Working Paper Series

The Effects of Technical Change on Labor Market Inequalities

WP 04-08

Andreas Hornstein
Federal Reserve Bank of Richmond

Per Krusell
Princeton University

Giovanni L. Violante
New York University

THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ▪ BALTIMORE ▪ CHARLOTTE

The Effects of Technical Change on Labor Market Inequalities*

Andreas Hornstein† Per Krusell‡ Giovanni L. Violante§

Federal Reserve Bank of Richmond WP 04-08
This Draft: December 20, 2004

Abstract

In this chapter we inspect economic mechanisms through which technological progress shapes the degree of inequality among workers in the labor market. A key focus is on the rise of U.S. wage inequality over the past 30 years. However, we also pay attention to how Europe did not experience changes in wage inequality but instead saw a sharp increase in unemployment and an increased labor share of income, variables that remained stable in the U.S. We hypothesize that these changes in labor market inequalities can be accounted for by the wave of capital-embodied technological change, which we also document. We propose a variety of mechanisms based on how technology increases the returns to education, ability, experience, and "luck" in the labor market. We also discuss how the wage distribution may have been indirectly influenced by technical change through changes in certain aspects of the organization of work, such as the hierarchical structure of firms, the extent of unionization, and the degree of centralization of bargaining. To account for the U.S.-Europe differences, we use a theory based on institutional differences between the United States and Europe, along with a common acceleration of technical change. Finally, we briefly comment on the implications of labor market inequalities for welfare and for economic policy.

Keywords: Inequality, Institutions, Labor Market, Skills, Technological Change.

JEL Classification: D3, J3, O3.

---

*Prepared for the *Handbook of Economic Growth* (Philippe Aghion and Steven Durlauf, Editors). We are grateful to Philippe Aghion for his suggestions on how to improve an early draft. We thank Stephan Fahr, Giammario Impullitti, Matthew Lindquist, and John Weinberg for comments, Hubert Janicki for research assistance and Eva Nagypal and Bruce Weinberg for providing their data. Krusell thanks the NSF for research support. Violante thanks the CV Starr Center for research support. Any opinions expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. For correspondence, our e-mail addresses are: `andreas.hornstein@rich.frb.org`, `pkrusell@princeton.edu`, and `gianluca.violante@nyu.edu`.

†Federal Reserve Bank of Richmond.

‡Princeton University, Institute for International Economic Studies, CAERP, CEPR, and NBER.

§New York University and CEPR.

1 Introduction

In this chapter we discuss the recent three decades of data on technology, productivity, and labor market outcomes. In particular, we explore the hypothesis that technological change has affected the labor market in various ways. We argue that (i) there is ample evidence indicating significant capital-embodied and/or skill-biased technological change and that (ii) this kind of technological change would plausibly lead to many of the transformations in the labor markets that we have observed. On the one hand, we are interested in possible implications of non-neutral technological change—of the kind we think we have experienced—on variables like wage inequality, unemployment, labor share, and unionization. On the other hand, we explore the possibility that the labor market can be used as an additional source of evidence of non-neutral technological change, a testing ground of sorts.

The past 30 years are particularly informative because they have contained rather important trend changes in several variables. We have seen a productivity slowdown common to all industrialized countries and common to almost all industries, together with continuing structural change away from manufacturing and toward services. An exception to this widespread productivity slowdown was the fast and accelerating productivity growth in the industries producing investment goods, in particular those producing equipment. Only very recently has there been a more widespread acceleration of productivity growth. Of course, in this context we are arguably in the midst of an “Information Technology Revolution.” We also discuss evidence of changes in the workplace—how production within firms is organized—possibly reflecting underlying changes in technology.

In the labor market, we have seen a sharp increase in wage inequality in the United States contrasting a roughly flat development in Europe, whereas we have witnessed a strong increase in European unemployment and no trend in U.S. unemployment.[1] The organization of labor markets seems to have changed too: for example, unions have lost prevalence during this period, and to the extent there have been unions, centralized bargaining has been

---

[1] Although the word inequality literally would suggest a zero-one classification—either there is inequality or there is equality—we will use the term loosely to reflect some measure of dispersion. That is, we will attach quantifiers such as “more” or “less” to the word.

replaced by decentralized bargaining in many sectors. Are all these developments consistent with basic economic theory and a short list of underlying technological driving forces? We argue that they are. To make our argument more convincing, we also put the past three decades in a historical perspective, going back as far as the early 20th century with data on technological change and the skill premium.

One distinctive feature of this literature is that the many different ideas have been presented in a wide variety of theoretical frameworks ranging from the neoclassical Ramsey-Cass-Koopmans growth model to the Schumpeterian endogenous growth model; from the traditional McCall search model to the Lucas-Prescott island economy; from the Mortensen-Pissarides matching model to the competitive directed search framework; and from the Bewley-Aiyagari incomplete-markets model to Arrow-Debreu economies with limited enforcement. We think two main reasons exist for the lack of a unified framework of analysis. First, this field of research is still relatively young; second, departing from the competitive model in studying labor markets is fairly natural, and many alternative frameworks exist that incorporate frictions. The main drawback of the lack of a unifying framework—we will repeat it often in the chapter—is that making structurally based quantitative comparisons between different mechanisms is difficult.

To us, these heterogeneous approaches pose a formidable challenge in the exposition. Our solution has been to give priority to presenting a range of ideas, using a variety of theoretical setups, rather than to discuss in great detail a few more specific frameworks. This approach has necessitated a summarization of some rather rich models in a few key equations, which misses some of the elegance and richness of the original frameworks. We hope, however, that our spanning a wide spectrum of ideas and macroeconomic effects of technological change helps paint a picture that is broader and that, at least in an impressionistic way, suggests that the main underlying hypothesis we are proposing is quite reasonable.

The presentation of the ideas in this chapter is organized into four parts. In the first part, Section 2, we review the main trends in the data on technological change and labor market inequalities. We then cover two kinds of theories that could account for the data.

In the second part of the chapter we cover "neoclassical" theory (i.e., models where wages directly reflect marginal productivity). We view the firm as hiring labor of different skill levels in a competitive and frictionless labor market. Wages, thus, will be influenced by technology in a very direct way. Similarly, the returns to education, ability, and experience, which we discuss in detail in Sections 3 and 4, respectively, will be directly tied to changes in technology. Therefore, within these kinds of theories, the shape of the production function of the firm is crucial. We then move beyond the production function of the firm or, rather, we attempt to go inside it. In particular, Section 5 explores the possibility that the organization of the workforce also has changed within firms. These transformations, of which there is some documentation, are arguably also a result of the kind of technological change we look at in this chapter. We point, in particular, to a recent literature that explores how firms are organized and how the IT revolution, by inducing organizational changes in the firm, had a substantial impact on wage inequality.

The second class of theories we cover, in the third part of the chapter, relies more on frictions in the labor market and deals more directly with how this market is organized. Here, technological change can still directly influence wages but there are new channels. For one, wages may not only reflect marginal productivity. Moreover, now unemployment is more in focus and is a function of technology, and since unemployment—through workers' outside option—may also feed back into wages, the picture becomes yet more complex. In the context of how wages are set, we furthermore argue in Section 6 that the importance of unions and their modus operandi are influenced by technology and, more generally, that labor income as a share of total income may respond to technological change in the presence of unions. An important point that we make in Section 7 is that "luck" can be a key part of wage outcomes for individuals active in a labor market with frictions, such as the search/matching frameworks, and that the "return to luck" can be greatly affected by technology as well. Finally, government participation in labor markets—labor-market "institutions," in the form of unemployment benefits, firing costs, and so on—likely interacts with technology in determining outcomes, and Section 8 completes the third part of this chapter by analyzing the interaction between technological shocks and labor market institutions in the context of the comparison between the United States and Europe.

The fourth and final part of the chapter asks the “So what?” question: given the significant transformations observed, is a government policy change called for? Our discussion here is very brief. It mainly points out that a basic element underlying any decisions on policy, namely, what the welfare outcomes of the changes in wages, unemployment, and so on, are for different groups in society, is studied only partially in the literature so far. Studies of changes in expected lifetime income of different groups exist, but it is reasonable to assume that risk matters too, especially with trend changes as large as those observed (at least to the extent they are hard to foresee and insure). In Section 9, we therefore cover some examples of more full-fledged attempts to look at the distribution of consumption and welfare outcomes of the changes in technology/labor market outcomes. Finally, Section 10 concludes the chapter.

2 A Look at the Facts

Before modelling the economic forces that connect changes in technology to labor market outcomes, it is useful to begin by summarizing how labor market inequalities and the aggregate technological environment evolved over the past three decades.

2.1 Labor Market Inequalities

In the late 1980s and early 1990s, an extensive body of empirical work has started to systematically document the changes in the U.S. wage structure over the past three decades. Levy and Murnane (1992) give the first overview of an already developed empirical literature. To date, Katz and Autor (1999) and, more recently, Eckstein and Nagypal (2004), offer the most exhaustive description of the facts. In between, numerous other papers have contributed significantly to our understanding of the data on wage inequality.[2]

---

[2] We refer the reader to the bibliographic lists in Levy and Murnane (1992), Katz and Autor (1999), and Eckstein and Nagypal (2004) for more details.

The typical data source used in the empirical work on the subject is the sequence of yearly cross-sections in the March *Current Population Survey (CPS)*. The other important data source is the longitudinal *Panel Study of Income Dynamics (PSID)*. In this section, we limit ourselves to stating the main facts and briefly commenting on them, omitting the details on the data sets, the sample selection, and the calculations that can be found in the original references. Unless otherwise stated, the data refer to a sample of male workers with strong attachment to the labor force, i.e., full-time, full-year workers.[3]

Observation 1 Wage inequality in the United States is today at its historical peak over the post World War II period. However, early in the century it was even larger. The returns to college and high school fell precipitously in the first half of the century and then rose again until now (Goldin and Katz, 1999).

In other words, the time series for inequality over the past 100 years is "U-shaped." Although the bulk of this chapter is devoted to interpreting the dynamics of the wage structure over the past three decades, it is useful to put the evidence in a historical perspective to appreciate that the high current level of inequality is not a unique episode in U.S. history. The rest of the facts characterize the evolution of inequality since the mid 1960s.[4]

Observation 2 Wage inequality increased steadily in the United States starting from the early 1970s. The 90-10 weekly wage ratio rose by 35 percent for both males and females in the period 1965-1995: from 1.20 to 1.55 for males, and from 1.05 to 1.40 for females. The increase in inequality took place everywhere in the wage distribution: both the 90-50 differential and the 50-10 differential rose by comparable amounts (Katz and Autor, 1999).

---

[3] Eckstein and Nagypal (2004) systematically document all the facts for males and females separately. Typically, measures of inequality in the literature refer to hourly or weekly wages, that is, they isolate the evolution of the "price" of certain labor market skills. The use of hourly or weekly wages then avoids the contamination of the data with endogenous labor supply decisions that, for example, is present in annual earnings.

[4] In Section 5, we return briefly to this historical pattern. In passing, we note that the data seems at odds with the so-called "Kuznets Hypothesis," i.e., the conjecture that income inequality first increases and then decreases as economies grow.

Qualitatively, the rise in inequality is present independently of the measure of dispersion and of the definition of labor income. For example, the standard deviation of log wages for males rose from 0.47 in 1965 to 0.62 in 1995, the Gini coefficient jumped from 0.25 to 0.34 (Katz and Autor, 1999), and the mean-median ratio rose from 1.00 to 1.18 over the same period (Eckstein and Nagypal, 2004). Inequality of annual earnings increased even more.[5]

Observation 3 The average and median wage have remained constant in real terms since the mid-1970s. Real wages in the bottom of the wage distribution have fallen substantially. For example, the 10th wage percentile for males declined by 30 percent in real terms from 1970 to 1990 (Acemoglu, 2002a).[6] On the contrary, salaries in the very top of the wage distribution have grown rapidly. In 1970, the workers in the top 1 percent of the wage distribution held 5 percent of the U.S. wage bill, whereas in 1998 they received over 10 percent (Piketty and Saez, 2003).

A large part of the absolute increase of top range salaries is associated with the surge in CEO compensation. Piketty and Saez (2003) document that in 1970 the pay of the top 100 CEOs in the United States was about 40 times higher than the average salary. By 2000 those CEOs earned almost 1,000 times the average salary.

We now list a set of facts on the evolution of *between-group* inequality, i.e. inequality between groups of workers classified by observable characteristics (e.g., gender, race, education, experience, occupation). For this purpose, it is useful to write wages w_{it} using the Mincerian representation

$$\ln w_{i,t} = X'_{i,t} p_t + \omega_{it}, \tag{1}$$

where X_{it} is a vector measuring the set of observable features of individual i at time t, p_t can be interpreted as a vector of prices for each characteristic in X, and ω_{it} is the residual unobserved component.

---

[5] The reason is, perhaps surprisingly, not a rise in the cross-sectional variance of hours worked, but rather a substantial increase in the wage-hours correlation over the past 30 years. See Heathcote et al. (2003) for an account of these facts.

[6] Note, however, that the wages of the 10th wage percentile have started to increase again since the late 1990s (Eckstein and Nagypal, 2004).

Observation 4 The returns to education increased slightly from 1950 to 1970, fell in the 1970s, increased sharply in the 1980s, and continued to increase, although at a slower pace, in the 1990s. For example, the college wage premium—defined as the ratio between the average weekly wage of college graduates (at least 16 years of schooling) and that of workers with at most a high school diploma (at most 12 years of schooling)—was 1.45 in 1965, 1.35 in 1975, 1.50 in 1985, and 1.70 in 1995 (Eckstein and Nagypal, 2004). If one estimates the coefficient on educational dummies in a standard Mincerian wage regression like (1), the finding is similar: the annual return to a college degree (relative to a high-school degree) was 33 percent in the 1980s and over 50 percent in the 1990s (Eckstein and Nagypal 2004).

We plot the college wage premium over the period 1963-2002 in Figure 1 (top panel).[7] Interestingly, if one slices up the college-educated group more finely into workers with post-college degrees and workers with college degree only, the rise in the skill premium is still very apparent. The return to post-college education relative to college education doubled from 1970 to 1990 (Eckstein and Nagypal 2004).

FIGURE 1

Observation 5 The returns to professional and white-collar occupations relative to blue-collar occupations display dynamics and magnitudes similar to the data stratified by education. For example, the professional-blue collar premium rose by 20 percent from 1970 to 1995 (Eckstein and Nagypal 2004).

Occupation is an interesting dimension of the wage structure that, until recently, received very little attention. For example, the "returns to occupation" appear large and significant, over and beyond returns to education. We discuss the theories of wage inequality that stress the changes in occupational structure in Section 7.

---

[7] Authors differ in their treatment of workers who have attended college for some years, but did not obtain a college degree. In Figure 1 (top panel), we have followed the bulk of the literature and assigned half of them to the numerator and half of them to the denominator (e.g., Autor et al. 1998).

Observation 6 The returns to experience increased in the 1970s and the 1980s and leveled off in the 1990s. For example, the ratio of weekly wages between workers with 25 years of experience and workers with 5 years of experience rose from 1.3 in 1970 to 1.5 in 1995 (Katz and Autor, 1999). An analysis by education group shows that the experience premium rose sharply for high-school graduates but remained roughly constant for college graduates (Weinberg, 2003b).

It is worth emphasizing that although entry of the baby-boomers into the labor market in the early 1970s had a significant impact on the experience premium, the dynamics described above are robust to this and other demographic effects. See for example, Juhn et al. (1993).[8]

Observation 7 Inequality across race and gender declined since 1970. The black-white race differential, for workers of comparable experience, fell from 35 percent in 1965 to 20 percent in 1990 (Murphy and Welch, 1992). The female-male wage gap fell from 45 percent in 1970 to 30 percent in 1995 (Katz and Autor, 1999).

We plot of the gender wage gap over the period 1963-2002 in Figure 1 (bottom panel). A unifying theory of the changes in the wage structure based on technological change should have something to say about gender as well as race. Admittedly, these two dimensions of inequality have been largely neglected by the literature. We return briefly to the gender gap in Section 4.

Observation 8 The composition of the working population changed dramatically over the past 40 years: in the period 1970-2000, women's labor force participation rate rose from 49 percent to 73 percent; college graduates rose from 15 to 30 percent of the male labor force and from 11 to 30 percent of the female labor force; professionals soared from 24 to 33 percent of the male labor force and from 8 to 28 percent of the female labor force (Eckstein and Nagypal, 2004).

---

[8] More recently, however, Card and Lemieux (2001) have argued in support of some "vintage effects" in the return to education. In particular, they argue that the college-high school premium is somewhat larger among the most recent cohorts of young workers entering the labor market.

We plot the relative supply of skilled workers and female workers over the time period 1963-2002 in Figure 1 (top and bottom panel, respectively).[9]

In terms of equation (1) , one can define the between-group component of wage inequality as the cross-sectional variance of $X_{it}'p_t$ and the within-group component as the variance of the residual ω_{it}. The fraction accounted for by observable characteristics, in turn, can be decomposed into what is caused by a change in the dispersion in the quantities of observable characteristics (X_{it}) for given vector of prices and into what is due to a change in the prices associated to each observable characteristic (p_t) for a given distribution of quantities.

Observation 9 Overall, changes in quantities and prices of observable characteristics (gender, race, education, experience) explain about 40 percent of the increase in the variance of log wages from 1963 to 1995. The price component is by far larger than the quantity component. Increasing *within-group inequality*, i.e., wage dispersion within cells of "observationally equivalent" workers, accounts for the residual 60 percent of the total increase. With respect to the timing, the rise in within-group inequality seems to anticipate that of the college premium by roughly a decade (Juhn et al. 1993).[10]

One can specify further the structure of the residual ω_{it} of equation (1), for example as

$$\omega_{it} = \phi_t \alpha_i + \varepsilon_{it},$$

where α_i is the permanent part of unobservable skills (e.g., "innate ability"), ϕ_t is its time-varying price, and ε_{it} is the stochastic component due to earnings shocks whose variance is also allowed to change over time. If one is prepared to assume that the distribution of innate ability in the population is invariant, then with the help of panel data one can separate the rise in the return to ability from the increase in the volatility of transitory earnings shocks.

Observation 10 Around one-half of the rise in residual earnings inequality is explained by the permanent components (e.g., a higher return to ability), with the rest accounted

---

[9] Skilled and unskilled labor are defined as in footnote 7.

[10] Notice that, typically, occupation is excluded from these regressions. Including occupation would reduce the fraction of unexplained wage variance.

for by transitory earnings shocks (Gottschalk and Moffitt 1994).[11]

Interestingly, the rise in the transitory component is not due to higher job instability or labor mobility (Neumark, 2000), but rather to more volatile wage dynamics, in particular faster wage growth on the job and more severe wage losses upon displacement (Violante 2002).

In Table 1 we report some key numbers on unemployment, wage inequality, and labor income shares for several OECD countries at five-year intervals from 1965 to 1995. We are particularly interested in the comparison between the United States and *continental* European countries (averaged in the row labeled Europe Average). For completeness, we include data for the United Kingdom and Canada, whose behavior falls somewhere between that of the United States and continental Europe.

TABLE 1

Observation 11 The time pattern of wage inequality over the past 30 years differs substantially across countries. The U.K. economy had a rise in wage inequality similar to that in the U.S. economy, except for the fact that the average real wage in the United Kingdom has kept growing (Machin 1996). Continental European countries had virtually no change in wage inequality, whereas over the same period they had large increases in their unemployment rates (roughly, all due to longer unemployment durations) and a sharp fall in the labor income share in GDP. On the contrary, in the United States both the unemployment rate and the labor share have remained relatively constant (Blanchard and Wolfers, 2000).

In 1965 the unemployment rate in virtually every European country was lower than in the United States. Thirty years later, the opposite was true: the U.S. unemployment rate rose only by 1.7 percent from 1965 to 1995, whereas the average unemployment rate increase of European countries was 8.4 percent.[12]

---

[11] Note that a rise in the return on ability does generate an increase in cross-sectional variation of wages because it multiplies individual ability in the log-wage Mincerian equation.

[12] Notice, however, that in the United States non-participation of the low-skilled males rose from 7 percent to 12 percent from the early 1970s to the late 1990s (Juhn, 1992 and Murphy and Topel, 1997).

The labor income share has declined only marginally in the United States—by 1.5 percentage points from 1965 to 1995—while on average it fell by almost 6 points in Europe. Wage inequality, measured by the percentage differential between the ninth and the first earnings deciles for male workers, rose only slightly in Europe, by 4 percent in the period from 1980 to 1995, and it even declined in some countries (Belgium, Germany, and Norway). Recall that, over the same period, earnings inequality surged in the United States: the OECD data show a rise of almost 30 percent, close to the numbers we reported earlier in this section.

Interestingly, the European averages hide much less cross-country variation than one would expect, given the raw nature of the comparison. For example, in 11 out of the 14 continental European countries, the increase in the unemployment rate has been larger than 6 percentage points, and in 9 countries the decline in the labor share has been greater than 5 percentage points.

Recently, Rogerson (2004) has argued that if one focuses on *employment* rate differences between the United States and Europe rather than on unemployment rate differences, a new set of insights emerges from the data. Employment rates in the United States start to increase relative to European employment rates twenty years before the divergence in unemployment rates. Moreover, the increase in European unemployment rates is correlated with the decline of European manufacturing employment.

2.2 Technological Change

The standard measure of aggregate technological change, total factor productivity (TFP), does not distinguish between the different ways in which technology grows. First, technology growth may differ across final-output sectors and second, it may have different effects on the productivity of different input factors. The recent experience of developed countries, however, seems to suggest that in the past 30 years technological change has originated in particular sectors of the economy and has favored particular inputs of production.

Arguably, the advent of microelectronics (i.e., microchips and semiconductors) induced a sequence of innovations in information and communication technologies with two features.

First, *sector-specific* productivity (SSP) growth substantially increased the productivity of the sector that produces new capital equipment, making the use of capital in production relatively less expensive. Second, *factor-specific* productivity (FSP) growth favored skilled and educated labor disproportionately. In other words, the recent technological revolution has affected the production structure in a rather asymmetric way.

Our assessment of the importance of SSP and FSP changes relies heavily on observed movements in relative prices. For SSP change, we rely on the substantial decline of the price of equipment capital relative to the price of consumption goods, a process that does not show any sign of slowing down. On the contrary, it shows an acceleration in recent years. For FSP change, we rely on the substantial increase in the wage of highly educated workers relative to less educated workers, the skill premium.

We first review the Solow growth accounting methodology for TFP within the context of the one-sector neoclassical growth model and then introduce SSP accounting and how it applies to the idea of capital-embodied technical change.[13] Next, we discuss how an acceleration of capital-embodied technical change might relate to the much-discussed TFP growth slowdown in the '70s and '80s; here, we discuss the possible relevance of the concept of General Purpose Technologies (GPTs). Finally, we explain the mapping between relative wages and FSP changes.

2.2.1 Total Factor Productivity Accounting

Standard economic theory views production as a transformation of a collection of inputs into outputs. We are interested in how this production structure is changing over time. At an aggregate *National Income and Product Accounts* (NIPA) level we deal with some measure of aggregate output, y, and two measures of aggregate inputs: capital, k, and labor, l. The production structure is represented by the production function, F: $y = F(k, l, t)$. Since the production structure may change, the production function is indexed by time, t. Aggregate total factor productivity changes when the production function shifts over time, i.e., when

---

[13] Our presentation is instrumental to the discussion of the impact of technological change on labor markets, and hence it is kept to the bare minimum. Jorgenson's (2004) chapter of this Handbook provides an exhaustive treatment of traditional and modern growth accounting.

there is a change in output which we cannot attribute to changes in inputs. More formally, the marginal change in output is the sum of the marginal changes in inputs, weighted by their marginal contributions to output (marginal products), and the shift of the production function, $\dot{y} = F_k \dot{k} + F_l \dot{l} + F_t$.[14] This is usually expressed in terms of growth rates as

$$\hat{y} = \eta_k \hat{k} + \eta_l \hat{l} + \hat{A}, \text{ with } \hat{A} = F_t / F, \tag{2}$$

where hats denote percentage growth rates, and the weight on an input growth rate is the elasticity of output with respect to the input: $\eta_k = F_k k/F$ and $\eta_l = F_l l/F$. Alternatively, if we know the elasticities, we can derive productivity growth as output growth minus a weighted sum of input growth rates.

Solow's (1957) important insight was that, under two assumptions, we can replace an input's output elasticity—which we do not observe—with the input's share in total revenue, for which we have observations. First, we assume that production is constant returns to scale, i.e., that if we are to double all inputs, then output will double, implying that the output elasticities sum to one: $\eta_k + \eta_n = 1$. Second, we assume that producers act competitively in their output and input markets, i.e., that they take the prices of their products and inputs as given. Profit maximization then implies that inputs are employed until the marginal revenue product of an input is equalized with the price of that input. In turn, this means that the output elasticity of an input is equal to the input's revenue share. For example, for the employment of labor, profit maximization implies that $p_y F_l = p_l$, which can be rewritten as $\eta_l = F_l l/F = p_l l / p_y y = \alpha_l$ (p_i stands for the price of good i). With these two assumptions, we can calculate aggregate productivity growth, also known as total factor productivity (TFP) growth, as

$$\hat{A} = \hat{y} - (1 - \alpha_l) \hat{k} - \alpha_l \hat{l}. \tag{3}$$

The Solow growth accounting procedure has the advantage that its implementation does not require very stringent assumptions with respect to the production structure, except con-

---

[14] The marginal change of a variable is its instantaneous rate of change over time; that is, if we write the value of a variable at a point in time as $x(t)$, then the marginal change is the time derivative, $\dot{x}(t) = \partial x(t) / \partial t$. Nothing is lost in the following if the reader interprets $\dot{x}(t)$ as the change of a variable from year to year; that is, $x(t) - x(t-1)$.

stant returns to scale, and it does not require any information beyond measures of aggregate output and input quantities and the real wage. This relatively low information requirement comes at a cost: this aggregate TFP measure does not provide any information on the specific sources or nature of technological change.

Given available data on quantities and prices for industry outputs and inputs, it is straightforward to apply the Solow growth accounting procedure and obtain measures of sector-specific technical change (see, for example, Jorgenson et al., 1987). Recently Jorgenson and Stiroh (2000) have documented the substantial differences in output and TFP growth rates across U.S. industries over the period 1958-1996. In particular, they point out that TFP growth rates in high-tech industries producing equipment investment are about three to four times as high as a measure of aggregate TFP growth. Also based on industry data, Oliner and Sichel (2000) and Jorgenson (2001) attribute a substantial part of the increase of aggregate TFP growth over the second half of the 1990s to one industry: semi-conductors.

2.2.2 Sector-Specific Productivity Accounting

The convincing evidence for persistent differences of SSP growth raises the potential of serious aggregation problems for the analysis of aggregate outcomes. We now discuss SSP accounting in a simple two-sector growth model that focuses on the distinction between investment and consumption goods. This approach provides a straightforward measure of SSP growth, and it keeps the aggregation problems manageable. Based on this approach, we present evidence of substantial increases of the relative productivity in the equipment-investment goods producing industries and stagnant productivity in the consumption goods industries since the mid 1970s.

Greenwood et al. (1997) use a two-sector model of the economy—where one sector produces consumption goods and the other new capital—to measure the relative importance of total-factor productivity changes in each of these sectors. Goods —consumption, c and new capital, x—are produced using capital and labor as inputs to constant-returns-to-scale technologies,

$$c = A_c F(k_c, l_c) \text{ and } x = A_x F(k_x, l_x); \tag{4}$$

and total factor inputs can be freely allocated across sectors,

$$k_c + k_x = k \text{ and } l_c + l_x = l. \tag{5}$$

Note that we have assumed that factor substitution properties are the same in the two sectors; that is, the functions relating inputs to outputs are the same. One can show that with identical factor substitution properties, the two-sector economy is equivalent to a one-sector economy with exogenous changes in the relative price of investment goods, $1/q$

$$y = c + x/q = A_c F(k, l). \tag{6}$$

In particular, the relative price of investment goods is the inverse of the relative productivity advantage of producing new capital goods:[15]

$$q = A_x / A_c. \tag{7}$$

The relative productivity of the investment goods sector is also called "capital-embodied" technical change, because q can be interpreted as the productivity level (quality) embodied in new vintages of capital.[16]

Accounting for quality improvements in new products is a basic problem of growth accounting.[17] This is especially true for our framework since we measure investment in terms of constant-quality capital goods. In a monumental study, Gordon (1990) constructed quality-adjusted price indexes for different types of producers' durable equipment. Building on Gordon's work, Hulten (1992), Greenwood et al. (1997), and Cummins and Violante (2002) have derived aggregate time series for capital-embodied technical change in the U.S. economy.[18] They use the property just described: that the constant-quality price of investment

---

[15] Jorgenson (2004), in this handbook, labels this methodology, where *relative productivity* growth is measured off the decline in relative prices, the "price approach" to growth accounting.

[16] Define investments in consumption units as $i = x/q$. Then, the aggregate resource constraint reads

$$c + i = A_c F(k, l),$$

and the law of motion for capital in efficiency units is $k' = (1 - \delta) k + iq$.

[17] See this Handbook's chapter by Bils and Klenow (2004) on the measurement of quality for an overview of the different approaches.

[18] Hulten's series strictly uses Gordon's data and therefore spans until 1983. Greenwood et al. extend

relative to consumption (precisely, nondurable consumption and services) reveals the extent of productivity improvements. Their main finding is that:

Observation 12 Productivity growth in the sector producing equipment investment has accelerated relative to the rest of the economy since the early to mid-1970s.

The solid line in Figure 2 shows the relative productivity of the equipment investment goods sector, q, for the period 1947-2000, normalized to 1 in the first year. This index grows at an annual rate of about 1.6 percent until 1975 and at an annual rate of 3.6 percent thereafter. In the 1990s, productivity growth embodied in capital has been spectacularly high, reaching an average annual rate just below 5 percent.

FIGURE 2

The measurement of SSP growth through changes in relative prices requires that the price measures used are appropriately adjusted for quality improvements, presenting a problem for the time period studied since, arguably, the IT revolution has caused large improvements in the quality of durable goods and has led to the introduction of a vast range of new items. Therefore, alternative ways of measuring capital-embodied productivity advancements have been proposed. Hobijn (2000) calculates the rate of embodied technical change by calibrating a vintage capital model. His findings are very similar to the price-based approach, both in terms of the average growth rate, and in terms of the timing of the technological acceleration. Bahk and Gort (1993) and Sakellaris and Wilson (2004) use plant-level data to estimate production functions and assess the productivity effects of new investments. They estimate the growth rate of capital-embodied technical change to be between 12 and 18 percent per year, much higher than the rest of the literature.

Gordon's index to 1992 by applying a constant adjustment factor to the National Income and Product Accounts (NIPA) official price index. Cummins and Violante update the series to 2000. Starting with Gordon's quality-adjusted price indexes for a variety of equipment goods from 1947 to 1983, they estimate the quality bias implicit in the NIPA price indexes for that period. Using the official NIPA series, they then extrapolate the quality bias from 1984 to 2000 for each equipment type and aggregate into an index for equipment and structure.

We calculate the rate of SSP change in the consumption goods sector based on the standard Solow approach. It is well known that the U.S. labor income share in GDP has been remarkably stable for the time period considered. We therefore choose a Cobb-Douglas parametric representation of the production function,

$$y = A_c k^{\alpha} l^{1-\alpha}, \tag{8}$$

with labor income share, $1 - \alpha = 0.64$ (Cooley and Prescott 1995). Conditional on observations for real GDP (in terms of consumption goods), the real capital stock, and employment, we can use this expression to solve for the SSP of the consumption sector A_c.[19] The common finding from this computation, as evident from the dashed line in Figure 2, is

Observation 13 Productivity in the sector producing consumption goods (precisely, nondurable and services) shows essentially no growth over the two decades 1975-1995.

The approach of Greenwood et al. (1997) defines aggregate output in terms of consumption goods. This is rather non-standard. The usual approach, especially as applied to the study of SSP, defines aggregate output growth as a revenue-weighted sum of sectoral output growth rates: a Divisia index (see, e.g., Jorgenson 2001, or Oliner and Sichel 2000). For this more standard approach, one can write aggregate TFP growth as the revenue-weighted sum of sectoral TFP growth. While the Divisia-aggregator approach is a definition with some desirable properties, the Greenwood et al. (1997) approach is based on a particular theory and requires certain identifying restrictions concerning the production structure. Hall (1973) shows that in multi-sector models a unique output aggregator, that is, a function that relates some measure of aggregate output to some measure of aggregate input, exists if

---

[19] It is important to adjust the capital and labor input measure for quality change. As pointed out above, quality adjustment of investment is useful so as to capture investment-specific technical change. The capital stock is then calculated as the cumulative sum of past undepreciated constant-quality investment. From our discussion of wage inequality it follows that the labor input needs to be adjusted for two reasons. First, the skill premium has been increasing since the mid-1970s, and thus the productivity of skilled labor, A_s, is increasing faster than the productivity of unskilled labor, A_u. Second, at the same time, the relative supply of skilled labor has been increasing, inducing large changes in the composition of the stock of labor. To account for quality changes, we use the labor input index computed by Ho and Jorgenson (1999). The dotted line in Figure 2 plots this quality index for labor which grows at an average rate of 0.8% per year.

certain separability conditions for the aggregate production possibility frontier are satisfied. The conditions for such an output aggregator to exist are, essentially, the ones imposed by Greenwood et al. (1997).[20] Given the definition of aggregate output in equation (6), SSP for consumption, or A_c, is then sometimes interpreted as neutral, or disembodied, aggregate technological change.

2.2.3 Reconciling the Acceleration in Investment-Specific Productivity Growth with the Slowdown in TFP: General Purpose Technology and Learning

The stagnation of aggregate TFP since the mid-1970s—evident from Figure 2—accounts for the phenomenon often referred to as a "productivity slowdown" in the growth accounting literature.[21] How can we reconcile the acceleration of investment SSP with a slowdown of consumption SSP? One interpretation builds on learning-by-doing (LBD). New investment goods do not attain their full potential as soon as they are introduced, but rather their productivity can stay temporarily below the productivity of older capital that was introduced some time ago. This feature is attributed to learning effects.[22]

These learning effects can be extremely important when the technological change is "drastic." Recent discussions suggest that the advent of microelectronics led to a radical shift in the technological paradigm, i.e., to a new "general purpose technology" (GPT). Bresnahan and Trajtenberg (1995) coined this term to describe certain major innovations that have the potential for pervasive use and application in a wide range of sectors of the economy. David (1990) and Lipsey et al. (1998) cite the microchip as the last example of such innovations that included, in ancient times, writing and printing and, in more recent times, the steam-engine and the electric dynamo.[23] Although it is hard to define the concept satisfactorily,

---

[20] For further details on this issue, see Hornstein and Krusell (2000).

[21] Since non-equipment investment represents more than three-fourths of GDP, the slowdown of consumption SSP change accounts for most of the slowdown of aggregate TFP change.

[22] The literature on learning effects is large. Lucas (1993) discusses the classic example of LBD related to the construction of the liberty ships of World War II. Bahk and Gort (1993) measure substantial LBD effects at the plant level. Irwin and Klenow (1994) present evidence of LBD in the production of semiconductors. Jovanovic and Nyarko (1995) document learning curves in several professions. Huggett and Ospina (2001) find evidence that the effect of a large equipment purchase is initially to reduce plant-level total factor productivity growth.

[23] Gordon (2000) offers a dissenting view on the issue of whether or not information technologies measure up to the great inventions of the past. In his view, the aggregate productivity impact of computers

given available data, we list as a "fact" the dominant view, which maintains that:

Observation 14 Technological change in the past 30 years displays a "general purpose" nature.

Though most of the evidence supporting this statement is anecdotal, there are some bits of hard evidence. Hornstein and Krusell (1996) document that the decline in TFP occurred roughly simultaneously across many developed countries. More recently, Cummins and Violante (2002) construct measures of productivity improvements for 26 different types of equipment goods. Using the sectoral input-output tables, they aggregate these indexes into 62 industry-level measures of equipment-embodied technical change, and document that their growth rate has accelerated by a similar amount in virtually every industry in the 1990s. Jovanovic and Rousseau (2004a) draw an articulated parallel between the diffusion of electricity in the early 20th century and the diffusion of information technologies (IT) eighty years later based on a variety of data. Their evidence supports the view that both episodes marked a drastic discontinuity in the historical process of technological change. Taken together, all these observations suggest that, similar to other past GPTs, IT has affected productivity in a general way over the past three decades.

There are two versions of the argument that IT are responsible for the observed productivity slowdown. According to one, the slowdown is real: when learning-by-doing is important in improving the efficiency of a production technique, abandoning the older, but extensively used technology to embrace a new method of production involves a "step in the dark" that can lead to a temporary slowdown in labor productivity (Hornstein and Krusell, 1996, Greenwood and Yorukoglu, 1997, and Aghion and Howitt, 1998, chapter 8).

An alternative, complementary version maintains that the slowdown is a statistical artifact due to mismeasurement: if the phase of IT adoption coincides with associated investments in organizational or intangible capital, as our Section 5 will suggest, then insofar as these investments are not included in the official statistics, measured TFP growth will first

---

and telecommunications equipment has been fairly small compared to, say, the telegraph, the railroad, or electricity.

underestimate and then overestimate "true" TFP growth (Hall, 2001, and Basu et al. 2003). The reason is that initially, when large investments in organizational capital are made, the "output side" of the mismeasurement is severe. Later, when the economy has built a significant stock of organizational capital, the "input side" of mismeasurement becomes dominant.

This explanation of the TFP slowdown is appealing, but extremely difficult to evaluate quantitatively because of the lack of direct evidence on how organizations learn. Using some theory, Hornstein (1999) argues that one key parameter is the fraction of knowledge that firms can transfer from the old to the new technology but also shows that the model's predictions vary significantly across plausible parameterizations. Atkeson and Kehoe (2002) build an equilibrium model to measure the dynamics of organizational capital during the "electrification of America". They criticize the Bahk and Gort (1993) view that organizational learning is reflected in an increase in the productivity of labor at the plant level: in an equilibrium model where labor is mobile, productivity is equalized across plants. Instead, they argue that when organizations learn, they expand in size. Thus, cross-sectional microdata on the size distribution of plants allows to identify the structural parameters of the stochastic process behind organizational learning.

Finally, Manuelli (2000) argues that, even in absence of learning effects, the anticipation of a future technological shock embodied in capital can result in a transitional phase of slowdown of economic activity. In the period between the announcement and the actual availability of the new technology, the existing firms prefer exercising the option of waiting to invest and the new firms prefer to delay entering. Consequently, output falls temporarily until the arrival of the new technology.[24]

2.2.4 Factor-Specific Productivity Accounting

In order to talk about changes in FSP, one possibility is to generalize the production function in equation (8) by disaggregating the contributions to production of the two labor inputs—skilled (e.g., more educated) and unskilled (e.g., less educated) labor. Suppose the aggregate

---

[24] We refer the reader to Hornstein and Krusell (1996) for a list of alternative explanations of the TFP slowdown that are not based on changes in technology.

labor input, l, is a CES function of skilled and unskilled labor, l_s and l_u, with FSPs A_s and A_u:

$$l = \left[(A_s l_s)^\sigma + (A_u l_u)^\sigma\right]^{1/\sigma},\ \sigma \leq 1. \tag{9}$$

Relative wage data can then be employed to understand the nature and evolution of FSP in the economy. With competitive input markets, the relative wages are a function of the relative FSP and the relative labor supply:

$$\ln\left(\frac{w_s}{w_u}\right) = \sigma \ln\left(\frac{A_s}{A_u}\right) - (1-\sigma)\ln\left(\frac{l_s}{l_u}\right). \tag{10}$$

The elasticity of substitution between skilled and unskilled labor here is $1/(1-\sigma)$. Katz and Murphy (1992) run a simple regression of relative wages on relative input quantities and a time trend to capture growth in the ratio $\frac{A_s}{A_u}$. They measure skilled labor input as total hours supplied to the market by workers with at least a college degree. Their estimate of the substitution elasticity—around 1.4 (or $\sigma = 0.29$)—indicates that a ten-percent increase in the relative supply of skilled labor implies a seven percent decline of the skill premium.[25] The estimated elasticity of substitution between factors, together with the growing skill premium, imply an increase in the relative FSP of skilled labor in excess of 11 percent per year. We conclude that the typical result of similar exercises on U.S. data is that:

Observation 15 Recent technological advancements have been favorable to the most skilled workers in the population. In other words, technical change has been *skill-biased.*

The "acceleration" in the rate of capital-embodied technical change, the "general purpose" nature of the new wave of technologies, and the "skill-biased" attribute of the recent productivity advancements are the three chief features of the new technological environment that seems to have emerged since the early to mid 1970s. The various economic theories that we are about to review in the rest of this chapter are built on various combinations of these features.

---

[25] The estimated input elasticity of about 1.4 is consistent with a large empirical literature on factor substitution that uses a wide array of data sets (time series as well as cross-section) and methods; see, e.g., Hamermesh (1993).

3 Skill-Biased Technical Change: Inside the Black Box

As we have just observed, the pattern of relative quantities of skills measured by education suggests that the behavior of the skill premium, that is, the increase in the wages of highly educated workers relative to those of less educated workers, should be attributed to a skill-biased labor demand shift, or to "skill-biased technical change." In the absence of a factor-bias in technological progress, the upward trend in the supply of skills documented in Figure 1 (top panel) would have reduced the skill premium.

Katz and Murphy (1992) were the first to use a production framework with limited substitution between skilled and unskilled labor to recover changes in relative FSP from changes in the skill premium. One should note a substantial drawback of the pure skill-biased technical change hypothesis: it is based on *unobservables* (relative FSP changes) that are measured residually from equation (10), so very much like TFP, it is a "black box". In this section we review the attempts to give some specific economic content to the notion of skill-biased technical change.

We start from the capital-skill complementarity conjecture advanced originally by Krusell et al. (2000). Next, we analyze models based on the Nelson-Phelps hypothesis: the adoption phase of a new technology requires skilled and educated workers. If one allows for an important role of FSP changes, then it is paramount to understand what economic forces induce these changes endogenously. In this context, we review the theory of "directed technical change" associated mainly to the work by Acemoglu (2002b, 2003b): exogenous spurts in the relative supply of skilled labor can induce the introduction of skill-biased technological advancements by affecting the incentives of the innovators.

3.1 Capital-Skill Complementarity

Krusell et al. (2000; KORV henceforth) argue that the dynamics of SSP that induced the substantial drop in the relative price of equipment capital is the force behind the rise in the skill premium. The decline in the price of equipment due to productivity improvements, especially those embodied in information and communication technologies, led to an increased use of equipment capital in production. KORV observe that, at least since Griliches (1969), various empirical papers support the idea that skilled labor is relatively more complementary to equipment capital than is unskilled labor. As a result, the higher capital stock increased skilled wages relatively more than unskilled wages. Consequently, the skill premium increased.

Thus, the key elements in KORV's analysis are: 1) separating the effect of equipment capital from that of other capital, mainly structures, 2) allowing equipment to have different degrees of complementarity with skilled and unskilled labor, 3) measuring the efficiency units of capital, especially the new technologies, correctly.[26]

KORV capture the differential complementarity between capital and skilled and unskilled labor using the following nested CES production function of four inputs: structures, k_s; equipment, k_e; skilled labor, l_s; and unskilled labor, l_u:

$$y = k_s^{\alpha} \left[\lambda \left[\mu \left(A_{k_e} k_e \right)^{\rho} + (1-\mu) \left(A_s l_s \right)^{\rho} \right]^{\sigma/\rho} + (1-\lambda) \left(A_u l_u \right)^{\sigma} \right]^{\frac{1-\alpha}{\sigma}}, \tag{11}$$

with $\rho, \sigma \leq 1$. Profit-maximizing behavior of a price-taking firm implies that the skill premium can be approximately written as

$$\ln\left(\frac{w_s}{w_u}\right) \simeq \sigma \ln\left(\frac{A_s}{A_u}\right) - (1-\sigma)\ln\left(\frac{l_s}{l_u}\right) + \lambda \frac{\sigma-\rho}{\rho} \ln\left(\frac{k_e}{l_s}\right)^{\rho}. \tag{12}$$

KORV estimate $\sigma = 0.4$ and $\rho = -0.5$, and thus that the skill premium increases with the stock of equipment capital.[27] They find that the relative productivity of skilled labor grows

---

[26] The quality-adjusted equipment capital stock is again based on the work of Gordon (1990) and subsequent updates, especially for IT.

[27] With this nested CES in 3 factors (equipment, skilled and unskilled labor) it is unclear how to define capital-skill complementarity. One possible, but slightly unorthodox, definition is that the skill premium rises with the stock of equipment. A more traditional definition involves comparing the Allen elasticities of substitution. The elasticity of substitution between equipment and unskilled labor is $1/(1-\sigma)$, while the elasticity of substitution between equipment and skilled labor is

$$\frac{1}{1-\sigma} + \left(\omega_e + \omega_s\right)^{-1} \left[\frac{1}{1-\rho} - \frac{1}{1-\sigma} \right],$$

where ω_e and ω_s are, respectively, the income shares of equipment and skilled labor. Thus, according to both definitions, the parameter estimates in KORV imply that equipment capital is more complementary with skilled labor compared to unskilled labor. See Ruiz-Arranz (2002) for a discussion of the various definitions of elasticity of substitution in production function with more than 2 inputs. Interestingly, Ruiz Arranz divides equipment into finer categories and finds that IT capital (defined as computers, communication equipment and software) is the subgroup with the largest degree of capital-skill complementarity.

at a modest 3 percent per year, a much more plausible number than the one estimated by Katz and Murphy (1992). Overall, KORV show that with their estimated parameters, the relative wage movements in the data can be quite closely tracked. This includes the decline of the wage premium in the 1970s, attributable to an acceleration in the growth of college enrollment due to the Vietnam war draft and the entry of the baby-boom cohorts.[28]

From equation (12) it follows that the skill premium can increase, even if the relative productivity of skilled labor remains constant and the relative supply of skilled labor increases, provided the equipment-skilled labor ratio trends upward fast enough. From this perspective, the results of KORV complement Katz and Murphy's (1992) work: when capital and skills are complementary in production, capital accumulation can explain a large fraction of the residual trend in skill-biased productivity growth.[29]

3.1.1 Further Applications of the Capital-Skill Complementarity Hypothesis

In KORV, the production structure is "centralized" through an aggregate production function. Jovanovic (1998) models an economy with vintage capital where production is decentralized into machine-worker pairs. Newer machines are more productive than older machines, and workers differ in their innate skill level. The pair's output is a multiplicative function of these two inputs. Jovanovic assumes perfect information (no coordination frictions), and hence the labor market equilibrium assignment displays "positive sorting" between skills and machines' productivity (Becker, 1973), i.e., capital-skill complementarity emerges endogenously.[30] An acceleration in the growth rate of technology embodied in

---

[28] Lee and Wolpin (2004) find evidence of capital-skill complementarity both in the goods-producing industries and services in the U.S. economy, and argue that it is an important ingredient to explain the pattern of relative wages and relative labor inputs across the two sectors, over the past 50 years.

[29] Acemoglu (2002a) argues that if the capital-skill complementarity hypothesis is valid, then in equation (10) the relative price of equipment should proxy the shift in the demand for skills and perform better than a linear time trend. However, he finds the trend is always more significant. First, as equation (12) shows, the right variable to add to the Katz-Murphy equation is not the relative price of equipment, but the equipment-skill labor ratio. Second, even using this latter variable one would be bound to find that the linear time trend is more significant because in an OLS regression the estimated coefficient on the time trend converges to its true value at a faster rate than the coefficient on the equipment-skill ratio. More importantly, the key insight of KORV is to give an economic content to the "skill-biased technical change" view, by replacing an unobservable trend with an observable variable.

[30] Holmes and Mitchell (2004) start from a more primitive level where production combines tasks of various complexity and the production factors can perform tasks at a given setup-cost per task. They show that

machines, that is, an increase of the relative productivity differences across vintages, has two effects: 1) for a given age range of machines, it widens the underlying distribution of job productivity differences, and since in equilibrium high-skilled workers are assigned to high-productivity machines, it magnifies the skill premium; 2) the faster rate of obsolescence shortens the optimal life of capital, that is, the range of operative vintages narrows, which tends to weaken inequality since the least productive workers are now matched with better machines. As we will see, these two counteracting forces will survive in the frictional economies of section 7, in spite of the different nature of the equilibrium assignment of workers to machines.

The capital-skill complementarity hypothesis has proved to be helpful to interpret the dynamics of the skill premium in other countries. Perhaps the most interesting example is Sweden. Lindquist (2002) documents that the facts to be explained in Sweden are qualitatively similar to the U.S. facts: between 1983 and 1999 the college premium rose by over 20% and the supply of skilled workers increased substantially. Sweden represents an especially interesting test case for the KORV model because Swedish labor market institutions are commonly believed to play a crucial role in wage setting, arguably making market forces less critical in determining relative wage movements. The main result of Lindquist (2002) is that capital-skill complementarity explains close to half of the dynamics of the skill premium.[31]

How can one reconcile the traditional strength of labor market institutions, such as unions and collective bargaining, in the Swedish labor market with the finding that market forces account for a large part of relative wage dynamics? One possibility is that institutions set the aggregate share of income going to labor in any given period—possibly extracting rents from firms. The distribution of these rents among workers is then determined by their individual outside options, which differ across skill levels and are affected by technical change. In section 8 we develop further this conjecture in the context of the decline in union membership in the United States, but the economic linkages between the dynamics of

---

under reasonable primitive assumptions on setup costs for capital, skilled labor and unskilled labor, the former two inputs display a form of complementarity.

[31] Lindquist uses the KORV specification for aggregate technology in equation (11) and estimates $\rho = -0.92$ and $\sigma = 0.31$.

institutions and technological progress are far from being well understood.

More international evidence in favor of the capital-skill complementarity model is offered by Flug and Hercowitz (2000). They estimate a strong effect of equipment investment on relative wages and employment of skilled labor using a panel data set for a wide range of countries around the world.

Recently, the capital-skill complementarity idea has been imported into the study of inequality at the business-cycle frequency. The skill premium is found to be close to *acyclical* in the United States: its contemporaneous correlation with output is positive, but not statistically different from zero. Lindquist (2004) argues that, since unskilled labor is relatively more pro-cyclical than skilled labor, a Cobb-Douglas production function in three inputs (capital, skilled labor, and unskilled labor) would predict a strongly pro-cyclical skill premium. Inspection of equation (12) suggests that introducing capital-skill complementarity in production counteracts this pro-cyclicality of the skill premium since, at impact, skilled hours respond more than the stock of equipment: the capital-skill complementarity effect is countercyclical and offsets the change in relative supply.[32]

In sum, the studies discussed in this section indicate that capital-skill complementarity is a quantitatively important ingredient in competitive theories of relative wage determination, within centralized as well as decentralized production structures and at high as well as low frequencies.

3.2 Innate Skills and the Nelson-Phelps Hypothesis

Nelson and Phelps (1966) argued that the wage premium for more skilled workers is not just the result of their having higher “static productivity”. Workers endowed with more skills, they contended, tend to deal better with technological change in the sense that their productivity is less adversely affected by the turmoil created by technological transformations of the workplace, or in that it is less costly for them to acquire the additional skills needed to use a new technology. Greenwood and Yorukoglu (1997) cite sources reporting that the skill

---

[32] Within a similar framework, Cohen-Pirani and Castro (2004) argue that capital-skill complementarity is important for understanding why the volatility of skilled labor (relative to GDP) has tripled after 1985.

premium also rose during the course of the first industrial revolution. In the context of the recent "IT revolution", Bartel and Lichtenberg (1987) provide evidence that more educated individuals have a comparative advantage at implementing the new technologies and Bartel and Sicherman (1998) argue that high-skilled workers sort themselves into industries with higher rates of technical change.

The theory has been formalized in various formats. Lloyd-Ellis (1999) embeds a race between the innovation rate and the "technological absorption rate" of workers (the maximum numbers of innovations that can be adopted per unit of time) in a general equilibrium model: at times when the innovation rate exceeds the absorption rate, wage inequality increases due to the fierce competition for scarce, adaptable labor. Galor and Moav (2000) formalize this hypothesis differently and assume that technological change depreciates the human capital of the unskilled workers faster than that of skilled workers (the "erosion effect"). Krueger and Kumar (2004) distinguish between workers with general education and workers with vocational skill-specific education and postulate that only the former type remains productive when new technologies are incorporated into production.

It is important to remark that this hypothesis, in all its versions, applies to educational skills as well as dimensions of skills that are not necessarily observable or correlated with education. Hence, it can potentially account for the rise in within-group (or residual) inequality. Ingram and Neumann (1999) offer some evidence on the increase in the return to certain categories of skills not fully captured by education. They match individual data on wages and occupations from the CPS with the skill content of several occupations, obtained from the *Dictionary of Occupational Titles* (DOT). DOT data contain information on how much each occupation requires of each of a wide range of skills such as verbal aptitude, reasoning development, numerical ability, motor coordination, and so on. Using factor analysis they group over 50 types of skills into four factors (intelligence, clerical skills, motor skills, and physical strength) and estimate that the return to "intelligence" has almost doubled from 1971 to 1998. Moreover, adding the quantity of this factor to a standard Mincerian wage regression weakens the implied increase in the returns to college education significantly.

The idea that the diffusion of IT may have raised the demand for adaptable skilled workers—thus, even within educational groups—has been formalized in various ways by Galor and Tsiddon (1997), Greenwood and Yorukoglu (1997), Caselli (1999), Galor and Moav (2000), and Aghion et al. (2002).

To illustrate the basic mechanism of such a theory, consider an economy where workers differ in their cost of learning the new technology.[33] Suppose that this economy starts in a steady-state equilibrium where production uses the "old" technology, $y_1 = A_1 k_1^{\alpha} l_1^{1-\alpha}$. The labor market is competitive; thus, in steady state, all workers are employed in the old sector and there is no wage inequality.

Suppose a new technology becomes available and the sector using this new technology can produce output with $y_0 = A_0 k_0^{\alpha} l_0^{1-\alpha}$ where $A_0 > A_1$. Because of the learning cost, labor is not perfectly mobile, and wages in the two sectors may differ. Capital, however, is free to move towards its more productive use, and factor-price equalization for capital yields

$$R_0 = R_1 \Rightarrow \frac{l_1}{k_1} = \left(\frac{A_0}{A_1}\right)^{\frac{1}{1-\alpha}} \frac{l_0}{k_0}. \tag{13}$$

It is straightforward to show that

$$w_0 = \left(\frac{A_0}{A_1}\right)^{\frac{1}{1-\alpha}} w_1 > w_1.$$

Therefore, in equilibrium, a premium emerges for those workers with low learning cost (i.e., high ability) who can adapt quickly and move to the new sector.

The skill premium increases due to two effects. With full mobility of labor, inequality would disappear. With no labor mobility and no capital mobility, the skill premium would reflect the productivity difference A_0/A_1. In this class of models, labor mobility is limited by the distribution of ability in the economy, but capital moves freely. Full mobility of capital induces a general equilibrium feedback that amplifies inequality: factor-price equalization requires capital to flow to the sector operating the new technology to equate marginal productivities of capital.[34] Thus, workers on the new technologies are endowed with more capital, which boosts their relative wages further.

---

[33] In the rest of this section, the exposition will be based mainly on the environment in Caselli (1999).

[34] One implication of this mechanism, evident from equation (13), is that a technological revolution should

In its typical version, the Nelson-Phelps hypothesis implies that the rise in the skill premium will be transitory: it is only in the early adoption phase of a new technology that those who adapt more quickly can reap some benefits. Over time there will be enough workers who learn how to work with the new technology to offset the wage differential. Note the difference with the KORV hypothesis, where the effect on the skill premium is permanent. Are new technologies and skills complementary in the whole production process or just in the adoption phase? Chun (2003) uses industry-level data for the U.S. to disentangle the impact of "adoption" and "use" of IT. He finds that the increase in the relative demand of educated workers from 1970 to 1996 in the U.S. is related significantly to both factors, but quantitatively the impact of use is twice as large.

3.2.1 Further Applications of the Nelson-Phelps Hypothesis

Aghion (2002) and Borghans and ter Weel (2003) emphasize that the Nelson-Phelps approach can explain why, in the 1970s, the college premium declined at the same time that the wage dispersion within college graduates increased. The idea is that in the early phase of IT diffusion in the 1970s, only educated workers with high ability adopted IT. Naturally, this higher return to ability increases within-group inequality. The contemporaneous acceleration in the growth of the supply of educated labor, due to exogenous factors, explains the relative fall in the average wage of college graduates.

Beaudry and Green (2003) compare the United States and Germany, highlighting an apparently puzzling feature of the data: the relative supply of skilled labor in the United States grew faster than in Germany, and yet the skill premium rose in the United States, but not in Germany. They outline a model that combines elements of Caselli (1999) and Krusell et al. (2000). Consider an economy where there are two technologies in operation and the "new" technology displays more capital-skill complementarity than the old one. An exogenous rise in the supply of skills increases the relative return to capital in the new sector. Capital then flows from the old to the new sector, and, ultimately, this higher capital

---

trigger a surge in the real rate of return on capital by a factor $(A_0/A_1)^{\frac{1}{1-\alpha}}$. Yearly U.S. long-term real interest rates were roughly 3 percent higher in the period 1980-1995 compared to the period 1965-1980. It is unclear whether this magnitude is quantitatively consistent with the observed increase in wage dispersion.

intensity can raise the relative wage of skilled labor if labor is not perfectly mobile because, as in Caselli's model, only skilled workers can quickly adapt. Thus, in the long-run, the country with the initial spur in the supply of skilled labor (the United States) finds itself with a larger skill premium.

In their original paper, Nelson and Phelps (1966) developed the concept of "technological gap", defined as the percentage difference between the technology operated by the typical machine in the economy and the one embodied in the leading-edge machine. They conjectured that a rise in the technological gap should be associated with a large skill premium because of the surge in the demand for educated workers needed to adopt the new, more productive technologies. Cummins and Violante (2002) use data on the quality-adjusted relative prices and quantities of equipment investment to construct a measure of the technological gap for the U.S. economy.[35] Figure 3 shows that the technological gap and the skill premium have moved largely in tune over the past half century, confirming—at least in the time-series dimension—even the most literal version of the Nelson-Phelps hypothesis.

FIGURE 3

Put differently, the size of the technological gap can be thought of a proxy for shifts in the relative demand of skilled workers.

3.3 Endogenous Skill-Biased Technical Change

In the literature we discussed so far, the sector bias and the factor bias of technical change were assumed to be exogenous. Over the past 20 years a substantial body of work in the field of growth theory has formalized the idea that the efforts of innovators are endogenous and

---

[35] Precisely, if q_t is the level of productivity embodied in the new investment at time t, then the average unit of productive capital in the economy at time t embodies a technology with productivity Q_t, defined as

$$Q_t = \sum_{j=0}^{\infty} (1-\delta)^j \, q_{t-j} \frac{i_{t-j}}{k_t},$$

where δ is the depreciation rate, i denotes investments and k the capital stock, both expressed in units of consumption. In other words, Q_t is the ratio between capital stock correctly measured in efficiency units (the numerator) and capital stock k not adjusted for quality. Then, the gap is defined as $\frac{q_t - Q_t}{Q_t}$.

respond to market incentives. The models belonging to the so-called "new growth theory" describe the endogenous determination of the *level* of innovative activity.

Recently, Acemoglu (1998, 2002b, 2003b) and Kiley (1999) have developed the idea that the composition, or *direction*, of innovations is also endogenous: if R&D activity can be purposefully directed towards productivity improvements of different inputs (capital, skilled labor, and unskilled labor), then it will be biased towards the factor that ensures the largest returns.

An important ingredient of this approach is that the returns to R&D targeted toward a given input are proportional to the total supply of that input, since "productivity" and "quantity" are complements in production. This creates a "market size" effect of R&D: productivity-improving resources are allocated to factor markets with large relative factor supplies.[36]

It is useful to see how this mechanism works within a simple model that represents a reduced form of the richer environments offered by Acemoglu and Kiley. Consider an economy with a given endowment of skilled and unskilled labor, l_s and l_u, and a production function (9) as in Section 2.2.4. Conditional on the FSPs, A_s and A_u, wages and employment are determined competitively, and the competitive equilibrium is Pareto-optimal. Now suppose that the Social Planner wants to maximize production subject to a given frontier of technological possibilities, that is, choices of A_s and A_u:

$$\max_{\{A_s,A_u\}} \left[(A_s l_s)^{\sigma} + (A_u l_u)^{\sigma}\right]^{1/\sigma}$$

$$s.t. \qquad \left[\lambda A_s^{\phi} + (1-\lambda)\, A_u^{\phi}\right]^{1/\phi} = 1.$$

Assume that the technological frontier is convex, that is, that $\phi > 1$. Rearranging the first-order conditions, one arrives at

$$\frac{A_s}{A_u} = \frac{\lambda}{1-\lambda}\left(\frac{l_s}{l_u}\right)^{\sigma/(\phi-\sigma)}, \tag{14}$$

which describes the optimal choice of skill-bias given the relative factor supply. The above equation shows that when skilled and unskilled labor are substitutes, $0 < \sigma \leq 1$, the skill

---

[36] It is useful here to draw a parallel with certain traditional endogenous growth models, where the scale effect determines the level of the growth rate. See, Jones (2004) for a survey of the models with scale effects.

bias is increasing in the relative supply of skills. This latter parametric condition implies that the marginal product of each innovation is increasing in its corresponding factor.

A surge in the relative endowment of skilled labor, like the one witnessed by the U.S. economy in the postwar period, induces the adoption of more skill-biased technologies in production. This force tends to counteract the direct relative supply effect on wage inequality. Can the endogenous skill bias be so strong in the long run as to overturn the initial supply effect?

To answer this question, we substitute the expression for the skill bias, (14), into the expression for the skill premium, (10), and obtain

$$\ln\left(\frac{w_s}{w_u}\right) \propto \frac{\sigma - \phi(1-\sigma)}{\phi - \sigma} \ln\left(\frac{l_s}{l_u}\right).$$

We see that the skill premium is increasing in the relative supply of skilled labor as long as $\phi \in \left(\sigma, \frac{\sigma}{1-\sigma}\right)$. Thus, theoretically, it is possible to explain a positive long-run relationship between the relative supply of skilled labor and the skill premium as the one depicted in Figure 1 (top panel).

One limitation of existing models of directed technical change, and also of most of the literature surveyed in this section, is that arguments for the skill-premium focus on the response of a relative price to exogenous changes in relative factor supplies. Whereas one can reasonably assume that "ability" is largely pre-determined with respect to the point in the life-cycle when agents start making economic decisions, education is not. One would expect that changes in returns to education as large as the ones we observed in the past 30 years would significantly affect the incentives to acquire education. However, it is an open question to what extent the observed changes in returns were predicted by the cohorts affected by these returns when they made their education decisions.[37]

Models of directed technical change, augmented by an endogenous supply of skills, can give rise to multiple steady states. If the innovators expect the supply of educated workers to rise, they will invest in skill-biased R&D which, in turn, will augment the returns to college and induce households to acquire human capital, fulfilling the innovators' expectations.

---

[37] See Abraham (2003) for a related analysis.

3.3.1 Sources of the Skill-Bias in Recent Times

Equation (14) shows that we can view the rise in the relative supply of educated workers as the driving force behind the recent skill-biased technical change. The change in relative supply was, according to Acemoglu (2002a) largely exogenous, at least initially, and a result of the high college enrollment rates of the baby-boom cohort and of the Vietnam war draft. The crucial issue, still unresolved, is whether the necessary parametric restrictions discussed earlier are plausible, and whether the initial shock is large enough.[38] What other changes in the economic environment can be listed as potential sources of skill-biased innovations?

First, there are possible interactions between capital-skill complementarity and the direction of technical change. Hornstein and Krusell (2003) have taken a first step at incorporating the idea of factor-biased innovations into the KORV explanation of the skill premium. Intuitively, an acceleration in capital accumulation due, for example, to an exogenous fall in the price of capital increases the returns to skill-biased innovations if capital is more complementary with skilled labor. Hence, capital-embodied productivity improvements can be the source of factor-biased technical progress. For a calibrated version of their model, Hornstein and Krusell find that a persistent decline of the relative price of capital results in a temporary, but very persistent, increase of the skill premium. In their model the skill premium not only increases because of capital accumulation (as in KORV), but it also increases because of the endogenously induced spur of skill-biased technical change.

Second, the increased openness to trade can play a role. Using a Schumpeterian growth model, Dinopolous and Segerstrom (1999) argue that if trade liberalization boosts the profitability for monopolistic suppliers by increasing the size of their markets, then resources shift from manufacturing to R&D activities. If, in turn, R&D is a skill-intensive sector, the skill premium rises. This model determines endogenously the level of R&D, but does not display endogenous factor bias in the equilibrium innovation rate. In Acemoglu (2003c), the direction of technical change is related to international trade. A natural assumption about

---

[38] In the richer model developed by Acemoglu (2002b), this parametric restriction requires that the elasticity of substitution between skilled and unskilled labor be larger than 2. Most of the empirical literature on factor substitutability, however, points at values around 1.5 (Hamermesh 1993).

factor endowments is that in the United States the ratio of skilled to unskilled labor is higher than in the rest of the world. After the U.S. economy opens to trade, the world prices are determined by the aggregate relative factor endowment, and thus skill-intensive goods become relatively more expensive. In the class of models with an endogenous factor bias, factors which produce goods with the highest relative price—and the highest expected profits—will be the target of the largest amount of innovative activity (the "relative price" effect). Thus trade opening induces skill-biased technical change. This mechanism can, under some conditions, explain also the increase in the skill premium in less-developed countries documented, for example, by Robbins (1996).

Third, Cozzi and Impullitti (2004) argue that government policy may also have contributed to the bias in technical change. In the 1980s, U.S. technology policy rapidly shifted its priority from security and defense to economic competitiveness in order to counteract the emerging dominance of Japan in the sectors producing high-tech goods.[39] Within a Schumpeterian growth model, they show that when the government reallocates its expenditures towards the (high-tech) manufacturing goods with the highest potential quality improvement, it creates a market-size effect that can lead to a rise in the innovation rate in those sectors and a net increase in the demand for skilled R&D workers and their wages.[40]

Although we have learned from the above analyses about possible channels influencing the skill premium, there is little work that allows us to quantify each of the channels. A careful calibration and evaluation of a model which incorporates these various channels would be an important first step in this direction.

---

[39] Japan's share of the high-tech goods markets rose from 7 percent to 16 percent during the period 1970-1990, while at the same time the U.S. share declined from 30 percent to 21 percent. In 1963 government spending on defense represented 1.37 percent of GDP. In 1980, it was down to 0.57 percent.

[40] Like the Dinopolous and Segerstrom (1999) model, strictly speaking, this is not a model of directed-technical change, since skilled labor works only in the R&D sector, and each manufacturing sector employs unskilled labor. However, in a version of the model with endogenous factor-bias and a structure of manufacturing where high-tech goods are produced by skilled labor and low-tech goods by unskilled labor, the shift in technology policy would have the same qualitative effect.

3.4 A Historical Perspective on the Skill Premium

In Section 2 we have observed that, over the last 100 years, wage inequality first declined and then increased, with the turning point somewhere around 1950. Can the theoretical models developed to interpret the increasing wage inequality for the second half of the 20th century also account for the declining wage inequality of the first half of the 20th century?

3.4.1 Capital-Skill Complementarity

Figure 4 plots the relative price of equipment together with the returns to one year of education (both college and high school) since 1929.[41] The pattern is rather striking and is broadly consistent with an explanation based on the capital-skill complementarity hypothesis. During the first half of the century, the price of capital increased which slowed the demand for educated labor and the skill premium. Then around mid-century it started to decline, fostering a strong demand shift in favor of educated labor.

FIGURE 4

This extension of the KORV analysis to the whole 20th century is yet to be performed formally.[42] Thus, before one fully subscribes to this explanation, it is worth discussing the key assumption behind the model. Is it an accurate historical assessment that the introduction of new capital goods has systematically increased the productivity of skilled labor relative to the productivity of unskilled labor? In other words, when can one date the birth of work organizations displaying capital-skill complementarity?

According to Goldin and Katz (1998), until the early 20th century there was no trace of skill-biased technical change; rather, the opposite bias was at work. The origins of capital-

---

[41] The relative price is computed from series available on the BEA website. In particular, compared to the series discussed previously in the chapter, there are no quality adjustments. As a result, the acceleration which occurred since the mid 1970s is less evident here. The series on the return to education for 1939, 1949, 1959, 1969, 1979, 1989, and 1995 are taken from Table 7 in Goldin and Katz (1999) and interpolated linearly for the missing years in between. The first datapoint for 1929 is obtained by linear interpolation from 1914.

[42] Admittedly, the evidence in Figure 4 is rather indirect. Looking directly at the stock of equipment (unadjusted for quality improvements), its average annual growth rate in the periods 1930-1950, 1950-1980, 1980-2000 is, respectively, 2.2%, 5.0% and 4.2%. However, when quality-adjusted, the growth rate of equipment from 1980-2000 is close to 8% (Cummins and Violante, 2002). See also Hornstein (2004) for a discussion of historical trends of U.S. capital-output ratios.

skill complementarity are associated with the introduction of electric motors and a shift away from assembly lines and toward continuous and batch processes. This development started in the second and third decades of the 20th century. In particular, the declining relative price of electricity and the consequent electrification of factories made it possible to run equipment at a higher speed. This, in turn, increased the demand for skilled workers for maintenance purposes. Since then, the introduction of new equipment, such as numerically controlled machines, robotized assembly lines, and finally computers, further increased the relative productivity of skilled labor. Thus, we conclude that based on anecdotal evidence, the period portrayed in Figure 4 is one where capital-skill complementarity became more important.

Mitchell (2001), in a related interpretation on the last century of data, emphasizes the technological aspects of optimal plant size. Mitchell documents a striking similarity between the historical path of wage inequality and the pattern of average plant size in manufacturing which rose over the 1900-1950 period and shrunk between 1950 and 2000, thus almost producing the mirror image of inequality at low frequencies. The time-path of plant size can be interpreted as an indicator of the magnitude of the fixed costs of capital and fits well with the evidence of Figure 4.

In Mitchell's model, production requires performing a large set of tasks with capital and two types of labor, skilled and unskilled. Entrepreneurs face a fixed cost to operate capital, skilled labor, and unskilled labor. Unskilled labor has a higher fixed cost and a lower variable cost than does skilled labor; e.g., unskilled labor is specialized and needs a certain amount of training to perform all the tasks, whereas skilled labor is naturally able to multi-task.[43]

The move from craft shops to assembly lines (1900-1950) induced a rise in the fixed cost: the optimal size of the plant rose and with a larger size, plants optimally employed more unskilled workers with large fixed cost, but low variable cost (wages). The demand for

---

[43] This idea is further developed in Holmes and Mitchell (2004). This paper develops a theory of the intrinsic difference between three key factors of production: capital, unskilled labor, and skilled labor. Based on this theory, the authors develop implications for: 1) how capital and skill intensity vary as a function of size of plants, 2) the micro-foundations of capital-skill complementarity, 3) the effect of trade on the skill premium and the historical relationship between the plant size-skill correlation and the skill-premium.

unskilled workers rose, weakening the skill premium. As an illustration of the importance of fixed costs for this type of production method, recall that all Ford plants had to be closed and redesigned when the "Model T" was discontinued (Milgrom and Roberts, 1990).

The shift toward more flexible, numerically controlled machines and IT capital (1950-2000) led firms to adopt a smaller scale of production and employ more highly skilled workers whose low fixed cost makes them preferable to unskilled workers in small plants. The increased demand for skilled labor thus raised the skill premium. Based on a calibration exercise, the model can account for two thirds of the movements in the skill premium.[44]

3.4.2 Directed Technical Change

The theory of directed technical change maintains that a growth in the relative supply of a factor of production should induce technical change biased in favor of that factor. Historically, there are two important episodes of largely "exogenous" spurs in relative factor supply.

First, there was an increase in the supply of unskilled labor in urban areas of England during the 19th century. A careful look at the nature of technological progress over this period supports the theory. Goldin and Katz (1998) argue that in the 19th century the wave of technological innovations substituted physical capital and raw labor for skilled artisan workers (Braverman, 1974 and Cain and Paterson, 1986). For example, automobile production began in artisanal shops where the car was assembled from start to finish by a small group of "all-around mechanics." Only a few decades later, the Tayloristic model of manufacturing would bring together scores of unskilled workers in large-scale plants to assemble completely standardized parts in a fixed sequence of steps for mass production.

Second, there was a surge in skilled labor (i.e., workers with literacy and numerical skills) due to the "high-school movement" of 1910-1940. As pointed out by Aghion (2002), with respect to this episode, the theory finds weaker support. On the one hand, as we discussed earlier, it appears that the first part of the 20th century indeed marked the beginning of

---

[44] Note that this model implies that the origin of capital-skill complementarity is to be located only around 1950, later than what was argued in Goldin and Katz (1998).

a transformation in production methods biased towards skilled labor (from assembly lines to continuous and batch production processes). On the other hand, there was a decline in the returns to high school and the returns to college were stable (see Figure 4). Why is it that this wave of skill-biased technical change, which was as strong as the one 50 years later, did not have a similar impact on the wage structure? This question remains unanswered to date.[45]

3.5 Technology and the Gender Gap

Here we explore briefly the interaction between the gender gap and the advancements of technological change, both in the market and in the household.

3.5.1 Technological Change in the Market

As evident from Figure 1 (bottom panel), since the mid-1970s the gender wage gap has closed substantially. Several studies have concluded that this is due to a rise in relative labor demand for women, as supply cannot have played a large role (Bertola et al. 1997). Was the recent technological revolution "gender-biased"?

Consider a simple model where jobs differ in their requirement of physical effort and all jobs are necessary for production of the final good. At the same time, men and women have two traits: physical ability and cognitive ability. The theory of comparative advantage then implies that men will be most efficiently assigned to jobs with high physical requirements and that women should work on jobs with a large fraction of cognitive tasks.

The arrival of a new technology, like computers, that increases productivity relatively more on jobs with high cognitive content therefore tends to raise the average wage of women more than it raises the average wage of men. Weinberg (2003a) tests this theory on microeconomic data for the United States and finds that the increase in computer use for women can explain up to 50 percent of the increase in the relative demand for female employment.

It is worth noting that the gender premium fell in spite of the fact that the female-male

---

[45] Institutions might have played a role in the 1940s. Goldin and Margo (1992) argue that the National War Labor Board operated an explicit policy of wage compression during that period.

relative supply ratio grew almost by a factor of 2 between 1960 and 2000, i.e., by as much as the growth in the relative supply of college-educated labor. In the perspective of the directed technical change literature, one is left to ponder whether rising female participation was also a force that led innovators to spend resources on capital goods complementary with cognitive skills rather than with physical skills in order to exploit women's comparative advantage. This hypothesis remains to be analyzed in detail.

3.5.2 Technological Change in the Household

The postwar period witnessed another form of technological revolution: one that did not take place in factories and plants, but rather in the household. Greenwood et al. (2004) argue that the decline in prices of household appliances (refrigerators, vacuum cleaners, washers, dishwashers, etc.) worked as "engines of liberations" for women: new and more productive capital in the households could free up potential hours to be supplied in the labor market. In particular, as household durables were introduced into the economy, the effective wage-elasticity of female labor supply increases, which, in turn, helps explaining the sharp rise in female market participation, even in the presence of not-so-large changes in the gender wage gap.[46]

4 Technical Change and the Returns to Experience

According to Card and DiNardo (2002), one of the most important challenges to the hypothesis that the recent changes in the wage structure are linked to technological progress is to explain the combination of the rise in the returns to labor market experience for the low-educated workers in the population and the flat, or declining, pattern of the experience premium for college graduates.

It turns out that the existing theoretical literature does not provide a unified answer to the question of how technological change affects the experience premium. Examples of the literature we review in this section include job-specific or technology-specific experience that,

---

[46] We refer the reader to Greenwood and Seshadri (2004) in this Handbook for a detailed analysis of this channel.

in principle, may be adversely affected by technological change, but that may also benefit from technological change if that change is of a 'general purpose' variety, that is, if it makes experience more widely applicable.[47] We also look at general labor-market experience as a vehicle to lower the cost of adapting to technological change.

4.1 Experience with General Purpose Technologies

An important feature of the recent technological developments that has not received much attention in the literature on inequality is its *general purpose nature*. Aghion et al. (2002) formalize the idea of "generality" of a technology and build a theoretical framework to understand how it affects various dimensions of wage inequality, such as the experience premium. They model generality in relation to human capital: a more general technology allows a larger degree of *transferability* of sector-specific experience across the different sectors of the economy. For example, the ability to use computers for word-processing or programming is useful in numerous sectors and jobs in the economy.[48] Given that actual technological change is uneven across sectors, transferability of experience then increases the value of experience, that is, the experience premium.

Consider a simple overlapping-generations (OLG) model with two-period lived agents, and two production sectors indexed by $i = 0, 1$. Each cohort of agents has measure one and works in both periods. Technological progress results in capital-embodied innovations that increase productivity by a factor $1 + \gamma$ occurring in each of the two sectors in alternation. Let "0" denote the new sector in the current period. Suppose, for simplicity, that production takes place with a fixed amount of capital, normalized to one: the production function in sector i (in the stationary transformation of the model) is $y_i = A_i^{\alpha} h_i^{1-\alpha}$, where A_i measures

---

[47] We will return to the issue on how technological change interacts with the accumulation of job/technology-specific knowlege in the frictional models of Section 7.3.

[48] A survey conducted by the U.S. Bureau of Labor Statistics emphasizes that

> "*...the technology, network systems, and software is similar across firms and industries. This is in contrast to technological innovations in the past, which often affected specific occupations and industries (for example, machine tool automation only involved production jobs in manufacturing). Computer technology is versatile and affects many unrelated industries and almost every job category*" (McConnell, 1996, page 5).

the efficiency of capital in sector i ($A_0/A_1 = 1+\gamma$), and h_{it} measures the effective labor input in sector $i = 0, 1$.

Young agents are always productive on the new technology, whereas old workers can productively move to the new sector only with probability σ. This captures the idea that young workers are more "adaptable" than old workers, possibly because of vintage effects in their schooling, or because the ability to learn declines with age. Moreover, assume that this "adaptability constraint" is binding, in the sense that: (1) the equilibrium fraction σ^* of old workers who moves equals σ; and, (2) there is not enough labor mobility (σ is sufficiently low) to offset the impact on wages of the sectoral productivity differential $1+\gamma$.

Newborn agents start working in the new sector with initial knowledge normalized to 1.[49] Agents accumulate η additional units of experience through learning-by-doing in the first period of work. The generality of the technology determines the degree of skill transferability for the old workers, τ_o, i.e., the fraction of accumulated knowledge η a worker can carry over if she moves to the leading-edge sector at the beginning of her second period of life. The entire knowledge η can be used if the worker stays in the old sector.

Aggregate human capital in the old sector h_1 is determined by old, non-adaptable workers, a fraction $1-\sigma$, who have accumulated $1+\eta$ units of experience. Human capital in the new sector is determined by the new cohorts that have one unit of experience, and old adaptable workers with transferable experience, that is, $h_0 = 1 + \sigma(1+\tau_o\eta)$. With competitive labor markets, the ratio between the prices of efficiency units of labor in the old and the new sector therefore is:

$$\frac{w_1}{w_0} = (1+\gamma)^{-\alpha}\left(\frac{h_0}{h_1}\right)^{\alpha} = (1+\gamma)^{-\alpha}\left[\frac{1+\sigma(1+\tau_o\eta)}{(1-\sigma)(1+\eta)}\right]^{\alpha}. \tag{15}$$

The steady-state experience premium, i.e. the average wage of old workers relative to the average wage of young workers, is therefore given by

$$x^* = \sigma(1+\tau_o\eta) + (1-\sigma)(1+\eta)\frac{w_1}{w_0}, \tag{16}$$

where one can see immediately that x^* is increasing in τ_o. That is, an increase in the generality of technological knowledge raises skill transferability and amplifies the experience

---

[49] Aghion et al. (2002) show that this is indeed the optimal choice of young cohorts, for general conditions.

premium of adaptable workers, who are able to transfer more of their cumulated skills. It also indirectly raises the experience premium of non-adaptable old workers by making effective adaptable labor input relatively more abundant in the economy, hence depressing the wage of young workers.

This result is particularly interesting in light of the fact that a version of this model that is based purely on the hypothesis that the rate of embodied technical change, γ, has accelerated would predict a decline in the experience premium. This is evident from the fact that the wage ratio w_1/w_0 is decreasing in γ: larger productivity differentials between the young and the old vintages represent a relative advantage to young workers who are more adaptable.

The more general model in Aghion et al. (2002) also features a flexible choice of capital. Another interpretation of generality of the technology offered in their paper is based on the *compatibility* of physical capital, i.e., the extent to which capital equipment embodying the old technology can be retooled—so as to embody the new leading-edge technology—and moved to the new sector. Under this interpretation, the arrival of a GPT, which increases the compatibility across vintages of capital, reduces the experience premium since it allows the transfer of more capital to the new sector where it benefits the young, inexperienced, but more adaptable workers.[50]

4.2 Vintage-Specificity of Experience

According to the GPT hypothesis, human capital becomes more transferable across sectors once the new technological platform has fully diffused throughout the economy. However, it is also reasonable to think that, at least in the transition phase, certain skills associated to the old way of producing quickly become obsolete. Or, put differently, human capital is vintage-specific. Thus, although in the final steady state skill transferability will be higher, it can undershoot during the transition.

---

[50] The model by Caselli (1999) outlined in section 3.2 has exactly this feature of capital mobility from the old technology to the new and more productive technology; thus, a version of that model where the young workers are those with the lowest learning cost would have the same counterfactual prediction for the experience premium.

To study the implications of vintage human capital for the experience premium, we can slightly modify the two-period OLG model in the previous section. To make this point starkly, consider the extreme case where old workers never find it profitable to move across sectors, so $\sigma^* = 0$, and suppose that when young workers join the new sector they lose a fraction $1-\tau_y$ of their initial knowledge (as before, normalized to 1). Modifying appropriately the equilibrium wage ratio (15), equation (16) for the experience premium becomes

$$x^* = \frac{(\tau_y + \eta)\, w_1}{\tau_y w_0} = \frac{1}{(1+\gamma)^{\alpha}} \left[\frac{\tau_y + \eta}{\tau_y}\right]^{1-\alpha}, \tag{17}$$

which shows that x^* is decreasing in the skill transferability rate for young workers, τ_y. The arrival of a new technology that makes the knowledge of its (young) users obsolete can widen the returns to experience.

In analyzing earlier equation (16) we argued that a rise in γ would depress the experience premium, which is a problem for the pure "acceleration hypothesis". Vintage human capital can overturn this result. Suppose, as in Violante (2002), that the degree of skill transferability is decreasing in the speed of technological improvements, i.e. $\tau_y = (1+\gamma)^{-\tau}$. Then, it is easy to see from (17) that as long as $\tau > \alpha/(1-\alpha)$, the experience premium will rise after a technological acceleration, since the loss of vintage-specific human capital incurred by young workers is larger than the productivity improvement embodied in physical capital.[51] In Section 7.3 we return to the role of vintage human capital and discuss the plausibility of the assumption that the extent to which skills are transferable depends on γ.

4.3 Technology-Experience Complementarity in Adoption

According to the standard technology adoption models, the adopters of the new technology are likely to be the young workers because they face a lower learning cost or a longer time horizon to recoup the adoption costs. Weinberg (2003b) challenges this view and argues that there is one other force that gives more experienced workers an advantage: complementarity between new technologies and skills, together with the fact that more experienced workers are

---

[51] Note that this large skill loss for young workers does not necessarily imply that it is not optimal for them to begin working in the new sector. Indeed, by working with the new technology in the current period they improve *future* skill transferability.

more skilled, should lead to the prediction that older workers will adopt the new technology. What force dominates? And what are the implications for the experience premium?

Weinberg looks at the empirical pattern of computer usage (i.e., adoption of one of the new recent technologies) over the life-cycle and shows that it differs dramatically between high-school graduates and college graduates (see Figure 5).

FIGURE 5

Among uneducated individuals the profile is hump-shaped and peaks around 30 years of experience, while for educated individuals it is downward-sloping. As expected, the adoption rates for college graduates are higher at any given age.

These data suggest that the answer to the first question above depends on the level of schooling: for low-educated workers, experience is a substitute for general education, and the more experienced workers are also more productive in the new technology. Workers with high education levels are all equally adaptable to the new technology, so, for such workers, additional experience has a small marginal return in adoption. Since the learning cost increases with age, the youngest are more likely to adopt the new technology.

Adding to this mechanism the assumption that new technologies are more productive yields that the adopters gain a wage increase, which is consistent with the different pattern of the experience premium for low and high education groups that we described in Section 2.

4.4 On-the-Job Training with Skill-Biased Technological Change

The models reviewed in this section treat the degree of skill transferability or adaptability of workers as exogenous. If old workers recognize that "new knowledge" is necessary for dealing with the transformed technological environment, then one should expect that they would be willing to forego some resources to acquire such skills through training.

Mincer and Higuchi (1991) advanced this hypothesis and found some supporting evidence from U.S. sectoral data: industries with faster productivity growth were also the ones with steeper experience profiles and lower job-separation rates. They interpreted these facts as reflecting the training channel in light of the findings of Lillard and Tan (1986) showing that the incidence of firm-specific on-the-job training is higher in sectors with high rates of productivity growth. Interestingly, Bartel and Sicherman (1998) document that the marginal impact of a rise in productivity growth on the likelihood of training (thus on the steepness of the wage profile) is stronger for low-educated workers, which is consistent with the pattern of the last 30 years mentioned in section 2.

The model developed by Heckman et al. (1998) explains the recent dynamics of the experience premium based precisely on this mechanism. To simplify the exposition, consider again a two-period OLG model where risk-neutral workers are endowed with a unit of human capital, work in both periods, and choose how much time to devote to on-the-job training and production in the first period. Training increases human capital in the second and final period. The problem of a worker at time t is:

$$\max_{\tau_t} \left\{ w_t \left(1-\tau_t\right) + \beta w_{t+1} h_{t+1} \right\}$$

$$s.t. \qquad h_{t+1} = \tfrac{A}{\theta} \tau_t^{\theta},$$

where τ_t is the fraction of the unitary time endowment spent in training, β is the discount factor, w_t is the wage rate at time t, and h_t is human capital at time t. We assume that production of human capital has decreasing returns in the time input. It is easy to see that optimal training and human capital are functions of expected wage growth:

$$\tau_t = \left(A\beta \frac{w_{t+1}}{w_t} \right)^{\frac{1}{1-\theta}} \text{ and } h_{t+1} = A^{\frac{1}{1-\theta}} \left(\beta \frac{w_{t+1}}{w_t} \right)^{\frac{\theta}{1-\theta}}.$$

The implied experience premium, that is, the wage of an experienced old worker relative to the wage of an inexperienced young worker at a given point in time, is then $x_t = h_t / \left(1-\tau_t\right)$.

In a stationary state where $w_t = w^*$ for any t, the optimal fraction of time spent in training is $\tau^* = \left(A\beta\right)^{1/(1-\theta)}$, and the corresponding steady-state experience premium is

$$x^* = \frac{1}{\theta} \left(\frac{A^{\frac{1}{1-\theta}} \beta^{\frac{\theta}{1-\theta}}}{1-\left(A\beta\right)^{\frac{1}{1-\theta}}} \right).$$

The steady-state experience premium is increasing in the productivity of training, A, and in the discount factor, β.

Suppose now that the economy undergoes a one-period transition toward a permanently higher level of skill-biased productivity. High-education (low-education) workers see their wage going up (down), i.e., $w_{t-1} = w_t = w^*$, $w_{t+n} = \bar{w}$ when $n > 1$, where for high-educated workers $\bar{w} > w^*$, and for low-educated workers $\bar{w} < w^*$. Since the two cases are perfectly symmetric, we solve for the transitional dynamics in the experience premium of the high-educated. Along the transition, in period t educated workers increase their investment in training since the anticipated rise in their wages increases the return to human capital accumulation, whereas in all future periods, i.e., $t + 1$ and higher, educated workers do not change their human capital investment decision since their anticipated wage change is not affected:

$$\tau_t = \left(A\beta \frac{\bar{w}}{w^*} \right)^{\frac{1}{1-\theta}} > \tau^*, \tau_{t+n} = \tau^* \text{ for } n \geq 1.$$

The implied sequence of experience premia for educated workers is given by

$$x_t = \frac{h^*}{1-\tau_t} > x_{t+1} = \frac{h_{t+1}}{1-\tau^*} > x_{t+2} = x^*. \tag{18}$$

The experience premium first rises from x^* to x_t and then falls gradually towards the steady state. For low-educated workers, the opposite pattern will hold. If one thinks of time $t - 1$ as 1965, i.e., the moment before the rise in inequality started, time t as 1975, and so on, this stylized model can qualitatively explain the rise in the experience premium for the less educated workers and the decline in the experience premium for the more educated in the 1980s.

The key force is the intertemporal substitution between working and training that the expected changes in wages bring along.[52] Also, as emphasized by Heckman et al. (1998), it is important to recognize that movements in earnings, $w(1 - \tau)$, can differ from movements in skill prices w when labor supply is endogenous. The major limit of the theory is probably that the mechanism depends crucially on the ability of agents to perfectly foresee changes in wage rates decades in advance.

---

[52] Dooley and Gottschalk (1984) also explore a mechanism based on human capital investment in order to explain the rising inequality within cohorts of young workers in the United States. They attribute the changes in expected wages to aggregate fluctuations in labor force growth: the baby-boom and, subsequently, the baby-bust.

5 Inside the Firm: the Organization of Work

Hayek (1945) argued that a fundamental problem of societies is how to use optimally the knowledge that is available but is dispersed across individuals. In frictionless markets, prices can solve this problem: they transmit knowledge about relative scarcity and relative productivity of resources. Since Coase (1937), it is well understood that frictions limit the efficiency of markets, and they divert certain transactions to occur within the boundaries of firms. Within the firm, the organization of work and production plays the role of the market as "information processor" to allow efficient use and transmission of knowledge.

It is therefore not surprising that the recent innovations that revolutionized the way in which information and communication take place have affected the workplace organization within firms and the boundaries of firms. Their impact on the wage structure is perhaps less clear. The maintained hypothesis in the literature is that the recent episodes of reorganization of production, especially in manufacturing, have favored adaptable workers who have general skills and who are more versed at multi-tasking activities. An alternative view, which we will develop later in this section, is that organizational change is not induced by technological change, but that the increased relative supply of skilled labor created the incentives to change the organization of production.

5.1 The Milgrom-Roberts Hypothesis: IT-Driven Organizational Change

Milgrom and Roberts (1990) were the first to emphasize the interaction between the diffusion of information technologies in the workplace and the reorganization of production. Their hypothesis builds on the idea that information technologies reduce a set of costs within the firm which triggers the shift towards a new organizational design. First, electronic data transmission through networks of computers reduces the cost of collecting and communicating data, and computer-aided design and manufacturing reduces the costs of product design and development. Second, there are complementarities among a wide group of strongly integrated activities within the firm (product design, marketing, and production), and pronounced non-convexities and indivisibilities in each activity.

As a result, as the marginal cost of IT declines, it is optimal to reorganize all activities to exploit this shock, and, due to non-convexities, organizational change can be sudden and drastic in nature. In particular, because of lower communication costs the layers in the hierarchical structure can be reduced, so that the organization of the firm becomes "flatter."[53] Workers no longer perform routinized, specialized tasks, but they are now responsible for a wide range of tasks within teams. These teams, in turn, communicate directly with managers. Because of the flexibility of IT capital, the scale of production decreases (recall the evidence in Mitchell (2001) on plant size), allowing greater production flexibility and product customization.

An elegant formalization of this hypothesis is contained in Bolton and Dewatripont (1994). They study the optimal hierarchical structure for an organization whose only objective is that of efficiently processing a continuous flow of information and show using their model that a reduction in communication costs leads to a flatter and smaller organization.

5.1.1 Implications for the Wage Structure

Although in their original papers neither Milgrom and Roberts nor Bolton and Dewatripont explore the implications of organizational change for the wage structure, a small but growing literature on IT-driven organizational change and inequality has developed since.

Lindbeck and Snower (1996) emphasize the "complementarity" aspect of the Milgrom-Roberts hypothesis. They consider a production function with two tasks and two types of workers. The Tayloristic model would assign one type of worker to each task, according to comparative advantages to exploit specialization. The alternative organization of production is the flexible model, where each type of worker performs both tasks. This more flexible organization is preferred when there are large informational complementarities across tasks. The introduction of IT capital amplifies these informational complementarities and makes the flexible organization more profitable. Moreover, firms increase the demand for skilled

---

[53] Rajan and Wulf (2003) use detailed data on job descriptions in over 300 large U.S. companies to document that the number of layers between the lowest manager and the CEO has gone down over time, i.e., organizations have become "flatter".

workers who are more adaptable and versed in multi-tasking, and the skill premium rises.

Möbius (2000) focuses on the "customization" aspect of organizational change. When products are standardized, demand is certain, and production tasks are perfectly predictable, inducing a high division of labor (the Tayloristic principle). New flexible capital allows firms to greatly expand the degree of product variety and customization in product markets. Larger variety implies a more uncertain demand mix because producers become subject to unpredictable "fad shocks" and producers therefore favor a flexible organization of production, with less division of labor. Once again, to the extent that the most skilled workers are also the most adaptable and versatile, the skill premium will increase.

The mechanism in Garicano and Rossi-Hansberg (2003) is based, instead, on the fall in the communication cost within the organization. Their paper has the particular merit of taking the literature on the internal organization of firms (e.g., Bolton and Dewatripont 1994) one step further by recognizing that organizational hierarchies and labor market outcomes are determined simultaneously in equilibrium. Consider an organization where managers perform the most difficult and productive tasks and workers specialize in a set of simpler tasks. Managers also spend a fraction of their time "helping" workers unable to perform their task, and by so doing, they divert resources away from their most productive activities. The fall in the cost of communication allows workers to perform a wider range of tasks, using a smaller amount of the manager's time. The implications for wage inequality are stark. First of all, since workers are heterogeneous in ability, and ability is complementary to the number of tasks performed, inequality among workers within the firm increases. Second, the pay of the manager relative to that of the workers rises because the manager can concentrate on the tasks with high return.

The previous papers have studied how IT-based advances have affected the organizational structure within firms. Saint-Paul (2001) addresses the spectacular rise in the pay of CEOs and a few other professions (e.g., sportsmen and performers) documented in Section 2 using a model where IT-based advances affect the organization of markets with frictions. Saint-Paul combines a model with "superstar" or "winner-take-all" effects (Rosen 1981) with the advent of information technology. In his model, human capital has two dimensions: productivity, i.e., the ability to produce units of output, and creativity, i.e., the ability to generate ideas that can spread (and generate return) over a segment of an economy, called a "network." The diffusion of information technology expands networks increasing the payoff to the most creative workers and widening the income distribution at the top. However, as networks become large enough, the probability that within the same network there will be somebody with another idea at least as good rises: superstars end up competing against each other, mitigating the inegalitarian effects of information technology. Under certain parametric assumptions, inequality first rises and then falls over time.

5.1.2 Empirical Evidence on the Complementarity between Technology, Organizational Change and Human Capital

Bresnahan et al. (2002) investigate the hypothesis that IT adoption, workplace reorganization, and product variety expansion (customization) are complementary at the firm level. Their view is that simply installing computers or communications equipment is not sufficient for achieving efficiency gains. Instead, firms must go through a process of organizational redesign. The combination of IT investments and reorganization represents a skill-biased force increasing the relative demand for more educated labor.

Their empirical analysis is based on a sample of over 300 large firms in the United States, and their definition of organizational change is a shift towards more decentralized decision making and more frequent teamwork. They find a significant correlation between IT, reorganization, and various measures of human capital.[54]

In a related paper, Caroli and Van Reenen (2001) argue that the existence of complementarities between organizational change and the demand for skilled labor leads to three predictions: 1) organizational change should be followed by a declining demand for less skilled workers; 2) in the vein of the directed technical change hypothesis (see next section), cheaper skilled labor should increase the occurrence of organizational change; and 3) organizational change should have a larger impact in workplaces with higher skill levels.

---

[54] See Brynjolfsson and Hitt (2000) for a survey on the empirical work documenting the causal link from adoption of information technology and organizational transformation within the firm.

They test these predictions combining two data sets, one for the United Kingdom and one for France, with information on changes in work organization, working practices, and the skill level of the labor force. Interestingly, they also have information on the introduction of new IT capital, so they can distinguish the effect of organizational change from that of skill-biased technical progress. They find some supporting evidence for all three predictions.

Baker and Hubbard (2003) offer an example where technological change not only affects the organizational design of firms but also the boundary of firms. In particular, they study how IT may have reduced the moral hazard problem in the U.S. trucking industry. Drivers may simply operate the trucks as employees of the dispatching company, or they may actually own the trucks they operate. If the dispatcher owns the truck, there is only limited assurance that the driver will operate in a way that preserves the value of the asset, since the dispatcher cannot perfectly monitor the driving operations. When this moral hazard problem is severe, decentralized ownership will be the outcome, that is, the driver owns the truck. Using detailed truck-level data, Baker and Hubbard show that with the introduction of a new monitoring technology—on-board computers linked to the company servers—the share of driver-ownership decreased significantly.

5.2 Directed Organizational Change

An alternative hypothesis to that put forth by Milgrom and Roberts is contained in several papers discussing the parallel change in the organization and in the pay structure of work. This view maintains that the driving force of organizational shifts is not technology, but rather the secular rise in the supply of skilled workers that created incentives to modify the organization of production: directed organizational change of sorts.

Acemoglu (1999) models a frictional labor market where firms must choose the amount of capital, k, when they are vacant, before meeting the worker. Consider a simple static version of Acemoglu's model. There are two types of workers, skilled and unskilled, where ϕ is the fraction of skilled ones. Skilled workers have productivity, h_s, and unskilled workers, h_u, which we normalize to $1 < h_s$. Output on each job is given by $y_i = h_i^{\alpha} k^{1-\alpha}$, where $i = s, u$. Wages and profits and are, respectively, a fraction ξ and $1 - \xi$ of output net of the cost of the capital installed k. The expected value of a firm choosing capacity k is

$$V(k) = (1-\xi)\left[\phi I^s \left(h_s^\alpha k^{1-\alpha} - k\right) + (1-\phi) I^u \left(k^{1-\alpha} - k\right)\right],$$

where I^i is an indicator variable that equals 1 if the firm accepts a match with a worker of type $i = s, u$ and 0 otherwise.[55] Suppose the firm chooses between two hiring strategies: a "pooling" strategy where it accepts all workers, $I^s = I^u = 1$, and a "separating" strategy where it only accepts skilled workers, $I^s = 1$, $I^u = 0$. Conditional on the hiring strategy, we can use the first-order condition to solve for the optimal choices of capacity, k^P and k^S. Substituting the capacity choice back into $V(k)$, the values of the two hiring strategies are:

$$\begin{aligned} V^P &= \kappa(1-\xi)\left[\phi h_s^\alpha + (1-\phi)\right]^{1/\alpha}, \\ V^S &= \kappa(1-\xi)\phi h_s, \end{aligned} \tag{19}$$

where κ is a constant depending only on α. Comparing these two values, we conclude that the payoff to the "separating" strategy, V^S, dominates the payoff of the "pooling" strategy, V^P, whenever

$$\left(\frac{1-\phi}{\phi^\alpha - \phi}\right)^{1/\alpha} < h_s. \tag{20}$$

Note that the left-hand side of this expression decreases in ϕ, the fraction of skilled workers. When the size of the skilled group is small, a "pooling" equilibrium arises where all firms invest the same amount of capital and search for both types of workers. As the relative size of the skilled group rises, the economy switches to a "separating" equilibrium where firms find it optimal to install more capital and accept exclusively skilled workers in their search process.[56] One can interpret the pooling and the separating equilibrium as different types of work organizations, displaying different degrees of segregation along the skill dimension within sectors. The switch from the low-segregation to the high-segregation organization stretches the wage structure and generates higher inequality.

In a related paper, Kremer and Maskin (1996) offer an alternative explanation for the rise in the degree of assortative matching in the workplace, using a frictionless assignment model.

---

[55] Here, for simplicity we assume that workers accept passively each job offer. We do not consider equilibria where firms randomize, i.e., where $I^i \in (0,1)$.

[56] In the more general version of the model, which is dynamic with free entry of firms, there are other firms who install a small amount of capital (unskilled jobs) and search exclusively for unskilled workers.

Their paper contains some suggestive evidence that the degree of sorting ("segregation") has risen within industries and plants. However, their model is based on an increase in the skill dispersion in the population, for which there is little evidence in the data.[57]

Thesmar and Thoenig (2000) embed a choice of organizational design into a Schumpeterian growth model. Firms can opt for a Tayloristic organization that has large product-specific set-up costs, with the benefit of a high level of productive efficiency. Alternatively, they can choose a new and more flexible organization that can be built with a lower initial fixed cost, but whose productivity level is lower.[58] As is common in this class of Schumpeterian models, there is an R&D sector, where product innovations are generated proportionately to the amount of skilled workers hired. The patent of each new product is then sold to a monopolistic producer who can choose optimally which organization of work to set up (Tayloristic or flexible) according to the volatility of the economic environment.

A rise in the supply of skilled workers will increase the innovation rate in the R&D sector: the higher the innovation rate, the shorter the product's life expectancy for a monopolistic producer, and the less profitable organizations with large fixed costs prove to be, compared to the more flexible production method. The model also produces a rise in segregation, since skilled workers tend to cluster into the R&D sector, as well as a rise in inequality as unskilled workers lose from the abandonment of the Tayloristic model since the production phase becomes less efficient.[59]

5.3 Discussion

The case examined by Baker and Hubbard (2003) is one where IT improves firms' monitoring ability of workers' effort. However, it is plausible that the trend towards a "flatter" organizational design where single-task routinized work is replaced by multi-tasking team-

---

[57] For example, Hoxby and Long (1998) report that the difference in the quality of education (measured by their wage) received by college students from institutions with varying degrees of selectivity has increased over time, but the increase is quantitatively small.

[58] This distinction between the Tayloristic firm and the new flexible firm is due to Piore and Sabel (1984).

[59] Duranton (2004) provides yet another framework for formalizing the concept of "skill segregation" in production and analyzes the implied wage structure in the economy. In his model, a rise in the relative supply of skilled workers can lead to higher segregation and more inequality.

work induces a *rise* in the cost of monitoring individual workers' effort. Firms would then, optimally, introduce incentive schemes (e.g., tournament contracts) with the result of increasing inequality in rewards. In other words, optimal contracts respond to technological and organizational changes that affect the extent of moral hazard within the firm. This line of research is largely unexplored at the moment.

All the models we surveyed in this section are qualitative in nature and, although they establish a logical link between organizational change and inequality, they do not provide any quantitative analysis. One of the main obstacles is that explicit models of organizations contain parameters and variables that are hard to observe, measure, and therefore calibrate (hierarchies, communication costs, number of tasks, etc.). Recently, several papers have started to measure, in various ways, "organizational capital" or "intangible capital" (see, e.g., Hall, 2001, McGrattan and Prescott, 2003, and Atkeson and Kehoe, 2002). A promising avenue for research would try to incorporate this measurement into models that link reorganization with changes in the stock of organizational capital and that relate the latter to the wage structure in order to perform a more rigorous quantitative analysis.

6 Technical Progress as a Source of Change in Labor Market Institutions

Throughout the chapter, up to this point, we have maintained a "competitive" view of the labor market and argued that skills are priced at their marginal product, potentially explaining large parts of the observed dynamics of inequality. However, the labor market displays very peculiar features compared to many other markets in the economy: a sizeable fraction of labor may be considered as under-employed in any given period (unemployment), individual workers often organize themselves into coalitions (unions), and wages frequently seem to be set through some explicit negotiation between firms and workers (individual and collective bargaining). These attributes of the labor market are, arguably, better captured by non-competitive models. We begin our departure from the purely competitive framework by introducing unions and collective bargaining.

Historically, unions and centralized bargaining have been key institutions in the determination of wages and other important labor market outcomes. Over the past 30 years, the economies of the United States and the United Kingdom experienced rapid deunionization. In the United States, in the late 1970s, 30 percent of male non-agricultural private-sector workers were unionized. By 2000, only 14 percent were unionized (Farber and Western, 2000). In the United Kingdom, union density among male workers was around 58 percent in the late 1970s and it has fallen uninterruptedly since to 30 percent today (Machin, 2000 and 2003). There is a variety of evidence that unions compress the structure of wages, even after controlling for workers' characteristics, and thus many economists suspect that their decline may have been an important factor in the increase in inequality in the Anglo-Saxon economies (see, e.g., Gosling and Machin, 1995, and DiNardo et al. 1996).

The existing literature has explored mainly two explanations for the decline in unions. The first generation of papers argued that an important force in the fall of unionization is the change in the composition of the economy away from industries, demographic groups, and occupations where union organization was comparatively cheaper and unions have been traditionally strong (Dickens and Leonard, 1985). However, Farber and Krueger (1992) estimate that compositional shifts can account for at most 25 percent of the decline in the United States and have played virtually no role since the 1980s. Machin (2003) reports that only around 20 percent of the U.K. union decline of the last two decades can be attributed to compositional change.

The second hypothesis is that the legal and political framework supporting union membership deteriorated in the 1970s and 1980s.[60] To date, this explanation seems to have gained rather broad acceptance, even though this view has limits as well. For example, the fall in union organizing activity precedes two key political events: the air-traffic controller strike of 1981 and Reagan's Labor Board appointments in 1983 (Farber and Western 2002). U.K. data also show that the fall in union membership pre-dates the first Thatcher government. Overall, we think that the forces behind rapid deunionization are not yet well understood.

---

[60] Some authors emphasized anti-union management practices (Freeman 1988). Others focused on changes in the composition of the National Labor Relation Board (Levy, 1985).

In most of continental Europe, unions are still strong, and there are no clear signs of decline in union coverage, but a marked change in union behavior has occurred over the past 30 years. Several indexes of coordination and centralization in unions' bargaining for Europe show a distinct trend towards more decentralized wage negotiations, especially in the Scandinavian countries, whose unionization rates are the highest (Iversen, 1998).

The standard explanation for the shift towards decentralized bargaining is based on the interaction between monetary policy and wage setting arrangements. With an independent national central bank, coordination in bargaining among unions is useful because it allows unions to internalize the implications of their wage claims on inflation. With the advent of the European monetary union and the institution of the European Central Bank within-country coordination proves less useful. However, the evidence in favor of this hypothesis is scant. First, monetary policy does not seem to Granger-cause centralization empirically (Bleaney, 1996). Second, we did not observe a substantial trend towards cross-border coordination in unions' bargaining.

Recently, a new hypothesis for deunionization and decentralization in unions' wage setting, based on skill-biased technological change, has been advanced by Acemoglu et al. (2001) and Ortigueira (2002). Their arguments rest on the view that unions are coalitions of heterogeneous workers which extract rents from employers and only exist insofar as members have an incentive to stay in the coalition and continue bargaining in a centralized fashion. The conjecture of these authors is that skill-biased technical change can dramatically alter such incentives.

6.1 Skill-Biased Technology and the Fall in Union Density

Here, we outline a reduced form model that conveys the basic trade-offs highlighted by Acemoglu et al. (2001). Suppose there are two kinds of workers: l_s of which are skilled and $l_u = 1 - l_s$ of which are unskilled. If employed in the competitive sector, these workers will receive wages equal to their productivity, h_s and $h_u < h_s$, respectively. We will think of skill-biased technological change as a rise in h_s relative to h_u.

Workers can also be employed in unionized firms and receive wages, w_s and w_u. A main characteristic of unions is that they compress wages. In our setup, this means that the wage gap between the unionized skilled and unskilled workers is smaller than the productivity gap, or

$$w_s - w_u = \kappa \left(h_s - h_u \right), \tag{21}$$

where $\kappa < 1$ is the degree of wage compression. This equation may arise for a variety of reasons. Collective decision-making within a union may reflect the preferences of its median voter, and if this median voter is an unskilled worker, he will try to increase unskilled wages at the expense of skilled wages. It is also possible that union members choose to compress wages because of ideological reasons or for social cohesion purposes. Or, in presence of idiosyncratic uncertainty, unions could offer insurance to their members by setting a flatter income profile. The empirical literature is broadly consistent with the notion that unions compress wages, though it does not distinguish among the various possible reasons for it (see Booth, 1995).

Union wages (w_s, w_u) must also satisfy some participation constraint for firms (who would otherwise either shut down or open a non-unionized plant). Suppose that this takes the form of non-negative profits:

$$h_s l_s + h_u \left(1 - l_s\right) + \Omega \left(h_s, h_u\right) - \left[w_s l_s + w_u \left(1 - l_s\right)\right] \geq 0, \tag{22}$$

where $\Omega \left(h_s, h_u\right) > 0$ is the additional contribution of unions to output, as a function of both types of labor.[61] This could be because unions, ceteris paribus, increase productivity (for example, Freeman and Medoff, 1984, and Freeman and Lazear, 1995, argue this). Or unions may encourage training (as in Acemoglu and Pischke, 1999).

Solving the wage compression and participation constraint equations (21) and (22) as equalities, we obtain the maximum wage that a skilled worker can be paid as a union member:

$$\overline{w}_s = h_s - \left(1 - l_s\right)\left(1 - \kappa\right)\left(h_s - h_u\right) + \Omega \left(h_s, h_u\right).$$

Intuitively, as w_s rises, w_u must increase too in order to satisfy the wage compression constraint (21) but since profits fall with labor costs, there is an upper bound to the wage of a

---

[61] As long as unions are sustainable, all workers, skilled and unskilled, will prefer to join the union.

skilled union member. Skilled workers will remain union members as long as what they are paid as union members exceeds their competitive salary,

$$\overline{w}_s \geq h_s. \tag{23}$$

From the no-quit condition (23) and the wage compression constraint (21), it follows that $\overline{w}_u \geq h_u$, so unskilled workers will always remain unionized. Observe that the slope of the maximum union skilled wage $\overline{w}_s$ as a function of the productivity of skilled workers h_s is:

$$\overline{w}_s'(h_s) = 1 - (1 - l_s)(1 - \kappa) + \Omega_1(h_s, h_u).$$

Since $\kappa < 1$, as long as the benefits of unionization, $\Omega(h_s, h_u)$, do not increase too rapidly in h_s (i.e., the benefits of unionization do not increase much with skill-biased technical change), we have $\overline{w}_s'(h_s) < 1$. Hence, there exists a cutoff level, h_s^*, such that $\overline{w}_s(h_s) < h_s$ for any $h_s > h_s^*$. This implies that once technical change takes h_s above h_s^*, the wage compression imposed by unions becomes unsustainable, and skilled workers will break away from unions.

Notice that skill-biased technical change is the cause of the deunionization and directly increases inequality. However, deunionization itself contributes to inequality as well. Before deunionization, the wage gap between skilled and unskilled workers is $w_s - w_u \leq \kappa(h_s - h_u)$, and widens smoothly with skill-biased technical change. It is only after deunionization that it jumps up discretely to $w_s - w_u = h_s - h_u$. Therefore, although deunionization is not the primary cause of the surge in wage inequality, it amplifies the original effect of these economic forces by removing the wage compression constraint imposed by unions.

6.2 Skill-Biased Technology and the Fall in Centralized Bargaining

In many European countries—in particular among the Scandinavian countries—the so-called "Ghent system" creates a fiscal-policy link among unions. Under this system, unemployment benefits are administered by the individual unions, but they are funded by the government through aggregate labor income taxation. Hence, not only does the net income of unions' members depend on their negotiated wage, but, through the equilibrium tax rate, also on the wage claims of other unions. Ortigueira (2002) outlines a model economy with this institutional feature, where there are two types of workers, skilled and unskilled, and two unions that can choose to coordinate their wage determination. Unemployment is generated through a frictional labor market with a standard matching function.

Under decentralized bargaining, unions take the tax as given. Ortigueira (2002) shows that there are two possible steady states: in one, unions expect a low tax, thus making moderate wage claims which, in turn, keep equilibrium unemployment and tax rate low, fulfilling the initial expectation; in the other steady state, unions expect a high tax rate, thus making strong wage claims that produce high unemployment and a high tax rate. This second equilibrium yields lower income and lower welfare for union members. Centralized bargaining avoids the coordination failure and the associated welfare losses that can arise in this "bad equilibrium," and hence it can be preferred by unions. Note, however, that the "good equilibrium" under decentralized bargaining is still the best outcome. It is the ex-ante uncertainty that the bad equilibrium could arise that makes coordination attractive.

However, consider what happens with the advent of a skill-biased technology that increases the demand for skilled workers sharply, reducing their unemployment incidence. When unemployment benefits are proportional to wages, the fact that skilled workers are much less likely to be unemployed decreases the social expenditures of the government. As a result, under decentralized bargaining, the equilibrium with high taxes and low welfare does not survive the advent of a skill-biased technology. This justifies the shift in unions' wage setting policies towards decentralization.[62]

6.3 Discussion

The testable implications that can be identified above are that (1) among the experienced workers, the most skilled leave the unions in response to technological improvements and that (2) among the new entrant cohorts, the most educated workers opt for non-unionized

---

[62] See also den Haan (2003) for a model with multiple steady states, one with low tax and unemployment rates and one with high tax and unemployment rates, applied to the U.S.-Europe comparison of labor market outcomes.

jobs. However, these implications are derived from theories of technology-induced deunionization that are rather exploratory; more sophisticated and rigorous models of unions (with endogenous membership and endogenous wage-compression mechanisms) are yet to be developed.

The recent empirical studies by Card (2001), for the United States, and Addison et al. (2004), for the United Kingdom, compare the unionization rate across several skill groups before and after the collapse in union density in these two countries (1973 and 1993 for the United States, and 1983 and 1995 for the United Kingdom). The common finding of these two papers, is that unionization declined most for the low- and middle-skill groups.[63] Taken at face value, this preliminary evidence is not favorable to the hypothesis discussed in this section. However, one has to be cautious in interpreting these results because this work does not control for unobserved heterogeneity.[64] Suppose that—as documented by Card (1996)—unobserved ability is higher among unionized workers with low observable skills. Given that unionized firms offer a compressed wage schedule, such a contract would attract the highest ability workers with low education and the lowest ability workers with high education. Moreover, assume that technological change induces a rise in the market return for innate ability, as discussed in section 3.2. Then, the theory suggests that one should observe exactly the cross-skill deunionization pattern documented from U.S. and U.K. data.

It should be mentioned that a technology-based theory of deunionization must also explain why union density did not fall (in fact, it expanded somewhat) in the public sector. Since the public sector is, by definition, sheltered from the international competition, it is reasonable to conjecture that the leap in competitive pressure faced by many manufacturing industries over the past 30 years eroded those rents that are, according to some researchers, at the heart of the existence of unions. A quantitative evaluation of the importance of this channel is yet to be performed.

---

[63] Note that wages in the union sector do not fully reflect skills. For this reason, these authors impute skill deciles to unionized workers based on what workers with similar observable characteristics (age, education, gender, race, etc.) would earn in the non-union sector.

[64] Card (2001) makes a rough adjustment for unobserved heterogeneity, based on Card (1996). A thorough analysis would require the use of longitudinal data, but both, Card (2001) and Addison et al. (2004) are restricted to repeated cross-sections.

Another avenue that so far has not been pursued is the analysis of deunionization in conjunction with the structural changes in workplace organization that occurred in the past 30 years. In Section 5, we argued that a distinct feature of the recent change in the production process, especially in manufacturing, is the switch from Tayloristic organizations, where workers repeatedly performed similar tasks around the conveyer belt, towards "flatter" organization built on teams where workers engage in multiple tasks and where the individual division of labor is much fuzzier. Union's wage setting arrangements, based on "equal pay for equal work", can be effective within a Tayloristic plant, but then become very inefficient in plants where production is organized through teams. There is no reason to assume that workers performing the same task will be equally productive, since they perform many other complementary operations simultaneously (see, e.g., Lindbeck and Snower, 1996).

7 Technological Change in Frictional Labor Markets

Most of the models presented so far feature an aggregate production technology, i.e., the production structure is centralized, and competitive labor markets. Constructing a frictional model of the labor market requires departing from both attributes and moving towards a decentralized production structure and a labor market with imperfect coordination between workers and firms in the matching process. This class of models gives rise to frictional equilibrium unemployment and "frictional equilibrium inequality". By frictional inequality, we mean wage dispersion that is purely an artifact of frictions and that, without frictions, would disappear. A useful way to think about this phenomenon is to introduce the concept of "return to labor market luck".

Throughout this chapter, we have discussed several models where technological progress produces a rise in the return to observable and unobservable *permanent* components of individual skills, such as educational attainment, age, and innate ability. These permanent factors greatly determine inequality of earnings among the population, but they are not by any means exhaustive. Earnings display a large *stochastic* component (e.g., events related to the luck of individuals, firms, or industries) that is responsible for their fluctuations around the permanent component.[65]

Gottschalk and Moffitt (1994) were the first to ask how much of the observed increase in inequality is attributable to a rise in earnings volatility and instability around its permanent component. They used a simple statistical model where log wages, w_{it}, for an individual i at time t—net of their predictable age profile—are assumed to be the sum of two orthogonal components, a fixed individual effect, α_i, and a stochastic ($i.i.d$) component, ε_{it}. Using the covariance structure of wages within a panel of U.S. males (constructed from PSID data), they reached the conclusion that the fraction of the total increase in cross-sectional inequality attributable to a surge in earnings volatility is between one third and one half.[66] One can interpret this fact as a rise in frictional inequality, or in the "return to labor market luck." The argument set forth is that the rapid diffusion of a new technology leverages the importance of these stochastic factors, raising the premium to workers with no observable distinguishing characteristics other than their good fortune.

Most of the work we review uses the random matching model of the labor market (see, e.g., Mortensen and Pissarides, 1998, or Pissarides, 2000). In this framework the existence of frictions creates a bilateral monopoly as a result of a meeting between a vacant firm and a worker. Wages are determined by bargaining over total output, so more productive firms tend to pay more, creating wage dispersion among ex-ante equal workers. We start by studying how technological change affects unemployment in this class of models. Next, we move to wage inequality. Random matching is a somewhat extreme characterization of frictions. In the last part of the section we contrast random search models to directed search models.

---

[65] A large empirical literature documents wage dispersion among observationally equivalent workers that cannot be fully reconciled with unobserved heterogeneity in permanent components. Abowd et al. (1999) document that firm effects still play a role, after controlling exhaustively for individuals' effects. Krueger and Summers (1988) found that a worker moving from a high to a low wage industry is subject to a wage loss roughly equal to the inter-industry differential.

[66] The subsequent literature on the subject demonstrated the robustness of this result to richer statistical models for the stochastic component of wages. See Haider (2001), Heathcote et al. (2003), and Meghir and Pistaferri (2004) for the United States and Blundell and Preston (1998) and Dickens (2000) for the United Kingdom.

7.1 Technological Progress and Frictional Unemployment

There is a sizeable literature trying to characterize how equilibrium unemployment reacts qualitatively to variations of the rate of technological change within a matching model à la Diamond-Mortensen-Pissarides (DMP) with vintage capital à la Solow (1960). Two distinct approaches emerge from the literature.

The first, that can be attributed to Aghion and Howitt (1994), argues that when new and more productive equipment enters the economy exclusively through the creation of new matches—because existing matches cannot be "upgraded"—it has a Schumpeterian "creative-destruction" effect: new capital competes with old capital by making it more obsolete and tends to destroy existing matches, because workers are better off separating from their old matches to search for the new firms endowed with the most productive technology. Thus, unemployment tends to go up as growth accelerates, due to a higher job-separation rate.

The second approach, due to Mortensen and Pissarides (1998), proposes an alternative view whereby the new technologies enter into existing firms through a costly "upgrading" process of old capital. In the extreme case where upgrading is free, we have the Solowian model of disembodied technological change, even though the carrier of technology is equipment. The separation rate is unaffected by faster growth and all the effects work through job creation. For small values of the upgrading cost, unemployment falls with faster growth, thanks to the familiar "capitalization effect": investors are encouraged to create more vacancies, knowing that they will be able to incorporate (and hence benefit from) future technological advances at low cost.[67]

Hornstein et al. (2003b) try to resolve the issue quantitatively. When they parameterize the model to match some salient features of the U.S. economy, they find that, in the vintage-matching model, the link between capital-embodied growth and unemployment does

---

[67] An interesting qualification to this result is provided by King and Welling (1995): if, unlike what is customarily assumed in this family of models, workers bear the full fixed search cost, then the capitalization effect leads to an increase in the number of searchers and to longer unemployment durations. See Pissarides (2000, chapter 3) for a detailed discussion on growth and unemployment in matching models of the labor market.

not importantly depend on to what parties—new matches or old ones—the benefits of the technological advancement accrue. The intuition for this "equivalence result" is that upgrading can be much better than creative destruction only if it is very costly for vacant firms to meet workers, but the data on the low average unemployment and vacancy durations imply that, in the model, this meeting friction is minor. That paper also shows that the same data on average unemployment duration impose severe restrictions on how much frictional wage inequality the model can generate. In the standard search model, high dispersion of wage opportunities makes workers very demanding and increases unemployment spells. Thus, a high wage dispersion could only coexist, in equilibrium, with long unemployment durations.

We now turn to the analysis of how technological progress impinges on frictional inequality in random matching models. In these models, however, the limits on the extent of wage inequality due to luck emphasized in Hornstein et al. (2003b) apply as well.

7.2 Technological Heterogeneity and the Returns to Luck

In a frictional labor market populated by ex-ante equal workers, an increase in technological heterogeneity can increase the return to luck. We explain this mechanism within a simple framework based on Aghion et al. (2002).[68] Consider an economy populated by a measure one of infinitely lived, ex-ante equal, and risk-neutral workers as well as by the same measure of jobs. Jobs are machines embodying a given technology. The technological frontier advances every period at rate $\gamma > 0$. The machines have a productive life of two periods. An age $j \in \{0, 1\}$ machine that is matched with a worker produces output, $y_j = (1+\gamma)^{-j} h$ (normalized relative to the age 0 machine), where h represents the skill level of the workers.

The labor market is frictional, i.e., workers separated from their jobs are randomly rematched with a vacant machine. To simplify, we assume that they always make contact with a machine. We postulate that, upon contact, the bilateral monopoly problem is solved by a rent sharing mechanism setting wages to be a constant fraction, ξ, of current output, y_j, where ξ is a measure of the bargaining power of workers.

It is easy to see that in an equilibrium where all job offers are accepted, the lucky half

---

[68] See also Manuelli (2000) and Violante (2002).

of the workers will be employed on new machines and the unlucky half on old machines. The variance of log wages is simply given by $var(\log w) = \gamma^2/4$, which is increasing in the rate of embodied technological change. Intuitively, in this economy all the heterogeneity is generated by technological differentials across machines. A technological acceleration (rise in γ) amplifies the productivity gaps between jobs. Since in this non-competitive labor market individual wages are linked to individual output, this acceleration then also raises wage dispersion even among ex-ante equal workers, i.e., it raises the return to luck.[69] As in Jovanovic (1998), however, if the scrapping age of capital is endogenous, the model would display an offsetting force. This force is due to the fact that, when the growth rate is higher, machines become obsolete faster, and firms scrap machines earlier. Therefore the equilibrium age range of machines in operation shrinks, compressing technological heterogeneity.

7.3 Vintage Human Capital with Frictions

A technological acceleration not only affects transitory residual wage inequality through its impact on the underlying distribution of job productivity differences. The technological acceleration may also affect the distribution of worker productivity differences if it interacts with the accumulation of job/technology-specific knowledge.[70] Violante (2002) extends the above model to include vintage human capital. Employed workers accumulate, through learning-by-doing, knowledge about the technology they are matched with. We normalize the amount of specific skills cumulated after every employment period to 1, so that the learning curve of the workers is concave, i.e., learning is faster for workers with lower initial skills. To keep the model tractable, we also assume that skills fully depreciate after two periods.

A worker on a machine of age i who moves on to a machine of age j next period can

---

[69] This increase in wage inequality is mirrored by a rise in wage instability along the lifetime of each worker: given a certain amount of labor turnover, larger cross-sectional productivity differences translate into more volatile individual wage profiles.

[70] The accumulation of job/technology-specific knowledge is also at the heart of the discussion of the experience premium in section 4.

transfer h_{ij} units of the accumulated skills to the new job :

$$h_{ij} = \min\left\{(1+\gamma)^{\tau(j-i+1)}, 1\right\}, \tag{24}$$

with $\tau > 0$ and $i, j \in \{0, 1\}$. The fraction of skills that can be transferred from an old to a newer machine is proportional to the *technological distance* between the two machines through a factor $\tau \geq 0$. The presence of the term γ in the transferability technology is crucial: the rate of quality improvement of capital-embodied technologies determines the degree by which new technology is different, more complex, and richer than the previous generation of machines. A higher γ reduces skill transferability in the economy.[71] Equation (24) and the depreciation assumption implies that we have three skill levels in the economy:

$$h_{01} = 1, \quad h_{00} = h_{11} = (1+\gamma)^{-\tau}, \quad h_{10} = (1+\gamma)^{-2\tau}, \tag{25}$$

and the corresponding wage rates (normalized relative to the wage on an age 0 machine) are $w_{ij} = \xi h_{ij} (1+\gamma)^{-j}$, $i, j \in \{0, 1\}$. Note that, given this simple expression for wages, the variance of log wages can be written as

$$var(\widetilde{w}) = \gamma^2 var(j) + var(\widetilde{h}) - 2\gamma cov(\widetilde{h}, j), \tag{26}$$

which is the sum of technological heterogeneity (the force discussed earlier), ex-post skill heterogeneity among workers, and the degree of assortative matching between skills and technologies measured by their covariance.[72]

One can prove that, for large enough γ, workers separate from firms every period.[73]

---

[71] The book by Gordon (1990) provides several examples of quality improvement in equipment requiring the performance of new tasks in the associated jobs. In the aircraft industry in the 1970s, new avionics were introduced that provided a safer but more complex navigation system. In the telephone industry, around the mid-1970s, electromechanical telephone switchboards were replaced by more sophisticated and flexible electronic equipment with larger programming possibilities. In the software industry, since the early 1980s, every new version of a software is equipped with new features. Those users who remain attached to an old version are often unfamiliar with many features of the new version.

[72] A rise in the degree of assortative assignment between workers' skill and machines' productivity is equivalent to a *fall* in the covariance component (recall that j is machine's age, which is inversely related to productivity) and a rise in the variance of wages.

[73] This result is related to the intertemporal trade-off intrinsic in the separation decision: choosing to remain on the old vintage improves the current wage (no vintage-specific skill is lost), but worsens future wages because in the next period the worker will have older knowledge, with low degree of transferability. As γ goes up, the expected future wage loss from holding old skills increases faster than the current wage gain, inducing the worker to optimally anticipate its separation decision.

Under this optimal separation rule, the equilibrium level of wage dispersion is given by

$$var(\widetilde{w}) = \gamma^2 \left[\frac{1}{4} + \frac{1}{2}\tau\left(\tau - 1\right)\right], \tag{27}$$

so it is increasing in γ whenever the variance is well defined (positive). In particular, the equilibrium displays $var(\widetilde{h}) = \frac{\gamma^2\tau^2}{2}$, and $cov(\widetilde{h}, j) = \frac{\gamma\tau}{4}$. The variance of skills is increasing in γ since a higher γ reduces the skill transferability of the bottom-end workers (h_{10}), while not affecting the skill level of the top end workers (h_{01}). The covariance between skills and age of technology is also increasing in γ, a force that restrains inequality because it worsens equilibrium sorting in the economy. The reason is that a larger γ reduces the skills of workers moving to the new technology relatively more than the skills of workers moving to old technologies.

A common criticism of this class of models is that the degree of churning in the labor market (i.e., labor mobility or job reallocation) has to rise in order to generate more volatile earnings, whereas the empirical literature documents no significant rise in labor mobility (Neumark, 2000).[74] This is a misconception. One way to unravel this issue exploits the equivalence between cross-sectional wage dispersion and individual wage instability in a model with ex-ante equal and infinitely lived agents. A technological acceleration has two effects. First, it curtails skill transferability, thereby increasing wage losses upon separation. Second, it reduces the average skill level of workers who find themselves, on average, on the steeper portion of a concave learning curve, which in turn implies higher wage growth on the job. Both these forces tend to raise individual earnings volatility, for any given level of labor mobility. Violante (2002) offers some evidence of wage losses upon separation and wage growth on the job being larger in the 1980s than in the 1970s and shows that a calibrated full-scale version of this model can account up to 90 percent of the rise in wage instability in the U.S. economy, while at the same time implying only a very modest rise in equilibrium labor turnover.

---

[74] The empirical literature on labor mobility contains partly opposing results: whereas Jovanovic and Rousseau (2004b) find a significant decline of labor mobility since the 1970s, Kambourov and Manovskii (2004) find that occupational mobility increased since the 1970s.

7.3.1 Occupation-Specific Human Capital

Occupation-specific experience may be one of the least transferable components of human capital, and a change in occupational mobility can have a big impact on the wage structure. Kambourov and Manovskii (2004) document an increase in occupational mobility in the United States from 16 percent in the early 1970s to 19 percent in the early 1990s.[75] Based on a calibration exercise, Kambourov and Manovskii argue that 90 percent of the rise in residual inequality (i.e., in both the permanent and the stochastic component) is due to increased occupational mobility.

The authors build a model of occupation-specific human-capital accumulation based on the equilibrium search framework of Lucas and Prescott (1974). At any one time workers can work in one occupation only. Workers choose their occupation based on their occupation-specific experience. When working in an occupation, workers increase their specific experience, and they lose some of this experience when moving between occupations. A worker's wage in a given occupation depends on the specific experience and the occupation's productivity.

The productivity of occupations is subject to shocks, and increased variability of these shocks directly increases wage variability. However, the total impact of occupational productivity shocks on wage inequality depends on the occupational choice response of workers. Workers in an occupation whose productivity declines choose to move in search of better occupations, and, by so doing, they dampen the effect of the shock on inequality. When the increased variance of productivity shocks is accompanied by decreased persistence –as conjectured by the authors– workers in occupations hit by moderately negative shocks may choose not to switch occupations because occupations which look profitable today may turn quickly into unproductive ones. This latter effect amplifies the direct effect of the initial shocks.[76]

---

[75] Kambourov and Manovskii use occupational data from the PSID at the three-digit level, including almost 1000 occupational groups.

[76] The model of Bertola and Ichino (1995), discussed in section 8, generates increased wage inequality through a similar mechanism.

7.3.2 A Precautionary Demand for General Skills

Gottschalk and Moffitt (1994) found that the transitory component of inequality is larger (and increased more) for low-education workers. Gould et al. (2001) model this phenomenon using a vintage human capital model where risk-averse workers choose their level of education. They study an economy where workers are ex-ante heterogeneous with respect to permanent innate ability, and the return to college education is increasing in ability. High-ability workers obtain a college education that provides them with general skills which do not depreciate as technology advances. Low-ability workers do not acquire general skills in college; rather, they acquire technology-specific experience through on-the-job learning. Here, we refer to workers with a college education as skilled and to workers without a college education as unskilled.

Gould et al. (2001) consider a shock to the economy that simultaneously increases the rate of embodied technological change and the ex-ante variance of technological progress across jobs.[77] This shock increases the "precautionary" demand for college education, since holding technology-specific skills becomes more risky. The lowest ability threshold for college graduates falls, and thus permanent inequality increases within skilled workers and falls within the group of unskilled workers. At the same time, the rise in the variance of embodied technological change means that "skill erosion" has a bigger impact on the relative wages of unlucky and lucky unskilled workers, so the increase in their wage variance is mostly determined by transitory components.

This mechanism relies on the assumption that the variance of technical progress is heteroskedastic in the sense that it rises with its mean. We know very little about this property: Cummins and Violante (2002) analyze the whole cross-industry distribution of equipment-embodied technical change for 62 industries in the United States from 1947 to 2000 and find little evidence of changing variance, although the mean grows substantially over the period. However, they document a rise in the cross-sectoral variance of the "technological

---

[77] This view of the past 30 years as being a period of high "turbulence" is also present in several models of the differential labor market performance between the United States and Europe, see Section 8.

gap" between average capital and leading-edge machines.[78] According to the transferability technology (24), the technological gap closely measures the degree of skill erosion of an average worker displaced in a given industry.

7.3.3 Explaining the Fall in Real Wages

Interestingly, in a set of model economies with vintage human capital (Helpman and Rangel, 1999, Gould et al. 2001, Violante, 2002, or Kambourov and Manovskii, 2004), during the transition to the new steady state, and notwithstanding the technological acceleration, the fall in the average skill level of the workforce can generate a temporary slowdown in average wage growth and a fall in the real value of wages at the bottom of the distribution—two facts that have been documented extensively for the period of interest.

To illustrate this point, let us return to the model from Section 7.3. Note that in an equilibrium where workers separate every period—as assumed—each skill type represents one fourth of all workers. The four skills types are reported in expression (25) . It is immediate to see that the normalized average log level of skills is $-\tau\gamma$, and thus it falls unambiguously when γ increases. This opens the interesting possibility that, in the model, the average wage could decrease along the transition following a technological acceleration.

Suppose that at time t the economy is in steady state with $\gamma = \gamma_L$ (and with the productivity of the new machine normalized to 1). The average log wage is then $\widetilde{w}_t = -\tau\gamma_L - \gamma_L/2$. Suppose now that γ rises to γ_H. Then, some simple algebra shows that in the next period, the average log wage is

$$\widetilde{w}_{t+1} = \frac{\gamma_H}{2} - \frac{\tau}{2}(\gamma_L + \gamma_H) = \widetilde{w}_t - \frac{\tau}{2}(\gamma_H - \gamma_L) + \frac{1}{2}(\gamma_L + \gamma_H).$$

Thus, despite the technological acceleration, the average wage decreases along the transition if $\tau > (\gamma_L + \gamma_H) / (\gamma_H - \gamma_L)$, that is, if τ or the increase in γ are large.

An alternative explanation for the fall in real wages –which does not depend on vintage human capital– is advanced by Manuelli (2000) within a frictional labor market model where workers have bargaining power and can seize a fraction of the firm's future stream of profits,

---

[78] See section 3.2.1 for a formal definition of the technological gap.

through wage negotiations. Consider what happens when it is announced that: 1) a new technology will be available in the future; but 2) the incumbent firms will be able to adopt it only with some probability (as in Greenwood and Jovanovic, 1999). Existing firms will anticipate a future increase in wages, driven by the new, more productive entrants. Hence, there will be a transitional phase before the arrival of the new technology, where the market value of the incumbent firms will fall and, with them, the wages they currently pay.

7.4 Random Matching vs. Directed Search as Source of Luck

So far, we have analyzed economies where the friction is due to random matching. Wong (2003) argues that models with random matching can have counterfactual implications.[79] It is well known that in a matching model with two types of workers (skilled and unskilled) and two types of firms (high-tech and low-tech), there can be multiple equilibria (Sattinger, 1995). There are equilibria with perfect sorting where skilled (unskilled) workers are matched with high- (low-) tech firms and equilibria that display some degree of "mismatch." In the latter class of equilibria, luck plays a role as skilled workers, ex-ante equal, can end up in jobs with different productivities. Suppose output is the product between efficiency of capital, z_i, where $i = l, h$ and $z_h > z_l$, and efficiency of labor, h_j, where $j = s, u$ and $h_s > h_u$, i.e., $y_{ij} = z_i h_j$. A wave of skill-biased technical change (or a capital-embodied technological acceleration) that increases the relative productivity of high-tech jobs (i.e., the ratio z_h/z_l) makes high-tech firms more picky in their choice of workers, as now the same skill differences translate in larger output differences. The equilibrium with mismatch is less likely to survive. When the economy switches to the equilibrium with perfect sorting, luck-driven inequality among ex-ante equal workers falls to zero.[80]

One of the key reasons why the model has this counterfactual prediction is that, due to random matching, prices (wages) have no signaling value. Shi (2002) analyzes exactly the same framework (a two-worker, two-firm economy) but he replaces Nash bargaining and

---

[79] See also Albrecht and Vroman (2002) for a similar environment.

[80] The argument in Wong (2003) regarding models with random matching is quite general; e.g., it applies in the model by Acemoglu (1999). From equation (20) of Section 5, note that as h_s rises (skill-biased technical change), the pooling, or mismatch, equilibrium is less likely to survive, so within-group inequality falls.

random matching with wage posting and directed search, following the alternative approach of "competitive search" (Moen, 1997). His conclusion is that random matching is not essential for technical progress to leverage the effect of luck in the labor market: directed search works equally well.

In this environment, skilled workers only apply to high-tech jobs, while unskilled workers apply to both types of jobs. Ex ante, every unskilled worker is indifferent between jobs, but inequality is generated ex post. Since high-tech firms always give priority to skilled applicants, unskilled workers applying for high-tech jobs are less likely to become employed than are unskilled workers applying for low-tech jobs. Therefore unskilled workers applying for high-tech jobs have to be offered higher wages than in low-tech jobs.

With free entry, a rise in the relative productivity of high-tech jobs (skill-biased technical change) induces the creation of more high-tech vacancies. More unskilled workers become attracted to the high-tech sector and in equilibrium their job finding probability in the high-tech sector falls, so wages rise. In the meantime, fewer unskilled workers stay in the low-tech sector, so their wages fall. In sum, wage inequality among ex-ante equal workers rises with the degree of skill bias in technology.

Can one conclude that directed search models are more suitable than random search models for studying problems where heterogeneity is crucial, such as wage inequality? The answer depends on the dimension of inequality studied. Directed search seems a more reasonable assumption when the trait determining heterogeneity is observable (e.g., education, general experience), whereas random matching fits better in the analysis of wage inequality when the source of heterogeneity is not directly observable (e.g., ability or vintage-specific skills).

8 Technology-Policy Complementarity: United States vs. Europe

A large portion of this chapter has been dedicated to the analysis of a number of different economic models designed to decipher the dynamics of the U.S. wage distribution over the past three decades, in light of changes in technology.

In this section we expand our viewpoint to include other dimensions of labor market inequality, which allows us to contrast the U.S. experience with the European experience. In Section 2 we documented that while wage inequality soared in the United States, both the labor share of income and the unemployment rate remained remarkably stable there. In sharp contrast, in most of the large continental European economies, the wage structure did not change much at all, while the labor share fell substantially and unemployment increased steadily. In particular, the increase in European unemployment largely reflects longer durations rather than higher unemployment incidence.

8.1 The Krugman Hypothesis

Why have we observed such different outcomes for two regions of the world standing at a similar level of development and, therefore, being subject to very similar aggregate shocks? Are we witnessing a sort of *devil's bargain*, i.e., a trade-off between inequalities: low unemployment can only be achieved by paying the price of soaring wage inequality? And, if so, what determines the position of each country along this trade-off?

In Table 2 we report, for the set of countries from Table 1, some indexes of the rigidity of various labor market institutions reproduced from Nickell and Layard (1999). The conclusion is unambiguous: compared to the United States, continental Europe has stricter employment protection legislation, more generous and longer unemployment benefits, less decentralized wage bargaining, and more binding minimum wage laws.

TABLE 2

The large majority of papers in the literature have taken the data exhibited in Table 2 as uncontroversial evidence that the reason for the observed differences can be found in the differences in labor market institutions between the United States and Continental Europe. Krugman (1994) was probably the first to provide a simple formalized model of this hypothesis. Simply put, the interaction between a severe technological shock and rigid European institutions have induced an adjustment through equilibrium *quantities* of labor (i.e., the employment distribution), whereas in the flexible U.S. labor market, the adjustment occurred through *prices* (i.e., the wage distribution).

Several authors have tried to test the Krugman hypothesis econometrically. The typical analysis is based on a cross-country panel of institutions and shocks, i.e., it allows for changing institutions over time, beyond aggregate shocks. A statistical model linking shocks and institutions to the dynamics of unemployment and wage inequality is estimated to evaluate the role of shocks and institutions, first separately and then interacted. The shocks considered are usually of technological nature and are measured through changes in measured TFP and changes in the labor share of income, possibly capturing a form of capital-biased technical change. In all cases the shock is assumed to be common across countries.

Blanchard and Wolfers (2000) argue that changing institutions alone have little explanatory power. The performance of the statistical model in explaining cross-country patterns of unemployment rates improves once shocks and institutions are interacted: an equal-size technological shock has differential effects on unemployment when labor market institutions differ. Bertola et al. (2001) provide further evidence for this view. Bentolila and Saint-Paul (1999) also study the evolution of the labor share across OECD countries since 1970. Using panel data techniques, they find that in the presence of institutions that promote wage rigidity, shocks that reduce employment also significantly reduce the labor share of income. One common problem in this empirical literature is that the results are, in general, not robust to the chosen specification.[81]

Another problem of this methodology is that the economic mechanism behind the interaction between technology and policy is not explicit. Consistent with the approach we have taken in the chapter so far, we will devote more space to quantitative analyses based on "structural" equilibrium models. In the rest of this section, we present the various frameworks the literature has explored to understand the interactions between technological progress and labor market institutions in shaping the various dimensions of inequality. We

---

[81] The recent results in Nickell and Nunziata (2002) seem to support an explanation of cross-country unemployment differentials largely based on changing institutions, with a common technological shock playing only a minor role.

have grouped these frameworks into six categories, according to the type of technological shock modeled: 1) a rise in microeconomic turbulence, linked to some fundamental change in technology, 2) a slowdown in total factor productivity, 3) an acceleration in the rate of capital-embodied productivity improvements, 4) skill-biased technical change, 5) a technological innovation whose adoption is endogenous, and 6) the structural transformation from manufacturing to services.

8.2 Rise in Microeconomic Turbulence

In Section 2 we have documented that roughly one-half of the rise in cross-sectional wage differentials in the United States is not associated to a higher return to permanent skills. Rather, it is due to increased wage "instability" over the workers' life time. In other words, transitory idiosyncratic shocks to labor productivity and wages have become more important over time (Gottschalk and Moffitt 1994). These larger temporary wage movements constitute important evidence that there has been a rise in the degree of microeconomic turbulence in the U.S. economy.

More evidence comes from the firm side. Campbell et al. (2001) show that the cross-sectional variability of individual stock returns has trended upward from 1962 to 1997. Chaney et al. (2003) and Comin and Mulani (2003) use Compustat firm-level data to demonstrate that the firm-level volatility of real variables, such as investment and sales, has gone up from 1970-1975 to 1990-1995. Overall, these papers provide snapshots, from very different angles, of an economy where idiosyncratic turbulence and volatility have risen to a high level.

Bertola and Ichino (1995) and Ljungqvist and Sargent (1998, 2003) argue that a rise in microeconomic turbulence that interacted with more or less rigid institutions can explain the U.S.-Europe dichotomy. Interestingly, the former authors identify wage rigidity and strict employment protection laws as the culprits, while the latter emphasize the generosity of unemployment benefits. Note, though, that one key premise behind these theories is that the surge in turbulence is common to the United States and Europe. We are not aware of any empirical work documenting trends in microeconomic instability in continental Europe.

Currently, this represents a limit for this class of explanations.

8.2.1 The Role of Wage Rigidity

The framework proposed by Bertola and Ichino (1995) is inspired by the Lucas and Prescott (1974) island-model of equilibrium unemployment. The economy is populated with a measure, L, of risk-neutral workers and a measure one of firms, indexed by $i \in [0,1]$. Each firm is subject to idiosyncratic productivity shocks that follow a two-state Markov chain taking values $\left(A^G, A^B\right)$, with $A^G > A^B$, and with transition probability, p, that the state (good, G, and bad, B) changes. When labor mobility is perfect, employment adjusts across good and bad firms to equalize wage differentials, and a unique market-clearing wage rate arises in equilibrium, i.e., there is no wage inequality.

Consider now the case where wages are flexible, but where workers have to pay a fixed moving cost, $\kappa > 0$, to change firms (this is the U.S.-like economy). In any period, workers observe the productivity level in all firms, but moving takes one period. Hence, when they start working, productivity might change. It is easy to see that the value functions of a worker in good- and bad-state firms, respectively, are

$$W^G = w^G + \frac{1}{1+r}\left[pW^B + (1-p)\,W^G\right], \tag{28}$$

$$W^B = \begin{cases} w^B + \frac{1}{1+r}\left[pW^G + (1-p)\,W^B\right] & \text{if staying,} \\ w^B - \kappa + \frac{1}{1+r}\left[pW^B + (1-p)\,W^G\right] & \text{if moving.} \end{cases} \tag{29}$$

If workers leave bad firms in equilibrium, the marginal worker has to be indifferent between staying in a B firm or moving, yielding

$$W^G - W^B = \frac{1+r}{1-2p}\kappa.$$

Using (28) and (29) together with this condition, one arrives at the expression for equilibrium wage inequality:

$$w^G - w^B = \frac{r+2p}{1-2p}\kappa. \tag{30}$$

On the one hand, the closer p is to 0, the more permanent are productivity changes. This justifies a large amount of wage-equalizing mobility, and hence there is smaller ex-post wage inequality across firms. On the other hand, the larger the degree of volatility in the economy (the closer p is to $1/2$), the riskier it is for a worker to move, as the new firm can quickly turn into the B state, and the cost κ is wasted. In this case, mobility will be low and the ex-post wage differential will increase.

Now consider the same experiment in a Europe-like economy where wages are rigid, i.e., where $w^B = w^G = w$, and where firing costs are prohibitively high, so that employment at every firm is constant at $\bar{l}$. To analyze this situation, Bertola and Ichino assume that firm i has a linear marginal revenue product $\pi\left(l^i\right) = z^i - \alpha l^i$, so that the marginal values for a firm in the G and B state of a unit of labor, respectively, are

$$\begin{aligned} V^G &= A^G - \alpha\bar{l} - w + \frac{1}{1+r}\left[pV^B + (1-p)V^G\right], \text{ and} \\ V^B &= A^B - \alpha\bar{l} - w + \frac{1}{1+r}\left[pV^G + (1-p)V^B\right]. \end{aligned}$$

In an equilibrium with free-entry, the hiring firm in the G state will have $V^G = 0$. Hence, the system above can be easily solved for $\bar{l}$ to give

$$\bar{l} = \frac{A^G - w}{\alpha} - \left(\frac{p}{r+2p}\right)\left(\frac{A^G - A^B}{\alpha}\right), \tag{31}$$

which shows that a rise in p that increases the degree of turbulence in the rigid economy will reduce average employment, i.e., it will increase the unemployment rate, $L - \bar{l}$. The reason is straightforward: when firms are constrained in their ability to shed labor in the face of a negative shock, they will be very cautious in hiring new workers even in the high-productivity state. Note, in fact, that the larger is the productivity differential $A^G - A^B$ across states, the higher will average unemployment in the economy be.

In conclusion, a similar increase in economic uncertainty induces more caution in workers' mobility and larger wage differentials in an economy with flexible wages whereas it leads to more caution in firms' hiring and lower average employment in an economy with rigid wages and costly layoffs. This result remains qualitative, as the authors did not try an exploration of the quantitative importance of their mechanism. In particular, it would be of interest to study by how much labor turnover needs to decline in order to generate a rise in wage inequality of the magnitude observed in the U.S. economy. Interestingly, as mentioned earlier, Jovanovic and Rousseau (2004b) document a substantial downward trend in labor mobility in the United States, from 50 percent in 1970 to 35 percent in 2000.

8.2.2 The Role of Welfare Benefits

Ljungqvist and Sargent (1998, 2003) propose an alternative mechanism based on the standard search model of unemployment (McCall 1970). Here, we present a stripped-down version of their argument. Consider an unemployed worker with skill level, h, who searches for a job, sampling wage offers every period from the stationary distribution, $F(w)$, with finite support $[\underline{w}, \bar{w}]$. Her skill level, when unemployed, decays at the geometric rate, δ, whereas, when employed, skills remain unchanged. Employment is an absorbing state (no exogenous breakup of jobs), and workers discount the future at rate r. Unemployment benefits are equal to b. The values of employment and unemployment for a worker of skills h are

$$W(w,h) = \frac{wh}{r}, \text{ and}$$

$$U(h) = b + \frac{1}{1+r}\int_{\underline{w}}^{\overline{w}} \max\{U(h'), W(w,h')\}\, dF(w)$$

$$s.t. \qquad h' = (1-\delta)h$$

respectively. The value of employment is simply the discounted present value of earnings, wh; the value of unemployment is given by the unemployment benefit plus the discounted future value of search with the lower skill levels $(1-\delta)h$. At the reservation wage, $w^*(h)$, the values of employment and unemployment are equalized, $U(h) = W(w^*(h), h)$. Standard algebra yields the following characterization of the reservation wage,

$$w^*\left(\frac{r+\delta}{1+r}\right) = r\frac{b}{h} + \frac{(1-\delta)}{1+r}\int_{w^*}^{\overline{w}} [1-F(w)]\, dw. \tag{32}$$

Ljungqvist and Sargent model the increased turbulence in the economy as a rise in the "skill obsolescence" parameter δ. The introduction of a new technological paradigm, or an acceleration in the rate of technological change, can lead to a higher rate of obsolescence, insofar as skills are at least partly technology-specific (recall our discussion in Section 7).

It is straightforward to show, through simple comparative statics, that w^* falls with δ: a worker aware that her skills will become obsolete faster during unemployment chooses optimally to reduce her time spent searching and decreases her reservation wage. As a result, the unemployment duration falls.

However, an increase in δ has an equilibrium effect on the distribution of workers across skills: the average skill level in the population falls, and one can show that the reservation wage declines in the skill level, i.e., $dw^*/dh < 0$. The key behind this result is that the unemployment benefits, b, do not depend on the *current* skill level, h, of the unemployed workers, whereas wage offers are naturally linked to h. A fall in h worsens the value of the average wage offer relative to the value of remaining unemployed with benefits b. Thus, both the reservation wage and unemployment duration increase.[82] The net effect of these two forces is qualitatively ambiguous, and only a quantitative analysis can determine which force is paramount. Note that it is easy to show that the derivative, dw^*/dh, is increasing (in absolute value) in b. Thus, in Europe-like economies with more generous benefits, the second effect tends to be stronger.

Ljungqvist and Sargent embed this simple mechanism in a much richer and detailed model. They calibrate the increase in turbulence to reproduce average earnings losses upon separation of the size estimated in the labor economics literature and show that in economies with generous welfare state (high b), the rise in microeconomic uncertainty brings about a surge in unemployment comparable to the one observed in continental Europe, with all the increase explained by longer durations, as the data suggest. In a "laissez-faire" economy with low b, the faster rate of skill obsolescence barely has any effect.

A related explanation is set forth by Marimon and Zilibotti (1999). In their model, unemployment insurance has the standard result of reducing employment, but it also helps workers find a suitable job. They construct two artificial economies which only differ by the degree of unemployment insurance and assume that they are hit by a common technological shock which enhances the importance of "mismatch". This shock reduces the proportion of

---

[82] One can easily generalize the model to allow b to depend on past earnings (thus, on *past* skills when employed) and the mechanism described would still be in place. This is what Ljungqvist and Sargent do.

jobs which workers regard as acceptable in the economy with unemployment insurance, and unemployment doubles in the Europe-like economy.

In the Ljungqvist-Sargent and Marimon-Zilibotti frameworks, the shock-policy interaction operates entirely through the *labor supply* side. These authors essentially argue that unemployment in Europe went up because, for the jobless, it was more beneficial to collect unemployment insurance than to work at a low wage, given that technological change made their skills obsolete (or made it difficult to use them on the current jobs).

8.3 Slowdown in Total Factor Productivity

A decline of TFP growth rates, such as measured for the United States and Europe after the mid-1970s (see Section 2) can reduce employment in a matching framework through the standard "capitalization effect." Consider the decision of a firm to create a job: the firm will compare the set-up cost with the discounted present value of profits. In a growing economy, where technical change is disembodied and benefits all firms equally, a productivity slowdown increases the "effective rate" at which profits are discounted and discourages the creation of new jobs (Pissarides 2000).

den Haan et al. (2001) evaluate this explanation quantitatively within the context of a standard matching model, à la Mortensen-Pissarides (1998). They find that for this channel to have a significant effect on unemployment, one needs to put restrictions on the shape of the cross-sectional distribution of firms' productivities. Since useful data to test these restrictions are scant, the mechanism remains largely unexplored.

Interestingly, in the same paper the authors argue that once the Ljungqvist and Sargent mechanism is embedded into a model with endogenous job destruction, the comparative statics for increased turbulence are reversed, i.e., unemployment falls. The reason is that as the speed of skill obsolescence rises, workers become more reluctant to separate, and job destruction falls.[83] This force dominates the effect described in the previous section.

---

[83] Recall that the original Ljungqvist and Sargent model is a standard search model where separations occur exogenously. Hence, workers are unable to respond to a negative shock hitting their job. In a matching model with wage bargaining, the workers can allow the firm to keep a larger fraction of output in order to avoid a separation in the face of a shock.

Ljungqvist and Sargent (2003) counter-argue that such an economic mechanism would be relevant only if every worker who separates (including those who quit voluntarily) were hit by faster skill obsolescence. In their view, a more reasonable assumption is that only the workers who suffer an exogenous layoff see their skills decreasing, in which case the original result in Ljungqvist and Sargent (1998) remains intact.

8.4 Acceleration in Capital-Embodied Technical Change

Several measures of embodied technical change suggest that the rate of technical change accelerated around the mid-1970s in the U.S. economy (see Section 2, especially Figure 2). A recent OECD study (Colecchia and Schreyer 2002) measures the decline in relative price for several high-tech equipment items across various countries in Europe from 1980 to 2000 and concludes that European countries experienced an acceleration quantitatively comparable to the United States. Jorgenson (2004, Table 3.5) measures the growth in the quality of the aggregate stock of capital across some OECD countries and finds that, even though the United States had the fastest average annual growth (1.5 percent from 1980 to 2001), Germany and Italy were quite close, with 1.3 percent and 1.1 percent annual growth rates, respectively.

Hornstein et al. (2003a) study precisely whether the interaction between an acceleration in capital-embodied growth, common between the United States and Europe, and certain labor market institutions whose strength differs between the United States and Europe, can explain the simultaneous evolution of the three dimensions of labor market inequalities quantitatively: the unemployment rate, the labor share, and wage inequality.

Their environment builds on the matching model with vintage capital developed by Aghion and Howitt (1994).[84] Consider a continuous-time economy populated by a stationary measure one of ex-ante equal, infinitely lived workers who supply one unit of labor inelastically. Workers are risk-neutral and discount the future at rate r. Production requires

---

[84] For expositional purposes, we simplify the framework in Hornstein et al. (2003a) substantially here. In particular, in the equilibrium of the original model, there are vacant firms with old vintages of machines, while here we make the standard assumption of matching models that all vacant firms embody the leading-edge technology.

one machine and one worker. Machines are characterized by their age, a, translating into match productivity, $e^{-\gamma a}$, where γ is the rate of technological progress embodied in capital.[85]

At any time firms can freely enter the market and post a vacancy at a cost, κ. Then they proceed to search for a worker in a frictional labor market governed by a standard constant-returns-to-scale matching function. Once matched, they produce and share output with the worker in a Nash fashion, with ξ denoting the bargaining power of the worker. At age $\bar{a}$ (determined endogenously), capital is scrapped and the job is destroyed.[86] Two key labor market policies are modeled explicitly: unemployment benefits b, and an employment protection system that combines a hiring subsidy T and an equal-size firing tax upon separation.

As is standard in this framework, it is possible to reduce the equilibrium of the model to two key equations—the job creation condition and the job destruction condition—in two unknowns, θ and $\bar{a}$. These equations, respectively, read

$$\begin{aligned} \kappa &= q(\theta)(1-\xi)\,S(0;\bar{a}) \\ e^{-\gamma\bar{a}} &= b + p(\theta)\xi S(0;\bar{a}) - (r-\gamma)\,T. \end{aligned} \tag{33}$$

Here, $q(\theta)$ and $p(\theta)$ are the meeting probabilities for firms and workers, respectively, expressed as a function of the vacancy-unemployment ratio, θ. We denote by $S(0,\bar{a})$ the "surplus" of a match of age 0, conditional on destruction taking place at age $\bar{a}$: the surplus is the value of the relationship for the parties (the discounted present value of the output stream), net of their outside options. Clearly the surplus is increasing in $\bar{a}$, as a longer match yields a bigger surplus.

The job-creation curve states that vacancies are created (and $q(\theta)$ falls) until the expected return of the marginal vacancy equals its cost, κ. The job-destruction curve states that, at age $\bar{a}$, the pair is indifferent between continuing operating the machine, which gives output

---

[85] As usual, we normalize all variables concerning a vintage a machine relative to the corresponding variable of the newest machine.

[86] Productivity improvements enter the economy only through new capital. This is the typical Schumpeterian "creative-destruction" mechanism, which is at the heart of unemployment in this class of models. As mentioned in section 7.1, Hornstein et al. (2003b) show that if one takes the view that technical progress can also benefit old machines, i.e., if old machines can be "upgraded" into new ones at a cost, then the model yields quantitatively similar results.

$e^{-\gamma \bar{a}}$, and separating, which yields the respective outside options for worker and firm (zero in equilibrium for the firm, because of free entry), net of the firing tax.

FIGURE 6

Figure 6 depicts the comparative statics of a rise in γ in a rigid economy (high b, high T) and in a flexible economy ($b = T = 0$) in the $(\theta, \bar{a})$ space.[87] Note that a low value for $\bar{a}$ corresponds to high separation rate and unemployment incidence, whereas a low value for θ corresponds to long unemployment durations. Thus, the two axes depict the two dimensions of equilibrium unemployment. To illustrate the result more sharply, we have chosen values for b and T in the rigid economy such that the initial equilibrium in the two economies is the same. This is possible since generous benefits and strict employment protection have offsetting effects on job destruction, while they are neutral on job creation, as evident from the equations in (33). The model is therefore consistent with an initial situation where, originally, the labor markets of the United States and Europe looked alike, as the data for the 1960s show. Figure 6 illustrates that a rise in γ has a dramatically different impact across the two economies, especially regarding the amplitude of the shift in the job destruction curve.

To understand intuitively the economic forces at work, it is useful to think of the acceleration in equipment-embodied technology as an “obsolescence shock.” As the rate of productivity growth of new capital accelerates, existing capital-worker matches—which have old vintages of capital—become obsolete faster. In the United States, this loss of economic value is to a higher extent borne by workers, whose wages fall in order to keep firms from scrapping capital and breaking up earlier to invest in better machines.

In Europe, however, labor payments are kept artificially high by generous unemployment benefits and by rents on firing costs, which make wages downwardly rigid. As a result, firms must bear the initial adjustment by destroying matches earlier and creating fewer jobs. The corresponding sharp increase in unemployment greatly improves the relative bargaining

---

[87] Once we recognize that $p'(\theta) > 0$, $q'(\theta) < 0$, and $S'(\cdot, \bar{a}) > 0$, understanding the slope of the two curves in (33) is immediate.

position of firms, which can now push workers closer to their outside option, thus reducing the labor share of output. Since the outside option is constant across all workers, this force also limits the rise in wage inequality that comes about with faster technical change because of larger productivity differentials across machines.[88] Thus, in response to a technological acceleration, an economy with rigid, European-like institutions would experience a higher unemployment rate, a more pronounced decline in the labor share, and a slower rise in wage inequality than would be observed in a more flexible economy.

Quantitatively, a permanent rise in the rate of capital-embodied productivity growth of the magnitude observed in the data can replicate a large fraction of the differential increase in the unemployment rate and of the capital share between the flexible U.S. economy and the rigid Europe-like economy (with the increase in unemployment taking place along the duration margin, as in the data). Wage inequality increases in the U.S.-like economy and declines in the rigid economy, but the changes generated by the model are rather small (recall our discussion of section 7.1).

8.5 Skill-Biased Technical Change

A number of explanations for the rise in wage inequality in the United States—many of which we have reviewed in Section 3—build on the idea of skill-biased technical change. Could this type of technological advancement, interacted with more rigid institutions, also be at the origin of the rise in European unemployment?

Mortensen and Pissarides (1999) explore this question in a model where the economy is populated by a finite number of types of workers, ex-ante different in their skill (productivity level). Skill is observable (e.g., education), and all workers endowed with the same skill level are segmented in their own labor market, which is modeled as frictional with a standard matching function governing the meeting process.

In this model, unemployed workers receive welfare benefits which are partly proportional to their wage (and skill level), and partly lump-sum. The equilibrium unemployment is

---

[88] This mechanism, which is based on technological heterogeneity and the existence of quasi-rents for workers, is the same as that analyzed in Section 7.2.

decreasing and convex in the skill level: low-skill markets have higher unemployment, as benefits represent a form of wage rigidity that is more binding at low levels of skills. As benefits become more generous, the convexity becomes more pronounced.

The skill-biased shock is introduced as a mean-preserving spread in the skill distribution, calibrated to match the rise in wage inequality in the economy with low benefits, like the United States. The model predicts a sharp surge in unemployment in the economy with generous benefits, due to the convex equilibrium relationship between unemployment and the skill level. A crucial ingredient of the Mortensen-Pissarides mechanism, which is present also in the Hornstein et al. (2003a) setup, is that welfare benefits are not fully proportional to wages and productivity; rather, they have a "flat", lump-sum component. If they were fully proportional, every skill market would just be a rescaled version of the highest-skill market, with the same unemployment rate. Hansen (1998) studies the institutional details of the welfare state in several European countries and argues that flat "social assistance" benefits are an important component of these welfare systems.[89]

Finally, the Mortensen and Pissarides model has the counterfactual implication that the rise in unemployment is concentrated among the low-skilled workers, whereas Nickell and Bell (1996) and Gottschalk and Smeeding (1997), among others, conclude that data from many European countries support the conclusion that unemployment rose proportionately across the entire skill spectrum.

8.6 Endogenous Technology Adoption

A careful look at Table 1 shows a non-monotonic evolution of the labor share in many European countries: the labor share rose between 1965 and 1980, only to decline sharply afterwards. In some countries this pattern is striking. In Portugal, for example, the labor share skyrocketed from 56 percent to 75 percent in the period 1965-1980, and then plunged

---

[89] Another key assumption is that markets are segregated across skills. This setup allows the model to avoid the criticism by Wong (2003) that skill-biased technical change can reduce the amount of mismatch and decrease within-group inequality. Mortensen and Pissarides describe their workers as differing in educational attainment. Given the observability of education, modeling firms as able to direct their search (and segregate the economy) seems appropriate.

to 68 percent by 1995. Blanchard (1997) and Caballero and Hammour (1998) proposed an explanation for these dynamics based on the idea that technological advancement responds to the relative cost of factor inputs.

The late 1960s and early 1970s witnessed a rapid evolution of capital-labor relationships in favor of labor in many European countries: "pro-labor" measures were introduced with the objective of consolidating unions' power, increasing the generosity and coverage of unemployment benefits, making economically motivated dismissals harder to justify.[90] The result was, in the language of Caballero and Hammour, an "appropriability shock" that shifted bargaining power away from capital.

In a model where the technological menu for capital-labor substitutability is fixed in the short run, but endogenous in the long run, one will observe an initial rise in the labor share as a result of such a shock. However, as time goes by, more and more firms respond to this institutional "pro-labor" push by introducing new technologies that substitute capital for labor. Therefore, in the long run the capital-labor ratio rises, and both the labor share and employment decline, as observed in the last two decades in Europe.

Why do the U.S. data not display the same pattern? According to Blanchard (1997), since the initial appropriability shock was much smaller, so was the response of capital. A natural question arises, if one follows this logic through: is it only a coincidence that the technological change away from unskilled labor was biased towards capital in Europe and toward skilled labor in the United States?

According to Acemoglu (2003a) the direction of the bias in technological innovations is endogenous (see also our discussion in Section 3) and institutional differences can be key in explaining different biases between the United States and Europe. Consider a flexible economy, like the United States, where firms can either produce with one unit of skilled labor with productivity h_s or one unit of unskilled labor with productivity h_u, where $h_u < h_s$. Output is $y_i = h_i$, with $i = u, s$, and the wage paid to the worker is simply a fraction, ξ, of output. Firms can also choose to pay a fixed cost, κ, and adopt a new technology that

---

[90] The "French May" in 1968 and the "Hot Italian Autumn" of 1969 are stark manifestations of the power of the labor movements in that period of European history.

increases output by a factor $1 + A < h_s/h_u$. Consider first an equilibrium where

$$\kappa > (1 - \xi) A h_s > (1 - \xi) A h_u,$$

so that it is not profitable for firms to implement the new innovation, and wage inequality is simply given by $w_s/w_u = h_s/h_u$. Suppose now that, due to technological progress, the cost of capital decreases to a new value, $\kappa' < \kappa$, such that it is always profitable to adopt it for skilled workers, but not for unskilled workers, i.e.,

$$(1 - \xi) A h_s > \kappa' > (1 - \xi) A h_u.$$

As a result of this adoption decision, wage inequality jumps to the higher level $(1 + A) h_s/h_u$ in the U.S. economy.

Consider now an alternative economy, like Europe, where, because of some institutional constraint, wages cannot fall below a fixed level, $\bar{w}$, where $\xi h_s > \bar{w} > \xi (1 + A) h_u$, so that the constraint is binding for the unskilled workers, even in the case of adoption, but never for the skilled workers.

Whenever the new cost level, κ', satisfies

$$A h_u > \kappa' > (1 - \xi) A h_u$$

in Europe, the new technology will be adopted also with unskilled workers; this is an effect of the minimum wage constraint. The intuition for this result is that, since firms in Europe pay a fixed wage $\bar{w}$ to the unskilled workers, whether or not they adopt the new technology, the institutional constraint makes the firm the residual claimant on output, once $\bar{w}$ is paid. The new technology increases output without changing the wage payment, and thus it may be optimal to adopt in an economy with wage rigidity and not to adopt in an economy with wage flexibility, with the obvious implication that inequality will not increase in Europe.[91]

Formalized models, where the direction of technical change is endogenous, are still in their infancy: in the case of this application to the U.S.-Europe comparison, one important

---

[91] This hypothesis is also consistent with the fact that, at least until the impressive productivity surge of 1995-2000, labor productivity grew faster in Europe (e.g., in France, Germany, Italy, and the United Kingdom) compared to the United States (Jorgenson 2004, Table 3.16).

extension would be verifying if this result survives when the "institutional wage rigidity" is endogenized so that it can respond to changes in the technological environment.

8.7 Sectoral Transformation

The standard approach to the U.S.-Europe differentials is built on comparing the diverging dynamics in the *unemployment* rate. Rogerson (2004) argues that the analysis of relative unemployment rates is misleading, and if one focuses instead on *employment-population ratios*, new insights surface.

In particular, Rogerson shows two new features of the data: 1) the relative deterioration of European employment starts as early as in the 1950s, whereas unemployment rates start diverging in the mid-1970s; 2) the deterioration of European unemployment is largely explained by the differential in manufacturing employment growth.[92]

These facts lead Rogerson to focus on the importance of the structural transformation occurring in the economy, i.e., the secular pattern of reallocation of resources across broad sectors of the economy: first from agriculture to manufacturing, and then from manufacturing to services. Expressed in terms of the shocks-institutions paradigm that we have highlighted in this section, the relevant shock is the transformation of modern economies into service-driven economies, and the relevant institutions are those which hampered the full development of a service sector in Europe.

Although this new approach is still in its infancy, and as such it lacks a quantitative assessment within a rigorous equilibrium model, it appears to be quite promising.

8.8 Discussion

Nickell and Layard (1999), in a widely cited piece in the most recent edition of the *Handbook of Labor Economics,* carefully review the empirical literature and conclude that time spent worrying about the effects of several labor market institutions on cross-country unemployment differentials is largely wasted, given these effects seem small and are often even

---

[92] The concept of "relative deterioration" refers to the difference between the U.S. variable (employment rate or unemployment rate) and its European counterpart.

ambiguous in sign. From the perspective of the research surveyed in this section, however, it seems that when institutional differences are studied in conjunction with technological change, the results are more encouraging.

Of course, much is still far from being well understood. First, once we recognize that the interaction of shocks and institutions is important, what are the key common shocks and the crucial institutional differences that can account for the facts? One would, for example, like to see a unified structural equilibrium framework where several shocks and institutions are jointly analyzed in order to investigate which shock-policy interaction is quantitatively important and which is not.

Second, in answering this question, more "discipline" is needed in the quantitative analysis. Often, the approach in the literature is to calibrate the shock by matching either the rise in wage inequality or the fall in the labor share. We maintain the view that changes in employment/unemployment, wage inequality, and income shares are intimately related and must be explained jointly: they are dimensions along which the model should be evaluated rather than calibrated. Thus, the shock should be calibrated, as much as possible, using independent observations. The use of data on technological change such as that for the relative price of equipment goods is such an example.

Third, it is important to note that we are not aware of any quantitative model of a rigid Europe-like economy that can generate a rise in equilibrium unemployment which is similar across all skill levels, which is what the data suggest.

Fourth, the literature is split between labor-supply models (Ljungqvist and Sargent; Marimon and Zilibotti) and labor-demand models (Bertola and Ichino; Caballero and Hammour; Hornstein et al.). Obviously, interpreting the European and U.S. labor market outcomes in terms of "labor demand" or "labor supply" is not mutually exclusive. In a theoretical framework with elements of vintage human capital and vintage physical capital, an embodied technological acceleration will also worsen the rate of skill obsolescence—exactly as in Ljungqvist and Sargent's paper. The next generation of investigations of the European (un-)employment puzzle should bring together supply and demand forces and allow a joint evaluation of their respective strength.

9 Welfare and Policy Implications

In traditional growth theory, technological progress is largely associated with productivity advancements, reflected in improvements in average wages, from which it would follow that there are welfare gains. While the first generation of growth models is based on the representative-agent assumption, the model economies we studied in this chapter are built on a heterogenous-agents model. By raising the wage differential between more and less skilled workers (between-group inequality) and by amplifying the amount of labor market uncertainty faced by ex-ante equal households in the economy (residual inequality), in these economies technological change can lead to welfare costs, at least for certain groups of workers, and it has first-order implications for policy. In what follows we give an account of some early work on the subject.

9.1 Lifetime Earnings Inequality

The majority of the empirical investigations on rising inequality in the United States focus on the cross-sectional distribution of wages and earnings. Friedman (1982) argues that data on cross-sectional inequality at a point in time are difficult to interpret, as they provide no information on the degree of economic mobility: the same distribution can be generated either by a "dynamic society" or by a "status society".

A better measure of inequality, which incorporates some of Friedman's concerns, is provided by the distribution of lifetime earnings. A stark example of the pitfalls implicit in making welfare and policy statements simply based on distributions at a point in time is provided by Flinn (2002). Flinn compares Italy and the United States and documents that, although the dispersion in cross-sectional yearly earnings inequality in the United States is several times larger than in Italy, the distribution of lifetime earnings in the United States is more compressed due to larger individual variability of labor income and shorter duration of non-employment experiences. In other words, in Friedman's language, Italy somewhat surprisingly looks more like a "status society" than does the United States.[93]

Two papers, so far, have studied the change in the distribution of lifetime earnings in the United States in the past three decades through the lenses of a structural model.[94] Heckman et al. (1998) solve a deterministic competitive OLG model with endogenous human capital accumulation to study the implications of the widening educational premium for lifetime-earnings inequality across cohorts.[95] Their model implies that the low-educated cohorts entering in the mid-1980s are those suffering the largest drop in lifetime earnings from skill-biased technical change: roughly 11 percent. At the same time, they calculate a rise in lifetime earnings of 6 percent for the college graduates in the same cohort.

Similarly, Bowlus and Robin (2003) use a search model with risk-neutrality, estimated on matched CPS data from 1977 to 1997, to study how changes in wage and employment dynamics over the past thirty years have affected the evolution of lifetime labor income inequality in the U.S. labor market. They find that the median worker suffered only a small decline in present value lifetime earnings, but that there is large heterogeneity across educational groups with lifetime earnings declining by over 25 percent for high-school graduates and increasing by almost 20 percent for college graduates.

These numbers are over twice as large as those in Heckman et al. (1998). One reason is that Heckman et al. model the acquisition of education and the costs associated with schooling explicitly. A large fraction of the changes in lifetime earnings is attributable to the surge in the returns to education: since education in reality is the outcome of a costly investment choice, the difference in earnings alone likely overstates the true welfare differential between the two groups in the analysis of Bowlus and Robin.

---

[93] Cohen (1999) performs a similar exercise between the United States and France and finds that, using annual wages, inequality in the United States is 60 percent greater than in France, but based on lifetime earnings, the difference reduces to 15 percent.

[94] See Aaronson (2003) for a measurement of changes in lifetime earnings inequality not based on a structural model.

[95] We have discussed a simple version of this model in Section 4.

9.2 Consumption Inequality

There is a definite gain in moving from studying hourly wages to lifetime labor income, if one wants to make inference on welfare. However, one important limit of the studies above is that they effectively assume complete insurance against those transitory income fluctuations that cancel out in the long run and thus do not affect lifetime income. With imperfect insurance against labor market risk, consumption is not determined only by purely permanent shocks that translate one-for-one into permanent income, but the degree of earnings variability and its persistence become important, too. In this sense, consumption is an even better measure of welfare than lifetime earnings.

The evidence based on *Consumption and Expenditure Survey (CEX)* data suggests that consumption inequality rose slightly during the first half of the 1980s (Cutler and Katz, 1992, and Johnson and Shipp, 1997) and has remained roughly stable thereafter (Krueger and Perri, 2002). Interestingly, Blundell and Preston (1998) document that in Britain, where the increase in wage inequality followed a pattern similar to the United States, the rise in consumption inequality was also strong until the early 1980s, but weaker afterwards. This path of consumption inequality is, at first sight, puzzling, especially since wage inequality keeps increasing in the 1990s, albeit at a slower pace. Three explanations for this puzzle have been provided so far.

Krueger and Perri (2002) developed the first formal model to solve this apparent puzzle. They consider an Arrow-Debreu economy with limited enforcement of contracts (Kocherlakota, 1996). In this economy, the degree of insurance market completeness is endogenous and responds to changes in income risk: as income shocks become larger and more persistent, the value of autarky declines, so agents are willing to enter more often into risk-sharing agreements. The central message of Krueger and Perri is that the rise of labor market inequality led to a development in financial markets—in particular the sharp expansion of consumer credit in the 1990s—and to a larger extent of risk sharing, limiting the rise in consumption inequality in this period.

Heathcote et al. (2003) offer an alternative interpretation for this pattern of rising and then flattening consumption inequality. Through a statistical decomposition of the rise in wage dispersion into permanent and transitory components, they conclude that the relative importance of the two components changes substantially over the sample period. From the late 1970s to around 1990 the permanent component increases sharply, but in the 1990s it ceases to grow, whereas there is a substantial increase in the variance of transitory shocks. A standard overlapping-generations model with "exogenously" incomplete-markets (Huggett 1996) predicts a trajectory for consumption inequality similar to the data: as the shocks become more transitory, they are easier to insure and tend to have a smaller impact on consumption. The finding that the first phase of the rise in inequality (1980s) had a more permanent nature than the second (1990s) is common to a number of empirical studies (Moffitt and Gottschalk, 1994, for the United States, and Dickens, 2000 and Blundell and Preston, 1998 for the United Kingdom). To our knowledge, there is no attempt to link this pattern of persistence with the nature of technical progress.

The third explanation is provided by Attanasio et al. (2003) who argue that once measurement error in the CEX data is properly taken into account, consumption inequality keeps rising also in the 1990s, and, hence, there is no puzzle.

9.3 Welfare Implications

Studying consumption inequality is a further improvement toward the understanding of the welfare costs of rising inequality, but a complete welfare analysis cannot abstract from leisure.

One approach that has been taken in the literature makes minimal assumptions regarding the structure of the underlying economic model. Krueger and Perri (2003), in an exercise similar in spirit to that in Attanasio and Davis (1996), estimate a stochastic process directly on consumption and leisure data from the CEX and use standard intertemporal preferences to compute the welfare costs of rising inequality. The computation of welfare losses "under the veil of ignorance," i.e., before the worker finds out whether she will be high- or low-skilled, yields numbers between 1 percent and 2 percent, with a difference in the welfare losses between the 90th percentile (net winners) and the 10th percentile (net losers) of just over 10 percent. To put this number in perspective, the estimate of Bowlus and Robin (2003) is 50 percent.

This approach is based entirely on revealed preferences, and has the advantage that no restrictive assumptions have to be made on the degree and the nature of market completeness. However, without a structural model (like those of Bowlus and Robin, 2003; and Heckman et al. 1998), strong faith must be placed in the reliability of the consumption and hours data from the CEX. In particular, if there are large transitory measurement errors, then one would overestimate the extent of economic mobility and underestimate the welfare losses coming from the change in the wage structure. Moreover, all that can be assessed through this methodology is the welfare cost of changes in consumption and leisure inequality, without knowing exactly what fraction of these changes are attributable to rising wage inequality rather than, for example, tax reforms or changes in financial and insurance markets that occurred over the same period.

A second approach, developed by Heathcote et al. (2003), builds on three steps: 1) an estimation of the dynamics of permanent and transitory components of individual wages over the period of interest, 2) a calibration of an OLG model with endogenous leisure and consumption choices and incomplete markets, 3) simulation of the model to compute the welfare costs of the changes in wage dynamics. This approach, thus, is fully structural, and, as such, it does not rely heavily on survey data on consumption and hours worked. Rather, welfare calculations are based on the changes in the model-generated consumption and leisure paths due exclusively to observed and well-measured changes in the wage process over the period. At the same time, it incorporates a realistic range of insurance avenues (a saving technology, labor supply, and social security) without going as far as assuming complete markets.

According to the calculations of Heathcote et al., (2003), welfare losses "under the veil of ignorance," although varied by cohort, average 2.5 percent across all cohorts, with a peak of 5 percent for the cohorts entered in the mid-1980s. The low-skill workers suffer a loss of 16 percent, and the high-skill workers enjoy a welfare gain of 13 percent. These numbers fall in between the estimates of Bowlus and Robin (2003) and those of Heckman et al. (1998).

Two main conclusions emerge. First, the welfare consequences of the observed rise in labor market risk are quite different across groups of workers: whereas the high-skill, high-educated workers are the winners, the low-skill, low-educated workers are the losers. Second, the ex-ante welfare loss from the rise in labor market risk in the United States is of the order of 2 percent of lifetime consumption, which is a rather large number.

9.3.1 Insurance and Opportunities in the Welfare Analysis of Wage Inequality

The quantitative studies on the welfare consequences of the recent rise in inequality point to a sizeable welfare loss. But does the absence of full insurance always imply a welfare decrease when risk increases? The answer is no. We have already mentioned the case studied by Krueger and Perri (2002) where, with endogenous market incompleteness, a rise in uncertainty can lead to more risk sharing in society and increase welfare. The same result can arise for different reasons in models where the extent of risk-sharing is limited exogenously (Bewley-Aiyagari economies). Consider, as do Heathcote et al. (2004), an economy populated by a measure one of infinitely-lived agents with preferences

$$U = E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_t^{1-\gamma} - 1}{1-\gamma} - \varphi \frac{h_t^{1+\sigma}}{1+\sigma} \right], \tag{34}$$

where $1/\gamma$ is the elasticity of intertemporal substitution and $1/\sigma$ is the Frisch elasticity of labor supply. Each agent i starts with zero wealth and faces log-normal productivity shocks to the efficiency units of labor ω_{it}. Shocks can be decomposed into two orthogonal components:

$$\ln \omega_{it} = \alpha_i + \varepsilon_{it}, \quad \text{with } \alpha_i \sim N\left(-\frac{v_\alpha}{2}, v_\alpha\right), \text{ and } \varepsilon_{it} \sim N\left(-\frac{v_\varepsilon}{2}, v_\varepsilon\right)$$

where α_i is the permanent-uninsurable component and ε_{it} is the transitory-insurable component. Note that the means have been normalized so that a rise in the variance of either component does not affect the average level of efficiency units.

After computing the allocations and substituting them into preferences (34), one can calculate the welfare gain of an increase in the two components of wage uncertainty—expressed as the equivalent consumption variation. The main finding is that one can obtain an (approximate) closed-form expression for the welfare gain: $\mathcal{W}$

$$\mathcal{W} = \frac{1}{\sigma}\Delta v_{\varepsilon} - \frac{(\gamma - 1) + \gamma(1 + \sigma)}{\gamma + \sigma}\Delta v_{\alpha}. \tag{35}$$

This expression only depends on two elasticity parameters (γ, σ) and on the change in the two variances $(\Delta v_{\alpha}, \Delta v_{\varepsilon})$. The key feature to note, in the above equation, is that the welfare gain is not always negative. For example, as $\gamma \to 0$ (risk-neutrality), the welfare gain is positive and proportional to the rise in overall inequality $(\Delta v_{\varepsilon} + \Delta v_{\alpha})$ through the Frisch elasticity.[96]

To understand this result, one has to keep in mind that there are two distinct effects of a rise in labor market uncertainty. On the one hand, "decreased insurance" induces a welfare loss. For example, as risk-aversion rises with γ or as the permanent-uninsurable component v_{α} expands, the second term becomes larger and the overall welfare gain tends to become negative. On the other hand, "improved production opportunities" induce a welfare gain. In presence of elastic labor supply (σ low), households supply more hours when they face a good productivity shock and enjoy leisure at times of low-productivity. When the variance of productivity shocks rises, this intertemporal behavior can improve households' welfare. The net effect depends on the parameterization of preferences and on the empirical assessment of what fraction of the rise in inequality is insurable.

9.3.2 Discussion

Economists have just started to tackle these issues, and many questions still lie ahead. One key area to explore is the role of the family in determining the welfare implications of the rise in wage inequality. Two offsetting forces are at work. First, there is positive assortative matching between spouses along the skill/education dimension. Second, shocks are imperfectly correlated between spouses (Hyslop 2001). While the first feature amplifies the surge in inequality and worsens welfare inequalities across families, the second establishes

---

[96] This qualitative result can be reproduced also starting from Cobb-Douglas preferences, albeit the expression in (35) is different.

a role for intra-family insurance in dampening the rise in labor market risk. Only a careful quantitative analysis can determine which force is dominant.

Finally, the current welfare studies abstract from some first-order "social" consequences of the rise in inequality and the fall in the wages of the unskilled, such as the decline in labor market participation for low-educated males (Murphy and Topel 1997), the rise in the crime rate (Kelly 2000), and the decline in the marriage rate (Gould and Paserman 2003).[97]

9.4 Brief Directions for Policy

Welfare losses originating from the rise in U.S. inequality in the past three decades are almost one hundred times larger than the standard estimates of the costs of business cycles (Lucas 2003). In this sense, policies that act by reallocating risk across agents (like social insurance policies) are a macroeconomic priority compared to policies that reduce the impact of aggregate risk (like monetary or fiscal stabilization policies). But among the myriads of possible government interventions, what are the right redistributive policies?

In Sections 3.2 and 7 we discussed two complementary views of the link between technology and inequality. The first of these views is that technological progress in the past three decades has been complementary to certain permanent individual characteristics, such as ability or education (technology-skill complementarity). The second view is that labor market history is scattered with shocks and stochastic events related to the luck of individuals, firms or industries that determine the degree of fanning out of the skill and earnings distributions among ex-ante equal workers. The rapid diffusion of a new technology amplifies the importance of these stochastic factors, increasing overall earnings instability (technology-luck complementarity).

The emphasis we placed on these two approaches is not just for classification purposes, since they have profoundly different policy implications. Insofar as we are interested in designing policies that reduce inequalities among households, models of technology-ability complementarity suggest that the intervention should be targeted early in the life of an

---

[97] Gould and Paserman argue that the higher male inequality in the United States increased the option value for single women to search longer for a husband.

individual, possibly during childhood when the key components of learning ability are being formed. Models of technology-luck complementarity seem to call for interventions that allow the disadvantaged (or unlucky) workers to rebuild their skill level after a shock, such as displacement due to skill obsolescence, has hit.

Examples of both types of policies are abundant in the U.S. economy.[98] In general, the most recent evaluations of programs entailing expenditures and treatment at early childhood report remarkable success. In contrast, the available evidence indicates that welfare-to-work and training programs directed toward adult workers are rather inefficient, as they generate only modest increases in permanent earnings levels (LaLonde et al. 1999).

According to Heckman (2000), the reason for the divergence in outcomes across these two classes of policies is twofold. First, investments in human capital at old ages are less efficient, since the elderly worker has less time to recoup the investment; second, "learning begets learning," so human capital, skills, and abilities acquired at a young age facilitate future learning.

In this sense, policymakers should have a life-cycle perspective: lifting the unskilled, displaced adults into skilled status is much easier and more efficient if the same workers have been developing their learning ability throughout childhood and youth, possibly with the help of government intervention. For the more mature low-skilled workers with limited learning ability who are subject to unavoidable wage losses due to biased technological change, targeted wage subsidies can be more effective than retraining programs.

10 Concluding Remarks

This chapter argues that labor market inequalities are shaped by technological change through a variety of economic mechanisms. Within the technology-labor market nexus, however, which of the specific mechanisms we evaluate are most likely to survive the test of

---

[98] Programs like the Perry Pre-School program and the Syracuse Pre-School program provide intense family development support to disadvantaged children at very young ages (from birth to 5 years). The Harlem program ensures frequent individual teacher-child sessions for children of age 3-5. Several programs for adult retraining of displaced workers were initiated throughout the United States under the Job Training Partnership Act of 1982 and the Economic Dislocation and Worker Adjustment Assistance Act of 1988.

time?

Before answering this question, it is useful to put things in perspective and recall that most of the statements we have made in this chapter are not all meant to represent general insights; rather, they allude to a particular historical episode. Specifically, technology has not always been skill-biased in the past: the transformation from artisanal workplaces to the factory in the 19th century had much the opposite effect (Goldin and Katz, 1998). Moreover, not all the drastic productivity advancements in the past were embodied in equipment: electricity was to a large extent embodied in new structures, as the electrification of production required a whole new blueprint for the plant (Atkeson and Kehoe, 2002). Even in reference to this particular historical episode, there are serious dissenting views on the overall impact of IT on the macroeconomy (e.g., Gordon 2000) and on the role of technology in explaining the observed changes in the U.S. wage structure (e.g., Card and DiNardo, 2002).

In returning to the original question, we identify three rather general categories that we find particularly interesting and plausible.

The first idea is *factor-specificity* of the recent technological advancements. In particular, the embodiment of productivity improvements in equipment capital goods, and the skill-bias of such productivity improvements. Whether in the Nelson-Phelps version of skills as a vehicle of adoption and innovation, or in the version of skills and capital as complementary in production, the skill-bias of the IT revolution is one of the most robust and pervasive in the literature. Skill-biased technical change and capital-skill complementarity are crucial to explain the climb of the skill premium, notwithstanding the continuous growth in the relative supply of skilled labor. A growing and promising avenue of research is on the endogenous determinants of the factor-bias in technological advancements (Acemoglu 2002b, 2003b).

The second idea is *vintage human capital.* The technological specificity of knowledge appears to be an important idea to explain some of the most puzzling aspects of the data such as the rise in within-group or “residual” inequality, the fall of the real wages at the bottom of the skill distribution, the growth in the returns to experience, and the slowdown of output growth in the aftermath of a technological revolution.

The third idea is the *interaction between technology and the organization of labor markets.* Radical technological developments, like those we have witnessed in the past three decades, are bound to interact deeply with the various aspects of the structure of labor markets, like the organization of production within the firm, labor unions, and labor market policies. Through this interaction, the literature has successfully interpreted the move from the Tayloristic to the flatter multi-tasking organizational design of firms, the decline of unionization, and the upward trend in unemployment rate in Europe. In particular, the comparison of the U.S. and European experiences seems a fruitful way of studying this channel.

These ideas are the building blocks of the most successful and influential papers in the first generation of models that we have surveyed in this chapter. Where will the literature go next? We argued in various parts of the chapter that one major weakness of this literature is the scarcity of rigorous quantitative evaluations of the theories proposed. Most of the papers reviewed are qualitative in nature. This is not too surprising, given the young vintage of the literature (which developed only starting from the mid 1990s), and given that, in any field, it naturally takes a long time before a handful of theoretical frameworks emerge as successful and begin to be used for a systematic quantitative accounting of the facts (e.g., the search and matching model in the theory of unemployment, and the neoclassical and the endogenous growth model in the theory of cross-country income differences). In this chapter we have highlighted some features that seem important for a successful theory of the link between technological change and labor market outcomes. Quantitative theory should be a priority within this field of research over the years to come.

References

[1] Aaronson, S. (2003), "The Rise in Lifetime Earnings Inequality Among Men," mimeo, Federal Reserve Board.

[2] Abowd J.M., F. Kramarz, and D.N. Margolis (1999), "High Wage Workers and High Wage Firms," *Econometrica* 67, 251-334.

[3] Abraham A. (2003), "Wage Inequality and Education Policy with Skill-biased Technological Change in OG Setting ," mimeo, Duke University.

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

[5] Acemoglu, D. (1999), "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence," *American Economic Review* 89, 1259-78.

[6] Acemoglu, D. (2002a), "Technical Change, Inequality and the Labor Market," *Journal of Economic Literature* 40, 7-72.

[7] Acemoglu, D. (2002b), "Directed Technical Change," *Review of Economic Studies* 69, 781-810.

[8] Acemoglu, D. (2003a), "Cross-Country Inequality Trends," *Economic Journal* 113, F121-149.

[9] Acemoglu, D. (2003b), "Labor- and Capital-Augmenting Technical Change," *Journal of the European Economic Association* 1, 1-37.

[10] Acemoglu, D. (2003c), "Patterns of Skill Premia," *Review of Economic Studies* 70, 199-230.

[11] Acemoglu, D., and J.-S. Pischke (1999), "The Structure of Wages and Investment in General Training," *Journal of Political Economy* 107, 539-72.

[12] Acemoglu, D., P. Aghion and G.L. Violante (2001), "Deunionization, Technical Change, and Inequality", *Carnegie-Rochester Conference Series on Public Policy* 55, 29-64.

[13] Addison, J.T., R.W Bailey and W.S. Siebert (2004), "The Impact of Deunionization on Earnings Dispersion Revisited", Departmental Working Paper 14172, Department of Commerce, University of Birmingham.

[14] Aghion, P. (2002), "Schumpeterian Growth Theory and the Dynamics of Income Inequality," *Econometrica* 70, 855-82.

[15] Aghion, P. and P. Howitt (1994), "Growth and Unemployment," *Review of Economic Studies* 61, 477-494.

[16] Aghion, P. and P. Howitt (1998), Endogenous Growth Theory, (MIT Press: Cambridge and London).

[17] Aghion, P., P. Howitt and G.L. Violante (2002), "General Purpose Technology and Within-Group Wage Inequality," *Journal of Economic Growth* 7, 315-345.

[18] Albrecht, J. and S. Vroman (2002), "A Matching Model with Endogenous Skill Requirements," *International Economic Review* 43, 283-305.

[19] Atkeson, A. and P.J. Kehoe (2002), "The Transition to a New Economy Following the Second Industrial Revolution," NBER Working Papers 8676

[20] Attanasio, O., E. Battistin and H. Ichimura (2003), "What Really Happened to Consumption Inequality in the US?" mimeo, Institute of Fiscal Studies.

[21] Attanasio, O. and S.J. Davis (1996), "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy* 104, 1227-62.

[22] Autor, D., L. Katz and A. Krueger (1998), "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113, 1169-1213.

[23] Bahk, B.H. and M. Gort (1993), "Decomposing Learning by Doing in New Plants," *Journal of Political Economy* 101, 561-583.

[24] Baker, G.P. and T.N. Hubbard (2003), "Contractibility and Asset Ownership: On-Board Computers and Governance in U.S. Trucking," *American Economic Review* 93, 1328-1353.

[25] Bartel A.P. and F.R. Lichtenberg (1987), "The Comparative Advantage of Educated Workers in Implementing New Technology," *Review of Economics and Statistics* 69, 1-11.

[26] Bartel, A. and N. Sicherman (1998), "Technological Change and the Skill Acquisition of Young Workers," *Journal of Labor Economics* 16, 718-755.

[27] Basu, S., J. Fernald, N. Oulton and S. Srinivasan (2003), "The Case of the Missing Productivity Growth: Or, Does information technology explain why productivity accelerated in the United States and not in the United Kingdom?," *NBER Macroeconomics Annual* 18 (MIT Press: Cambridge and London, 9-63.

[28] Beaudry, P. and D.A. Green (2003), "Wages and Employment in the United States and Germany: What Explains the Differences?" *American Economic Review* 93, 573-602.

[29] Becker, G. (1973), "A Theory of Marriage", *Journal of Political Economy* 81, 813-846.

[30] Bentolila, S. and G. Saint-Paul (1999), "Explaining Movements in the Labor Share," *Contributions to Macroeconomics* 3 (1).

[31] Bertola, G. and A. Ichino (1995), "Wage Inequality and Unemployment: United States vs. Europe," in B. Bernanke and J. Rotemberg, eds., *NBER Macroeconomics Annual* 10 (MIT Press: Cambridge and London), 13-54.

[32] Bertola, G., F.D. Blau and L.M. Kahn (1997), "Swimming Upstream: Trends in the Gender Wage Differential in 1980s," *Journal of Labor Economics* 15, 1-42.

[33] Bertola, G., F.D. Blau and L.M. Kahn (2001), "Comparative Analysis of Labor Market Outcomes: Lessons for the United States from International Long-Run Evidence," in: A.B. Krueger and R.M. Solow, eds., The roaring nineties: Can full employment be sustained? (Century Foundation Press: New York), 159-218.

[34] Bils, M. and P. Klenow (2004), "Measuring Quality Change and Externalities," in: P. Aghion and S. Durlauf, eds., Handbook of Economic Growth, (North-Holland: Amsterdam).

[35] Blanchard, O. (1997), "The Medium Run," *Brookings Papers of Economic Activity (Macroeconomics)* 2, 89-141.

[36] Blanchard, O. and J. Wolfers (2000), "The Role of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence," *Economic Journal* 110, C1-33.

[37] Bleaney, M. (1996), "Central Bank Independence, Bargaining Structure, and Macroeconomic Performance in the OECD Countries," *Oxford Economic Papers* 48, 20-38.

[38] Blundell, R. and I. Preston (1998), "Consumption Inequality and Income Uncertainty," *Quarterly Journal of Economics* 113, 603-640.

[39] Bolton, P. and M. Dewatripont (1994), "The Firm as a Communication Network," *Quarterly Journal of Economics* 109, 809-839.

[40] Booth, A. (1995), The Economics of the Trade Union, (Cambridge University Press: Cambridge, UK).

[41] Borghans, L., and B. ter Weel (2003), "The Diffusion of Computers and the Distribution of Wages," mimeo, Maastricht University

[42] Bowlus, A.J. and J.-M. Robin (2003), "Twenty Years of Rising Inequality in U.S. Lifetime Labor Income Values," forthcoming *Review of Economic Studies.*

[43] Braverman, L. (1974), Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century. (Monthly Review Press: New York).

[44] Bresnahan, T.F. and M. Trajtenberg (1995), "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics* 65, 83-108.

[45] Breshnahan, T.F., E. Brynjolfsson and L.M. Hitt (2002), "Information Technology, Workplace Organization and the Demand for Skilled Labor: Firm-Level Evidence," *Quarterly Journal of Economics* 117, 339-376.

[46] Brynjolfsson, E. and L.M. Hitt (2000), "Beyond Computation: Information Technology, Organizational Tranformation and Business Performance," *Journal of Economic Perspectives* 14, 23-48.

[47] Caballero, R.J. and M.L. Hammour (1998), "Jobless Growth: Appropriability, Factor Substitution, and Unemployment," *Carnegie-Rochester Conference Series On Public Policy* 48, 51-94.

[48] Cain, L. and D. Paterson (1986), "Biased Technical Change, Scale, and Factor Substitution in American Industry, 1850-1919," *Journal of Economic History* 46, 153-64.

[49] Campbell, J., M. Lettau, B. Malkiel and Y. Xu (2001), "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk," *Journal of Finance* 56, 1–43.

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

[51] Card, D. (2001), "The Effects of Unions on Wage Inequality in the U.S. Labor Market," *Industrial and Labor Relations Review* 54, 296-315.

[52] Card, D. and J. DiNardo (2002), "Skill Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," NBER working paper 8769.

[53] Card, D. and T. Lemieux (2001), "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis," *Quarterly Journal of Economics* 116, 705-46.

[54] Caroli, E. and J. Van Reenen (2001), "Skill-Biased Organizational Change? Evidence From a Panel of British and French Establishments," *Quarterly Journal of Economics* 116, 1449-1492.

[55] Caselli, F. (1999), "Technological Revolutions," *American Economic Review* 89, 78-102.

[56] Chaney, T., X. Gabaix and T. Philippon (2003), "The Evolution of Microeconomic and Macroeconomic Volatility," mimeo, NYU Stern.

[57] Chun (2003), "Information Technology and the Demand for Educated Workers: Disentangling the Impacts of Adoption versus Use," *Review of Economics and Statistics* 85, 1-8.

[58] Coase (1937), "The Nature of the Firm," *Economica* 4, 386-405

[59] Cohen, D. (1999), "Welfare Differentials Across French and U.S. Labor Markets: A General Equilibrium Interpretation," CEPREMAP working paper 9904.

[60] Cohen-Pirani, D., and R. Castro (2004), "Why Has Skilled Employment Become so Procyclical after 1985?," mimeo Carnegie-Mellon University.

[61] Colecchia, A. and P. Schreyer (2002), "ICT Investment and Economic Growth in the 1990s: Is the United States a Unique Case? A Comparative Study of Nine OECD Countries," *Review of Economic Dynamics* 5, 408-442.

[62] Comin, D. and Mulani (2003), "Diverging trends in macro and micro volatility: Facts," mimeo NYU.

[63] Cooley, T.F. and E.C. Prescott (1995), "Economic Growth and Business Cycles," in: T.F. Cooley, ed., Frontiers of Business Cycle Research (Princeton University Press: Princeton), 1-38.

[64] Cozzi G., and G. Impullitti (2004), "Technology Policy and Wage Inequality", mimeo NYU

[65] Cummins, J. and G.L. Violante (2002), "Investment-Specific Technological Change in the U.S. (1947-2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics* 5, 243-284.

[66] Cutler, D.M. and L.M. Katz (1992), "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980's," *American Economic Review* 82, 546-51.

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

[68] den Haan W.J. (2003), "Temporary Shocks and Unavoidable Transitions to a High-Unemployment Regime," mimeo London Business School.

[69] den Haan, W., C. Haefke and G. Ramey (2001), "Shocks and Institutions in a Job Matching Model", NBER working paper 8463.

[70] Dickens, R. (2000), "The Evolution of Individual Male Earnings in Great Britain: 1975-1995," *Economic Journal* 110, 27-49.

[71] Dickens, W.T., and J.S. Leonard (1985), "Accounting for the Decline in Union Membership: 1950-1980," *Industrial and Labor Relations Review* 38, 323-34.

[72] DiNardo, J., N.M. Fortin and T. Lemieux (1996), "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica* 64, 1001-1044.

[73] Dinopolous and Segerstrom (1999), "A Schumpeterian Model of Protection and Relative Wages, " *American Economic Review* 89, 450-472.

[74] Dooley and Gottschalk (1984), "Earnings Inequality among Males in the United States: Trends and the Effect of Labor Force Growth," *Journal of Political Economy* 92, 59-89.

[75] Duranton, G. (2004), "The Economics of Production Systems: Segmentation and Skill-Biased Change," *European Economic Review* 48, 307-36.

[76] Eckstein, Z. and E. Nagypal (2004), "U.S. Earnings and Employment Dynamics 1961-2002: Facts and Interpretations," mimeo Northwestern University.

[77] Farber, H. and A.B. Krueger (1992), "Union Membership in the United States: The Decline Continues," NBER working paper 4216.

[78] Farber, H.S. and B. Western (2000), "Round Up the Usual Suspects: The Decline of Unions in the Private Sector, 1973-1998," Princeton University, Industrial Relations Sections, working paper 437.

[79] Farber, H.S. and B. Western (2002), "Ronald Reagan and the Politics of Declining Union Organization," *British Journal of Industrial Relations* 40, 385-401.

[80] Flinn, C. (2002), "Labour Market Structure and Inequality: A Comparison of Italy and the U.S." *Review of Economic Studies* 69, 611-45.

[81] Flug, and Hercowitz (2000), "Equipment Investment and the Relative Demand for Skilled Labor: International Evidence," *Review of Economic Dynamics* 3, 461-485.

[82] Freeman, R.B (1988), "Contraction and Expansion: The Divergence of Private Sector and Public Sector Unionism in the United States," *Journal of Economic Perspectives* 2, 63-88.

[83] Freeman, R.B. and E.P. Lazear (1995), "An Economic Analysis of Works Councils," in: J. Rogers and W. Streeck, eds., Works Councils: Consultation, Representation, and Cooperation in Industrial Relations, National Bureau of Economic Research Comparative Labor Markets Series, (University of Chicago Press: Chicago and London), 27-50.

[84] Freeman, R.B. and J.L. Medoff (1984), "What Unions Do: Evidence, Interpretation, and Directions for Research," Harvard Institute of Economic Research Discussion Paper: 1096.

[85] Friedman, M. (1982), Capitalism and Freedom, (Chicago University Press: Chicago).

[86] Galor, O., and O. Moav (2000), "Ability Biased Technological Transition, Wage Inequality Within and Across Groups, and Economic Growth," *Quarterly Journal of Economics* 115, 469-97.

[87] Galor, O., and D. Tsiddon (1997), "Technological Progress, Mobility, and Economic Growth," *American Economic Review* 87, 362-382.

[88] Garicano, L. and E. Rossi-Hansberg (2003), "Organization and Inequality in a Knowledge Economy," mimeo Stanford.

[89] Gittleman, Maury, and Mary Joyce (1996); "Earnings Mobility and Long-Run Inequality: An Analysis Using Matched CPS Data," *Industrial Relations* 35, 180-197.

[90] Goldin, C. and L.M. Katz (1998), "The Origins of Technology-Skill Complementarity," *Quarterly Journal of Economics* 113, 693-732.

[91] Goldin, C. and L.M. Katz (1999), "The Returns to Skill in the United States Across the Twentieth Century," NBER working paper 7126.

[92] Goldin, C. and R. Margo (1992), "The Great Compression: The Wage Structure in the United States at Mid-Century," *Quarterly Journal of Economics* 107, 1-34.

[93] Gordon, R.J. (1990), The Measurement of Durable Good Prices, (University of Chicago Press: Chicago).

[94] Gordon, R.J. (2000), "Does the 'New Economy' Measure Up to the Great Inventions of the Past?" *Journal of Economic Perspectives* 14, 49-74.

[95] Gosling, A. and S. Machin (1995), "Trade Unions and the Dispersion of Earnings in British Establishments, 1980-90," *Oxford Bulletin of Economics and Statistics* 57, 167-84.

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

[97] Gottschalk, P. and T.M. Smeeding (1997), "Cross-National Comparisons of Earnings and Income Inequality," *Journal of Economic Literature* 35, 633-687.

[98] Gould, E.D., O. Moav and B.A. Weinberg (2001), "Precautionary Demand for Education, Inequality and Technological Progress," *Journal of Economic Growth* 6, 285-315.

[99] Gould, E.D. and M.. Paserman (2003), "Waiting for Mr. Right: Rising Inequality and Declining Marriage Rates," *Journal of Urban Economics* 53, 257-281.

[100] Greenwood, J. and B. Jovanovic (1999), “The IT Revolution and the Stock Market,” *American Economic Review* 89, 116-122.

[101] Greenwood, J. and A. Seshadri (2004), “Technological Progress and Economic Transformation,” in: P. Aghion and S. Durlauf, eds., Handbook of Economic Growth, (North-Holland: Amsterdam).

[102] Greenwood J. and M. Yorukoglu (1997), “1974,” *Carnegie-Rochester Conference Series on Public Policy* 46, 49-96.

[103] Greenwood, J., Z. Hercowitz and P. Krusell (1997), “Long-Run Implications of Investment-Specific Technological Change,” *American Economic Review* 87, 342-362.

[104] Greenwood, J., A. Seshadri and M. Yorukoglu (2004), “Engines of Liberation,” *Review of Economic Studies*, forthcoming.

[105] Griliches, Z. (1969), “Capital-Skill Complementarity," *Review of Economics and Statistics* 5, 465-68.

[106] Haider, S. (2001), “Earnings Instability and Earnings Inequality of Males in the United States: 1967-1991,” *Journal of Labor Economics* 19, 799-836.

[107] Hall, R.E. (1973), “The Specification of Technology with Several Kinds of Output,” *Journal of Political Economy* 81, 878-892.

[108] Hall, R.E. (2001), “The Stock Market and Capital Accumulation,” *American Economic Review* 95, 1185-1202.

[109] Hamermesh, D.S. (1993), Labor Demand, (Princeton University Press: Princeton, NJ)..

[110] Hansen, H. (1998), “Transition from Unemployment Benefits to Social Assistance in Seven European OECD Countries,” *Empirical Economics* 23, 5-30.

[111] Hayek, F.A. (1945), “The Use of Knowledge in Society,” *American Economic Review* 35, 519-530.

[112] Heathcote, J., K. Storesletten and G.L. Violante (2003), "The Macroeconomic Implications of Rising Wage Inequality in the U.S.," mimeo NYU.

[113] Heathcote, J., K. Storesletten and G.L. Violante (2004), "Insurance and Opportunites: The Welfare Analysis of Wage Dispersion," mimeo NYU.

[114] Heckman, J.J. (2000), "Policies to Foster Human Capital," *Research in Economics* 54, 3-56.

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

[116] Helpman, E. and A. Rangel (1999), "Adjusting to a New Technology: Experience and Training," *Journal of Economic Growth* 4, 359-383.

[117] Ho, M.S. and D.W. Jorgenson (1999), "The Quality of the U.S. Workforce," mimeo Harvard University.

[118] Hobjin, B. (2000), "Identifying Sources of Growth," mimeo, Federal Reserve Bank of New York.

[119] Holmes, T. and M.F. Mitchell (2004), "A Theory of Factor Allocation and Plant Size," NBER Working Papers 10079.

[120] Hornstein, A. (1999), "Growth Accounting with Technological Revolutions," *Federal Reserve Bank of Richmond Economic Quarterly* 85 (3), 1-24.

[121] Hornstein, A. (2004), "(Un)Balanced Growth," *Federal Reserve Bank of Richmond Economic Quarterly* 90 *Fall*, 25-45.

[122] Hornstein, A. and P. Krusell (1996), "Can Technology Improvements Cause Productivity Slowdowns?", *NBER Macroeconomics Annual* 11 (MIT Press: Cambridge MA), 209-259.

[123] Hornstein, A. and P. Krusell (2000), "The IT Revolution: Is it Evident in the Productivity Numbers?" *Federal Reserve Bank of Richmond Economic Quarterly* 86 *Fall*, 49-78.

[124] Hornstein, A. and P. Krusell (2003), "Implications of the Capital-Embodiment Revolution for Directed R&D and Wage Inequality," *Federal Reserve Bank of Richmond Economic Quarterly* 89 *Fall*, 25-50.

[125] Hornstein, A., P. Krusell and G.L. Violante (2003a), "Vintage Capital in Frictional Labor Markets", mimeo NYU.

[126] Hornstein, A., P. Krusell and G.L. Violante (2003b), "A Quantitative Study of the Replacement Problem in Frictional Economies," mimeo NYU.

[127] Hoxby C.M. and B.T. Long (1998), "Explaining Rising Income and Wage Inequality Among the College-Educated," mimeo harvard University.

[128] Huggett, M. (1996), "Wealth Distribution in Life-Cycle Economies," *Journal of Monetary Economics* 38, 469-94.

[129] Huggett, M. and S. Ospina (2001), "Does Productivity Growth Fall After the Adoption of New Technology?" *Journal of Monetary Economics* 48, 173-95.

[130] Hulten, C.R. (1992), "Growth Accounting When Technical Change is Embodied in Capital," *American Economic Review* 82, 964-80.

[131] Hyslop, D. (2001), "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings," *American Economic Review* 91, 755-77.

[132] Ingram B. and G. Neumann (1999), "An Analysis of the Evolution of the Skill Premium, " mimeo University of Iowa.

[133] Irwin, D.A. and P.J. Klenow (1994), "Learning-by-Doing Spillovers in the Semiconductor Industry," *Journal of Political Economy* 102, 1200-1227.

[134] Iversen, T. (1998), "Wage Bargaining, Central Bank Independence and the Real Effects of Money", *International Organization* 52, 469-504.

[135] Johnson, D. and S. Shipp (1997), "Trends in Inequality Using Consumption-Expenditures in the U.S. from 1960 to 1993," *Review of Income and Wealth* 43, 133-52.

[136] Jones C. (2004), "Growth and Ideas" in: P. Aghion and S. Durlauf, eds., Handbook of Economic Growth, (North-Holland: Amsterdam).

[137] Jorgenson, D.W. (2001), "Information Technology and the U.S. Economy," *American Economic Review* 91, 1-32.

[138] Jorgenson, D.W. (2004), "Accounting for Growth in the Information Age," in: P. Aghion and S. Durlauf, eds., Handbook of Economic Growth (North-Holland: Amsterdam).

[139] Jorgenson, D.W. and K.J. Stiroh (2000), "U.S. Economic Growth at the Industry Level," *American Economic Review* 90, 161-67.

[140] Jorgenson, D.W., F. Gollop and B. Fraumeni (1987), Productivity and U.S. Economic Growth (Harvard University Press: Cambridge MA).

[141] Jovanovic, B. (1998), "Vintage Capital and Inequality," *Review of Economic Dynamics* 1, 497 - 530.

[142] Jovanovic, B. and Y. Nyarko (1995), "A Bayesian Learning Model Fitted to a Variety of Empirical Learning Curves," *Brookings Papers on Economic Activity (Microeconomics)*, 247-99.

[143] Jovanovic, B. and P. Rousseau (2004a), "General Purpose Technologies," in: P. Aghion and S. Durlauf, eds., Handbook of Economic Growth, (North-Holland: Amsterdam).

[144] Jovanovic, B. and P. Rousseau (2004b), "Specific Capital and the Division of Rents," mimeo NYU.

[145] Juhn, C. (1992), "Decline of Male Labor Market Participation: The Role of Declining Market Opportunities", *Quarterly Journal of Economics* 107, 79-121.

[146] Juhn, C., K. Murphy and B. Pierce (1993), “Wage Inequality and the Rise in Returns to Skill,” *Journal of Political Economy* 101, 410-442.

[147] Kambourov, G. and I. Manovskii (2004), “Occupational Mobility and Wage Inequality”, mimeo University of Pennsylvania.

[148] Katz, L. and D. Autor (1999), “Changes in the Wage Structure and Earnings Inequality,” in: O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, vol. 3 (North-Holland: Amsterdam), 1463-1555.

[149] Katz, L. and K. Murphy (1992), “Changes in Relative Wages, 1963-1987: Supply and Demand Factors,” *Quarterly Journal of Economics* 107, 35-78.

[150] Kelly, M. (2000), “Inequality and Crime,” *The Review of Economics and Statistics* 82, 530-539.

[151] Kiley, M.T. (1999), “The Supply of Skilled Labour and Skill-Biased Technological Progress,” *The Economic Journal* 109, 708-24.

[152] King, I. and L. Welling (1995), “Search, unemployment, and growth,” *Journal of Monetary Economics* 3, 499-507

[153] Kocherlakota, N. (1996), “Implications of Efficient Risk Sharing without Commitment,” *Review of Economic Studies* 63, 595-609.

[154] Kremer, M. and E.S. Maskin (1996), “Wage Inequality and Segregation by Skill,” NBER Working Papers 5718.

[155] Krueger A.B. and L.H. Summers (1988), “Efficiency Wages and the Inter-Industry Wage Structure,” *Econometrica* 56, 259-93.

[156] Krueger, D., and Kumar (2004); “Skill-specific rather than General Education: A Reason for US-Europe Growth Differences?,” *Journal of Economic Growth* 9, 167-207.

[157] Krueger, D. and F. Perri (2002), "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory", mimeo Stanford.

[158] Krueger, D. and F. Perri (2003), "On the Welfare Consequences of the Increase in Inequality in the United States," in: M. Gertler and K. Rogoff, eds., *NBER Macroeconomics Annual* 18, (MIT Press: Cambridge, MA).

[159] Krugman, P. (1994), "Past and Prospective Causes of High Unemployment," *Economic Review* (Federal Reserve Bank of Kansas City) 79 (4), 23-43.

[160] Krusell, P., L. Ohanian, J.-V. Rios-Rull and G.L. Violante (2000), "Capital Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica* 68, 1029-1053.

[161] LaLonde, R.J., J.J. Heckman and J. Smith (1999), "The Economics and Econometrics of Active Labor Market Programs," in: O. Ashenfelter and D. Card, eds., Handbook of Labor Economics, Vol 3A (North-Holland: Amsterdam), 1865-2097.

[162] Lee, D. and K. Wolpin (2004), "Intersectoral Labor Mobility and the Growth of the Service Sector", mimeo NYU.

[163] Levy, P.A. (1985), "The Unidimensional Perspective of Reagan's Labor Board", *Rutgers Law Journal* 16, 269-390.

[164] Levy, F. and R. J. Murnane (1992), "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations", *Journal of Economic Literature* 30, 1333-1381.

[165] Lillard, L.A. and H.W. Tan (1986), "Training: Who Gets It and What Are Its Effects on Employment and Earnings?" (Santa Monica: RAND Corporation Report R-3331-DOL/RC).

[166] Lindbeck, A. and D.J. Snower (1996), "Reorganization of Firms and Labor Market Inequality," *American Economic Review, P&P* 86, 315-321.

[167] Lindquist, M.J. (2002), "Capital-Skill Complementarity and Inequality in Sweden," mimeo Uiversity of Stockholm.

[168] Lindquist, M.J. (2004), "Capital-Skill Complementarity and Inequality Over the Business Cycle," *Review of Economic Dynamics* 7, 519-540.

[169] Lipsey, R.G., C. Bekar and K. Carlaw (1998), "The Consequences of Changes in GPTs," in: E. Helpman, ed., General Purpose Technologies and Economic Growth (MIT Press: Cambridge MA), 193-218.

[170] Ljungqvist, L. and T.J. Sargent (1998), "The European Unemployment Dilemma," *Journal of Political Economy* 106, 514-50.

[171] Ljungqvist, L. and T.J. Sargent (2003), "European Unemployment and Turbulence Revisited in a Matching Model," mimeo NYU.

[172] Lloyd-Ellis, H. (1999), "Endogenous Technological Change and Wage Inequality," *American Economic Review* 89, 47-77.

[173] Lucas, R.E. Jr. (1993), "Making a Miracle," *Econometrica* 61, 251-72.

[174] Lucas, R.E. Jr. (2003), "Macroeconomic Priorities," *American Economic Review* 93, 1-14.

[175] Lucas, R.E. Jr. and E.C. Prescott (1974), "Equilibrium Search and Unemployment," *Journal of Economic Theory* 7, 188-209.

[176] Machin, S. (1996), "Wage Inequality in the UK," *Oxford Review of Economic Policy* 12, 47-64.

[177] Machin, S. (2000), "Union Decline in Britain," *British Journal of Industrial Relations* 38, 631-45.

[178] Machin, S. (2003), "New Workplaces, New Workers: Trade Union Decline and the New Economy", forthcoming in H. Gospel and S. Wood, eds., The Future of Unions ,Volume 1 (Routledge: London).

[179] Manuelli, R. (2000), "Technological Change, the Labor Market, and the Stock Market", NBER Working Papers 8022.

[180] Marimon, R. and F. Zilibotti (1999), “Unemployment vs. Mismatch of Talents: Reconsidering Unemployment Benefits,” *Economic Journal* 109, 266-91.

[181] McCall, J.J. (1970), “ Economics of Information and Job Search,” *Quarterly Journal of Economics* 84, 113-26.

[182] McConnell, S. (1996), “The Role of Computers in Reshaping the Workforce,” *Monthly Labour Review* 119 (August), 3-5.

[183] McGrattan, E., and E.C. Prescott (2003), “Taxes, Regulations, and the Value of U.S. Corporations: A General Equilibrium Analysis,” Staff Report 309, Federal Reserve Bank of Minneapolis.

[184] Meghir, C. and L. Pistaferri (2004), “Income Variance Dynamics and Heterogeneity”, *Econometrica* 72, 1-32.

[185] Milgrom, P. and J. Roberts (1990), “The Economics of Modern Manufacturing: Technology, Strategy,and Organization,” *American Economic Review* 80, 511-528.

[186] Mincer, J. and Y. Higuchi (1991), “Wage Structures and Labor Turnover in the United States and Japan,” *Journal of the Japanese and International Economies* 2, 97-133.

[187] Mitchell, M.F. (2001), “Specialization and the Skill Premium in the 20th Century,” Staff Report 290, Federal Reserve Bank of Minneapolis.

[188] Möbius, M. (2000), “The Evolution of Work”, mimeo Harvard.

[189] Moen, E.R. (1997), “Competitive Search Equilibrium,” *Journal of Political Economy* 105, 385- 411.

[190] Mortensen, D.T. and C.A. Pissarides (1998), “Technological Progress, Job Creation, and Job Destruction,” *Review of Economic Dynamics* 1, 733-53.

[191] Mortensen, D.T. and C.A. Pissarides (1999), “Unemployment Responses to ‘Skill-Biased’ Technology Shocks: The Role of Labour Market Policy,” *Economic Journal* 109, 242-65.

[192] Murphy, K. and R. Topel (1997), "Unemployment and Nonemployment," *American Economic Review P&P* 87, 295-300.

[193] Murphy, K. and F. Welch (1992), "The Structure of Wages," *Quarterly Journal of Economics* 107, 285-326.

[194] Nelson, R.R. and E.S. Phelps (1966), "Investment in Humans, Technological Diffusion, and Economic Growth," *American Economic Review* 56, 69-75.

[195] Neumark, D. (2000), "Changes in Job Stability and Job Security: A Collective Effort to Untangle, Reconcile and Interpret the Evidence," in: D. Neumark, ed., On the Job: Is Long-Term Employment a Thing of the Past? (Russell Sage Foundation: New York), 1-27.

[196] Nickell, S. and B. Bell (1996), "Changes in the Distribution of Wages and Unemployment in OECD Countries," *American Economic Review P&P* 86, 302-08.

[197] Nickell, S. and R. Layard (1999), "Labor Market Institutions and Economic Performance," in: O. Ashenfelter and D. Card, eds., Handbook of Labor Economics, vol. 3C (North Holland: Amsterdam), 3029-84.

[198] Nickell, S. and L. Nunziata (2002), "Unemployment in the OECD since the 1960s: What Do We Know?" mimeo, Bank of England.

[199] OECD Employment Outlook (1996), OECD: Paris.

[200] Oliner, S.D. and D.E. Sichel (2000), "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives* 14, 3-22.

[201] Ortigueira, S. (2002), "The Rise and Fall of Centralized Wage Bargaining," mimeo Cornell.

[202] Piketty, T. and E. Saez (2003), "Income Inequality in the United States, 1913-1998," *Quarterly Journal of Economics* 118, 1-39.

[203] Piore, M.J. and C. F. Sabel.(1984), The Second Industrial Divide, (Basic Books: New York).

[204] Pissarides, C.A. (2000), Equilibrium Unemployment Theory, (MIT Press: Cambridge MA).

[205] Rajan, R. and J. Wulf (2003), "The Flattening Firm: Evidence from Panel Data on the Changing Nature of Corporate Hierarchies," mimeo Chicago GSB.

[206] Robbins D. (1996), "Evidence on Trade and Wages in Developing World," OECD Technical Papers 119.

[207] Rogerson, R. (2004), "Two Views on the Deterioration of European Labor Market Outcomes," *Journal of the European Economic Association* 2, 447-455.

[208] Rosen, S. (1981), "The Economics of Superstars," *American Economic Review* 71, 845-58.

[209] Ruiz-Arranz, M. (2002). "Wage Inequality in the U.S.: Capital-Skill Complementarity vs. Skill-Biased Technological Change," mimeo, Harvard University.

[210] Saint-Paul, G. (2001), "On the Distribution of Income and Worker Assignment under Intrafirm Spillovers, with an Application to Ideas and Networks," *Journal of Political Economy* 109, 1-37.

[211] Sakellaris, P. and D.J. Wilson (2004), "Quantifying Embodied Technical Change", *Review of Economic Dynamics* 7, 1-26.

[212] Sattinger, M. (1995), "Search and the Efficient Assignment of Workers to Jobs," *International Economic Review* 36, 283-302.

[213] Shi, S. (2002), "A Directed Search Model of Inequality with Heterogeneous Skills and Skill-Based Technology," *Review of Economic Studies* 69, 467-91.

[214] Solow, R. (1957), "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics* 39, 312-320.

[215] Solow, R. (1960), "Investment and Technological Progress," in: K. Arrow, S. Karlin, and P. Suppes, eds., Mathematical Methods in the Social Sciences (Stanford University Press: Stanford CA), 89-104.

[216] Thesmar, D. and M. Thoenig (2000), “Creative Destruction and Firm Organization Choice”, *Quarterly Journal of Economics* 115, 1201-37.

[217] Violante, G.L. (2002), “Technological Acceleration, Skill Transferability and the Rise in Residual Inequality,” *Quarterly Journal of Economics* 117, 297-338.

[218] Weinberg, B.A. (2003a), “Computer Use and the Demand for Women Workers,” *Industrial and Labor Relations Review* 53, 290-308.

[219] Weinberg, B.A. (2003b), “Experience and Technology Adoption,” mimeo, Ohio State University.

[220] Wong, L.Y. (2003), “Can the Mortensen-Pissarides Model with Productivity Changes Explain U.S. Wage Inequality?” *Journal of Labor Economics* 21, 70-105.

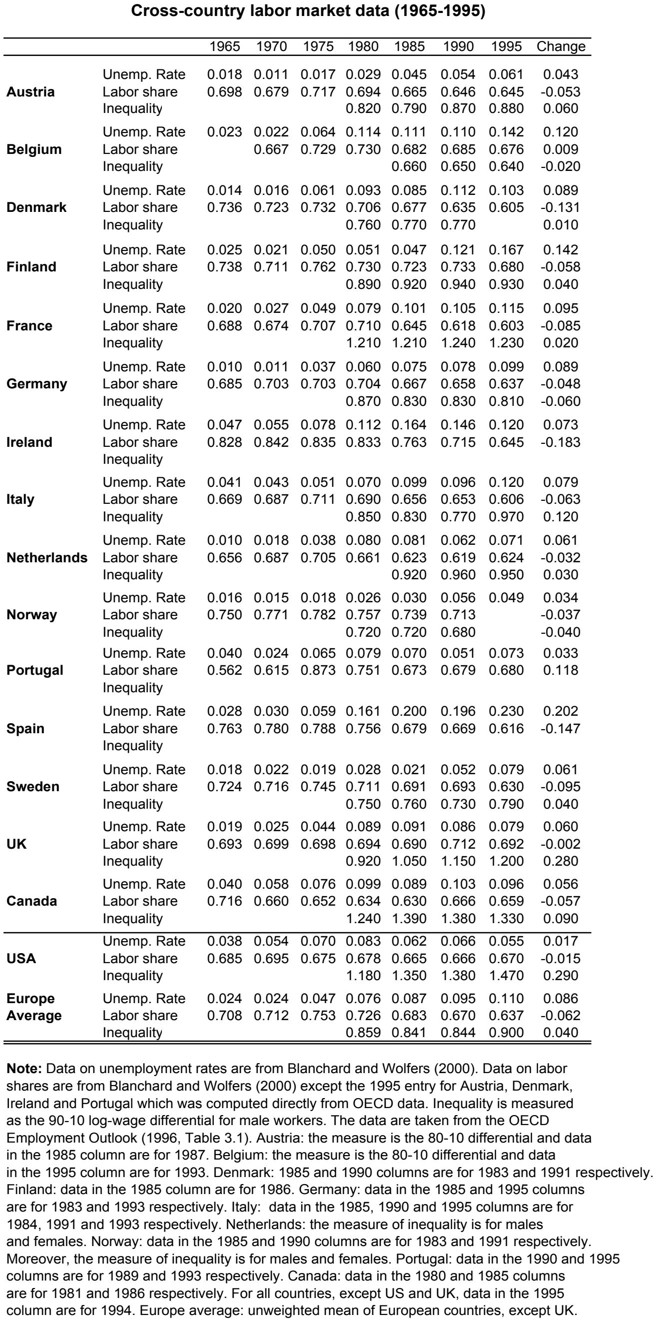

Cross-country labor market data (1965-1995)

		1965	1970	1975	1980	1985	1990	1995	Change
Austria	Unemp. Rate	0.018	0.011	0.017	0.029	0.045	0.054	0.061	0.043
	Labor share	0.698	0.679	0.717	0.694	0.665	0.646	0.645	-0.053
	Inequality				0.820	0.790	0.870	0.880	0.060
Belgium	Unemp. Rate	0.023	0.022	0.064	0.114	0.111	0.110	0.142	0.120
	Labor share		0.667	0.729	0.730	0.682	0.685	0.676	0.009
	Inequality					0.660	0.650	0.640	-0.020
Denmark	Unemp. Rate	0.014	0.016	0.061	0.093	0.085	0.112	0.103	0.089
	Labor share	0.736	0.723	0.732	0.706	0.677	0.635	0.605	-0.131
	Inequality				0.760	0.770	0.770		0.010
Finland	Unemp. Rate	0.025	0.021	0.050	0.051	0.047	0.121	0.167	0.142
	Labor share	0.738	0.711	0.762	0.730	0.723	0.733	0.680	-0.058
	Inequality				0.890	0.920	0.940	0.930	0.040
France	Unemp. Rate	0.020	0.027	0.049	0.079	0.101	0.105	0.115	0.095
	Labor share	0.688	0.674	0.707	0.710	0.645	0.618	0.603	-0.085
	Inequality				1.210	1.210	1.240	1.230	0.020
Germany	Unemp. Rate	0.010	0.011	0.037	0.060	0.075	0.078	0.099	0.089
	Labor share	0.685	0.703	0.703	0.704	0.667	0.658	0.637	-0.048
	Inequality				0.870	0.830	0.830	0.810	-0.060
Ireland	Unemp. Rate	0.047	0.055	0.078	0.112	0.164	0.146	0.120	0.073
	Labor share	0.828	0.842	0.835	0.833	0.763	0.715	0.645	-0.183
	Inequality								
Italy	Unemp. Rate	0.041	0.043	0.051	0.070	0.099	0.096	0.120	0.079
	Labor share	0.669	0.687	0.711	0.690	0.656	0.653	0.606	-0.063
	Inequality				0.850	0.830	0.770	0.970	0.120
Netherlands	Unemp. Rate	0.010	0.018	0.038	0.080	0.081	0.062	0.071	0.061
	Labor share	0.656	0.687	0.705	0.661	0.623	0.619	0.624	-0.032
	Inequality					0.920	0.960	0.950	0.030
Norway	Unemp. Rate	0.016	0.015	0.018	0.026	0.030	0.056	0.049	0.034
	Labor share	0.750	0.771	0.782	0.757	0.739	0.713		-0.037
	Inequality				0.720	0.720	0.680		-0.040
Portugal	Unemp. Rate	0.040	0.024	0.065	0.079	0.070	0.051	0.073	0.033
	Labor share	0.562	0.615	0.873	0.751	0.673	0.679	0.680	0.118
	Inequality								
Spain	Unemp. Rate	0.028	0.030	0.059	0.161	0.200	0.196	0.230	0.202
	Labor share	0.763	0.780	0.788	0.756	0.679	0.669	0.616	-0.147
	Inequality								
Sweden	Unemp. Rate	0.018	0.022	0.019	0.028	0.021	0.052	0.079	0.061
	Labor share	0.724	0.716	0.745	0.711	0.691	0.693	0.630	-0.095
	Inequality				0.750	0.760	0.730	0.790	0.040
UK	Unemp. Rate	0.019	0.025	0.044	0.089	0.091	0.086	0.079	0.060
	Labor share	0.693	0.699	0.698	0.694	0.690	0.712	0.692	-0.002
	Inequality				0.920	1.050	1.150	1.200	0.280
Canada	Unemp. Rate	0.040	0.058	0.076	0.099	0.089	0.103	0.096	0.056
	Labor share	0.716	0.660	0.652	0.634	0.630	0.666	0.659	-0.057
	Inequality				1.240	1.390	1.380	1.330	0.090
USA	Unemp. Rate	0.038	0.054	0.070	0.083	0.062	0.066	0.055	0.017
	Labor share	0.685	0.695	0.675	0.678	0.665	0.666	0.670	-0.015
	Inequality				1.180	1.350	1.380	1.470	0.290
Europe Average	Unemp. Rate	0.024	0.024	0.047	0.076	0.087	0.095	0.110	0.086
	Labor share	0.708	0.712	0.753	0.726	0.683	0.670	0.637	-0.062
	Inequality				0.859	0.841	0.844	0.900	0.040

Note: Data on unemployment rates are from Blanchard and Wolfers (2000). Data on labor shares are from Blanchard and Wolfers (2000) except the 1995 entry for Austria, Denmark, Ireland and Portugal which was computed directly from OECD data. Inequality is measured as the 90-10 log-wage differential for male workers. The data are taken from the OECD Employment Outlook (1996, Table 3.1). Austria: the measure is the 80-10 differential and data in the 1985 column are for 1987. Belgium: the measure is the 80-10 differential and data in the 1995 column are for 1993. Denmark: 1985 and 1990 columns are for 1983 and 1991 respectively. Finland: data in the 1985 column are for 1986. Germany: data in the 1985 and 1995 columns are for 1983 and 1993 respectively. Italy: data in the 1985, 1990 and 1995 columns are for 1984, 1991 and 1993 respectively. Netherlands: the measure of inequality is for males and females. Norway: data in the 1985 and 1990 columns are for 1983 and 1991 respectively. Moreover, the measure of inequality is for males and females. Portugal: data in the 1990 and 1995 columns are for 1989 and 1993 respectively. Canada: data in the 1980 and 1985 columns are for 1981 and 1986 respectively. For all countries, except US and UK, data in the 1995 column are for 1994. Europe average: unweighted mean of European countries, except UK.

Table 1: Data on the evolution of the labor share, the unemployment rate, and wage inequality across OECD countries from 1965-1995.

Cross-country institutions data (1984-1995)

	Labor Standards	Employment Protection	Union Density	Bargaining Centralization	Ratio of min. to avg. wage	Benefit Repl. Rate	Benefit Duration
Austria	5	16	46.2	17	0.62	0.50	2.0
Belgium	4	17	51.2	10	0.60	0.60	4.0
Denmark	2	5	71.4	14	0.54	0.90	2.5
Finland	5	10	72.0	13	0.52	0.63	2.0
France	6	14	9.8	7	0.50	0.57	3.0
Germany	6	15	32.9	12	0.55	0.63	4.0
Ireland	4	12	49.7	6	0.55	0.37	4.0
Italy	7	20	38.8	5	0.71	0.20	0.5
Netherlands	5	9	25.5	11	0.55	0.70	2.0
Norway	5	11	56.0	16	0.64	0.65	1.5
Portugal	4	18	31.8	7	0.45	0.65	0.8
Spain	7	19	11.0	7	0.32	0.70	3.5
Sweden	7	13	82.5	15	0.52	0.80	1.2
UK	0	7	39.1	6	0.40	0.38	4.0
Canada	2	3	35.8	1	0.35	0.59	1.0
USA	0	1	15.6	2	0.39	0.50	0.5
Europe Average	5.15	13.77	44.52	10.77	0.54	0.61	2.38

Note: Data are taken from Nickell and Layard (1999), Tables 6, 7, 9, 10. Labor standards are summarized in an index whose max value is 10 and refers to labor market standards enforced by legislation. The employment protection index ranges from 1 to 10. Union density is measured as a percentage of all salary earners. Centralization is an index where 17 corresponds to the most centralized regime. Benefit duration is in years. Europe average: unweighted mean of European countries, except UK.

Table 2: Data on various labor market institutions across OECD countries. Averages for the period 1985-1995.

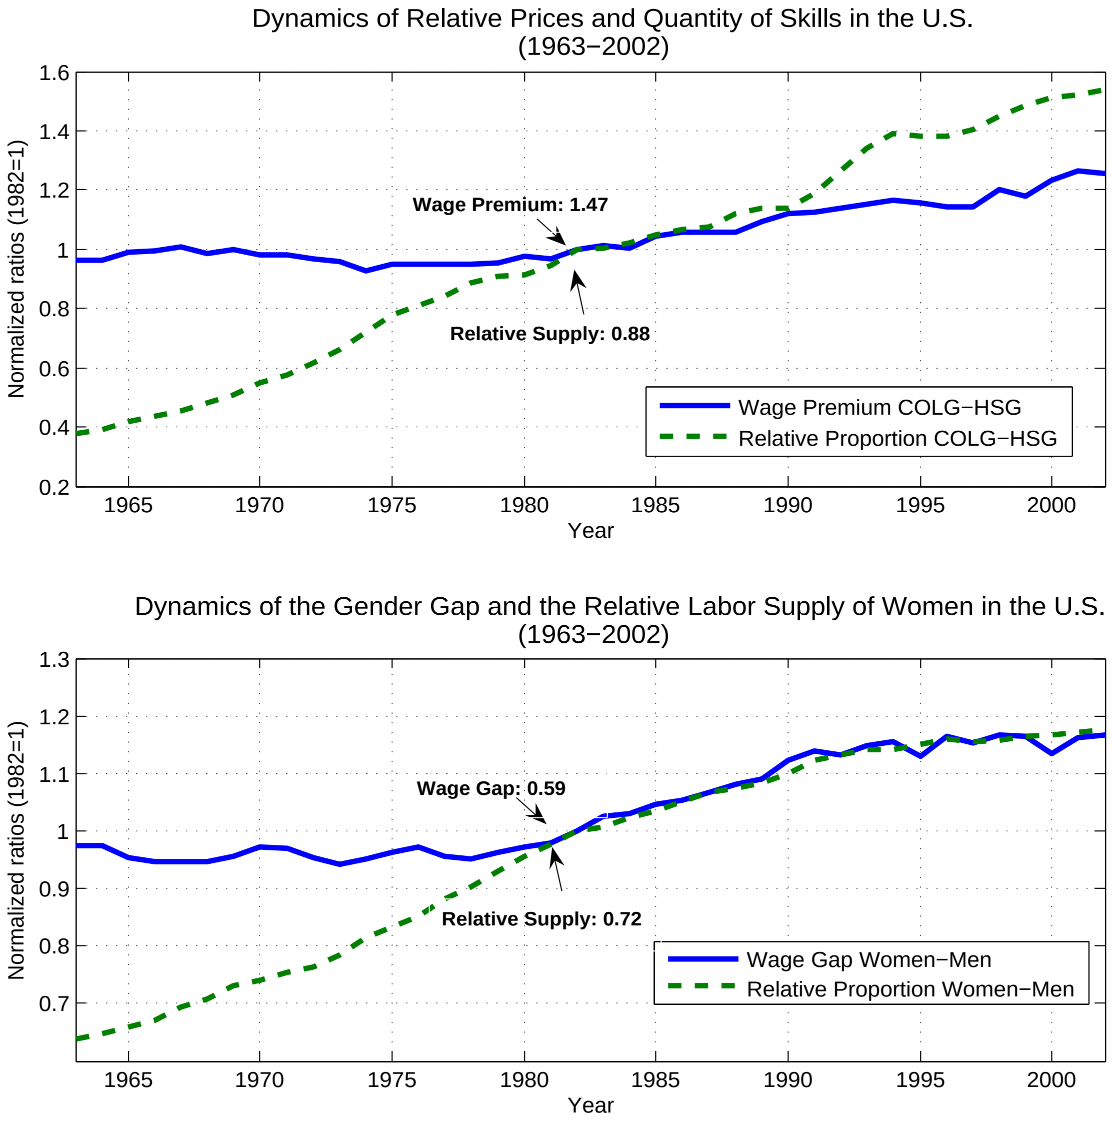

Figure 1: The top panel depicts the evolution of the skill premium (average wage of college graduates relative to the wage of high-school graduates) and of the relative quantity of skilled workers, from 1963-2002. The bottom panel depicts the evolution of the gender gap (average wage of female workers relative to the wage of male workers) and of the relative quantity of female workers, over the same period of time.

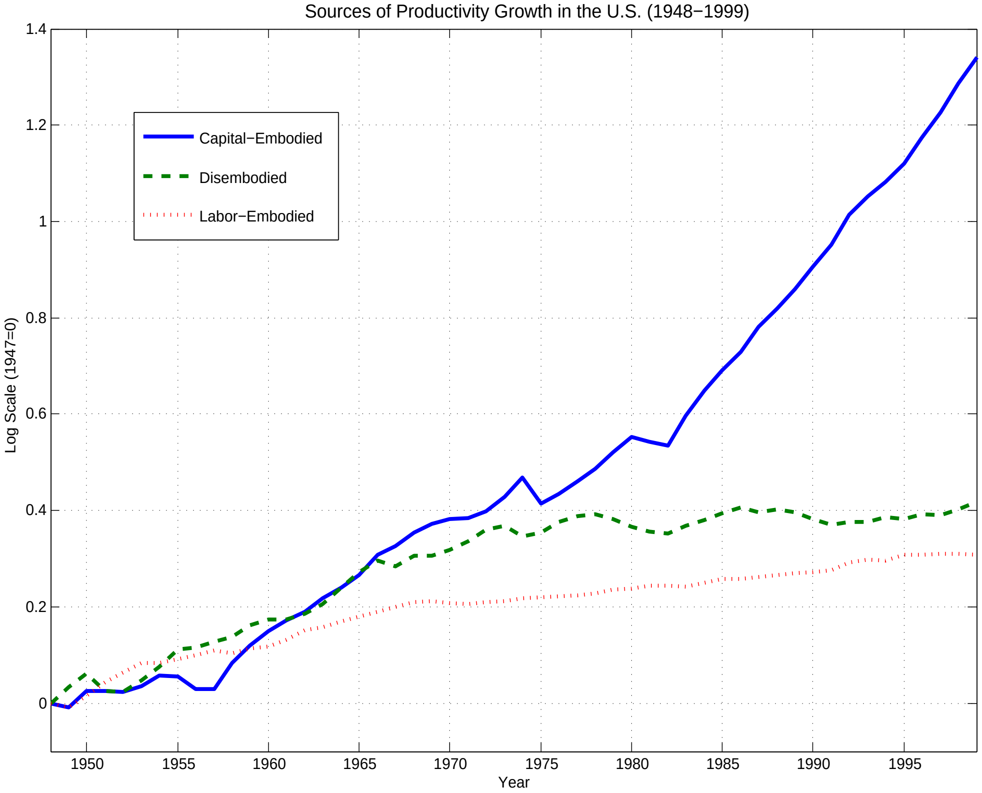

Figure 2: The figure depicts the dynamics of three sources of productivity growth in the post-war U.S. economy: disembodied, capital-embodied, and labor-embodied. Source: Cummins and Violante (2002).

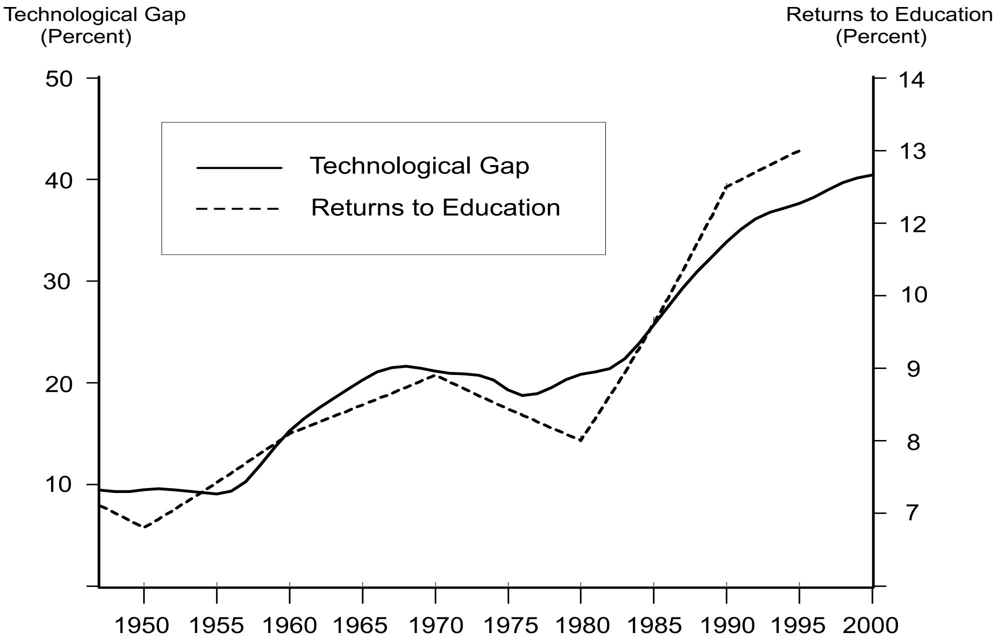

Figure 3: The figure illustrates the joint dynamics of the returns to education and the technological gap (1947-2000) in the U.S. economy. The figure is reproduced from Cummins and Violante (2002).

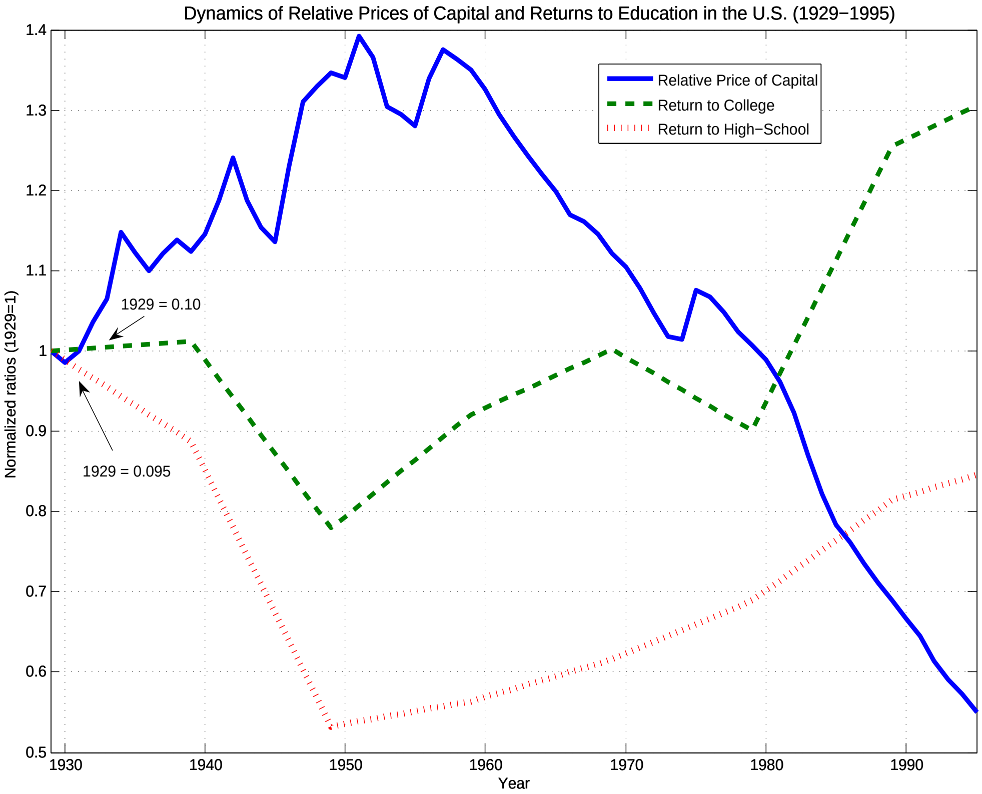

Figure 4: The figure depicts the dynamics of the relative price of capital and the returns to education from 1929-1995 in the U.S. economy. Source: Cummins and Violante (2002) and Goldin and Katz (1999).

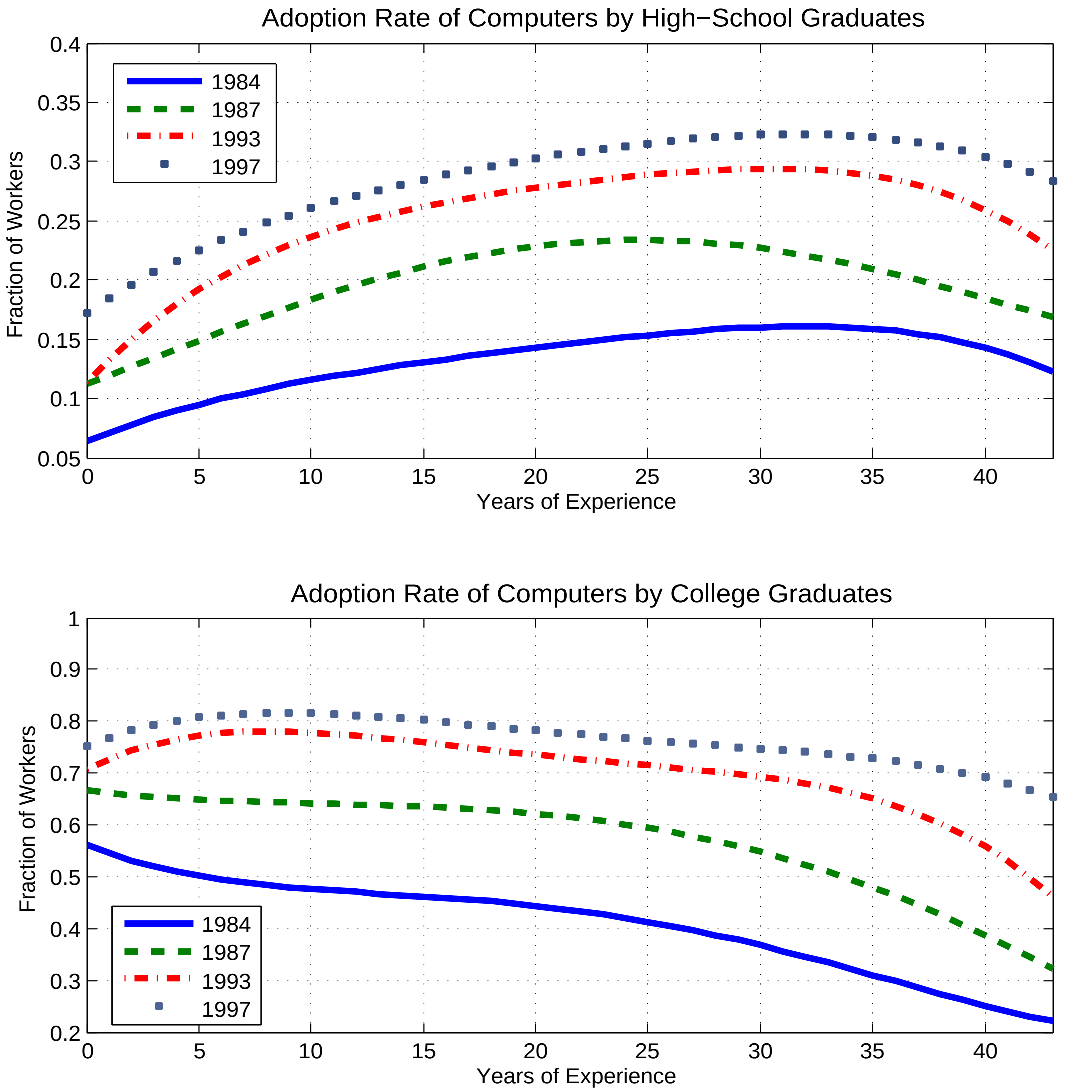

Figure 5: The top panel depicts the experience profile of the adoption rate of computers for U.S. high-school graduates for 1984, 1987, 1993, and 1997. The bottom panel plots similar experience profiles for college graduates. The figure is reproduced from Weinberg (2003b).

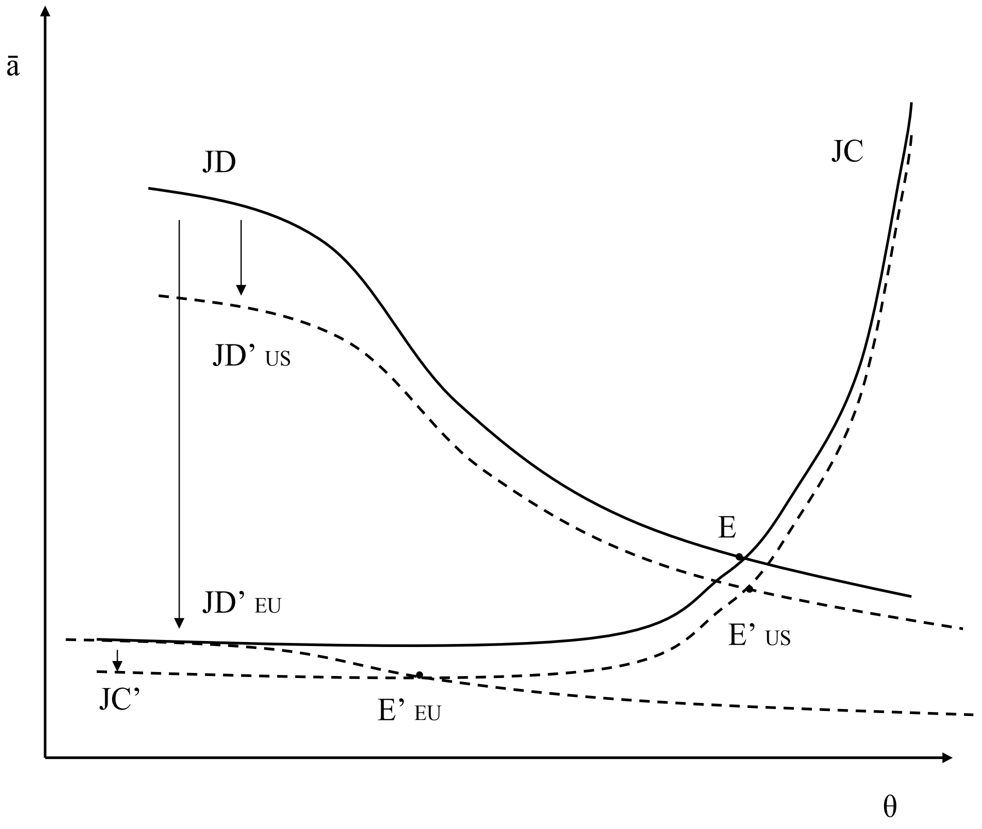

Figure 6: The figure illustrates graphically the equilibrium comparative statics of the model by Hornstein, Krusell and Violante (2003a). Following an acceleration in the rate of capital-embodied technical change, both the job-creation (JC) and the job-destruction (JD) curves shift. The amplitude of the shift is regulated by institutions, and hence it differs between the flexible economy (US) and the rigid economy (EU).