the first adaptability constraint is binding, but n_{10} can take any value between 0 and $\sigma(n_{01} + n_{11})$. The relative labour supply x_0/x_1 is thus indeterminate.

(b) $v_{10} > v_1$, or equivalently: $\omega_0/\omega_1 > \Omega$. In this case the adaptability constraint is binding for workers of type $i = 0, 1$ as any of these workers with the chance of moving to the new technology will take that opportunity, hence $n_{i0} = \sigma(n_{0i} + n_{1i}), i = 0, 1$.This, together with the market clearing and stationarity conditions satisfied by the labour flows n_{ij},and substituting for the n_{ij}'s in the above equations defining x_0 and x_1, leads to the relative labour supply: $x_0/x_1 = \sigma(1+\sigma\tau)/(1-\sigma) = \chi$.

(c) $v_{i1} > v_{i0}$. In this case workers of type i prefer to be operating old capital, and therefore $n_{i0} = 0$, so that $x_0 = 0$.

5.2.3 Equilibrium Within-Group Wage Inequality

We can now easily determine the wage distribution in stationary equilibrium by the intersection between the labour demand schedule in (1) and the labour supply schedule defined by (a)–(c) (see Figure 4). Let Φ denote the relative wage ω_0/ω_1 corresponding to the intersection between the relative demand schedule and the vertical line $x_0/x_1 = \chi$, that is: $\Phi = \frac{1+\gamma}{(1+\eta)^{1-\alpha}}\left[\frac{1-\sigma}{\sigma(1+\sigma\tau)}\right]^{\alpha}$. Then, using the fact that the maximum wage is earned by the individuals working on the leading-edge two periods in a row, that $\omega_0/\omega_1 = \max(\Omega, \Phi)$, and that $\Omega < 1$, we can express the ratio R_ω between the maximum and the minimum wage as:

$$\begin{aligned} R_\omega &= \max\left[\frac{\omega_{00}}{\omega_0}, \frac{\omega_{00}}{\omega_1} = \frac{\omega_{00}}{\omega_0}\frac{\omega_0}{\omega_1}\right] \\ &= \max\left[1+\tau, (1+\tau)\Phi\right] \end{aligned}$$

In particular this measure of *within-group wage inequality* decreases with adaptability σ, but *increases both with the rate of embodied technical progress* as parametrized by γ and *with the transferability of knowledge by adaptable workers* as parametrized by τ. Both variables, in turn, are likely to have increased during the acceleration phase in the diffusion of new Communication and Information technologies: (i) for example, McHugh and Lane (1987), Gordon (1990), Greenwood and Yorukoglu (1997) and Krusell et al. (2000) show there has been an acceleration in the rate of embodied technical change since the mid-1970s, e.g as measured by the decline in the quality-adjusted price of equipment goods; (ii) the "general nature" of the technological wave in Communication and Information implies that the acceleration phase of that wave should be accompanied by an increased *similarity* between successive vintages of capital, which in turn is likely to increase the degree of *skill transferability* across technologies.

Thus the approach based on the notion of adapting to major technological change can shed light, not only on the observed evolution of the college premium, but also on the increase in residual wage inequality. Furthermore, it does so in a way that can be made consistent with at least three puzzling facts: first, to the extent that residual wage inequality has to do with the stochastic nature of workers' adaptability to the newest vintage more than with innate ability, the rise in within-group inequality induced by the diffusion of a new GPT, should primarily affect the transitory component of income, in line with the empirical work of Blundell and Preston (1999); second, the increase in residual wage inequality should be mirrored by a rise in individual life-time wage instability as documented by Gottschalk and Moffitt (1993), as individual luck in adapting faster to a new sector, will obviously vary over time; third, if the economy comprises several educational groups of workers with more educated workers being more able to adapt or transfer recently acquired knowledge to the newest vintages, then one can more easily conceive of the possibility that a fall in the education premium be accompanied by a rise in residual inequality, as appears to have been the case in the US during the mid and late 70s and also possibly in the late 1990s. For example, an increase in the relative supply of educated labour that would occur whilst the Information revolution hits the economy, would have exactly that effect: namely, to temporarily reduce the education premium, meanwhile the continuing increase in the speed of embodied technical progress and in the transferability of recently acquired knowledge induced by the new GPT would generate a sustained rise in within-group inequality. The alternative theories of within-group inequality based upon market-size effects and unobserved ability do not seem to provide an equally convincing explanation for the diverging patterns of the between-group and the within-group components of wage inequality.

6 Policy Implications

We have suggested in this paper that the rise in wage inequality that has been observed in many countries over the past thirty years may be attributable to the introduction of computer technology, which raised the pace of technological change and thus increased the premium placed on adaptability in labour markets, but which in the long run is not skill-biased in the usual sense of the term. We conclude by discussing briefly some of the policy conclusions that might follow if this suggestion were confirmed by subsequent empirical investigations.

According to the standard view that regards skilled labour as nothing but an input, along with unskilled labour and capital, to the technology for producing goods, to reduce the wage-inequality created by skill-biased technological change would require either an increase in the relative supply of skilled labour or an increase in the relative supply of capital, depending on the exact form of the aggregate production function. The current view cautions that either of these remedies might simply accelerate the process of rising inequality. More skilled labour is likely to speed up the (skill-intensive) process of technological change underlying the rise in inequality. Likewise, according to our argument in Howitt and Aghion (1998), faster capital accumulation will also tend to accelerate the rate of technological change through a scale effect. Moreover, our analysis of within-group inequality suggests that it ought to be higher within more educated groups, to the extent that education enhances the ability to transfer the lessons of experience from one leading-edge technology to another. Thus an increase in education, even if it reduces the education premium, may nevertheless increase the overall degree of wage inequality.

As to the question of what sort of increase in education is most likely to favour a reduction in inequality, the analysis suggests that general education, which teaches fundamental analytical and problem-solving skills, and which fosters creativity and an open attitude to novel intellectual challenges, may reduce inequality more than education in narrow technology-specific skills. One of the reasons why adaptability to new technologies is random is that the technologies themselves evolve in unpredictable ways. Investing in skills that are complementary with current technologies is increasingly risky as the pace of uncertain change accelerates. Investing in fundamental skills and in appreciation for new ideas fosters the adaptability that has been increasingly rewarded by labour markets in recent years.

In any event, the analysis suggests that increasing inequality is not an inevitable consequence of technological change, and that it is likely to be transitory. Our view of the causes of between-group wage inequality implies an education premium which should eventually go back to more normal levels, as people become increasingly used to computer technology and as the technology becomes increasingly easy to use. Similarly our view of the causes of within-group inequality is one that relies heavily on the assumption that we are in a transitional period with innovations coming more rapidly than usual and being more than usually associated with a common underlying general purpose technology. Eventually the transition will end and the degree of residual inequality should revert to normal. This suggests that transitory programs of direct assistance to those with low earnings may be a more efficient way of dealing with high levels of inequality than permanent structural remedies.

References

Acemoglu, D. (1998): "Why Do Technologies Complement Skills? Direct Technical Change and Wage Inequality", *Quarterly Journal of Economics, 113, 1055-1090.*

Acemoglu, D. (1999): "Patterns of Skill-Premia", mimeo MIT.

Acemoglu, D (2000): "Technical Change, Inequality, and the Labour Market", mimeo MIT.

Acemoglu, D., P. Aghion and G. Violante (2000): "Deunionization, Technical Change, and Inequality", mimeo UCL.

Aghion, P. and P. Howitt (1992): "A Model of Growth Through Creative Destruction", *Econometrica*, 60, pp.323-351.

Aghion, P. and P. Howitt (1998): *Endogenous Growth Theory.* Cambridge: MIT Press.

Aghion, P., P. Howitt and G. Violante (2000): "Technology, Knowledge, and Within-Group Inequality", mimeo UCL.

Autor, D., L. Katz and A. Krueger (1998): "Computing Inequality: Have Computers Changes the Labour Market?", *Quarterly Journal of Economics, 113, pp. 1169-1214.*

Becker, Gary S. (1964): *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education.* New York: NBER.

Benhabib, Jess, and Mark M. Spiegel (1994): "The Role of Human Capital Economic Development: Evidence from Cross-Country Data." *Journal of Monetary Economics,* 34, 143-73.

Berman, E., J. Bound and Z. Griliches (1994): "Changes in Demand for Skilled Labour Within US Manufacturing: Evidence from the Annual Survey of Manufactures", *Quarterly Journal of Economics*, 109:2, pp.367-97.

Blundell, R. and I. Preston (1999): "Inequality and Uncertainty: Short-Run Uncertainty and Permanent Inequality in the US and Britain", mimeo UCL.

Bresnahan, T.F, and M. Trajtenberg (1995): "General Purpose Technologies: Engines of Growth?", *Journal of Econometrics,* 65, pp.83-108.

Caballero, R. and M. Hammour (1997): "The Macroeconomics of Specificity", mimeo MIT.

Card, D (1996): "The Effects of Unions on the Structure of Wages: A Longitudinal Analysis", *Econometrica,* 64, pp. 957-979.

Comin, D (1999): "An Uncertainty-Driven Theory of the Productivity Slowdown: Manufacturing", mimeo Harvard University.

David, P.A. (1990): "The Dynamo and the Computer: A Historical Perspective on the Productivity Paradox", *American Economic Review*, 80:2, pp.355-61.

Dinopoulos, Elias, and Paul Segerstrom (1999): "A Schumpeterian Model of Protection and Real Wages." *American Economic Review,* 89, 450-72.

Fortin, N.M. and T. Lemieux (1997): "Institutional Changes and Rising Wage Inequality: Is there a Linkage?, *Journal of Economic Perspectives*, 11:2, pp.75-96.

Freeman, R.B. (1993): "How Much Has Deunionization Contributed to the Rise in Male Earnings Inequality?", in *Uneven Tides: Rising Inequality in America.* S. Danziger and P. Gottschalk, eds. New York: Russell Sage Foundation, pp.133-63.

Galor, O. and O. Moav (1999): "Ability Biased Technological Transition, Wage Inequality Within and Across Groups, and Economic Growth", forthcoming in the *Quarterly Journal of Economics.*

Garcia-Peñalosa, C. and B. Koebel (1998): "International Trade, Material Prices and Wages", mimeo Nuffield College, Oxford.

Goldin, C and L. Katz (1999): "The Returns to Skill across the Twentieth Century United States", Mimeo Harvard.

Gordon, R.J. (1990): *The Measurement of Durable Good Prices*, NBER Monograph Series, University of Chicago Press.

Gottschalk, P. and R. Moffitt (1994): "The Growth of Earnings Instability in the U.S. Labour Market", *Brookings Papers of Economic Activity*, 2, pp.217-272.

Greenwood, J. and M. Yorukoglu (1997): "1974", *Carnegie-Rochester Series on Public Policy*, 46, pp.49-95.

Heckman, J., L. Lochner and C. Taber (1998): "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labour Earnings with Heterogeneous Agents", *Review of Economic Dynamics*, 1, pp.1-58.G

Helpman, E., editor (1998): *General Purpose Technologies and Economic Growth* , The MIT Press.

Hornstein, Andreas, and Per Krusell (1996): "Can Technology Improvements Cause Productivity Slowdowns?" In *NBER Macroeconomics Annual, 1996*, edited by Ben Bernanke and Julio Rotemberg, 209-59. Cambridge, MA: MIT Press.

Howitt, P (1998): "Measurement, Obsolescence, and General Purpose Technologies", in *General Purpose Technologies and Economic Growth,* E. Helpman ed., The MIT Press.

Howitt, Peter (1999): "Steady Endogenous Growth with Population and R&D Inputs Growing." *Journal of Political Economy,* 107, 715-30.

Howitt, Peter, and Philippe Aghion (1998): "Capital Accumulation and Innovation as Complementary Factors in Long-Run Growth." *Journal of Economic Growth,* 3, 111-30.

Jones, C. (1995): "R&D Based Models of Economic Growth", *Growth of Political Economy*, 103, pp.759-784.

Katz, L. and D. Autor (2000): "Changes in the Wage Structure and Earnings Inequality" in O. Ashenfelter and D. Card, eds, *The Handbook of Labour Economics,* Volume 3, Amsterdam: Elsevier.

Krusell, Per, Lee Ohanian, José-Victor Ríos-Rull, and Gianluca Violante (1994): "Capital-Skill Complementarity and Inequality." unpublished.

Lucas, Robert E. Jr. (1988): "On the Mechanics of Economic Development." *Journal of Monetary Economics,* 22, 3-42.

Machin, S. (1996a): "Wage Inequality in the UK", *Oxford Review of Economic Policy*, 12:1, pp.47-64.

Machin, S. (1996b): "Changes in the Relative Demand for Skills", in *Acquiring Skills, Market Failures, their Symptoms and Policy Responses*, A.L. Booth and D.J. Snower eds., Cambridge: Cambridge University Press.

Machin, S. (1997): "The Decline of Labour Market Institutions and the Rise in Wage Inequality in Britain", *European Economic Review*, 41:3-5, pp.647-57.

Machin, S. and J. Van Reenen (1998): "Technology and Changes in Skill Structure: Evidence from Seven OECD Countries", Berkeley Center for Labour Economics, Working Paper..

McHugh, Richard, and Julia Lane (1987): "The Role of Embodied Technical Change in the Decline of Labor Productivity." *Southern Economic Journal,* 53, 915-24.

Murphy, K. and F. Welch (1992): "The Structure of Wages", *Quarterly Journal Of Economics,* 107, pp. 255-285.

Nelson, R. and E. Phelps (1966); "Investment in Humans, Technological Diffusion, and Economic Growth", *American Economic Review*, 61, pp.69-75.

Rivera-Batiz, Luis A., and Paul M. Romer (1991): "Economic Integration and Endogenous Growth." *Quarterly Journal of Economics,* 106, 531-56.

Rubinstein, Y. and D. Tsiddon (1999): "Coping with Technological Progress: The Role of Ability in Making Inequality so Persistent", mimeo Hebrew University.

Shleifer, Andrei (1986): "Implementation Cycles." *Journal of Political Economy,*94, 1163-90.

Stokey, N (1996): "Free Trade, Factor Returns, and Factor Accumulation", *Journal of Economic Growth,* 1, pp. 421-447.

Violante, G.L. (1996): "Equipment Investment and Skill Dynamics: A Solution to the Wage Dispersion Puzzle?", mimeo University College London.

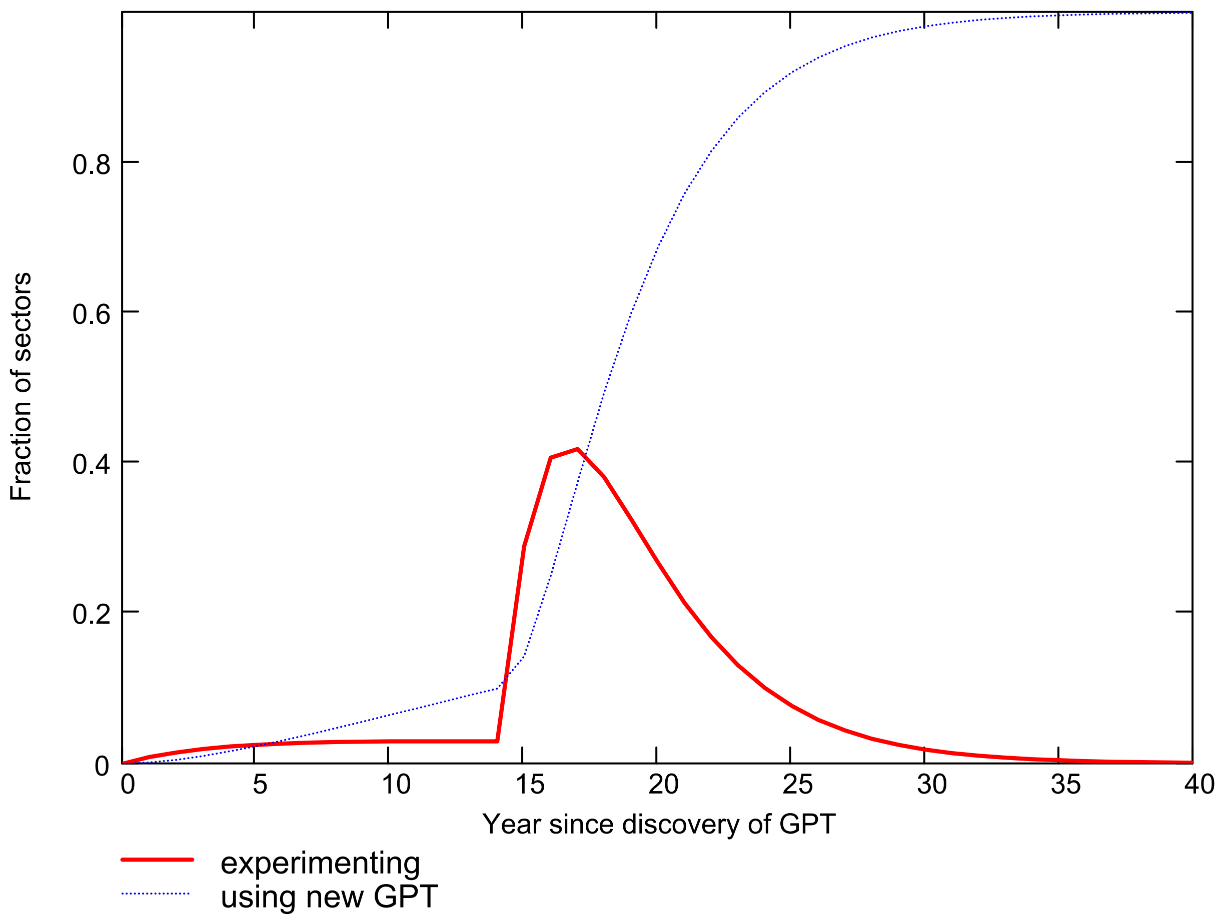
0.8
0.6
0.4
0.2
0
Fraction of sectors
0
5
10
15
20
25
30
35
40
Year since discovery of GPT
experimenting
using new GPT

Figure 1

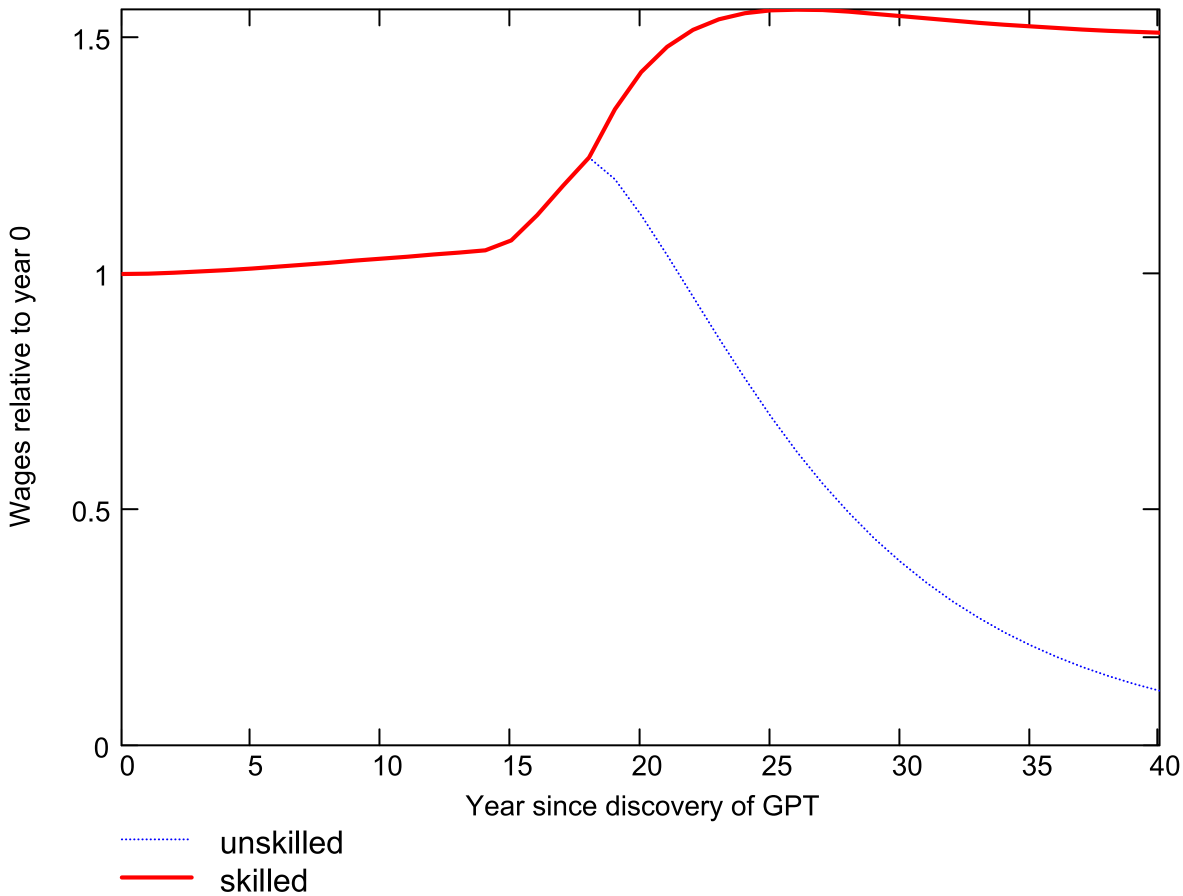

1.5
1
0.5
0
0
5
10
15
20
25
30
35
40
Wages relative to year 0
Year since discovery of GPT
unskilled
skilled

Figure 2

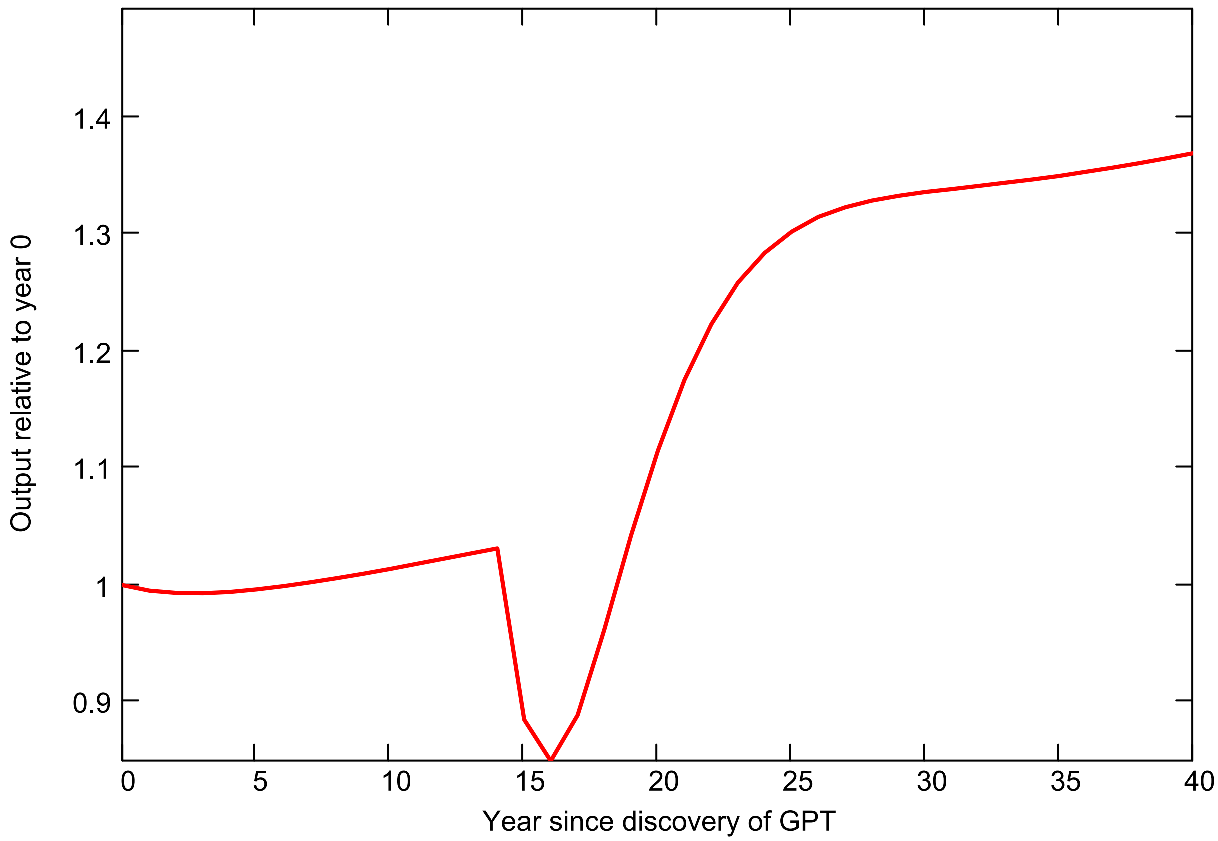

1.4
1.3
1.2
1.1
1
0.9
0
5
10
15
20
25
30
35
40
Output relative to year 0
Year since discovery of GPT

Figure 3