Wage Inequality and Technological Change: A Nelson-Phelps Approach.

Philippe Aghion, Peter Howitt and Gianluca Violante

December 28, 2001

1 Introduction

The sharp increase in wage inequality that has taken place since the early 1980s in developed countries, especially in the US and the UK, has sprung intense debates among economists. The rapidly growing literature on the subject reflects substantial progress in narrowing down the search for robust explanations, in particular by emphasizing the primary role of (skill-biased) technological progress; yet this literature leaves some important puzzles still open to further inquiry.

The first puzzle concerns the evolution of wage inequality *between* educational groups; although the relative supply of college-educated workers increased noticeably within the past 30 years, the wage ratio between college graduates and high-school graduates rose substantially in countries like the US and the UK between the early 1980s and the mid-1990s. In the US, for example, Autor et al. (1998) show that the ratio of "college-equivalents" (defined as the number of workers with a college degree plus half the number of workers with some college education) to "non-college equivalents" workers (defined as the complementary set of workers) 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 parallel to these movements in relative supply, the ratio between the average weekly wages of college- and high-school graduates went up by more than 25 percent during the period 1970-1995, although it had fallen by 0.11% a year on average during the previous period.

The second puzzle is that wage inequality has also increased sharply *within* educational and age groups: in particular Machin (1996a) finds that the *residual* standard deviation in hourly earnings 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 third 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).

In this paper we develop an explanation for these puzzles which is based on the idea that technological change is skill-biased, not only in the usual sense of enhancing educated workers' productivity in producing goods and services under given technological conditions but also in the sense of raising the reward to adaptability. The argument follows Galor and Tsiddon (1997) and Galor and Moav (2000) in building on the idea of Nelson and Phelps (1966) that skills are not just an input to the production of goods and services but also a factor in the creation and absorption of new technological knowledge.[1]

The paper is organized as follows. In Section 2 we review and discuss the main existing explanations for the increase in wage inequality. In Section 3 we develop our own explanation which combines adaptability considerations and the idea that an important driving force behind the evolution of the wage structure is the arrival of new "General Purpose Technologies". Finally, Section 4 concludes by mentioning a few potential extensions and policy implications of our approach.

2 Alternative explanations of rising wage inequality

2.1 Trade

One possible explanation, inspired by Heckscher-Ohlin theory, is that the observed upsurge in wage inequality was the direct consequence of *trade liberalization*; in a nutshell, a globalization boom should drive up the demand for skilled labor in developed countries, where skilled labor is cheap relative to developing countries, and it should drive down the relative demand for unskilled labor which is relatively expensive in developed countries. Unfortunately, trade liberalization failed to be supported by the evidence. First, as argued by Krugman and others, how could trade liberalization 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,

[1] While our explanation of the facts is similar to that of Galor and Moav, it differs in placing emphasis on the *generality* of technological change as well as its pace, in generating a steady-state increase in within-group wage inequality that is not just transitional, and in basing the explanation of within-group inequality on luck rather than on ability.

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 labor 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 et al. (1994) for the US and Machin (1996b) 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. Finally, the Heckscher-Ohlin theory would predict that the ratio of skilled to unskilled employment should have gone down in skill-intensive industries in developed economies, which again did not happen.

2.2 Institutions

A second explanation would link wage inequality to institutional change. The primary candidate is the *deunionization* trend experienced by the US and the UK since the Reagan-Thatcher period. 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. The argument is simply that unionization is often positively correlated with wage compression,[2] so that deunionization should naturally lead to an increase in wage inequality.[3] DiNardo, Fortin and Lemieux (1996) and Card (1996) examine the effects of deunionization on the wage distribution and find that it can explain a significant part of the rise in *male* inequality. However, the attempt to attribute the increase in wage inequality to deunionization faces at least three important shortcomings: First, "timing" considerations: in the UK the rise in wage inequality started in the mid-seventies but 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. Second, the empirical studies above find no impact of deunionization on the rise of female inequality. Third, most of the impact of deunionization on inequality takes place through a declining middle region of the wage distribution, so deunionization fails to explain the fall

[2] 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.

[3] For example, Card (1996).

in the lower deciles and the sharp surge in the upper deciles. [4]

Changes in direct government intervention in setting the minimum wage is another potential candidate: the nominal minimum wage was fixed at \$3.35 an hour for all the 80's, so in real terms it has declined over the whole period. DiNardo et al. (1996) and Lee (1999) examine empirically the role of the declining minimum wage and conclude that while one can detect large effects in the 50-10 wage differential for males, both rising inequality for women and the increase in the male college-high school premium cannot be attributed to changes in minimum wage regulations.

2.3 Technology

A number of empirical studies have pointed to a significant impact of *skill-biased technical change (SBTC)* on the evolution of wage inequality. For example, using R&D expenditures and computer purchases as measures of technical progress, Berman et al. (1994) found that these two factors could account for as much as 70% of the move away from production to non-production labor over the period 1979–1987. Murphy and Welch (1992) find that the share of college labor 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 et al. (1998) and assuming an elasticity of substitution of 1.4 between skilled and unskilled labor, 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). Krusell et al. (2000) argued that skill-biased technical change can be interpreted as an increased growth in the stock of capital equipment since capital and skills are complementary in the aggregate production function.

This is only a starting point, however, as we still need to understand why we observed an *accelerated* increase in wage inequality following a sharp increase in the supply of skilled workers in the early 1970s. The existing literature provides two main answers to this puzzle.

On the one hand, Katz and Murphy (1992) argue that the observed accelerated increase in the college premium during the 1980s was the combined result of: (i) secular skill-biased

[4] 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 et al (2000) on deunionization and skill-biased technical change and Acemoglu (1999) on trade liberalization and skill-biased technical change.

technical change at a constant pace over the past fifty years; (ii) the temporary fall in the college premium caused by the baby-boom driven increase in the relative supply of skilled labor in the early 1970s; before moving back to its secular path, the college premium was bound to increase at an accelerated rate.

An alternative view is that there has been a true *acceleration* in the pace of skill-biased technical change since the 1970s. This view was first put forward by Krusell et al. (2000), who pointed to the increased rate of decline in the relative price of production equipment goods since the mid-1970s,[5] which they interpreted as an increase in the pace of capital-embodied technological progress. However, the "acceleration view" begs two questions: What caused the increased rate of capital-embodied technological progress? And, how do we reconcile a technological acceleration with the fact that measured total factor productivity growth was much slower in the two decades following 1975 than during the previous two decades?

Acemoglu (1998, 2000) provided the first compelling explanation for the rise in the pace of skill-biased technological change: his main idea is that the increased relative supply of college-educated workers in the 1970s was responsible for a shift in the *direction* of technological 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 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 labor. Whereas Acemoglu takes

[5] See McHugh and Lane (1987) and Gordon (1990).

as given the overall level of innovation in the economy, Aghion and Howitt take 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.

This market size explanation is quite appealing, especially since it appears 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 labor supply in the late 1960s. On the other hand, this explanation raises two issues which we should like to discuss briefly.

Issue 1: *Historical Perspective.* The above story can account for the dynamic pattern followed by the skill premium in the US after the 'baby boom' increase in skilled labor supply in the early seventies, but it does not explain why the rise in wage inequality occurred around this time in contrast with other historical episodes in which similar increases in the supply of educated labor have not been followed by any noticeable increase in wage inequality. For example, in a recent paper on "The Returns to Skill across the Twentieth Century in the United States", Goldin and Katz (1999) show that in spite of a substantial increase in the relative supply of educated labor between 1900 and 1920 following the so-called "high-school movement", the wage ratio between white collars and blue collars fell continuously during the first half of the century and especially during the 1920s and the 1940s. Moreover, even though they mention a "strong association between changes in the use of purchase in electricity and shifts in employment toward more educated labor," 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 really 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* 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.

Issue 2: *Productivity Slowdown.* In a highly influential paper, Jones (1995) points out that OECD countries have experienced substantial increases in the average duration of schooling and in R&D levels during the past fifty years, yet there has been no apparent payoff in terms of faster growth; if anything, measured *productivity growth has slowed*, especially between the mid 1970s and the early 1980s.[6] These findings appear to be at odds with R&D-based models of growth that predict that the innovation rate should significantly increase when the supply of skilled labor s increases.[7] The Acemoglu model is actually more subtle in the sense that it predicts a change in the direction — not the speed — of technological change. Yet the growth rate, as derived from the above model, should still increase following a (discontinuous) increase in relative skilled-labor supply, which is still at variance with Jones-style evidence, at least up until the mid 1990s. To reconcile the market size explanation with this evidence, Acemoglu (2000) invokes the existence of decreasing returns in R&D aimed at skill-biased technical progress. Now, even if individual researchers may experience decreasing returns in their R&D activities, it is not clear why the *whole* economy should: the exception would be if individual innovations were more like secondary discoveries induced by an economy-wide fundamental breakthrough, which becomes more and more incremental over time.

An alternative way to reconcile the acceleration theory with the evidence on measured productivity is to argue that the recent wave of technological change marked he birth of a new technological paradigm. A natural explanation for the aggregate productivity slowdown is that it takes some time before producers and developers economy-wide become fully acquainted with the new technological platform. This brings us naturally to the notion of *General Purpose Technology*[8], which we develop in the next section.

3 General purpose technologies and the premium to adaptability

In this section we sketch a pair of models that show in more detail how the Nelson-Phelps idea of skill as a measure of adaptability to technological change together with the notion

[6]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; 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.

[7]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 until the late 1990s. But something more is needed to reconcile the above arguments with the fact that productivity growth actually appeared to fall.

[8]See Bresnahan- Trajtenberg (1995), and the papers in the Helpman (1998) volume.

of General Purpose Technology, can help to account for several observed 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.[9]

In the long run, we assume that everyone will adapt to a given technology. Our reason for making this assumption is that in the long run technological progress automates skills. Ballpoint pens are easier to use than straight pens, modern photocopiers take less skill than lithograph machines, operating a car skillfully requires less training than riding a horse skillfully, electronic calculators are less demanding than mental addition, and so forth. Even the ability of fast-food clerks to read is being replaced by computer graphics. Thus in the long run we assume there is no skill bias to technological change. In the short run, however, some people catch on faster than others, and they earn a premium for their adaptability. The introduction of a new GPT can enhance this premium in several ways, as we explain below.

3.1 Between-group inequality

Our first model is a simplified version of the one involving social learning presented in Aghion and Howitt (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 others in a similar situation. In order for a firm to begin experimenting with a new technology it must first find another firm that has used the technology successfully in solving problems similar to the ones this firm faces; it can then use the other firm's experience as a "template" for its own experimentation. For a long time, improvements in knowledge will take place slowly because there are so few successful firms from which to draw a template; but eventually a point will be reached when enough other firms are using the new technology that experimentation snowballs. This quickly raises the demand for skilled labor and thereby raises the skill premium.

We begin by assuming that there are two technologies that can be used in any sector; an

[9] Gould, Moav and Weinberg (2001) also emphasize the interplay between education and random adaptability in explaining between- and within group inequality. In their model, less able and therefore less educated workers acquire technology-specific skills through on the job training, whereas more able workers acquire general skills through education; within-group inequality among low ability workers then results primarily from the random obsolescence of existing technologies, whereas within group inequality among the high ability workers results from changes in the ability composition among them.

old technology and a new one that has just arrived in the form of a GPT. Aggregate final output is produced by labor 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 labor produces intermediate goods using a one-for-one technology, so that $x(i)$ also denotes the labor demand flow in sector i. The total labor force L is divided into skilled and unskilled workers. Whereas old sectors are indifferent between skilled and unskilled workers, experimentation with and implementation of the new GPT can be done only by skilled labor.

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:

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

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

In each sector i, moving from the old to the new GPT requires two steps. First, as indicated above, a firm in that sector must acquire a template on which to base experimentation; second, the firm must succeed in its experimentation with the new GPT. Let n_0 denote the fraction of sectors that have not yet acquired a template, n_1 the fraction of sectors that are currently experimenting on the new GPT, and $n_2 = 1 - n_0 - n_1$ the fraction 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 given by:

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

Now, suppose that for a templated firm to succeed in implementing the new GPT, it must employ at least H units of skilled labor per period. We can think of this labor as being used in 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:

$$\begin{aligned} \dot{n}_1 &= \lambda(n_2)(1 - n_1 - n_2) - \lambda_1 n_1 \\ \dot{n}_2 &= \lambda_1 n_1 \end{aligned}$$

with initial condition $n_1(0) = 0$, $n_2(0) = 0$. The time path of n_0 is 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.[10] 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, because it starts and ends at zero.

The transition process from the old to the new GPT can 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 labor 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 labor market will remain unsegmented, with the real wage being the same for skilled and unskilled labor and determined by the labor-market-clearing equation:

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

where x_O, x_N, and H denote the labor demands respectively by an old manufacturing sector, a sector using the new GPT, and an experimenting sector.[11]

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

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

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

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 labor force, the labor market may 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 labor market clearing conditions. The skilled wage is determined by the labor-market clearing equation for skilled labor:

$$L_s = n_1 H + n_2 x_N,$$

while w_u is obtained from the market-clearing equation for unskilled labor, namely:[12]

$$L - L_s = (1 - n_2) x_O.$$

Figure 2 depicts the time-path of real wages, assuming a relatively large cost H of experimentation.[13] The skill premium, here measured by the ratio (w_s/w_u), starts increasing around when the diffusion of the new GPT across sectors accelerates, as a result of the increased demand for skilled labor 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.[14]

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:

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

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

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

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

[14] This simple model of GPT diffusion and between-group wage inequality can be extended easily to accommodate 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 and 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:

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

This explanation of increased inequality between skill groups is also consistent with the observed dynamic pattern of the skill 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 labor 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. This mechanism provides an explanation for the first puzzle we raised in our Introduction.

With regard to the comparison between the recent period and the early 1900s (the *historical-perspective* issue), the diffusion of the electric dynamo did not result in a comparable increase in the skill premium because that earlier GPT was not nearly as skill-biased as recent information technologies, even in the transitional sense used in the Nelson-Phelps framework. For example, whereas workers operating steam engines needed to know how to maintain and repair their own engines, the maintenance of new electrical machinery only required firms to hire a limited number of skilled workers specialized in that task. Thus the appearance of bottlenecks that segment the labor 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 easy to reconcile also with the *productivity-slowdown* issue. 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 labor 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 the diffusion process. In addition, the explanation appears also to be consistent with the observed deceleration of between-group wage inequality during the late 1990s; 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

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.

the benefits of a GPT that is getting more familiar and more user-friendly.

3.2 Within-group inequality

Existing explanations for the rise in residual wage inequality[15] 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, 1994). In the remaining part of this section, we shall argue, using the Nelson-Phelps framework, that the diffusion of a new general purpose technology can account for the evolution of within-group wage inequality in a way that is consistent with these and other puzzling facts.

The basis of this argument is the model of Aghion, Howitt and Violante (2001)[16], which is built on prior work of Violante (1996, 2001). The model, like that of the preceding section, involves differential rates of adaptability, which in this case are entirely random, among a group of *ex ante* identical workers with the same educational background. The diffusion of a new GPT raises (within-group) inequality for two reasons. First the rise in the speed of technical change associated with the new GPT increases the market premium to those workers who are lucky enough to adapt quickly to new versions of the technology. Second, to the extent that the new GPT generates a wave of secondary innovations that are closely related to one another, its diffusion especially benefits workers who are lucky enough to adapt quickly several times, who can profit from transferring recently acquired knowledge to the task of working with the latest innovations.

[15] See Acemoglu (1998), Rubinstein and Tsiddon (1999) and Galor and Moav (2000) for models of within-group inequality based on differences in innate ability.

[16] The model in this subsection is a stripped-down version of the framework developed in Aghion, Howitt and Violante (2001). In that paper, we assume an overlapping-generations structure and capital-embodied technological progress. We also allow for two dimensions of generality of the GPT - - not just the transferability of skills τ but also the compatibility of capital across successive technologies.

3.2.1 Basic Framework

We construct an infinite-horizon discrete-time model with sequential productivity-improving innovations in a one-good economy. Each period a new technology arrives. At any date t some fraction of workers can work with technology t, which has just arrived. The rest are all able to work with technology $t-1$, which arrived last period. Thus no one will find it profitable to employ a technology that is more than one period old.

Production relations Aggregate output using each technology is given by the same Cobb-Douglas production kernel. Thus the economy's GDP at each date t is:

$$Y_t = K_{0t}^{\alpha} \left(A_t x_{0t}\right)^{1-\alpha} + K_{1t}^{\alpha} \left(A_{t-1} x_{1t}\right)^{1-\alpha}, \qquad 0 < \alpha < 1,$$

where x_{0t} and K_{0t} are the labor and capital inputs working with technology t (that is, in "sector 0"), x_{1t} and K_{1t} are the inputs working with technology $t-1$ (in "sector 1"), and A_t is the (labor-augmenting) productivity parameter of technology t. Each new technology is $(1+\gamma)$ times more productive than the previous one, so:

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

Capital is perfectly mobile between sectors. Therefore in a competitive equilibrium the marginal product of capital must be the same in each sector:

$$\partial Y_t / \partial K_{0t} = \partial Y_t / \partial K_{1t},$$

which can be written, using the Cobb-Douglas production function, as:

$$\alpha \left(K_{0t}/A_t x_{0t}\right)^{\alpha-1} = \alpha \left(K_{1t}/A_{t-1} x_{1t}\right)^{\alpha-1}.$$

This capital-market equilibrium condition implies that the capital stock per efficiency unit of labor must be the same in each sector:

$$K_{0t}/A_t x_{0t} = K_{1t}/A_{t-1} x_{1t} \equiv k_t.$$

Because the capital stock per efficiency unit is the same in each sector, therefore the marginal product of labor will be higher in sector 0, where it works with the new technology, than in sector 1, where it works with the old technology, by the growth factor $1+\gamma$. That is:

$$\begin{array}{lll} \partial Y_t / \partial x_{0t} = (1-\alpha) \left(K_{0t}/A_t x_{0t}\right)^{\alpha} A_t & = (1-\alpha)\, k_t^{\alpha} A_t \\ \partial Y_t / \partial x_{1t} = (1-\alpha) \left(K_{1t}/A_{t-1} x_{1t}\right)^{\alpha} A_{t-1} & = (1-\alpha)\, k_t^{\alpha} A_t / \left(1+\gamma\right) \end{array}.$$

Adaptation, Learning and Transferability At each date t a given fraction σ_t of all workers is randomly given the capacity to adapt to (that is, to work with) the new technology in that period. In equilibrium they will all take advantage of this opportunity by working in sector 0.[17]

Experience with a technology produces learning by doing that enhances the productivity of a worker who remains working with the same technology as last period, by the factor $1 + \eta$. In order to ensure that workers in equilibrium always choose to work with the most recent technology when they are able to adapt, we assume that:

$$0 < \eta < \gamma.$$

If the worker switches to working with a technology that is one period newer, then some of this learning by doing can be transferred, and the worker's productivity is thus enhanced by the factor $1 + \tau\eta$, where τ is a parameter measuring the generality of the technology, with:

$$0 < \tau < 1$$

We assume however that learning cannot be transferred if the worker switches to working with a technology that is more than one period newer. That is, skills acquired by experience eventually (in this case after 2 periods) become obsolete. Moreover, we assume that once a technology is one period old, the benefits of learning by doing on that technology can be shared by all who work with it, even those who did not work with it last period. That is, eventually the skills involved in working with a particular technology become publicly available.

It follows from this description of adaptability, learning and transferability that different workers will supply different amounts of labor input depending upon their work experience. Suppose, as a normalization, that a worker having no experience working with a technology the same as, or one period older than, the one he or she is currently working with supplies one unit of labor input. Then the amount supplied by each worker will be given by the following table:

[17] As we shall see, to work in sector 1 would give the worker a lower current wage, because $\gamma > \eta$. If we suppose, as seems reasonable, that the probability of being adaptable next period is not enhanced by working in sector 1 this period, then working in sector 1 would also yield a lower continuation value, since it would require the worker to give up the possibility of transferring knowledge to the leading edge if he or she were to be adaptable next period.

TABLE 1
Labor units supplied by each worker

	Working this period in sector	
	0	1
Worked last period in sector 0	$1+\tau\eta$	1
Worked last period in sector 1	$1+\eta$	$1+\eta$

Thus every worker employed this period in sector 1, working with the old technology, supplies $1+\eta$ units of labor input no matter where he or she was employed last period, because all such workers benefit equally from the learning by doing of those who worked on the technology last period when it was new. Workers who go from the old sector 1 last period to the new sector 0 this period are skipping a technology so, in accordance with our assumption that learning by doing becomes obsolete in two technological generations, they supply only the basic one unit of labor input. Workers who go from sector 0 to sector 0 are able to adapt to the latest technology two periods in a row, and they are able to supply $1+\tau\eta$ units of labor input because they can transfer the fraction τ of their learning by doing from last period's leading-edge technology to this period's.

3.2.2 Equilibrium Within-Group Wage Inequality

In a competitive labor market each worker will receive a wage equal to his or her supply of labor input, as given by Table 1 above, times the marginal product of labor in his or her sector (sector 0 if the worker is lucky, in the sense of being able to adapt to the latest technology, or sector 1 if unlucky). Thus there will be three separate wages observed at each date t:

$$\begin{aligned}
w_{0t}^{0} &= (\partial Y_t/\partial x_{0t})\,(1+\tau\eta) &&= (1-\alpha)\,k_t^{\alpha}A_t\,(1+\tau\eta)\\
w_{0t}^{1} &= (\partial Y_t/\partial x_{0t}) &&= (1-\alpha)\,k_t^{\alpha}A_t\\
w_{1t} &= (\partial Y_t/\partial x_{1t})\,(1+\eta) &&= (1-\alpha)\,k_t^{\alpha}A_t\,(1+\eta)\,/\,(1+\gamma)
\end{aligned}$$

where w_{0t}^{i} denotes the wage of a someone working with the new technology (in sector 0) who worked in sector i last period, and w_{1t} is the wage of everyone working with the old technology (in sector 1).

The three wage rates are listed above in decreasing order of size. Thus the highest wage w_{0t}^{0} will be earned by those workers who have been lucky enough to be adaptable two periods in a row, while the lowest wage w_{1t} will be earned by those who are unlucky this period.

Our measure of inequality is the ratio R between the maximum and minimum wage:

$$R_\omega = w_{0t}^0 / w_{1t} = (1 + \tau\eta)(1 + \gamma) / (1 + \eta) > 1$$

It follows that this measure of within-group wage inequality increases with both the rate of technical progress as parametrized by γ and the transferability of knowledge 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. As mentioned before, 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 technical change since the mid-1970s, e.g. as measured by the decline in the quality-adjusted price of equipment goods; likewise, the "general nature" of the technological wave in communication and information implies that the acceleration phase of that wave was probably accompanied by an increased similarity between successive vintages, which in turn should have increased the degree of skill transferability across technologies.

The implied surge in inequality is purely "residual" as workers are ex-ante equal. Furthermore, the model's residual wage inequality is linked to the stochastic nature of workers' adaptability to the newest vintage more than with innate ability, thus the model predicts that the diffusion of a new GPT should primarily affect the transitory component of income, since individual luck in adapting quickly to a new sector will obviously vary over time, in line with the empirical work of Blundell and Preston (1999), and Gottschalk and Moffitt (1994)

4 Concluding Remarks

In this paper we have argued that the notion of human capital shaped by Nelson and Phelps (1966), as labor *adaptable* to the new technologies, is crucial in understanding the labor market experience of the U.S. economy in the past three decades. The arrival of a new General Purpose Technology and the implied rise in the demand for labor adaptable to the new technological platform produced a surge in the return to adaptability.

Insofar as formal education is correlated with adaptability, this mechanism can explain the dynamics of the educational premium, as argued in Section 3.1. 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, as argued in Section 3.2. Moreover, if the economy comprises several educational groups of workers, with more educated workers being more able to transfer recently acquired knowledge to the newest vintages, then a fall in the education premium could easily 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 labor that would occur when the information revolution hits the economy might temporarily reduce the education premium, but 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 continue to sustain a rise in within-group inequality. The alternative theories of within-group inequality based upon market-size effects and unobserved innate ability do not seem to provide an equally convincing explanation for these divergent patterns of between-group and within-group inequality.

What policy lessons can we learn from the labor market experience of the past thirty years? Rapid and major technical change transmits benefits to the whole society in the long-run. However, in the short run it can create large redistributions. Adaptability is the key skill that allows workers to reap the benefits of technological improvements form the early start and prevents situations of skill obsolescence and prolonged exclusion from the labor market.[18] This makes the case for the creation of a more favorable institutional environment for the acquisition of general skills for the young (formal education) and for retraining programs later in life ("continuous learning") for those workers whose jobs are made obsolete by technological progress.

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 (1998): *Endogenous Growth Theory.* Cambridge: MIT Press.

Aghion, P., P. Howitt and G. Violante (2001): "General Purpose Technology and Wage Inequality", mimeo UCL.

[18] Murphy and Topel (1997) documents the rise in nonparticipation and long-term unemployment among prime-aged males in the US and argues that this phenomenon is strictly linked to the surge in wage inequality.

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

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.

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.

DiNardo John, Nicole Fortin, and Thomas Lemieux, (1996): "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semi-parametric Approach," *Econometrica*, 64 , 1001-1044.

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 (2000): "Ability Biased Technological Transition, Wage Inequality, and Economic Growth", *Quarterly Journal of Economics, 115,* pp. 469-97.

Galor, O. and D. Tsiddon (1997): "Technological Progress, Mobility, and Economic Growth," *American Economic Review*, 87, pp. 363-82.

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.

Gould, E., O. Moav and B. Weinberg (2001): "Precautionary Demand for Education, Inequality, and Technological Progress", forthcoming in the *Journal of Economic Growth.*

Greenwood, J. and M. Yorukoglu (1997): "1974", *Carnegie-Rochester Series on Public Pol-*

icy, 46, pp. 49-95.

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

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

Jones, C. (1995): "R&D Based Models of Economic Growth", *Journal 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.

Katz, L. and K. Murphy (1992): "Changes in Relative Wages: Supply and Demand Factors", *Quarterly Journal of Economics,* pp. 35-78.

Krusell, P., L. Ohanian, J.-V. Ríos-Rull and G. Violante (2000): "Capital-Skill Complementarity and Inequality", *Econometrica, 68:5, pp. 1029-1054.*

Lee, David (1999), "Wage Inequality in the U.S. During the 1980s: Rising Dispersion or Falling Minimum Wage?," *Quarterly Journal of Economics*, 114, 977-1023.

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.

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

Murphy, Kevin and Robert Topel, (1997): "Unemployment and Nonemployment," *American Economic Review Papers and Proceedings*, 87 , 295-300.

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.

Rubinstein, Y. and D. Tsiddon (1999): "Coping with Technological Progress: The Role of

Ability in Making Inequality so Persistent", mimeo Hebrew University.

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

Violante, G.L. (2001): "Technological Acceleration, Skill Transferability and the Rise in Residual Inequality," forthcoming, *Quarterly Journal of Economics*

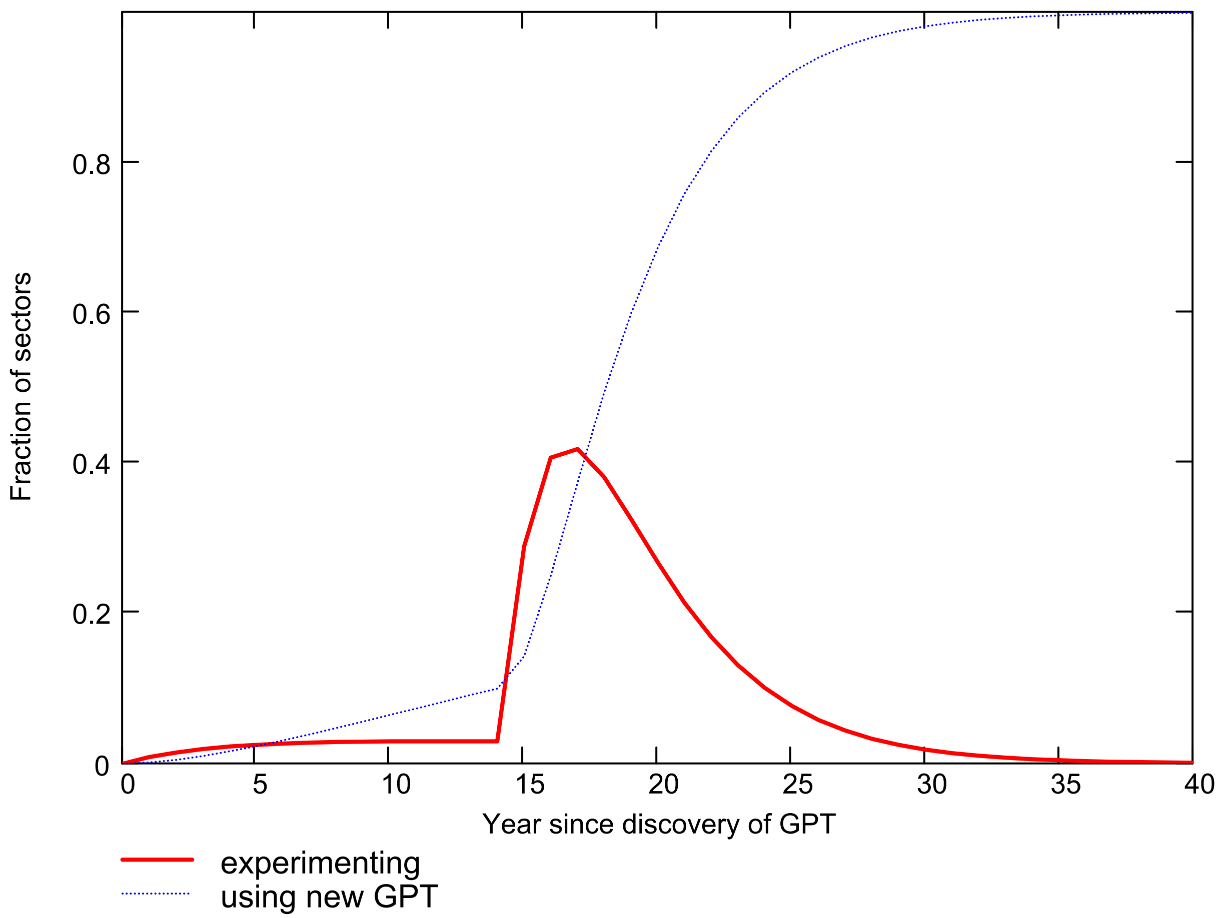

Fraction of sectors
0.8
0.6
0.4
0.2
0
0
5
10
15
20
25
30
35
40
Year since discovery of GPT
experimenting
using new GPT

Figure 1

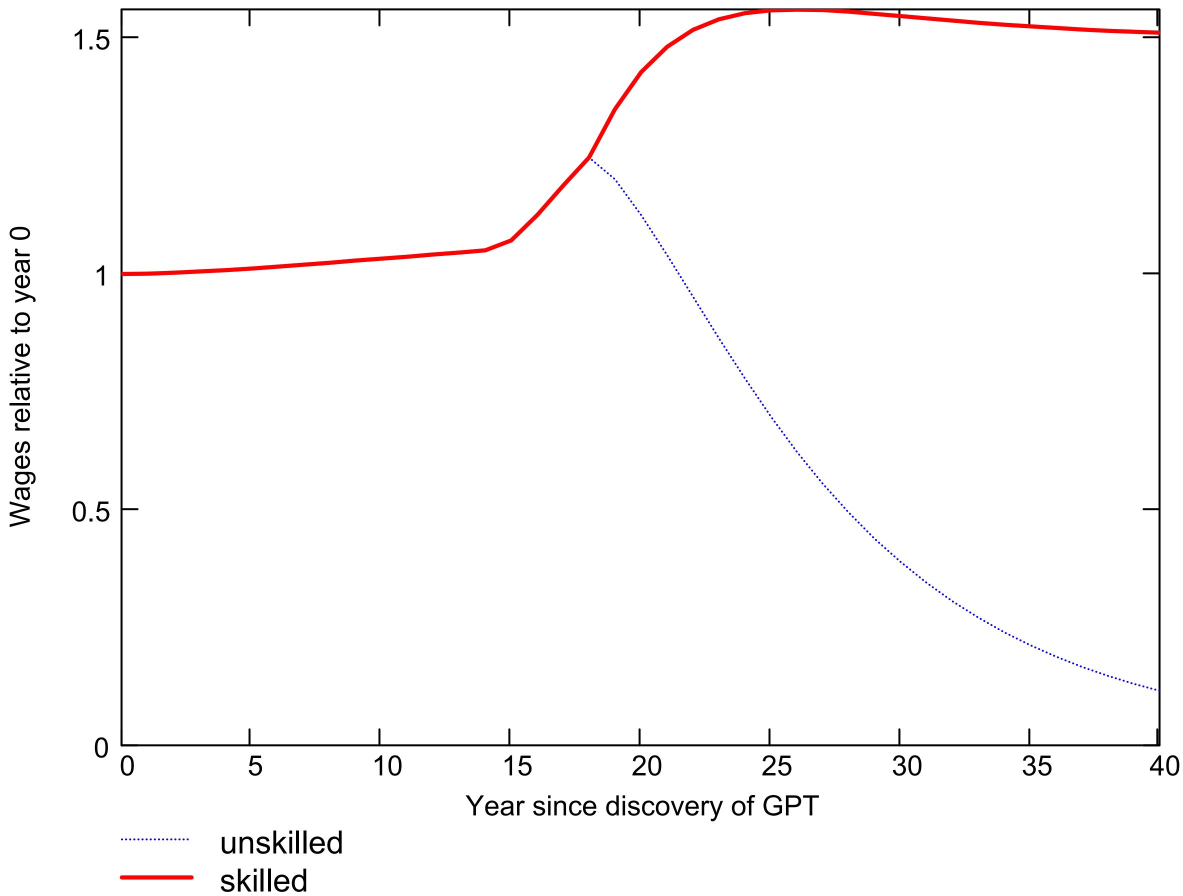

Wages relative to year 0
0
0.5
1
1.5
0
5
10
15
20
25
30
35
40
Year since discovery of GPT
unskilled
skilled

Figure 2

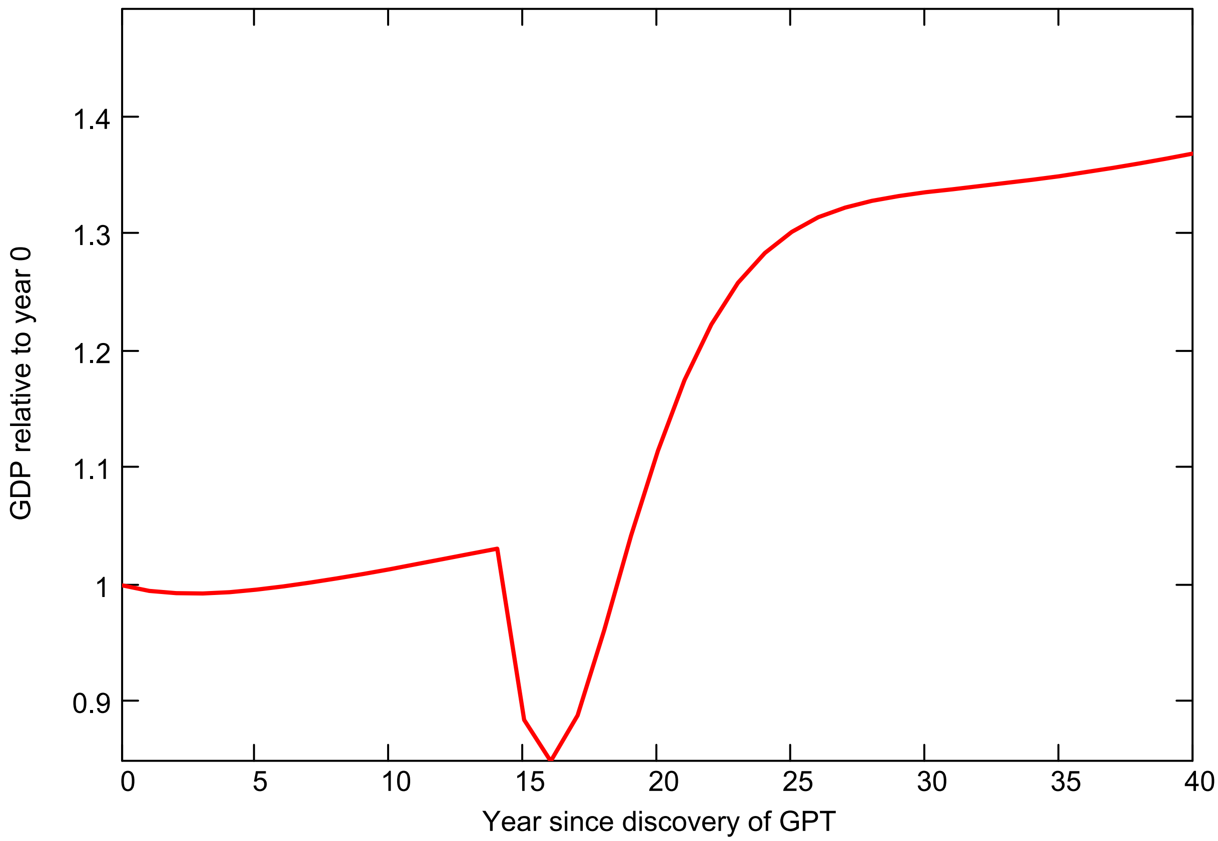

1.4
1.3
1.2
1.1
1
0.9
GDP relative to year 0
0
5
10
15
20
25
30
35
40
Year since discovery of GPT

Figure 3