

Inequality and Economic Growth: The Perspective of the New Growth Theories

Philippe Aghion, Eve Caroli, and Cecilia García-Peñalosa.¹

1. Introduction

THE QUESTION of how inequality is generated and how it reproduces over time has been a major concern of social scientists for more than a century. Yet the relationship between inequality and the process of economic development is far from being well understood.

First, on the effect of inequality on growth in market economies: the conventional textbook approach is that inequality is good for incentives and therefore good for growth, even though incentive and growth considerations might (sometimes) be traded off against equity or insurance goals. On the other

hand, development economists have long expressed counterarguments, although not in a formalized way. For example, Michael Todaro's book *Economic Development* provides four general arguments why "greater equality in developing countries may in fact be a condition for self-sustaining economic growth," namely: (a) disaving and/or unproductive investment by the rich; (b) lower levels of human capital held by the poor; (c) demand pattern of the poor being more biased towards local goods; and (d) political rejection by the masses.² Recently, the view that inequality is growth-enhancing has been further challenged by a number of empirical studies, often based on cross-country regressions of GDP growth on income inequality. They all find a negative correlation between the average rate of growth and a number of measures of inequality.³

An interesting case study is that of South Korea and the Philippines during the past thirty years, discussed by Roland Benabou (1996). In the early 1960s, these two countries looked quite similar with regard to major macroeconomic

¹Aghion: University College London and EBRD. Caroli: INRA-LEA and CEPREMAP, Paris. García-Peñalosa: GREQAM, Marseille. This survey draws heavily from joint work with Patrick Bolton, Peter Howitt, Abhijit Banerjee, and Gian Luca Violante. We are particularly grateful to Tony Atkinson for many insightful remarks and for drawing our attention to numerous relevant pieces of work. We also benefited from the comments of Beatriz Armendariz, Roland Benabou, Steve Bond, Theo Eicher, Juan Antonio García-Diez, Jon Temple, Steve Redding, Andrea Richter, the editor (John McMillan), and the referees. Finally, we wish to thank the "Cost of Inequality" group of the McArthur Foundation and the School of Public Policy at UCL for invaluable intellectual and financial support. An earlier version of this survey was given as the Raffaele Mattioli Lecture in L. Bocconi University.

² See Todaro (1997) pp. 165–66. We are grateful to Tony Atkinson for pointing this out to us.

³ This literature is comprehensively reviewed by Benabou (1996).

indicators (GDP per capita, investment per capita, average saving rates), although they differed in the degree of income inequality. In the Philippines the ratio of the income share of the top 20 percent to the bottom 40 percent of the population was almost twice as large as in South Korea. Over the following thirty-year period, fast growth in South Korea resulted in a fivefold increase of the output level, while that of the Philippines barely doubled. That is, contrary to what the standard argument predicts, the more unequal country grew more slowly.

The first part of this paper is concerned with providing new theoretical insights on the effects of inequality on growth. In contrast to the arguments provided by Todaro, it will try to reconcile the above empirical findings with existing microeconomic theories of incentives, concentrating on credit market imperfections and moral hazard. We show that, in some instances, greater inequality may reduce an economy's rate of growth.

The immediate implication of our analysis is that redistribution can foster growth. However, the growth process is unlikely to leave inequality unchanged. The question then arises of whether this feedback creates a virtuous circle, in which redistributive policy can be used to reduce inequality, which in turn would accelerate growth and thereby automatically induce further reductions in inequality. Or, on the contrary, does growth initiate a vicious circle because it spontaneously increases inequality, therefore calling for permanent redistributive efforts? The early literature on the evolution of income inequality over the process of development used to be dominated by the so-called *Kuznets hypothesis*.⁴ Using both cross-country data

and time series, Simon Kuznets (1963) found an inverted U-shaped relation between income inequality and GNP per head. This result was interpreted as describing the evolution of the distribution of income over the transition from a rural to an industrial economy: income inequality should increase during the early stages of development (due to urbanization and industrialization) and decrease later on (as industries would already attract a large fraction of the rural labor force). And indeed, in the US the share of total wealth owned by the 10 percent richest households rose from 50 percent around 1770, to about 75 percent around 1870, and then receded back to 50 percent in 1970.

Up to the 1970s, the Kuznets hypothesis seemed to account for the experience not only of the US but also of most of the OECD countries,⁵ where there appeared to be a virtuous circle: lower inequality would foster growth, which in turn would reduce inequality. However, the downward trend in inequality experienced by these economies during the twentieth century has reversed sharply in recent times. In particular, the past fifteen years have witnessed a significant increase in wage inequality. For example, during the 1980s, the ratio of the 90th to the 10th percentile of the male wage distribution increased by 27 percent in the UK and by 18 percent in the US.⁶

In light of the recent evidence, new theories are needed to understand the

⁵ The evidence that growth generates higher inequality in less-developed countries is less clear-cut. The relevance of the Kuznets curve for developing countries was challenged more than twenty years ago by authors such as Irma Adelman and Cynthia Morris (1973). Sudhir Anand and Ravi Kanbur (1993) present an account of this literature.

⁶ See OECD *Employment Outlook* (1993), pp. 61–62.

⁴ See Kuznets (1955) and (1963).

impact of growth upon inequality. Economic growth during the past twenty years has been closely associated with three phenomena: trade liberalization, technical change, and the emergence of new organizational forms. The second part of the survey will analyze the way in which these three factors might affect inequality, and, in particular, how they may account for the absence of a virtuous circle between inequality and growth.

The two parts of the survey are concerned with different concepts of inequality. The first part concentrates on *wealth* inequality, while the second focuses on *wage* inequality. This choice of focus is by no means arbitrary. First, when looking at the effects of inequality on growth, we are primarily interested in the ways in which “distribution” can affect aggregate output and growth through its impact on individual investments in human or physical capital. What is relevant then is the distribution of wealth, no matter whether this wealth results from the accumulation of labor earnings or capital income. Second, when looking at the effects of growth on inequality, one may want to distinguish a priori between changes in labor earnings and other sources of income. Concentrating on the former rather than on household income allows us to abstract from changes in redistributive policies, in interest rates, or in patterns of household formation.⁷

2. *Is Inequality Always Good for Growth?*

2.1 *New Evidence and Old Theories*

The 1990s have witnessed a resurgence of interest in the determinants of economic growth. The development of

endogenous growth theory and the availability of comparable data on national incomes and growth rates for a large cross-section of countries has permitted the empirical analysis of the causes of national differences in growth rates. Within this vast literature, several studies have examined the impact of inequality upon economic growth. The picture they draw is impressively unambiguous, since they all suggest that greater inequality reduces the rate of growth.⁸ Such a result comes as a surprise both with regard to traditional theories in the field and to the channels through which inequality might affect the growth process.

The evidence on inequality and growth follows what has by now become the standard approach to cross-country comparisons of the determinants of growth. The average annual rate of growth of per capita GDP over the period 1960–85 is regressed on a set of explanatory variables at the start of the period (i.e. around 1960) in order to assess their relative contribution to growth. The various studies dealing with the impact of inequality are consistent in both the data⁹ and the methodology used. The underlying theory is that wealth inequality determines investment in physical or human capital, which in turn affects the long-run growth rate. Unfortunately, the absence of data on the distribution of wealth for a sufficient number of countries forces researchers to use proxies in empirical studies. The most common approach is to use data on income inequality as a proxy for wealth inequality. It is generally argued that this is unlikely to be a major problem since both measures of

⁸ See Alesina and Rodrick (1994), Perotti (1992, 1993, 1996), and Persson and Tabellini (1994).

⁹ The data used are the World Bank data on personal income distribution and the Summers-Heston data set on national incomes.

⁷ See François Bourguignon (1998) and Gary Burtless (1999).

distribution generally vary together in cross-sections. Alternatively, the true distribution of wealth is sometimes proxied by the distribution of land.¹⁰

Reduced-form equations are then estimated, in which the average rate of growth of GDP is regressed on a measure of inequality around 1960. They all give consistent results. Alberto Alesina and Dani Rodrik (1994) regress the average growth rate over 1960–85 on the Gini coefficient of income and of land around 1960. The estimated coefficients imply that both variables have a negative impact on growth, even when controlling for the initial per capita income and for the primary-school enrollment rate in 1960. Greater inequality in the distributions of income and land thus appears to slow down economic growth. Symmetrically, equality seems to be growth-enhancing. Torsten Persson and Guido Tabellini (1994) regress the average growth rate of GDP over the same period on the income share accruing to the third quintile of the income distribution for a cross-section of developed and developing countries. This variable represents the income share of the middle class and is thus considered to be a measure of equality in the underlying distribution. Its impact on growth is positive, significant, and robust to the introduction of other explanatory variables. Persson and Tabellini also find a positive effect of their measure of equality on growth using time series data for nine developed economies over the period 1830–1985. Similar results are obtained by Roberto Perotti (1996) for a larger cross-section of countries. Table 1, column 1, reports a reduced-form equation in which the average rate of growth is regressed on the most standard regressors in the growth literature

¹⁰ Perotti (1996) is an example of the first approach, while Alesina and Rodrik (1994) use the second one.

TABLE 1
THE IMPACT OF INCOME INEQUALITY ON GROWTH

Dependent Variable	(1)	(2)	(3)
	GR	GR	MTax
Constant	-0.18 (-1.37)	0.004 (0.47)	0.164 (1.13)
GDP	-0.002 (-1.77)	-0.004 (-2.39)	-0.021 (-1.50)
MSec	0.031 (4.05)	0.004 (0.38)	
FSec	-0.025 (-3.06)	0.001 (0.10)	
PPPI	-0.002 (-0.30)	-0.0005 (-0.07)	
MID	0.118 (2.84)		-0.096 (-0.19)
MTax		0.090 (3.61)	
N.Obs.	67	49	49
R ²	0.30	0.22	0.30
Estimation	OLS	2SLS	2SLS

Source: Perotti (1996), Table 4, p. 160, and Table 8, p. 170.

Note: t-statistics in parentheses. The variables are GR: average rate of growth of per capita GDP 1960–85; GDP: per capita GDP 1960; MSec: average years of secondary schooling in the male population; FSec: average years of secondary schooling in the female population; PPPI: purchasing power parity value of the investment deflator relative to the US in 1960; MID: combined income share of the third and fourth quintile; MTax: average marginal tax rate over the period 1970–85.

(per capita GDP; average years of secondary schooling in the male and the female populations; and the value of the investment deflator as a proxy for market distortions) and on a measure of the size of the middle class, MID, defined as the income share of the third and fourth quintiles.¹¹

¹¹ This cross-country evidence has been recently criticized by Kristin Forbes (1998). Using panel data for a cross-section of countries, Forbes finds a positive relationship between inequality and growth. Her analysis, however, presents three problems. First, the Arellano-Bond Generalized Method of Moments used results in excessively

In addition to the evidence that a more equal income distribution is beneficial to growth, the empirical literature also provides insights as to the channels through which inequality affects economic growth. Perotti (1992) emphasizes the role of credit constraints. Using the loan-to-value ratio for domestic mortgages as a proxy for credit availability, he finds that greater credit availability has a positive and significant effect on the growth rate. Moreover, as the income share of the lowest two quintiles decreases—that is, as inequality rises—the impact of credit availability on growth increases. Similarly, the negative effect of inequality on physical capital investment is enhanced by credit frictions. Other empirical studies have put forward the role of macroeconomic volatility as the transmission mechanism between inequality and growth. Income inequality is found to be positively correlated with volatility, measured by the standard deviation of the annual rate of growth of GDP.¹² Cross-country regressions also find that greater volatility of the growth rate consistently reduces the average rate of growth during that period. This is partly due to its deterring effect on physical and human capital investment.¹³

small standard errors when the sample is small, casting doubt over the significance of her coefficients (see Richard Blundell and Stephen Bond 1998). Second, the assumed lag structure—inequality today affects growth in five years' time—is ad hoc. Thirdly, in order to obtain a positive and significant coefficient, Forbes needs to restrict the data on inequality to the so-called “high-quality” subset compiled by Deininger and Squire (1996). Atkinson and Andrea Brandolini (1999) show that the criteria used to construct this subset are inappropriate and that there is no reason to exclude other countries from the analysis.

¹² See Ricardo Hausmann and Michael Gavin (1996) and Richard Breen and García-Peñalosa (1998).

¹³ See Garey Ramey and Valerie Ramey (1995), Hausmann and Gavin (1996), and Inter-American Development Bank (1995).

An obvious question then arises: can the negative impact of inequality on growth be reduced by redistribution? William Easterly and Sergio Rebelo (1993) examine the impact of fiscal policy on growth for a large cross-section of developed and developing countries. Using measures of redistribution such as the marginal and average tax rates and different types of social spending, they find that redistribution has, if anything, a positive effect on growth rates. These results are striking in that they contradict the traditional view, discussed below, that redistribution is harmful for growth. More recent work by Perotti (1996) displays similar results. He estimates a two-stage least squares growth regression in which fiscal policy variables are endogenously determined by inequality. Redistribution, measured by the marginal tax rate, appears to have a positive and significant impact upon economic growth, as we can see in column 2 of Table 1. Casual evidence from developing countries also points to the fact that redistribution in the form of land or education reform has played an important role in fostering economic growth.¹⁴

Cross-country growth regressions have been the subject of criticism due to their ad hoc specification and the fragility of many of the results.¹⁵ However, recent empirical evidence provides, at worst, enough reasons to cast doubt over the validity of traditional theories of the influence of inequality and redistribution upon growth. These theories maintained that inequality should, if at all, have a *stimulating* effect on capital accumulation and growth.

¹⁴ This has been especially so in the case of the East Asian economies. See World Bank (1993).

¹⁵ Jonathan Temple (1999) discusses in detail the problems of cross-country regressions and provides an assessment of why they are nevertheless a powerful tool.

The view that wealth inequality should be growth-enhancing is based on three arguments. First comes Kaldor's hypothesis that the marginal *propensity to save* of the rich is higher than that of the poor. If the growth rate of GDP is directly related to the proportion of national income that is saved, more unequal economies are bound to grow faster than economies characterized by a more equitable distribution of income. Joseph Stiglitz (1969) formalized this argument in a Solow growth model, showing that with a linear saving function, aggregate behavior is independent of the distribution. François Bourguignon (1981) went one step further and showed that with a convex savings function, aggregate output does depend on the initial distribution and is higher along the more unequal steady-state. When combined with an AK production function, this leads to the prediction that more unequal economies will grow faster.

A second reason why inequality may enhance growth has to do with *investment indivisibilities*: investment projects, in particular the setting up of new industries or the implementation of innovations, often involve large sunk costs. In the absence of a broad and well-functioning market for shares, wealth obviously needs to be sufficiently concentrated in order for an individual (or a family) to be able to cover such large sunk costs and thereby initiate a new industrial activity.¹⁶

Lastly, the idea that there is necessarily a trade-off between productive efficiency and equality is based on *incentive* considerations, first formalized by James Mirrlees (1971). Namely, in a moral hazard context where output depends on the *unobservable* effort borne by agents or "employees," rewarding

the employees with a constant wage independent from (the observable) output performance will obviously discourage them from investing any effort. On the other hand, making the reward too sensitive to output performance may also be inefficient from an *insurance* point of view when output realizations are highly uncertain and the employees are *risk averse*.¹⁷

The basic incentive argument carries over to the aggregate economy when agents are identical and/or capital markets are perfect, as shown by Rebelo (1991). In a Ramsey-Cass-Koopmans growth model with perfect capital markets, the rate of growth of individual consumption is given by

$$g = \frac{r - \rho}{\sigma},$$

where ρ is the intertemporal discount rate, r the *after-tax* interest rate and σ the intertemporal elasticity of substitution. If all agents have the same preference parameters, this expression is also the aggregate rate of growth. By making the after-tax rate of interest smaller, greater taxation reduces the return to saving, thus lowering the incentives to accumulate capital and hence the rate of growth.

The traditional view in economic theory, then, is that there is a fundamental *trade-off* between productive efficiency (and/or growth) and social justice. Redistribution has both a direct and an indirect effect on growth. On the one hand, it reduces differences in income and wealth, and hence lowers the rate of growth. On the other, it has a negative effect as income redistribution that is financed through an income tax diminishes the incentives to accumulate wealth.¹⁸

¹⁷ Canice Prendergast (1999) surveys the literature on the provision of incentives, and discusses the trade-off between risk and incentives.

¹⁸ An example of this early literature is the book by Arthur Okun (1975), *Equity and Efficiency: The Big Tradeoff*.

¹⁶ This issue has been recently emphasized by policy advisers to transition economies in Central and Eastern Europe and the former Soviet Union.

Overall, the view that inequality is *necessary* for accumulation and that redistribution harms growth is at odds with the empirical evidence summarized above. A possible explanation, put forward by Alesina and Rodrik (1994) and Persson and Tabellini (1994), combines political economy arguments with the traditional negative incentive effect of redistribution. These authors maintain that inequality affects taxation through the political process when individuals are allowed to vote in order to choose the tax rate (or, equivalently, vote to elect a government whose program includes a certain redistributive policy). If inequality determines the extent of redistribution, it will then have an indirect effect on the rate of growth of the economy. In general, we would expect that in very unequal societies, more voters prefer high redistribution than in more equal societies. If redistribution reduces the incentives to invest, and hence the growth rate, then more equal societies would grow faster.

Although it accounts for the negative correlation between inequality and growth found by reduced-form equations, the political economy approach is not fully supported by the data. It implies that greater inequality increases the extent of redistribution, which in turn has a direct negative effect on economic growth. As pointed out earlier, redistribution is found to have a positive rather than negative influence on growth. Moreover, when measures of redistribution such as tax rates or the extent of social spending are regressed on measures of inequality, the coefficients are either insignificant or have a sign opposite to what the theory predicts, as seen in column 3 of Table 1.¹⁹

The evidence regarding the impact of inequality and redistribution upon eco-

omic growth has still to be accounted for by economic theory. In the remainder of this section, we analyze the effects of inequality on growth in economies in which wealth or human capital endowments are heterogenous across individuals and *capital markets are imperfect*. We argue that there are at least three reasons why inequality may have a direct negative effect on growth:

- (a) inequality reduces investment opportunities;
- (b) inequality worsens borrowers' incentives;
- (c) inequality generates macro-economic volatility.

Redistribution to the less endowed, by reducing inequality, can therefore be growth-enhancing in such an economic environment.

2.2 *The Opportunity-Enhancing Effect of Redistribution*

One cornerstone of neoclassical economics is the assumption that there are diminishing returns to capital. It is precisely this assumption that drives the familiar convergence results, both at the cross-country level (as in the Solow growth model) and for individuals (as shown by Stiglitz 1969, and, more recently, by Robert Tamura 1991). The convergence results rely crucially on the existence of perfect capital markets. As Stiglitz (1969) first pointed out, when there are decreasing returns to capital and capital markets are imperfect, individual wealth will not converge to a common level and the aggregate level of output may be affected by its distribution.

This subsection reconsiders Stiglitz's arguments in the context of the recent literature on endogenous growth. The crucial difference between Stiglitz's approach and ours is the concept of the aggregate production function. Stiglitz follows the standard growth model in

¹⁹ See Perotti (1996) and Peter Lindert (1996).

assuming that aggregate output is produced by the *aggregate stock of capital*. That is, $y_t = f(k_t)$, where k_t denotes aggregate capital, defined as the sum of the capital held by individuals, $k_t = \sum_i k_{i,t}$. Once we bring in capital-market imperfections, it is hard to maintain this approach. When a bank refuses to lend funds, it does so to a specific agent with a particular investment project. It hence becomes more appropriate to think of individuals as producers themselves, rather than simply rentiers. Aggregate output is then the sum of the output generated by each production unit, $y_t = \sum_i y_{i,t}$, where agent i 's output is in turn a function of her own capital stock, $y_{i,t} = f(k_{i,t})$. In general, $\sum_i f(k_{i,t})$ and $f(\sum_i k_{i,t})$ are not equal. When individuals are limited in their borrowing capacity, the distribution of wealth affects their production possibilities. This in turn has an impact on the aggregate level of output and, in an endogenous growth model, also on its rate of growth.

In what follows, we will present a simple model in which we assume away all possibilities of borrowing and lending. Using such an extreme form of capital-market imperfections, we are able to show that when the individual production function is concave, greater inequality in the distribution of wealth results in a lower rate of growth. Redistributing wealth from the rich (whose marginal productivity of investment is relatively low, due to decreasing returns to individual capital investments) to the poor (whose marginal productivity of investment is relatively high, but who cannot invest more than their limited endowments), would enhance aggregate productivity and therefore growth. In other words, *redistribution creates investment opportunities* in the absence of well-functioning capital markets, which in turn increases aggregate productivity and growth.

This idea was first put forward in an influential paper by Oded Galor and Joseph Zeira (1993).²⁰ The particular formulation we use in this subsection is borrowed from Benabou (1996), who considers a simple endogenous growth model in which growth is driven by externalities in the accumulation of physical (or human) capital.²¹ In this model, individual production generates spillovers that increase the level of technology available to all production units. The presence of knowledge spillovers across individuals implies that individual production functions differ from the aggregate production function.

More formally, suppose that when individual i invests an amount of physical or human capital $k_{i,t}$ at date t , production takes place according to the technology

$$y_{i,t} = A_t k_{i,t}^\alpha,$$

where $0 < \alpha < 1$. A_t is the level of human capital or technical knowledge available in period t , and it is common to all individuals. The level of technology is endogenous, as the economy exhibits both learning-by-doing and knowledge spillovers. Learning-by-doing means that the more an agent produces one period, the more she learns, and hence the greater the level of knowledge available in the next period. The presence of spillovers implies that the learning done by one individual affects the level of technology of all other agents in the economy. These

²⁰ This paper examines the impact of the distribution of wealth on aggregate output through investments in human capital. Galor and Zeira assume that there are both credit market imperfections and indivisibilities in human capital investments. Parental wealth determines whether or not an individual invests in education, which in turn determines the bequest to her offspring, and hence the next generation's investment opportunities. The initial distribution of wealth then affects both aggregate output and the long-run distribution of wealth and skills.

²¹ This type of model is discussed by Robert Barro and Xavier Sala-i-Martin (1995), ch. 4, and Aghion and Peter Howitt (1998), ch. 1.

two assumptions are captured by the following equation,

$$A_t = \int y_{i,t-1} di = y_{t-1}.$$

That is, the accumulation of knowledge results from past aggregate production activities.

As a result of learning-by-doing, growth depends on individual investments. The rate of growth between period $t-1$ and period t is given by $g_t = \ln(y_t/y_{t-1})$, that is,

$$g_t = \ln \frac{\int A_t k_{i,t}^\alpha di}{A_t} = \ln \int k_{i,t}^\alpha di,$$

It can then be expressed simply as

$$g_t = \ln E_t[k_{i,t}^\alpha],$$

where $E_t[k_{i,t}^\alpha]$ is the mathematical expectation over the output generated by individual investment levels at date t . The rate of growth therefore depends on the distribution of individual capital investments.

To see how investments are determined, consider an economy with only one good that serves both as capital and consumption good. There is a continuum of overlapping-generation families, indexed by $i \in [0,1]$. Each individual lives for two periods. The utility of an individual i born at date t is given by $U_i^t = \log c_{i,t} + \rho \cdot \log c_{i,t+1}$, where $c_{i,t}$ and $c_{i,t+1}$ denote current and future consumption respectively. Individuals differ in their initial endowments. In order to abstract from intergenerational transfers and bequest decisions, suppose that initial endowments are randomly determined at birth. Let the endowment of individual i upon birth at date t be given by

$$w_{i,t} = a \cdot \varepsilon_{i,t},$$

where a is a constant and $\varepsilon_{i,t}$ is an identically and independently distributed random variable, with mean $\frac{1}{a}$.

Individual i can either consume her endowment, or invest it into the pro-

duction of the future consumption good. Production of the future consumption good (i.e., of the good available at date $t + 1$) is assumed to take place according to the technology $y_{i,t} = A_t k_{i,t}^\alpha$, with the level of technology given by $A_t = y_{t-1}$.

In the absence of capital market imperfections, all individuals choose to invest the same amount of capital $k_{i,t} \equiv k_t^*$, no matter what the initial distribution of human capital or “wealth” across individuals (see Benabou 1996). The reason is that the opportunity cost of investing is the rate of interest, both for lenders and borrowers. Hence all individuals wish to invest up to the point where the marginal product of capital is equal to the rate of interest. Those whose wealth is above this level lend, those whose wealth is below it borrow. As a result, aggregate output and growth are not affected by the distribution of wealth.

Conversely, when capital markets are highly imperfect and therefore credit is scarce and costly, equilibrium investments under laissez-faire will remain unequal across individuals with heterogeneous endowments. Consider the extreme situation in which borrowing is simply not possible and agents are constrained by their wealth, $k_{i,t} \leq w_{i,t}$. In this case, individual investments are simply a constant fraction of their wealth $k_{i,t} = s \cdot w_{i,t}$. Thus, in contrast to the perfect capital-market case, when credit is unavailable equilibrium investments will differ across individuals, being an increasing function of their initial endowments in human capital. Individual outputs will then be given by $y_{i,t} = (s \cdot w_{i,t})^\alpha$, and the rate of growth is determined by the distribution of endowments,

$$g_t = \alpha \ln s + \ln \int_0^1 (w_{i,t})^\alpha di.$$

We can now ask whether more inequality is good or bad for growth. Because

of decreasing returns with respect to individual capital investments k_i (in other words, the fact that $\alpha < 1$ and therefore the production function $f(k_i)$ is concave) greater inequality between individual investments for a given aggregate capital stock will reduce aggregate output.²² Therefore, the more unequal the distribution of individual endowments, and hence investments, the smaller current aggregate output and therefore the lower the growth rate in the above model.

There is now a role for suitably designed redistribution policies in enhancing aggregate productive efficiency and growth. Consider an ex-ante redistribution of endowments, which consists of taxing highly endowed individuals directly on their endowments, and then using the revenues from this tax to subsidize the less endowed. Thus, the post-tax endowment of individual i can be simply defined by

$$\hat{w}_{i,t} = w_{i,t} + \beta(w_t - w_{i,t}), \quad 0 < \beta < 1,$$

where w_t is the average endowment. Those with above-average wealth pay a tax of $\beta(w_{i,t} - w_t)$, while those with below-average wealth receive a net subsidy, $\beta(w_t - w_{i,t})$. Because it is a lump-sum tax, it does not change the returns to $k_{i,t}$, and hence it only affects the incentives to invest insofar as it changes the current wealth of the individual.

As the tax rate β increases and the distribution of disposable endowments becomes more equal across individuals,

²² This follows from a standard result in expected utility theory. Let X and Y be two random variables such that Y is obtained from X through a sequence of mean-preserving spreads. If the utility function u is concave, expected utilities are such that $Eu(Y) \leq Eu(X)$.

Then, since the expectation over $w_{i,t}^\alpha$ is given by

$$E_t(w^\alpha) = \int_0^\infty w^\alpha \cdot f_t(w) dw,$$

where $f_t(w)$ is the density function over individual endowments at date t , the growth rate is reduced by a mean-preserving spread.

investments by the poorly endowed will increase while investments by the rich will decrease. However, as we already argued, because the production technology exhibits *decreasing* returns with respect to individual capital investments, we should expect redistribution to have an overall *positive* effect on aggregate output and growth. The rate of growth becomes

$$g = \alpha \ln s + \ln \int_0^1 (w_{i,t} + \beta(1 - w_{i,t}))^\alpha di.$$

Now consider the term under the integral sign. As β increases, the heterogeneity among individual investment levels (which are proportional to $[w_{i,t} + \beta(1 - w_{i,t})]$) decreases, and therefore so does the aggregate efficiency loss due to the unequal distribution of w_i . In the limiting case where $\beta = 1$, the term under the integral sign is constant across individuals i , and the highest possible growth rate is achieved.

The implication of the foregoing analysis is that, when credit is unavailable, redistribution to the poorly endowed, that is, to those individuals who exhibit the highest marginal returns to investment, will be growth-enhancing. Note that this “opportunity creation effect” of redistribution does not rely on incentive considerations: even if one could force the poor to invest *all* their initial endowments rather than maximize intertemporal utility as in the above analysis, redistributing wealth from the richest to the poorest individuals would still have an overall positive effect on aggregate productivity and growth, because of decreasing returns to individual investments.

2.3 *The Positive Incentive Effect of Redistribution: Questioning the Traditional Argument*

Our modeling of capital-market imperfections in the previous subsection

was somewhat extreme, as we simply assumed away all possibilities of borrowing and lending. Using such a reduced form representation of credit market imperfections, we were able to show that redistribution creates investment opportunities. This result holds whatever the source of credit market imperfections. Capital market imperfections can be due to traditional incentive problems—such as ex-ante moral hazard—but also to repayment enforcement problems—stemming from contract incompleteness or ex-post moral hazard. In this subsection we consider in detail the implications of the former. We will illustrate that in the presence of ex-ante moral hazard, greater inequality reduces aggregate incentives to accumulate wealth. This will enable us to challenge the view that the incentive effect of redistribution should always be negative.

In two related papers, Abhijit Banerjee and Andrew Newman (1993) and Aghion and Patrick Bolton (1997) introduce moral-hazard considerations as the explicit source of credit-market imperfections, and then examine the impact of redistribution on the output level and growth rate. In both papers, the source of moral hazard is the presence of limited liability; that is, the fact that a borrower's repayment to his lenders cannot be greater than his wealth. To understand the consequences of limited liability, think of an agent with no wealth who borrows in order to invest in a risky project. Suppose that the probability of success of the project depends on the (costly) effort exerted by the individual. If the project succeeds, the individual keeps the output minus the amount borrowed, while if it fails the individual incurs no loss as she has disbursed no funds. Consequently, the optimal amount of effort that she chooses to exert will be less than the amount her lenders would like her to.

What is important in order to find moral hazard is that effort is increasing in the wealth of the individual. As we will see, this occurs because the more an individual needs to borrow in order to get production started, the less incentives she has to supply effort, in that she must share a larger fraction of the marginal returns from her effort with lenders. An immediate consequence of this result is that redistributing wealth towards borrowers will have a *positive* effect on their effort *incentives*. Whenever this positive incentive effect more than compensates the potentially negative incentive effect on lenders' efforts, then such a redistribution will indeed be growth-enhancing based on incentive considerations only.

To illustrate the argument, we will now use a simple model of inequality and growth with credit market imperfections, based on Aghion and Bolton (1997). Specifically, we assume again the existence of a continuum of nonaltruistic, overlapping-generation families, indexed by $i \in [0,1]$. The utility of individual i in generation t is

$$U_t^i = c_{i,t} - h(e_{i,t}),$$

where $c_{i,t}$ denotes individual i 's second-period consumption (for simplicity we assume that individuals consume only when old), $e_{i,t}$ is the nonmonetary effort incurred by individual i when young and $h(e_{i,t}) = A_t e_{i,t}^2 / 2$ denotes the nonmonetary cost of effort. The parameter A_t still measures productivity on the current technology. The endowment of individual i is taken to be an idiosyncratic proportion of average knowledge at date t , that is, $w_{i,t} = \varepsilon_{i,t} \cdot A_t$.

The production technology involves an extreme form of U-shaped average cost curve with respect to capital investments, namely:

- the production activity requires a *fixed* and indivisible capital outlay equal to $k_{i,t} = \phi \cdot A_t$;

- conditional upon the required investment $\varphi \cdot A_t$ being made at date t , the output from investment in this technology is uncertain and given by

$$y_{i,t} = \begin{cases} \sigma \cdot A_t & \text{with probability } e_{i,t} \\ 0 & \text{with probability } 1 - e_{i,t}. \end{cases}$$

We assume that second-period outcomes $y_{i,t}$ are independently identically distributed across individuals of the same generation.

The source of capital market imperfection will be moral hazard with limited wealth constraints (or limited liability); in other words, the assumptions are that:

- individual efforts e_i are not observable;
- a borrower's repayment to her lenders cannot exceed her second period output $y_{i,t}$.

We can now analyze how effort choices are affected by the wealth of the individual. Consider first the effort decision of an individual who does not need to borrow, that is, for whom $w_i \geq \varphi A$. The problem she faces is²³

$$\max_e \left\{ e \cdot \sigma A - A \frac{e^2}{2} \right\},$$

which gives the first-best level of effort, $e^* = \sigma$.

An agent with initial endowment $w_i < \varphi A$ needs to borrow an amount $b_i = \varphi A - w_i$ in order to invest. Let r be the unit repayment rate. The individual chooses her effort to maximize the expected second-period revenue net of both repayment to the lenders and effort cost, namely

$$\max_e \left\{ e(\sigma A - r(\varphi \cdot A - \varepsilon_i \cdot A)) - A \frac{e^2}{2} \right\}.$$

The resulting level of effort is

$$e(r, w_i) = \sigma - r \left(\varphi - \frac{w_i}{A} \right),$$

which is less than the first-best effort e^* , and is decreasing in r and increasing in w_i . That is, for a given interest rate, the lower a borrower's initial wealth, the *less* effort she will devote to increasing the probability of success of her project.

Before turning to the analysis of redistribution, let us make two important remarks. First, individuals with initial wealth $w_i \geq \varphi A$ (in other words, the lenders), will systematically supply the first-best level of effort because they remain residual claimants on all returns from such effort: $e_i(w_i \geq \varphi A) = e^*$. Second, when analyzing the relationship between initial wealth and effort, we have treated the repayment schedule r as given. However, because the risk of default on a loan increases with the size of the loan (the probability of success $e(r, w)$ decreases when w decreases), the unit repayment rate r may vary with w to reflect the change in default risk. Aghion and Bolton (1997) show that even once this effect is taken into account, effort is still increasing in w_i .

The growth rate of the economy is given by $g_t = \ln(y_t/y_{t-1})$. Assuming, as in the previous subsection, that there is learning-by-doing, i.e. $A_t = y_{t-1}$, we can express it as

$$\begin{aligned} g_t &= \ln \frac{\int_0^1 e_i \cdot \sigma A_t di}{A_t} \\ &= \ln \sigma + \ln \int_0^1 e_i di, \end{aligned}$$

where $e_i \leq \sigma$.

If either (a) or (b) were violated, then the first-best effort, $e^* = \sigma$, would automatically be elicited from *all* individuals no matter what their human-capital endowments are. The growth rate would then be unaffected by the distribution of endowments and always be equal to

²³ For simplicity, we suppress time subscripts.

$g = \ln \sigma^2$. This corresponds to nothing but the case of *perfect* capital markets—that is, capital markets that do *not* suffer from incentive problems. When there are incentive problems, the more unequal the distribution of wealth is—that is, the larger the number of individuals with wealth below the threshold level φA —the lower the aggregate level of effort will be. Consequently, inequality has a negative effect on both the income level and the growth rate.

We now have all the elements we need to analyze the incentive effects of redistribution. Because individuals with initial wealth $w_i \geq \varphi A$ supply the first-best effort $e^* = \sigma$, raising a lump-sum tax $t_i < w_i - \varphi A$ on the endowment of each such individual and then distributing the total proceeds among borrowers will:

- i. not affect the effort e^* supplied by the wealthy, whose *after-tax* endowments remain strictly above the required fixed cost φA ;
- ii. increase the effort supplied by any subsidized borrower.

Such tax-subsidy scheme will then have an unambiguously positive *incentive* effect on output and growth, as efforts either increase or remain constant as a result of redistribution.

We have just put the traditional incentive-distribution trade-off upside-down, since we have shown that in the context of an imperfect credit market with moral hazard, redistribution enhances growth. An interesting question then arises quite naturally: is redistribution *self-sustaining*, in the sense that a one-time redistribution of wealth would have permanent effects on aggregate output and growth? Or, at the opposite, should redistribution policies be maintained on a more permanent basis? The above model would call for permanent redistribution policies, but to a large extent this conclusion follows from the as-

sumption that individual shocks on initial wealth are i.i.d. over time and across individuals.

The conclusion that redistribution policies must be sustained over time in order to have long-lasting effects on aggregate output was obtained by Aghion and Bolton (1997) in a model of inequality and growth in which individuals' wealth endowments are inherited from their own parents. The wealth distribution is therefore endogenous and varies over time, but under the above technological assumptions it can be shown to converge towards a unique *invariant* distribution.²⁴ Compared to the first-best distribution that would emerge in the absence of credit-market imperfections (i.e. in the absence of moral hazard), this (second-best) invariant distribution involves a higher fraction of the population lying *below* any level of wealth at any point in time. This, in turn, implies both that aggregate effort—and therefore aggregate output—is strictly lower than in the absence of moral hazard in credit; and that social mobility is also strictly reduced by the existence of credit market imperfections, in the sense that it takes longer for any individual borrower to eventually become a net lender.

²⁴ More specifically, the high degree of decreasing returns experienced by a rich individual on his own project leads to an ever-increasing supply of investable funds as the economy grows. This, in turn, pushes the gross interest rate downwards towards 1 whatever its initial level, and consequently the wealth distribution converges to a unique invariant distribution which is independent from initial conditions. On the other hand, Banerjee and Newman (1993) and Piketty (1997) obtain a multiplicity of invariant distributions, so that a one-time redistribution can have permanent effects on aggregate output. For example, in Piketty (1997) high-initial interest rates can be self-reinforcing through higher credit rationing to the poor and lower accumulation of funds by the rich. Interest rates can thus remain high in the long run. The resulting multiplicity of equilibrium interest rates also generates a multiplicity of invariant wealth distributions.

A suitable redistribution policy can then be designed which, by improving borrowers effort incentives, will increase both social mobility and aggregate output. However, because the invariant distribution of wealth under *laissez-faire* was unique—and therefore history-independent—a one-time redistribution cannot have permanent effects on aggregate output and social mobility. Lastly, it is worth pointing out that in this whole literature on inequality and economic development, both the technology and the institutional framework (in particular, the degree of capital market imperfections) are given once and for all and remain independent of the level of development of the economy. Our sense is that redistribution policies should be sustained over time, although in a way that accounts for the evolution of capital markets and credit institutions.²⁵

2.4 *Macroeconomic Volatility*

Empirical evidence suggests that macroeconomic volatility may be another channel through which inequality might affect economic growth. Several explanations have been put forward to account for the correlation between (high) inequality and macroeconomic volatility. Alesina and Perotti (1996) maintain that causality runs from high inequality to political instability, which in turn results in macroeconomic volatility. The approach we take in this subsection, based on Aghion, Banerjee, and Thomas Piketty (1997), postulates a direct effect of inequality on macroeconomic fluctuations. Inequality, however, takes the form of unequal access to investment opportunities across individuals, which, together with a high degree of capital market imperfection, can generate persistent credit cycles.

Specifically, we consider a dynamic economy in which only a fraction of the active population has access to high-yield investment opportunities. There are a number of reasons why access to investment opportunities may be restricted. Particular skills, ideas, or connections may be required, and often there may be crucial information that can only be acquired by those already in the business. Investment indivisibilities are another potential cause. Individuals may also differ in their attitudes towards risk, hence only those with little risk-aversion will be willing to undertake risky projects rather than work under a riskless employment contract. It is this inequality of access to investments and the consequent separation of investors and savers that will give rise to volatility.

Consider an economy where there are two production technologies: a traditional technology and a high-yield technology. Two crucial assumptions are needed for inequality to affect volatility:

1. Inequality of access to investment: Only a fraction of savers can directly invest in high-yield projects, whereas all individuals can invest in the low-yield technology.
2. Credit market imperfections: Because of incentive compatibility considerations, an investor with wealth w can borrow only a limited amount, νw , where $\nu < \infty$.

Now assume that all individuals in the economy save a constant fraction of their wealth, s . What do the saving and the investment functions look like? The total supply of funds in period t is a fraction s of the aggregate level of wealth in period $t - 1$. Savings at t are therefore independent of any variable in that period. The total demand for investment in the high-return project at time t is proportional to the wealth of those who have access to

²⁵ See Barro (1999).

the high-yield investment, and thus is completely determined by the previous period's income and by the (exogenously given) credit multiplier. There is therefore no market-clearing mechanism that will equalize the supply of funds and the demand for investment in the more productive technology. Consequently the economy will experience either "idle" savings (i.e. a fraction of savings are not invested in high-yield projects) or unrealized investment opportunities.

The link between inequality and volatility hence stems from the fact that those who invest and those who save are not the same individuals. During booms, investors' net wealth increases and therefore so does their borrowing capacity, *vw*. Investors can thus accumulate debt during booms, thereby increasing the demand for investable funds. The interest rate is given by the marginal product of capital. Since all funds are invested in the high-yield technology, interest rates are high during booms. Eventually, the accelerated increase in their debt repayment obligations ends up squeezing the borrowing capacity of investors who have access to the high-yield technology. In the meantime, savings keep on increasing due to the rise in wealth during the boom. This results in a fraction of savings becoming idle, which then have to be invested in the traditional technology. At this point the economy experiences a slump: the marginal product of capital falls and interest rates drop. This in turn allows the investors to progressively reconstitute their borrowing capacity, and so eventually the economy will reenter a boom. If the fraction of the population with high-yield investment possibilities is small enough and/or the credit multiplier low enough, there will be continuous oscillations of the investment level. Such volatility of investment in turn implies that there are unexploited produc-

tion possibilities and hence the long-run growth rate is lower than it could be.²⁶

The government has two structural policy options to try to move the economy out of the above cyclical equilibrium into a situation in which all savings are invested in the high-return production technology. One is to reduce the borrowing constraints, thus increasing the credit multiplier and ensuring that there is sufficient demand for funds. This is, however, a hard policy to implement unless the government is willing to lend to individuals itself. Moreover, if the credit constraint is the result of a moral hazard problem, such as that examined in subsection 2.3, it would not be possible to increase the credit multiplier without generating adverse incentive effects. A second structural policy consists in reducing the degree of inequality of access to investment. By increasing the fraction of savers that can directly invest in high-yield projects, the economy can move to a permanent-boom situation and thus increase its growth rate. Structural reforms such as investing in infrastructure or in human capital, or reducing the bureaucratic obstacles faced by entrepreneurs that wish to set up a firm, would reduce entry barriers and promote growth.

Structural policies may be hard to implement though, especially in the short-run. An alternative would be to transfer the idle funds from savers to investors.

²⁶ Recent work by Aghion, Philippe Bacchetta, and Banerjee (1998) considers a small open economy extension of the same framework and develops a model of financial crises. The real interest rate is fixed at the international market-clearing level and the transmission variable becomes the price of non-tradeable goods in terms of the tradeable good. When investors are constrained in their borrowing, the accumulation of debt results in an increase in the relative price of non-tradeables relative to tradeables. At some point, the credit market constraint becomes binding and brings about the collapse in the price of nontradeables, that is, a financial crisis.

This policy ensures that all savings are invested in the high-yield technology. The problem is that it transfers resources from those that are worse off to those that are better off. Yet Aghion, Banerjee, and Piketty (1997) show that this policy does not entail negative distributive consequences for savers. The higher level of income trickles down to savers for two reasons: first, the interest rate is higher, so (poor) lenders are better off; second, as more capital is invested in the high-yield technology, the productivity of labor, and thus the wage rate, are also higher.

2.5 Discussion

The main conclusion we can draw from this section is that when agents are heterogeneous and capital markets imperfect, greater inequality may have a negative impact on growth. Moreover, the traditional argument that redistribution is detrimental to incentives and growth is strongly challenged. The mechanisms we have discussed in the above subsections crucially rely on the presence of credit market imperfections, and hence on the limited degree of financial development in an economy. Financial constraints have been found to play a crucial role in the entrepreneurial process. The results obtained by David Evans and Boyan Jovanovic (1989) point at the importance of liquidity constraints in the building up of enterprises in the US. According to the data, capital is essential for starting a business and liquidity constraints tend to exclude those with insufficient funds at their disposal. The problem is likely to be accentuated in other countries. Indeed, if we take the UK and the US as the benchmark for measuring financial development, we find that most industrial economies (and, of course, all the so-called emerging market economies) suffer from in-

sufficiently developed capital markets.²⁷ This means that we can expect such mechanisms to be at work not only in less-developed countries, but also in many developed economies.

Our theoretical framework suggests that when capital markets are highly imperfect and the production technology exhibits diminishing returns to capital, inequality in the distribution of wealth is bad for growth. Redistribution from the rich to those who are poorly endowed with physical or human capital creates investment opportunities, thus fostering growth. A possible interpretation of the opportunity-enhancing model is in terms of education investments. Investments in human capital are characterized by strong diminishing returns. Moreover, borrowing in order to make such an intangible investment is usually expensive, particularly in developing countries,²⁸ and hence family wealth becomes a major determinant of the size of the investment. If we view $k_{i,t}$ as an investment in education and g as the rate of growth of human capital (which in turn determines the rate of output growth, as argued by Robert Lucas 1988), then our analysis predicts a negative relationship between wealth inequality and the rate of growth. These are the conclusions obtained by several papers that analyze how inequality affects growth through agents' investments in education.²⁹

²⁷ Rafael La Porta, Florencio Lopez-de-Silanes, Andrei Shleifer and Robert W. Vishny (1997) provide measures of the size of capital markets. The debt to GNP ratio is of similar magnitude in France and Germany as in the UK and the US (around 1), but it is only 0.55 in Italy. However, the equity to GNP ratio is 1 in the UK, 0.58 in the US, and much lower in the other three countries (0.23 in France, 0.13 in Germany, and 0.08 in Italy).

²⁸ See World Development Report (1990) and George Psacharopoulos (1986).

²⁹ For example, Galor and Zeira (1993), Perotti (1993), and García-Peñalosa (1995).

Second, wealth inequality is detrimental to growth, as borrowers tend to underinvest in effort when effort is unobservable and there is limited liability. Two remarks deserve to be made at this point. First, the mechanism we have described does not rely on the presence of fixed costs. Piketty (1997) shows that the results we obtained in subsection 2.3 also hold when the returns in the event of success are a continuous function of the amount invested. Second, the ex-ante moral hazard formulation we have presented can be applied not only to investments in physical capital, but also to human capital investments and, more interestingly, to research activities. Consider for a moment an economy where the sources of growth are technological improvements brought about by innovations. Suppose also that innovations result from individual research efforts, and that there are decreasing returns to such efforts. Then, the analysis in subsection 2.3. implies that redistribution from lenders to borrowers improves the borrowers' incentives to provide the adequate level of effort in the handling of investment projects. Here again, redistribution through a lump-sum tax/transfer system may be growth enhancing. This does not mean that the traditional negative incentive effect has disappeared. If redistribution were financed through distortionary (ex-post) taxation, there would be two incentive effects which conflict with one another: the standard effect whereby taxation reduces net returns and therefore lenders' incentive to invest, and moral hazard with limited liability which decreases the effort exerted by borrowers. Then, whether redistribution raises or reduces the rate of growth depends on whether the standard effect of taxation on those individuals with large wealth endowments is smaller or greater than the positive im-

pact on the effort of those with small endowments.

In both the opportunity and incentive enhancing models, inequality arises from the initial wealth distribution. It can then be reduced, and growth thereby fostered, by means of taxes and subsidies on individuals' endowments. Redistribution takes the form of standard fiscal policy. This is no longer the case when the main source of inequality does not lie in the distribution of wealth, but rather in the social and institutional environment that affects access to investment projects. In this case, inequality gives rise to a spurious dynamics in which fast growth is associated with high volatility, thus endangering its own sustainability. Redistribution must then take the form of structural policies aiming at restructuring core financial institutions. Such policies are obviously far more difficult to devise and implement than a mere tax-subsidy mix. Enhancing growth by means of redistributive policies may thus prove highly demanding, depending on the extent to which inequality originates from the distribution of wealth or from institutional characteristics.

Finally, note that our analysis in this first part of the survey questions the idea—also associated with the Kuznets hypothesis—that for poor countries inequality is more likely to be positively correlated with growth. On the contrary, as our previous analysis suggests, the less developed the credit markets and the larger the separation between borrowers and investors, the bigger the scope for redistributive policies aimed at creating opportunities, improving borrowers' incentives, and reducing macroeconomic volatility.

Now, if reducing inequality fosters growth in particular in economies at an early stage of development, an important question arises: is there any virtuous circle? In other words, does economic

development bring about a steady decrease in inequality which in turn would reduce the need for sustained redistributive policies at a later stage of development? This question refers to the other half of the Kuznets hypothesis, namely the half that deals with inequality and growth in richer countries. However, as we already argued in the introduction, many OECD countries have experienced a sharp increase in income and wage inequality over the past twenty years. The next section considers and discusses three candidate explanations for this phenomenon—namely, trade liberalization, skill-biased technical change, and organizational change—all of which are closely related to the growth process.

3. *Does Growth Increase Earnings Inequality?*

3.1 *Recent Trends in Income and Earnings Inequality*

If one thing is clear from the analysis of times-series data over the last twenty years, it is that there are no universal trends. Anthony Atkinson (1996) illustrates the diversity of experiences in a number of industrial economies, all at similar levels of development and experiencing similar rates of growth. Since 1970, the Gini coefficient of (gross) household income has increased sharply in the US and the UK, remained roughly constant in Germany, and decreased in France, Italy, and Canada. Such differences do not come as a surprise. The determinants of the degree of income inequality in a country include social and political forces as well as economic ones.³⁰ In particular, government transfers are the second larg-

³⁰ See Atkinson (1996, 1997) and Piketty (1996) for illuminating surveys of the relevant empirical and theoretical literature on the determinants of income inequality.

est source of household income,³¹ suggesting that even if growth matters in shaping the distribution of income, policy choices also play a crucial role.

In seeking to understand the impact of growth and technical change on inequality, it therefore becomes more appropriate to concentrate on the distribution of earnings rather than of income. The traditional approach, emphasized by development economists, has concentrated on how labor-saving technological change would affect capital and wage income. However, in what follows, we are going to examine only the distribution of labor earnings.³² There are two reasons for doing this. One is simply that labor income is the main source of personal and household income, and hence its distribution has major implications for income inequality. The other is that, in developed economies, the increase in both borrowing and saving across all income groups makes it difficult to assess the distributional consequences of an increase in the interest rate at the expense of wages.³³

General interest into questions of how growth affects inequality has been revived in the last few years by the widening in earnings distribution experienced by a number of the world's most developed economies. Recent empirical studies³⁴ have pointed to a substantial

³¹ In 1993, social security benefits accounted for 14 percent of household income in the UK (Atkinson 1997, p. 305), and government transfers reduced the Gini coefficient from 52 percent to 37 percent (Atkinson 1996, figure 2.7).

³² We will use the terms wage distribution and earnings distribution interchangeably, even if this is not entirely accurate. Data on earnings includes labor incomes other than wages, such as those of self-employed workers.

³³ See Atkinson (1996), section 2.

³⁴ Such as those by Kevin Murphy and Finnis Welch (1992), Chinhui Juhn, Murphy, and Brooks Pierce (1993), and Machin (1996a). Gottschalk and Smeeding (1997) review this literature in an earlier issue of this journal.

TABLE 2
BETWEEN-GROUP CHANGES IN WAGES IN THE UK AND THE US

<i>Differential</i>	Late 1960s– Early 1970s		Late 1970s– Early 1980s		Late 1980s– Early 1990s		
	Year	Ratio	Year	Ratio	Year	Ratio	
<i>Educational Differentials (males)</i>							
UK	University/no-qualification	1974	1.64	1980	1.53	1988	1.65
US	College/High school	1969	1.49	1979	1.37	1987	1.51
<i>Occupational Differentials (manufact.)</i>							
UK	Non-manual/manual	1970	1.35	1980	1.31	1990	1.49
US	Non-production/production	1970	1.56	1980	1.53	1990	1.64
<i>Age Differentials</i>							
UK	40–49/21–24	1974	1.27	1980	1.29	1990	1.36
US	45–49/20–24	1970	1.76	1980	1.88	1990	2.14

Source: Machin (1996a), Table 4, p. 52.

increase in earnings inequality in several OECD countries during the past twenty years. This has been the case in the US, the UK, Canada, New Zealand, Australia, but also in Austria, Belgium, Japan, and to a lesser extent Sweden. Although some countries have escaped this rise—earnings inequality remained quite stable in Finland and France over the period, while it declined in Germany and Italy, at least until the beginning of the 1990s—most developed economies saw the secular trend of stability or even reduction in earnings inequality reversed in the 1980s.

The most spectacular rise has undoubtedly taken place in the UK and North America. The ratio of the 90th to the 10th percentile of the male wage distribution rose from 2.53 to 3.21 in the UK between 1980 and 1990, and from 4.76 to 5.63 in the US over 1980–89.³⁵ In both countries, these changes can be decomposed into three main elements:

- (a) An increase in *educational* wage differentials, that is, in wage inequality across different educational cohorts.

Between 1980 and 1988, the wage ratio of university graduates to workers with no qualification increased by almost 8 percent in the UK, and the wage ratio of college to high school graduates rose by over 10 percent in the US over 1979–87 (see Table 2).

- (b) An increase in *age-related* wage differentials. Over the decade between 1980 and 1990, the ratio of wages of older to younger workers rose by 5.4 percent in the UK and by 13.8 percent in the US (see Table 2).
- (c) An increase in within-group wage inequality, that is inequality which is not accounted for by the above between-group changes. Yona Rubinstein and Daniel Tsiddon (1998) present evidence of an increase in wage inequality within educational groups in the US. Wage dispersion has increased since 1970, for both high school and college graduates. Moreover, they find that dispersion is generally greater for more educated workers.

Further evidence is provided by Stephen Machin (1996a) who looks at the residual standard deviation obtained from simple human-capital

³⁵ See OECD *Employment Outlook* (1993).

equations where the logarithm of real hourly earnings is regressed on age (and its square) and the number of years of schooling. His estimates on British and American data yield similar results: the residual standard deviation in real hourly earnings increased by 23 percent in the UK and by 14 percent in the US between 1979 and 1993.

The US and the UK are the only countries that have experienced large and simultaneous increases in the returns to both experience and education, as well as within-group inequality. However, many other OECD countries—except Germany, Finland and Italy—have experienced some rise in at least one of these three components over the 1980s. As reviewed by Peter Gottschalk and Timothy Smeeding (1997), returns to experience have sharply increased in Australia, Canada, France, Israel, and the Netherlands. Evidence regarding the returns to education is more mixed but they have increased in Israel and Sweden, while within-group inequality also displays an upward trend in a number of countries. Overall, Gottschalk and Smeeding conclude that: “what we observe is a diversity of experiences but with almost all countries experiencing some increase in earnings inequality” (Gottschalk and Smeeding 1997, p. 652).³⁶

In order to explain earnings patterns, we need to understand the behavior of the supply and demand for different kinds of labor. As Jan Tinbergen (1975) pointed out, observed relative wages are the outcome of a “race” between the forces increasing the supply of skills—mainly education, but also experience—and those increasing the de-

mand for skills required by firms—technical change. Up to the late 1970s, this race was won by supply forces, thus resulting in falling relative wages.³⁷ However, the recent trends we have just described and the absence of a marked change in labor supply trends³⁸ indicate that there has been an *acceleration* of the rate of growth of the relative demand for skills, which in turn has increased the skill premium. Then, the question is not *why* has demand been steadily increasing—the only answer to this is technical change—but rather *what* has caused this acceleration.³⁹

The timing of this change has coincided with two events. One of them is the rapid growth of imports from developing countries, caused by cheaper and faster transport and telecommunications, a switch to export-oriented policies in many developing economies, and the rapid growth of East-Asian countries. The other is the rapid diffusion of computers in the work place. Hence, two competing explanations have been proposed for this structural change in the relative demand for skilled labor: the *impact of trade* with the rapidly growing East Asian economies, and *skill-biased technological change*. More recently, a third explanation has been put forward that emphasizes the role of *organizational change* within firms. It argues that the specific way in which workers interact and learn in the workplace is likely to be crucial in determining their productivity, and hence wages. The purpose of the second part of the

³⁷ Tinbergen, writing in 1975, stated that “in the past 70 years that race has been *won by education*” thus resulting in falling relative wages, but his extrapolation of income and education trends predicted that by 1990 the race “*will be lost by education*” (Tinbergen 1975, p. 103).

³⁸ See George Johnson (1997), who documents the absence of a substantial change in the rate of growth of skilled labor.

³⁹ See Adrian Wood (1998).

³⁶ See also Steven Davis (1992) for a comparative study of the evolution of relative wages in thirteen countries, among which four middle-income economies are included.

paper is to examine the validity of these hypotheses.

Our analysis is organized as follows. Section 3.2 evaluates the trade hypothesis, and discusses a growing empirical literature that has tried to sort out how much of the observed rise in the skill premium can be explained by trade versus skill-biased technical change. These studies claim to present evidence that trade can only explain a small fraction of the change in wage inequality. However, once we take into account the role of imports as intermediate production factors, the trade hypothesis can account for the steady rise in wage inequality across educational groups. Section 3.3 considers the role of skill-biased technical change as a possible explanation for the observed increase in earnings dispersion. First, in subsection 3.3.1, we analyze the effect of *disembodied technical change*, and in particular we argue that the acceleration in the diffusion of new technologies can result in episodes of increasing wage inequality across skill groups. Subsection 3.3.2 focuses on *embodied technical change*. We introduce learning-by-doing and intersectoral mobility considerations to provide an explanation for the observed increase in wage inequality within educational groups. Lastly, section 3.4 focuses on the impact that changes in the *internal organization of firms* have had on the productivity gap between individuals with different skills or experience.

3.2 International Trade

3.2.1 The Heckscher–Ohlin Model

The argument that trade is responsible for the increase in wage inequality stems largely from Heckscher–Ohlin theory.⁴⁰ According to it, countries spe-

cialize in the production of those commodities that use intensively the factors of production they are abundantly endowed with. Developing countries that are abundant in unskilled labor but scarce in skilled labor tend to export goods that are intensive in the former. Developed countries export skill-intensive commodities, such as computer software, and import labor-intensive manufactures and primary products. Under such specialization conditions, a globalization boom causes predictable inequality trends. In the poor country—where abundant unskilled labor is cheap and scarce skilled labor is expensive—the trade boom drives up the demand for unskilled labor and drives down the demand for skilled labor, thus erasing some earnings inequality. In the rich country—where (relative to the poor country) unskilled labor is expensive and skilled labor is cheap—the trade boom drives up the demand for the latter and drives down the demand for the former.

The argument is best illustrated with the aid of a simple figure, taken from a paper by Adrian Wood and Cristobal Riddo-Cano (1996). Suppose there are two production factors: skilled labor and unskilled labor. There is a large number of goods of different skill intensities. Now consider two countries. Country A has a lower relative endowment of skilled labor than country B. Figure 1 depicts the relationship between the ratio of the skilled to the unskilled wage, denoted ω , and the relative labor supply or fraction of the labor force that has skills, denoted s . The given skill ratios are represented by the two vertical lines, where SA is the supply of skills in country A and SB that in country B.

When the two countries start trading, the skill-rich economy will experience an increase in the price of skill-intensive

⁴⁰ Adrian Wood was one of the first to put forward this hypothesis in his book *North-South Trade, Employment and Inequality* (1994).




Figure 1. Trade and Relative Wages

Source: Wood and Ridao-Cano (1996), Figure 2, p. 3.

commodities, and hence an increase in the relative demand for skilled labor. Similarly, country A will experience an increase in the price of the labor-intensive goods and a reduction in the demand for skills. That is, the demand curve becomes flatter, represented by the $D'D'$ line. If wages are flexible, this change in demand results in a higher skill premium in country B and a lower skill premium in country A than under autarchy. Thus, earnings inequality increases in the rich economy. Alternatively, if the wage of unskilled labor is to some extent rigid, the fall in the relative demand for unskilled labor would manifest itself in a rise in unskilled unemployment. Trade, and the resulting shift in demand, is often seen as the cause of the increase in the skill premium in the UK and the US, and of

higher unemployment among the unskilled in the rest of Europe, where labor market rigidities have maintained the level of unskilled wages.⁴¹

3.2.2 An Empirical Test

Conventional trade theory then predicts that the increase in exports of goods intensive in unskilled labor, due to improved transport and rapid growth in East Asia in the past two decades, would cause the rise in skilled-wage differentials observed in the UK and the US. A growing empirical literature has tried to sort out how much of the increase in the skill premium can be explained by trade with newly industrializing countries, as opposed to alternative

⁴¹ This distinction is, however, somewhat oversimplified, as documented by Stephen Nickell and Brian Bell (1996) and Nickell (1997).

explanations such as skill-biased technical change. The empirical tests have mainly concentrated on a specific implication of the model we have just developed. Trade between high-skill and low-skill economies should cause, in the former, an increase in the demand for (domestic) skill-intensive commodities at the expense of the demand for domestic unskilled intensive commodities. This, in turn, would induce a reallocation of labor *between* low-skill and high-skill industries. In contrast, skill-biased technical change would induce a shift in labor demand towards skilled labor *within* all industries.

Eli Berman, John Bound, and Zvi Griliches (1994) test this hypothesis by decomposing the variations in the share of white-collar workers in employment.⁴² Let $E_i = X_i/N_i$ be the share of white-collar workers (X) in total manufacturing employment (N) of industry i , and $S_i = N_i/N$ be the share of industry i 's employment in total employment. Then the aggregate change in the share of white-collar workers in employment, ΔE , can be expressed as

$$\Delta E = \sum_{i=1}^I \Delta S_i \bar{E}_i + \sum_{i=1}^I \Delta E_i \bar{S}_i,$$

where the bar over a variable indicates that the variable is kept constant at its start-of-the-period value. The first term then reflects changes in the share of employment between industries, as caused by an increase in the relative demand addressed to white-collar intensive industries due, for example, to trade liberalization. The second term captures the within-industry component, that is, skill upgrading likely to be due to skill-biased technical progress.

Berman, Bound, and Griliches find

⁴² This blue/white-collar distinction proves to be highly correlated with a high-school/college educational classification.

that, in the United States, the within component is, by far, the most important one. It accounts for 70 percent of the rise in the white-collar share in employment between 1979 and 1987. Using a similar method for the UK between 1979 and 1990, Machin (1996b) reports that 82 percent of the increase in the nonmanual share is due to within-industry shifts. Hence, only a minor part of the shift away from manual/blue-collar workers to non-manual/white-collars is due to between-industry changes.

Several other studies arrive at similar conclusions. For instance, George Borjas, Richard Freeman, and Lawrence Katz (1992) obtain that a maximum of 15 percent of the growth of the college-noncollege wage differential in the US is due to imports. A crucial link in the argument is that, for relative wages to increase, there should be a fall in the prices of less skill-intensive goods relative to those of skill-intensive goods in rich countries. Empirical studies find little evidence that the relative prices of less skill-intensive goods have fallen in either the United States or Europe during the 1980s, as the trade arguments would require.⁴³ Paul Krugman (1995) summarizes this research as indicating that the effect of East Asian imports on industrialized countries' labor markets has been small.

3.2.3 *Trade in Intermediate Goods*

While the above-mentioned evidence points to skill-biased technical change as an unavoidable part of the wage inequality story, the fact that the shift in labor demand has taken place mainly

⁴³ Matthew Slaughter (1998) reviews the evidence on price changes in the US, and Robert Lawrence (1996) and Damien Neven and Charles Wyplosz (1998) provide evidence on Europe. However, Wood (1998) argues that this evidence cannot be interpreted as refuting the trade hypothesis.

within industries (and in both traded and nontraded good sectors) does not suffice to disregard trade liberalization when trying to explain the observed increase in wage inequality between skilled and unskilled labor. The argument crucially relies on the assumption that traded goods are primarily final goods, and, consequently, trade liberalization would shift demand from unskilled-intensive to skill-intensive goods. If, on the other hand, traded goods were mainly *inputs* into further production, the implications would be very different.

A lower price of intermediate goods would shift the demands for other inputs, increasing that of complementary production factors and reducing the demand for substitutes. If unskilled labor were more substitutable for intermediate goods than skilled labor, cheap intermediate goods would increase the relative demand for skilled workers, shifting the within-industry labor demands in all industries that use such inputs, irrespectively of whether they themselves produce traded or untraded goods and of whether they are skill—or unskilled-labor intensive. In this scenario, *trade liberalization that reduces the price of intermediate goods would shift relative labor demands but not the demand for final goods*. Consequently, one cannot interpret the evidence of an absence of between-industry shifts as refuting the role of trade.

Recent work by Martin Falk and Bertrand Koebel (1997) and by Koebel (1997) lends support to this hypothesis, as they find evidence that unskilled labor is more substitutable to material inputs than skilled labor in the manufacturing and construction sectors in Germany. Falk and Koebel estimate a system of input demand functions for five factors of production: materials, capital, unskilled labor (no qualifica-

tions), skilled labor (high school degree or equivalent), and high-skilled labor (university degree). Each factor demand is estimated as a function of the prices of all five inputs, the level of output and a time trend. The main aim of the exercise is to obtain cross-price elasticities of demand, which will tell us whether two inputs are complements or substitutes. The cross-price elasticity is defined as $\varepsilon_{ij} = \partial i^* / \partial p_j \cdot p_j / i^*$, where i^* is the amount of factor i demanded and p_j is the price of factor j . A value of $\varepsilon_{ij} > 0$ implies that goods i and j are substitutes, while $\varepsilon_{ij} < 0$ means that they are complements. The system is estimated for biannual German data, covering the period 1977–94. The estimated degree of substitutability between material inputs and unskilled labor in the manufacturing and construction sectors is large, with elasticities of substitution respectively equal to 0.14 and 0.34. The demand for high-skilled labor is increased by a lower price of material inputs in some sectors and is unaffected in others. Work in progress by García-Peñalosa, Koebel, and Cameron shows that, even if there are no direct complementarities between materials and skilled labor, the fact that unskilled workers and material inputs are substitutes suffices for a reduction in the price of material inputs to increase the relative demand for high-skilled labor.

Such relative price changes actually occurred in Germany and in the US during the 1980s. The price of material inputs relative to that of unskilled labor fell by an average 2.4 percent per year in Germany between 1977 and 1994, and by 1.3 percent per year in the US over the period 1979–87 (see respectively Koebel 1997, and Ana Revenga 1992). Given the elasticities of substitution reported above, this substantial reduction in the relative price of materials may have been a cause of the upward

shift in the demand for high skilled labor experienced in both countries.⁴⁴

The extent to which the reduction in the relative price of materials in developed countries is due to trade liberalization has, to our knowledge, not been investigated. Nevertheless, Trudy Owens and Wood (1997) provide some insights about import and export patterns of different countries. They divide traded goods into manufactures and a broad definition of primary goods, which includes processed primary products such as gasoline. Primary products represent less than 30 percent of total exports in developed countries and the East-Asian “high-performing” economies. For all other countries, primary commodities account for between 45 and 90 percent of all exports. Although the differences are less marked for imports, Owens and Wood show that the share of primary products in total imports tends to increase with a country’s income level. The large share of primary products in exports from poorer to richer countries suggests that increased openness and the reduction of trade barriers may have largely contributed to the decline in the relative price of material inputs by industrialized economies.

Evidence that intermediate goods have played an important role in the change in relative wages is provided in a recent paper by Robert Feenstra and Gordon Hanson (1999). They examine the relative influence of technology and trade on relative wages in the US during the 1980s. They measure technology by the share of computers in total capital, and trade by the degree of “out-

sourcing,” defined as the proportion of intermediate goods that are imported. Their results imply that, although computer use is important, outsourcing can explain between 15 and 40 percent of the change in the relative wage of nonproduction labor.

3.3 *Technical Change*

A growing literature argues that the shift in the relative labor demand has been caused by technological change. Now, if technological change is to generate an increase in wage inequality, it must be because technological change is *biased* toward certain skills or specializations, in the sense that it reveals and enhances new differences in abilities among workers across or within educational cohorts.⁴⁵ We will argue in the next two subsections that this may come as a result of both disembodied and embodied technical change.

3.3.1 *Disembodied Technical Change*

3.3.1.a *A Basic Explanation Based on General Purpose Technologies*

The rise in the wage premium for skilled workers raises a theoretical puzzle. Although technological change can exert an upward pressure on the demand for skilled workers and thereby increase their wage premium over unskilled workers, education should eventually lead to an expanded supply of skilled labor and thereby to a fall in the wage differential.⁴⁶ In what follows we

⁴⁵ See Juhn, Murphy, and Pierce (1993) and Piketty (1996).

⁴⁶ This approach is best captured in the work of Theo Eicher (1996), who examines the interaction between endogenous human capital accumulation and technological change. Skilled labor is assumed to be an essential input in education, research, and in the absorption of innovations into production. The absorption of bursts in technological change then requires the withdrawal of skilled labor from research and education which subsequently increases the cost of human capital

⁴⁴ In Germany, the 1980s witnessed a rapid decline in the level of employment of the least qualified workers, some increase for those with intermediate qualifications, and a fast rise in employment of highly-educated workers, together with a moderate increase in the skilled–unskilled wage ratio (5 percent between 1980 and 1994). See Falk and Koebel (1997), p. 21.

examine how the adoption and implementation of new technologies may result in a situation in which the wage premium grows even when the relative supply of skilled workers is increasing.

One of the features that stand out in our review of the evidence on the dynamics of wage inequality is the episodic nature of upswings and downswings. A natural explanation for this pattern is the acceleration in the diffusion of new “general purpose technologies” (GPTs).⁴⁷ Let us pause for a moment and explain in more detail what we have in mind here.

A GPT is a technological invention (or breakthrough) that affects the entire economic system. However, whilst each GPT raises aggregate output and productivity in the long run, it also causes cyclical fluctuations while the economy adjusts to it. Examples of GPTs include the steam engine, the electric dynamo, the laser, and the computer. As argued by economic historians such as Paul David (1990), there are several reasons to believe that the diffusion of a new GPT to the entire economy should be *nonlinear*. For example, the existence of *strategic complementarities* (or *network externalities*) between the various sectors of the economy may generate temporary lock-in effects, of the kind already emphasized by Andrei Shleifer (1986) in his paper on implementation cycles. Shleifer’s central idea is that a firm chooses not to implement the new GPT as long as no one else does, unless forced to do so by some “exogenous” factor such as the continuous rise in labor costs.

Externalities can also arise in the

transmission of skills from parent to offspring. If these externalities are specific to a particular sector, they will affect the distribution of skills during the phase of implementation of a GPT, and hence the rewards received by workers. Galor and Tsiddon (1997) argue that there are two components that determine individual earnings: sector-specific skills inherited from parents and individual ability. The arrival of a new GPT erodes the stock of specific human capital, enhancing mobility across sectors. This generates a larger concentration of individuals with high levels of ability and human capital in the technologically advanced sectors, and earnings inequality rises. As existing technologies become more accessible, the effect of parental specific skills starts to dominate. Mobility is diminished and inequality declines. The cyclical nature of technological progress (i.e., inventions followed by innovations) together with an external effect of parental choice of sector of employment give rise to cyclical movements of the wage differential between skilled and unskilled workers.⁴⁸

Another potential source of non-linearity in the diffusion of new GPTs which we formalize below lies in the phenomenon of *social learning*. That is, the way most firms learn to use a new technology is not to discover everything on their own but to learn from the experience of other firms in a similar situation. For a firm to learn from other firms, the problems to be solved before

investment and research. Higher rates of technological change are thus accompanied by a higher relative wage but lower relative supply of skilled labor.

⁴⁷This analysis draws on Aghion and Howitt (1998), ch. 8.

⁴⁸Rubinstein and Tsiddon (1998) provide empirical support for this approach. They find that, in the US, parent’s education plays a crucial role in determining the wage of an individual. Controlling for education of the child, the impact of parental education on wages is greater in the 1980s than in the 1970s. However, the return to college education for individuals whose parents did not have a college degree is the same in the two decades. The correlation between the parent’s and the child’s education has also increased from the 1970s to the 1980s.

the technology can successfully be implemented must bear enough resemblance to the problems solved by others so that it is worthwhile trying to use the procedures of those successful firms as a “template” on which to prepare for adoption. Thus, the fact that at first no one knows how to exploit a new GPT means that almost nothing happens in the aggregate. Only minor improvements in knowledge take place for a long time, because successful implementation in any sector requires firms to make independent discoveries with little guidance from the successful experience of others. But if this activity continues for long enough, a point will eventually be reached when almost everyone can see enough other firms using the new technology to make it worth their while experimenting with it. Hence, even though the spread of a new GPT takes place over a long period of time, most of the costly experimentation through which the spread takes place may be concentrated over a relatively short subperiod, during which there is a cascade or *snowball effect* resulting in an *accelerated* demand for skilled labor. This in turn will cause the skill premium to rise.

More formally, suppose that aggregate output is produced by “labor” according to the constant return technology:

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

where $A(i) = 1$ in sectors where the old GPT is still used, and $A(i) = \gamma > 1$ in sectors that have successfully innovated, while $x(i)$ is manufacturing labor used to produce the intermediate good in sector i . The total labor force L is actually divided into skilled and unskilled workers. Whilst old sectors, that is those with $A(i) = 1$, can indifferently use skilled and unskilled workers, the experimentation

and the implementation of the new GPT requires skilled labor.

For simplicity, we do not detail the supply side of the labor market, but simply assume that the fraction of skilled workers is monotonically increasing over time, for example as a result of schooling and/or training investments. The supply of skilled labor at time t is then given by:

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

where $s < 1$ is the initial fraction of skilled workers and β is a positive number measuring the speed of skill acquisition.

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

We assume that in each sector i , moving from the old to the new GPT requires two steps. First, a firm in that sector must acquire a “template” on which to base experimentation. Second, the firm must use this template to discover how to implement the GPT in that particular sector. Let n_0 denote the fraction of sectors that have not yet acquired a template; n_1 denote the fraction of sectors that have acquired a template but are still experimenting; and $n_2 = 1 - n_0 - n_1$ the fraction of sectors that have succeeded in making the transition to the new GPT.

A sector will acquire a template if a firm in that sector either makes an independent discovery or finds its template by “imitation,” that is, by observing at least k “similarly located” firms that have made a successful transition to the new GPT. The Poisson arrival rate of independent discoveries to such a sector is $\lambda_0 < 1$. The Poisson arrival rate of opportunities to observe m similarly




Figure 2. Fraction of Sectors that Have Adopted the GPT

located firms is assumed to equal unity. The probability that such an observation will pay off (in other words, the probability that at least k among the m similar firms will have successfully experimented the new GPT) is given by the cumulative binomial:

$$\varphi(m, k, n_2) = \sum_{j=k}^m \binom{m}{j} n_2^j (1 - n_2)^{m-j},$$

since n_2 is the probability that a randomly selected firm will have succeeded in implementing the new GPT. The flow of sectors that acquire a template and can thus start *experimenting* on the new GPT, will be equal to n_0 times the flow probability of each sector making the transition: $\lambda_0 + \varphi(m, k, n_2)$.

Now we assume that for an experimenting firm to actually succeed in implementing the new GPT, it must employ at least H units of skilled labor per period. We can think of this labor as being used in formal R&D, informal R&D, or in an experimental start-up firm. In any case it is not producing current out-

put. Instead, it allows the sector to access a Poisson process that will deliver a workable implementation of the new GPT with an arrival rate of λ_1 . Thus the flow of new sectors that can implement the new GPT will be the number of experimenting sectors n_1 , times the success rate per sector per unit of time λ_1 .

We can summarize the discussion to this point by observing that the evolution over time of the two variables n_1 and n_2 is given by the autonomous system of ordinary differential equations:

$$\dot{n}_1 = [\lambda_0 + \varphi(m, k, n_2)](1 - n_1 - n_2) - \lambda_1 n_1$$

$$\dot{n}_2 = \lambda_1 n_1$$

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

Figure 2 depicts the solution to the above system.⁴⁹ Not surprisingly, the time-path of n_2 follows a *logistic* curve, accelerating at first and slowing down as n_2 approaches 1, with the maximal

⁴⁹ Figure 2 displays the behavior of n_1 and n_2 in the case where $\gamma_0 = 0.005$, $\lambda_1 = 0.3$, $m = 10$, and $k = 3$.




Figure 3. Wages Relative to the Initial Wage

growth rate occurring somewhere in the middle. Likewise the path of n_1 must peak somewhere in the middle of the transition, in as much as it starts and ends at zero. If the arrival rate λ_0 of independent discoveries is very small then both n_1 and n_2 will remain near zero for a long time. The number of sectors engaging in experimentation peaks sharply in year 20 due to social learning.

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

labor and determined by the labor market clearing equation:

$$(1 - n_2) \cdot x_0 + n_2 \cdot x_N + n_1 \cdot H = L$$

$\underbrace{\hspace{2em}}$
 labor
 demand by
 an old sector

$\underbrace{\hspace{2em}}$
 labor
 demand by
 a new sector

$\underbrace{\hspace{2em}}$
 labor demand by
 an experimenting
 sector

However, in the later phase of transition, where the fraction of new sectors has grown sufficiently large that it can absorb the whole skilled labor force, the labor market will become segmented, with skilled workers being exclusively employed (at a higher wage) by new sectors whilst unskilled workers remain in old sectors. Let w_u and w_s denote the real wages respectively paid to unskilled and skilled workers. We now have $w_s > w_u$, since the two real wages are determined by two separate labor market clearing conditions. The skilled wage is determined by the expression $L_2 = n_1 \cdot H + n_2 \cdot x_N$, while w_u is obtained from $L_1 = (1 - n_2) \cdot x_0$, with $L_1 = L - L_2$.

TABLE 3
THE PRODUCTIVITY SLOWDOWN

	Average annual rate of growth of TFP			
	1950–60	1960–73	1973–79	1979–89
US	1.56	1.65	0.00	0.75
UK	0.50	1.47	0.64	0.52

Source: Wolff (1996), Table 1.

Note: The data is for the whole economy, rather than just manufacturing.

Figure 3 depicts the time-path of real wages⁵⁰ in the benchmark case of the previous subsection. The skill premium (w_s/w_u) starts increasing sharply in the year $n = 21$ when social learning accelerates the flow of new sectors in the economy, and the premium keeps on increasing although more slowly during the remaining part of the transition process. Since everyone ends up earning the same (skilled) wage, standard measures of wage inequality first rise and then fall.

3.3.1.b *Skill-Biased Technical Change and the Productivity Slowdown*

The main argument put forward against the skill-biased technical change hypothesis is that we have not observed an increase in the rate of productivity growth since the early 1980s. Many industrial economies have experienced a substantial reduction in the rate of growth of total factor productivity since 1973.⁵¹ Table 3 reports the results obtained by Edward Wolff (1996). In the US, total factor productivity grew at an average annual rate of 1.61 percent between 1950 and 1973, and at 0.47 percent over the 1973–89 period. In the UK, its rate of growth fell from 1.05

percent in the first period to 0.56 percent in the second. This *productivity slowdown* questions the validity of a hypothesis that attributes the increase in wage differentials to an acceleration of technological progress.

However, several counterarguments can be put forward. To start with, standard calculations of total factor productivity do *not* provide an appropriate measure of the rate of technical change. Empirical calculations are based on a Cobb-Douglas production function of the form $Y_t = A_t K_t^\alpha L_t^{1-\alpha}$, where A_t is the technology parameter (or total factor productivity), K_t is the stock of capital, and L_t is an aggregate measure of the labor input, usually the size of the labor force times average years of education. Total factor productivity is then measured as the residual growth rate, once the rate of growth of capital and labor have been accounted for: $g_A = g_Y - \alpha g_K - (1 - \alpha)g_L$, where g_i denotes the growth rate of factor i . The first problem with this approach is that the rate of growth of this residual can be explained by factors such as the age of capital, the capital labor ratio and the rate of output growth, indicating that it is a poor measure of the level of technology. Second, given that manufacturing and services have different capital intensities, there is a measurement problem if the importance of the share of the service sector in total output is growing but the parameter α is treated as constant over time. Thirdly, the technological parameter A_t affects all types of workers equally, and hence it is not an appropriate measure of biased technical change.

More direct measures of technical change are therefore needed in order to test for the existence of skill-biased technical change. The empirical literature has proposed two measures: R&D expenditures and computer use. If

⁵⁰ The same parameter values as for Figure 2 are used, together with $\beta = 0.05$ and $s = 0.25$.

⁵¹ Wolff (1996) and Griliches (1994), for example, provide evidence on this.

research is an important source of the level of technology, then lagged R&D expenditures can be treated as a proxy for the rate of technical change. The same applies to computers, since they are one of the major technological breakthroughs of the past twenty years.

The skill-biased technical change hypothesis is supported by estimates of equations relating the employment share of nonmanual workers to the two above measures of technical change.⁵² Using US manufacturing data, Berman, Bound and Griliches (1994) show that both computers (as a share of total investment in 1974) and R&D expenditures have a positive and significant impact on the increase in share of non-production workers in the total wage bill⁵³; these two factors account for 70 percent of the move away from production labor over 1979–1987. Extending the study to nonmanufacturing, David Autor, Lawrence Katz, and Alan Krueger (1997) corroborate the importance of technical change in accounting for the increase in skilled workers as a proportion of the wage bill. A similar analysis for a number of OECD countries shows that R&D expenditures have a positive and significant impact on the share of nonmanual workers in employment.⁵⁴ Similarly, at the plant level, Timothy Dunne, John Haltiwanger, and Kenneth Troske (1996) show, using US data, that changes in the R&D stock have a positive and significant impact on the secular change in skill shares. Mark Doms, Dunne, and Troske (1997) also find that technology indicators are associated with a higher proportion of skilled workers when plant-level cross-sectional data is used. Similar results

are obtained with British, French and Spanish plant level data.⁵⁵ Overall, empirical evidence from a number of OECD countries, indicates that more technologically advanced industries (or plants) are more likely to have increased their relative use of skilled workers in the 1980s.

The same variables that have a positive impact on the employment share of skilled workers, namely R&D expenditures and computer use, are found to be associated with higher levels of productivity at the micro level.⁵⁶ A question then arises: how come that this is not reflected in aggregate measures of productivity? The GPT model proposed in the previous section provides a potential explanation for this puzzle. The implementation of a new GPT may induce a temporary productivity slowdown. It involves diverting a substantial fraction of physical and human resources into the risky experimentation of the new GPT; also, it leads to the proliferation of (secondary) product innovations whose output contributions cannot be measured at once by the available statistics and which also have limited externalities on the overall R&D process. Therefore, a faster pace of innovation may result in a temporary productivity slowdown during the time interval where the new GPT is being *experimented* by the various sectors of the economy.⁵⁷

⁵² See Machin (1996b) for the UK, Emmanuel Duguet and Nathalie Greenan (1997) for France, and Victor Aguirregabiria and Cesar Alonso-Borrega (1997) for Spain.

⁵³ See Griliches (1986), Griliches and Jacques Mairesse (1984), Erik Brynjolfsson and Lorin Hitt (1998).

⁵⁴ A similar effect is obtained by Galor and Tsiddon (1997). During the phase of implementation there is a reduction in the concentration of high-ability workers in the technologically advanced sectors, which diminishes the likelihood of technological breakthroughs. Acemoglu (1998) provides an alternative explanation, based on the idea that researchers can target their effort to innovations

⁵² Lucy Chennells and John Van Reenen (2000) review the empirical literature on this topic.

⁵³ All regressions include controls for the level of output and the stock of physical capital.

⁵⁴ Machin and Van Reenen (1998).

3.3.2 Embodied Technical Change and Within-Group Inequality

Skill-biased technological change as analyzed in subsection 3.3.1 cannot account for the fact that around 60 percent of the total increase in wage inequality over the past twenty years is *within* groups of individuals with apparently the same education level and the same number of working years. Focusing on a theory that explains only the rise in between-group inequality is thus not sufficient.

Now, one might go on and argue that this empirical observation reflects measurement problems. In particular, existing studies do not properly discriminate between schools and universities with different levels of standing. As pertinent as such a criticism may be, it cannot entirely account, however, for the magnitude of the within-cohort effects; nor does it account for the fact that the wage gap between workers initially hired by the same firm at the same wage, is also increasing over time.

Gianluca Violante (1996) develops a new explanation for the observed increase in wage dispersion *within* educational groups. His explanation is based on the notion of vintage-specific skills primarily acquired by workers through learning-by-doing. More specifically, suppose that technological knowledge is embodied in equipments of different

vintages. Workers are ex-ante identical (i.e. they all have the same educational background), but are randomly matched with machines of different vintages. At any point in time, workers can either remain on the same job and improve their skills on the current machine through learning-by-doing or move to newer machines. New machines are more productive than old ones, but leaving an old machine involves a (partial) loss of skills. As a result, workers become increasingly heterogenous as their specific labor-market histories unfold, involving different patterns of accumulation or transfer of skills on the job and between jobs. This, in turn, will generate an increased variance in productivity, and thus in wages, among workers with different matching and vintage histories.

The following formalization, based on recent work by Aghion, Howitt, and Violante (1998), might help develop more precise intuitions on the effects of technical change and education policy on *within* cohorts wage inequality.

We consider the simplest possible case, where new technologies are embodied in capital goods which only last for two periods but do not depreciate after one period. Each period a new technology arrives, and it is adopted by those firms that acquire new capital to replace their old and obsolete equipment. The new technology at date t allows the production of final output according to:

$$Y_t = K_t^\alpha (A_t \cdot X_t)^{1-\alpha},$$

where X_t is the number of efficiency units of labor working with technology t , K_t is the amount of capital used by technology t , and A_t is the technological parameter. Profit-maximization by firms will lead to an equilibrium amount of capital which is proportional to the technological level A_t , and for notational simplicity

that complement either skilled or unskilled labor. Since research is a fixed cost, the returns to R&D depend on the number of workers that will be able to use the new technology. For as long as the number of skilled workers was small, R&D targeted products that were complementary with unskilled workers and hence technical change reduced the skill premium. The expansion of education since the 1960s made it profitable to invent machinery to be used by skilled rather than unskilled workers. Technical change *became* skill-biased and the wage ratio started to increase even though there was no change in the aggregate rate of productivity growth.

we shall assume that $K_t^* = A_t$, so that the equilibrium level of output by technology t can be expressed as:

$$Y_t = A_t \cdot X_t^{1-\alpha}. \quad (1)$$

Each new technology is γ times more productive than the previous one, where $\gamma > 1$ measures the rate of technical progress. That is:

$$A_t = \gamma \cdot A_{t-1}$$

for all t .

Since capital goods last only for two periods and do not depreciate after one period, at any point in time only two technologies are in operation: the leading-edge technology t which operates according to equation (1), and the old technology $t-1$ which operates according to:

$$Y_t^O = A_{t-1} \cdot X_{t-1}^{1-\alpha}. \quad (2)$$

Not all workers can move to the leading edge, and this is the primary source of within-cohort inequality; more specifically, the labor market is *frictional* in the sense that a randomly selected fraction σ of all workers at most can relocate at once to the leading edge.

Let 0 (respectively 1) denote the leading edge (respectively the old) technology at any time in steady-state, and X_0 and X_1 the amount of efficiency units of labor employed respectively by new and old firms in steady-state. We denote η the rate of learning-by-doing fully enjoyed by those workers that remain on the same technology for two periods; τ the part of acquired knowledge that a worker moving to the new leading edge can transfer with him; and ζ the spillover of acquired knowledge to new workers. We then have:

$$X_0 = (1 + \tau\eta) \cdot x_{00} + x_{10}$$

$$X_1 = (1 + \eta) \cdot x_{01} + (1 + \zeta\eta)x_{11}$$

where x_{ij} is the steady-state labor flow from the i -technology ($i = 0,1$) last period to the j -technology ($j = 0,1$) this period.

Workers are paid their marginal productivity. The wage of each individual will then depend on two things: the technology she is currently operating, and the technology that she operated the previous period. To see how relative wages are affected by technical change, consider first the simplest case, in which there is full appropriability by firms but no transferability, i.e. $\zeta = 1$ and $\tau = 0$. In this case, wages are vintage-specific and not worker specific. There are, thus, only two wages, which depend on the current technology used by the worker but not on her past work experience. That is, the wage for those operating the new technology is $w_{00} = w_{10} = w_0$, and for those using the old technology is $w_{01} = w_{11} = w_1$, where

$$w_0 = (1 - \alpha) \cdot \gamma \cdot A_1 \cdot X_0^{-\alpha}$$

$$w_1 = (1 - \alpha)(1 + \eta)A_1 \cdot X_1^{-\alpha}$$

There are two possible situations. One is that the relocation constraint is not binding, that is, not all workers using the old technology want to move and operate the new one. Equilibrium would then imply that workers relocate to newer lines in such a way that all individuals are indifferent between working with either technology. This in turn will result in *complete equality of lifetime earnings across workers*. The more interesting case is the one in which the constraint imposed by σ is binding. It is then possible to show⁵⁸ that the wage ratio between the two types of workers is given by

$$\frac{w_0}{w_1} = \left(\frac{1 - \sigma}{\sigma} \right)^\alpha \frac{\gamma}{(1 + \eta)^{1 - \alpha}} > 1.$$

Two parameters affect the ratio of wages, η and γ . A higher rate of learning-by-doing reduces the wage ratio by

⁵⁸ Relative employment is given by $X_1/X_0 = (1 + \eta)(1 - \sigma)/\sigma$, where σ is the fraction of workers using the leading-edge technology, $(1 - \sigma)$ the number using the old one, and $(1 + \eta)$ the productivity of the latter.

making the productivity of those working with the old technology relatively higher. On the other hand, faster technological change as measured by a higher γ will result in more unequal earnings between workers on the leading edge and workers on the old technology.

Let us now allow for the *transferability* of skills by workers who move to the new leading-edge technology that is, for $\tau > 0$. One can then show that *an increase in the degree of knowledge transferability τ has the unambiguous effect of increasing wage inequality* when measured either by the ratio between the top and the bottom wage or by the variance of log-wages. This is not too surprising, since *for given labor flows x_{ij} 's*, more transferability tends to provide an even greater advantage to those workers who adapt quickly to the new leading edge. (What is less obvious is the fact that this direct effect of transferability always dominates the indirect effect working through the relocation decision of workers, i.e. through the x_{ij} 's).

This result yields two potentially testable implications. First, if we interpret the transferability parameter τ as measuring the *general level of education* within the whole cohort of workers, then our analysis suggests that *a higher degree of within-cohort inequality should be observed within more highly educated cohorts*. Recent empirical work by Rubinstein and Tsiddon (1998) and Amanda Gosling, Machin, and Costas Meghir (1998) appears to confirm this first prediction of our model. Second, if we define τ as a measure of the *generality* of the (current leading-edge) technology, then we obtain the prediction that (within-cohort) wage inequality should increase when the economy enters a phase of implementation of a new GPT.

We can now ask what would be the effects of increasing education spending on earnings inequality. The basic model of skill-biased technical progress outlined in subsection 3.3.1. implied that education would have the unambiguous effect of reducing wage inequality between skilled and unskilled workers. However, when technical progress is embodied in particular production lines, the answer to this question will depend on the particular kind (or design) of public support to education. To begin with, any education or training policy that increases the growth rate γ will automatically *increase* earnings inequality. More interesting are the effects of education policies aimed at increasing the adaptability of workers to new lines. Think of a situation in which there are two levels of transferability, $\underline{\tau}$ and $\bar{\tau}$, corresponding to high-school graduates and college graduates, respectively. The overall degree of wage dispersion is determined by both within-group and across-group inequality. If there are decreasing returns to labor for both types of workers, a policy of education expansion that increases the number of workers with high adaptability (thus decreasing their marginal product) would reduce the extent of across-group inequality. Let us thus refer to this effect as a *decreasing returns effect*. However, as argued above, the educational group with a higher τ exhibits greater within-group wage dispersion due to a *transferability effect*. Whether the decreasing return effect or the transferability effect dominates will obviously depend on the parameters of the model. What is important is that policies that increase the average level of education of the workforce may *increase* earnings inequality because they will place the majority of workers in sections of the labor market where previous work experience has a

strong effect on current wages, and where, consequently, wage dispersion is high.

A further effect of transferability occurs if we allow τ to vary across individuals. Following Galor and Omer Moav (1998), suppose that the degree of transferability depends on the *innate ability* of the worker, so that the more able a worker is, the less knowledge depreciates as she moves to the leading-edge technology. Galor and Moav examine this effect in a model in which technical change is skill-biased and education decisions are endogenous. They find that an increase in the rate of technical change increases both the return to ability and to education, bringing about a rise in the skill premium and in wage inequality within educational-groups (as differences in ability now result in greater differences in productivity and therefore wages).

The other parameter of interest in our exposition of vintage effects and embodied technical change is ζ . The degree of appropriability of knowledge by firms is not an invariant technological parameter, but rather the outcome of the particular structure and internal organization chosen by firms. The environment in which firms are operating can therefore affect the extent to which they are willing to provide channels for spillovers of knowledge accumulated by workers, and hence affect the dispersion of earnings. The next section addresses this issue in detail and examines the extent to which the increase in wage inequality can be attributed to organizational change.

3.4 *Organizational Change*

The preceding analysis of skill-biased technical progress and wage inequality may appear simplistic in its representation of firms as “production functions” in which skills *and* current wages

are both being taken as given once a technology has been adopted. The reality is indeed more complex. First, skills are only partly determined by technological requirements. They are to a substantial extent firm-specific, that is, specific to the particular type of *organizational* form chosen by enterprises. Second, as a result of this firm-specificity, wages are not entirely market determined, but instead are the outcome of complex bargaining processes⁵⁹ that are also affected by the organizational structure of firms.

In recent years, a growing literature has emphasized the impact of organizational change upon rising wage inequality.⁶⁰ The underlying idea is that, as changes in organization take place, the productivity gap between individuals with different skill levels increases. In this subsection we discuss the impact of *organizational change* on wage inequality. In what follows we will put forward two hypotheses:

1. In recent years, organizational change has been called for and made possible by the development of new technologies. Hence, it does not stand as an autonomous event but rather as a response to other shocks. In particular, it can be understood as being itself the result of skill-biased technical change.⁶¹
2. Whether organizational considerations magnify or dampen the effect of skill-biased technical change on wage inequality depends on the relative importance of across-firm effects compared to within-firm effects, and on whether employment flexibility is higher within or between firms.

⁵⁹ See Lars Stole and Jeffrey Zwiebel (1996).

⁶⁰ See for example Kremer and Maskin (1996), Acemoglu (1996) or Assar Lindbeck and Dennis Snower (1996).

⁶¹ See David Thesmar and Mathias Thoenig (1998).

3.4.1 *What is Organizational Change?*

Recent developments in the structure of European and US companies have been characterized by the following trends:⁶²

- (a) A move towards flatter organizations with both decentralization of decision making within firms (creation of independent profit centers with greater flexibility, and authority being allocated to units' managers) and a decrease in the number of hierarchical layers accompanied by a wider *span of control* at each layer.⁶³
- (b) The development of collective work in the form of work teams, workers involvement groups and quality circles.
- (c) A shift away from hierarchical into more "organic" structures characterized by the replacement of vertical communication channels by *horizontal* (cross-department) *channels*, less hierarchy and a reduction in specialization (i.e., of the extent to which a particular agent can be identified with any particular task).
- (d) A growing segregation of workers by skills, leading to a higher homogeneity of firms' skill/employment structure. Economic activity has shifted away from firms like General Motors which use both high and low-skilled workers to firms such as Microsoft and McDonald's whose workforces are much more homogenous. For example, in the US, the correlation between wages of production workers in the same manufacturing plant

rose from 0.76 in 1975 to 0.80 in 1986.⁶⁴

These changes in the organization of firms tend to be adopted in clusters, as shown by Casey Ichniowski and Kathryn Shaw (1995). Moreover, they also go hand in hand with technological change. Recent work by Nathalie Greenan (1996) for France and by Erik Brynjolfsson and Lorin Hitt (1998) for the US shows that the use of newly developed information technologies are positively and strongly correlated with the adoption of new forms of workplace organization. The reasons for joint technical and organizational change are twofold. First, the spreading of new technologies heavily relying on the management of information calls for new organizational forms. Decentralization of decision making and intensification of communication are necessary to capture the full benefit of the upcoming technological opportunities. Brynjolfsson and Hitt (1998) provide evidence on that point. Using Fortune 1000 data on technical change as well as a recent survey of organizational practices in the US, they estimate a Cobb-Douglas production function with physical capital, labor, and a measure of information-technology capital interacted with some index of decentralization. This last term has a strongly positive and significant impact on productivity, thus implying that "reorganized" firms get a higher return on the information technology investments.

The second reason why technical and organizational changes should be correlated is that the spread of new technologies makes organizational change possible. For example, the flattening of firms' structure is encouraged by a higher efficiency of the *monitoring technology* which allows any given principal to increase his span of control and

⁶² These four trends are identified in recent work by, respectively, Michael Piore (1994), Paul Osterman (1994), Elisabeth Scott, K. O'Shaughnessy, and Peter Cappelli (1996), and Kremer and Maskin (1996). Caroli (1999) surveys this literature.

⁶³ By *span of control* we mean the number of downstream agents or units that are being supervised or monitored by a given layer.

⁶⁴ See Kremer and Maskin (1996), p. 1.

hence to reduce the number of intermediate supervisory layers.⁶⁵ The move towards less hierarchical (more organic) firms is also linked to the development of new technologies, insofar as these reduce the cost of direct, lateral *communication*. A major comparative advantage of a nonhierarchical structure over a hierarchical one lies in the former's ability to process new information more quickly and thereby to respond faster to changes in demand. This comparative advantage is enhanced when communication costs are reduced.⁶⁶ Finally, as technical change takes place, skill homogeneity tends to increase within firms. Michael Kremer and Eric Maskin (1996) show that, in such circumstances, segregated equilibria will emerge, characterized by "assortative matching." Three ingredients are necessary for this to happen: the production process requires two tasks which are complementary, these tasks are unequally sensitive to skills, and the various skill levels are imperfect substitutes. Under such conditions, when skill dispersion reaches a certain threshold, cross-matching of workers with different skill levels is replaced by assortative matching: high and low skilled workers are employed by different firms in which the labor force is perfectly homogenous.

Technical progress can then be seen as having two effects on wage inequality: a direct one, examined in subsection 3.3, and an indirect one through induced changes in firms' organizational structures that may amplify or dampen

the direct effect. Given that technical progress tends to be skill-biased, one could conjecture that organizational change will necessarily increase wage inequality. We will challenge this view in the next subsection, arguing that the consequences of organizational change heavily depend on the choice made by firms regarding human resources flexibility.

3.4.2 *Organizational Change and Wage Inequality: The Role of Internal versus External Flexibility*

The main argument in favor of the view that organizational change necessarily fosters wage inequality is that organizational change itself is *skill-biased*. In a similar way as technical change, it shifts upwards the relative demand for skilled workers. Using US data, Timothy Bresnahan, Brynjolfsson, and Hitt (1998) show that decentralization in workplace organization has a positive impact on firms' investments in human capital, independent of that of information technologies. Similarly, Caroli and John Van Reenen (1999) show that both technical and organizational change reduce the share of unskilled manuals within British and French enterprises. One reason for this is that nonhierarchical firms rely on direct, horizontal communication among workers and on task diversification as opposed to specialization. They hence require *multi-skilled* agents, who can both perform varied tasks *and* learn from other agents' activities. If educated workers have a relative advantage at multi-skilling, they will get a premium in such organizations.⁶⁷ This is likely to be the case, since education is, by far, the main provider of the kind of general knowledge that is required for a worker to be truly multi-skilled.

⁶⁵ Greater efficiency in monitoring may be due to new monitoring technologies which make it possible for a principal to directly monitor more agents, or to an increase in total factor productivity, which increases workers' incentive to exert effort and thus requires a lower degree of monitoring (Acemoglu and Newman 1997).

⁶⁶ See Roy Radner (1993), Bolton and Mathias Dewatripont (1994), Timothy Van Zandt (1997).

⁶⁷ See Scott, O'Shaughnessy, and Cappelli (1996).

A second reason why organizational change may be at the root of rising inequality has to do with skill segregation. As within-firm skill structure tends to become more homogenous, the productivity gap across firms increases. Assume that the economy is initially in a situation where all types of workers are employed together and that it then moves towards a segregated equilibrium. Segregation takes place between high-skill (and sometimes high physical capital, as argued by Daron Acemoglu 1996) enterprises and mixed or low-skill firms. Since the former are more productive than the latter, all other things being equal, wage inequality is bound to increase across firms within an industry. This suggests that the spreading of more organic (and highly homogenous) firms may be partly responsible for the rise in wage inequality.

However, strong arguments point in favor of a reverse relationship. As decision rights are delegated to lower layers of the hierarchy, firm owners need to avoid moral hazard (or free-riding problems) in teams. The outcome in terms of wage inequality will heavily depend on the choice of the delegates. The owner may choose to concentrate the delegated rights on a small number of "team leaders." This unevenness in authority allocation will, in turn, result in an increased wage differential between the team leader and his co-team members,⁶⁸ all of whom are likely to be drawn from the same educational cohort. On the contrary, it may be optimal to delegate those rights to a large number of team members. In this case, incentive considerations may lead to a

rise in wages for workers in the lower part of the occupational structure and hence, to a narrowing of the corresponding wage gap. This prediction is actually in accordance with the results displayed by Peter Cappelli and Kermit Daniel (1995) on US data. They show that the ratio of supervisors to blue collar workers' wages is significantly lower in those firms which have introduced Total Quality Management schemes.⁶⁹

When supervisory layers are removed, whether the wage differential between top managers (or top supervisors) and the downstream teams should increase or decrease will also depend to a large extent on the nature of tasks or decision rights that are being transferred upstream and downstream. If even a few complex tasks are being transferred downstream, top managers will have to deal with simpler problems and the move toward decentralization is likely to translate into a smaller revenue gap between the top management and downstream units.⁷⁰ Finally, segregation leads to more homogenous skill structures inside firms which, in turn, mechanically reduces wage inequality within any given enterprise.

Overall, *organizational change seems to have an ambiguous impact upon wage inequality*. The move toward more organic firms may increase or reduce inequality, depending on whether multi-skilling (and hence education) requirements overcome or not the consequences of despecialization and decision rights delegation induced by the flattening of organizations. Similarly, rising homogeneity of firms' skill structure will result in higher or lower inequality depending on the relative importance of across-firm effects—which

⁶⁸ This is most likely to be the case whenever: (a) employees are highly responsive to monetary incentives and do not only value the private benefits of control (see Aghion and Jean Tirole 1997); (b) there is *little* job rotation between current team leaders and other team members.

⁶⁹ Such schemes typically involve transferring some decision making down to line workers.

⁷⁰ See Caroli, Greenan, and Dominique Guellec (1997).

enhance wage inequality—as compared to the within-firm effects, which tend to reduce it.

Our conjecture here is that the outcome of organizational change in terms of inequality heavily depends on the type of flexibility chosen by firms in the management of human resources. If they relied on *external* flexibility, firms would react to any new need for human capital by firing unskilled workers and hiring skilled ones from the outside labor market. Such a strategy is bound to boost wage inequality across groups since it shifts the demand for skills upwards. On the contrary, if firms relied on *internal* flexibility, they would choose to promote workers from the bottom end of the occupational or skill structure up to some higher layers. This may occur through formal training or rotation across jobs. The difference between both strategies is particularly clear when comparing the US and the UK on the one hand and, Japan and Germany on the other hand. In the former, flexibility is mainly external⁷¹ while, in the latter, human capital accumulation inside firms is an important factor and internal flexibility thus plays a crucial role.⁷² As we have seen, wage inequality has sharply increased in the UK and the US since the early 1980s, while it has only moderately increased in Japan, and has slightly fallen in Germany. A possible cause of these trends is the difference in the choices of these countries regarding employment flexibility. More generally, we conjecture that flexibility strategies are important determinants of wage patterns in that they condition the impact of organizational change upon wage inequality.

⁷¹ See Lisa Lynch (1993), Erika Groshen and David Levine (1998), and David Soskice (1993).

⁷² See Kauzo Koike (1988) and David Marsden (1990).

3.5 Discussion

We have seen in this section that growth and economic development do not necessarily entail a reduction in inequality, as the recent experience of many OECD countries shows. Attempts to explain the increase in earnings inequality have focussed on trade, technological change and organizational change. All three factors are key components of the growth process and all of them have the effect of widening the earnings distribution. Nevertheless, our discussion implies that technological change is the most important factor, since both trade liberalization and organizational change only affect earnings inequality insofar as they are associated with technical change.

We have argued in this section that the key influence of trade liberalization goes through imported material goods. In particular, recent evidence of a *substitution* between (imported) material inputs and unskilled labor across various industries in Germany, has led us to suggest that by lowering the cost of purchasing physical input—or by improving the variety (or mix) of physical inputs—trade liberalization may well be partly responsible for the observed change in the *direction* of technical change towards increasingly skill-demanding technologies. Such a change may occur in all sectors, including those that do not directly trade with the South but yet are substituting material input for unskilled labor. In this view, skill-biased technical change and trade liberalization end up being *complementary* with the trade effect being carried over by technical change.

Technical progress is by itself is a crucial source of inequality whenever it is not neutral, that is, if it affects differently the productivity of the various types of labor. We have seen that in the

case of GPTs, the diffusion process generates a rise and then a decrease in wage inequality, thus giving rise to a “temporary” Kuznets’ curve during the transition from the old to the new GPT. We have also shown that, when embodied in capital vintages, technical progress enhances within-cohort inequality, first because old technologies get obsolete faster and second, because of a rising gap between the rate of productivity growth and that of learning-by-doing. However, this “growth effect” on inequality may be reversed by education. If education increases workers’ mobility between vintages, as the rate of technical change rises more workers reallocate to new vintages and the productivity of those remaining on old lines increases, due to the existence of decreasing returns. The overall impact on inequality will of course depend on the relative size of both effects.

Finally, technical change also affects relative labor demands and wage inequality through its impact on organizational change. More organic structures, characterized by horizontal communication and task diversification, develop. Organizations become flatter and more homogenous with regard to skills. By fostering changes in workplace organization that are themselves skill biased, technical change magnifies its direct impact upon the relative demand for skills. However, the impact of organizational change on wage inequality is ambiguous. On the one hand, organic firms require multi-skilling, which increases the returns to skills. On the other hand, as decision rights are delegated downstream, firm owners need to avoid moral hazard. One way to do so is to raise wages in lower rank occupations, which reduces in turn occupational wage differentials. The distributional impact of organizational change thus remains undecided and is highly sensitive

to the type of flexibility (internal or external) underlying human resources strategies.

Overall, technological change appears both as the major source of economic growth *and* as the main vector through which the growth process is likely to affect the distribution of earnings. It is therefore at the core of the relationship from growth to inequality. However, the extent to which the growth process actually induces rising inequality depends on the institutional characteristics of each country. In particular, labor market institutions are crucial. Deunionization as well as a sharp decrease in the minimum wage are found to have noticeably contributed to the rise in wage inequality in the US and the UK over the last decade.⁷³ More generally, whether wage bargaining is centralized or not influences patterns of relative wages and unemployment rates across groups of workers with different skills.⁷⁴ So, the characteristics of labor market institutions may either magnify or dampen the impact of technological changes upon wage inequality. This suggests that, if greater equality is to be a target of economic policy, it has to be tackled directly, since market forces by themselves will, most likely, not do it all.

4. Conclusions

We have analyzed the relationship between inequality and economic growth from two directions. The first part of the survey has examined the effect of inequality on growth, showing

⁷³ See John DiNardo, Nicole Fortin, and Thomas Lemieux (1996), Fortin and Lemieux (1997) for the US, and Machin (1997) for the UK.

⁷⁴ See Gottschalk and Smeeding (1997) for an extensive discussion of that point. Atkinson (1997) also discusses the role of labor market institutions when assessing the impact of forces other than supply and demand on wages.

that when capital markets are imperfect, there is not necessarily a trade-off between equity and efficiency. It therefore provides a theoretical framework which explains two recent empirical findings, namely, the negative impact of inequality and the positive effect of redistribution upon growth. In the second part, we have analyzed several mechanisms whereby growth may increase wage inequality, both across and within education cohorts. Technical change, and in particular the implementation of “General Purpose Technologies,” stands as a crucial factor in explaining the recent upsurge in wage inequality.

Our analysis calls for further empirical evidence. As far as the impact of inequality on growth is concerned, the evidence arises mainly from cross-country regressions. It is well known that these are subject to a number of limitations. A thorough test of the theories we have discussed would therefore require other types of evidence, such as time-series analysis and controlled experiments that test the microeconomic foundations of our analysis. In particular, experiments that increase the endowment of less well-off individuals and follow their subsequent economic decisions would help us quantify the effect of credit market constraints upon the incentives to invest or exert effort and, by extension, upon growth.

On the effect of growth on earnings inequality, the recent empirical literature has focused on the experience of developed economies. However, the mechanisms we have explored may also be at work in less-developed countries. In these countries, trade liberalization is likely to have had conflicting effects on the distribution of earnings. On the one hand, standard trade theory implies that trade liberalization should result in a reduction in the skill premium. On the other hand, trade flows bring in new

technologies and ideas that enhance the productivity of all workers, but especially that of skilled workers. Clearly, the notion of “skilled” and “unskilled” workers differs across countries. In less-developed economies, those at the top of the earnings distribution often have no more than a secondary degree. Yet, they have skills that will be enhanced by the arrival of new technologies, thus increasing their wage relative to that of uneducated workers. An empirical investigation of the evolution and the determinants of income inequality in these countries stands as a question to be tackled.

An important aspect of our analysis of biased technical change is that its effect on earnings inequality is nonlinear. In the case of disembodied technical change, the arrival of a new GPT increases the skill premium because of the high demand for skilled “experimentation” labor during the first stages of social learning. The skill premium starts tapering off thereafter as most sectors have made the transition to the new GPT and the supply of skilled labor keeps increasing. Similarly, in the case of embodied technical change, the arrival of a GPT would initially raise the transferability of knowledge (because of the generality of the current leading-edge technology), hence increasing within-cohort wage inequality. This increase would halt once the new technology is so widely spread that all workers have had some experience with it. These two mechanisms thus tend to generate a kind of alternative Kuznets curve, with inequality first rising and then falling (or, in the case of within-cohort inequality, leveling) during the transition to a new technological paradigm. These theories would consequently be validated by the data if we could observe a decrease in inequality fifteen or twenty years after the arrival

of the GPT. Some work is starting to be done with the most recent US data, which seems to point out a reduction in the skill premium since the mid 1990s. The crucial question for the next round of empirical research is, then, whether the increase in across- and within-cohorts inequality is beginning to slow down.

The main policy implication that emerges from the first part of the survey is that, when capital markets are imperfect, there is scope for redistributive policies which are also growth-enhancing. How to redistribute hence becomes a crucial issue. Given that capital market imperfections are at the root of the relationship between inequality and growth, transfers or subsidies to borrowers are an important policy tool. This is particularly relevant in the case of investments in human capital. Increased access to education would also reduce inequalities between dynasties, as it would diminish the effect of family wealth upon individuals' investment possibilities, thereby increasing ex-ante equality. The question is whether it goes as far as reducing ex-post inequality.

Once skill-biased technical change is taken into account, ex-post inequality may actually be increased by rising educational levels. In the case of disembodied technical change, education does narrow the differential between skilled and unskilled workers and has therefore the direct effect of reducing wage inequality. However, increasing the supply of skills has a counteracting impact on wage inequality because it is itself a cause of skill-biased technical change. Had there not been an initial mass of skilled labor, the (nonlinear) implementation of new GPTs and changes in the direction of technical change towards skill-demanding technologies would not have taken place. The *long-run* effect of a greater stock of educated labor is consequently ambiguous. When technical

change is embodied in capital goods, the effect of education can again be to either increase or decrease wage inequality. On the one hand, education increases the mobility of workers across lines, helping them to catch up from old to new vintages, and thus reducing the variance of wages across workers with similar initial skills. On the other hand, education also increases the transferability of skills across lines, thereby enhancing the comparative advantage of those workers who have the opportunity to reallocate quickly to leading-edge technologies. This means that even in the *short run* the impact of education on the dispersion of wages is ambiguous.

The impact of skill-biased technical change on inequality may be somewhat mitigated by some aspects of organizational change. In particular, policies aimed at enhancing internal labor market flexibility would increase firms' incentives to train and promote their existing workers (rather than competing in the hiring of already-trained workers) in response to technological change. This is likely to reduce the extent to which technological change increases the dispersion of wages. However, other aspects of organizational change, such as decentralization of decision making and skill segregation would have the opposite effect.

Overall, our discussion in this survey points to an important efficiency role for *sustained* redistribution. Indeed, as suggested by the second part of the survey, a one-time reduction in after-tax inequality that would foster investment incentives and growth in the short run would result in a (maybe temporary) upsurge in inequality as a consequence of the accelerated technical progress it induces. In other words, the absence of a durable virtuous circle à la Kuznets calls for permanent redistribution policies in order both to control the level of

inequality and to foster social mobility and growth. The details of how such policies should be designed and implemented constitute a whole research area still to be explored.

REFERENCES

- Acemoglu, Daron. 1996. "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence," CEPR Discussion Paper 1459.
- . 1998. "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *Quart. J. Econ.*, 113:4, pp. 1055–89.
- Acemoglu, Daron and Andrew Newman. 1997. "The Labor Market and Corporate Structure," MIT Working Paper 97/08.
- Adelman, Irma and Cynthia T. Morris. 1973. *Economic Growth and Social Equity in Developing Countries*. Stanford: Stanford U. Press.
- Aghion, Philippe; Philippe Bacchetta, and Abhijit Banerjee. 1998. "Capital Markets and the Instability of Open Economies," mimeo, University College, London.
- Aghion, Philippe; Abhijit Banerjee, and Thomas Piketty. 1997. "Dualism and Macroeconomic Volatility" mimeo, University College, London.
- Aghion, Philippe and Patrick Bolton. 1997. "A Trickle-Down Theory of Growth and Development with Debt Overhang," *Rev. Econ. Stud.*, 64:2, pp. 151–62.
- Aghion, Philippe and Peter Howitt. 1998. *Endogenous Growth Theory*. Cambridge: MIT Press.
- Aghion, Philippe; Peter Howitt, and Gianluca Violante. 1998. "Technology, Knowledge and Inequality," mimeo, University College, London.
- Aghion, Philippe and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *J. Polit. Econ.* 105:1, pp. 1–29.
- Aguirregabiria, Victor and Cesar Alonso-Borrego. 1997. "Employment Occupational Structure, Technological Capital and Reorganisation of Production," Universidad Carlos III de Madrid, Working Paper 97–12.
- Alesina, Alberto and Roberto Perotti. 1996. "Income Distribution, Political Instability, and Investment," *Europ. Econ. Rev.*, 40:6, pp. 1203–28.
- Alesina, Alberto and Dani Rodrik. 1994. "Distributive Politics and Economic Growth," *Quart. J. Econ.*, 109:2, pp. 465–90.
- Anand, Sudhir and S. M. Ravi Kanbur. 1993. "The Kuznets Process and the Inequality–Development Relationship," *J. Devel. Econ.*, 40:1, pp. 25–52.
- Atkinson, Anthony B. 1996. "Seeking to Explain the Distribution of Income," in *New Inequalities: The Changing Distribution of Income and Wealth in the United Kingdom*. John Hills, ed. Cambridge: Cambridge U. Press, pp. 19–48.
- . 1997. "Bringing Income Distribution in from the Cold," *Econ. Journal*, 107:441, pp. 297–321.
- Atkinson, Anthony B. and Andrea Brandolini. 1999. "Promise and Pitfalls in the Use of Secondary Data-Sets: A Case Study of OECD Income Inequality," mimeo, Nuffield College, Oxford.
- Autor, David H.; F. Lawrence Katz, and Alan Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *Quart. J. Econ.*, 113:4, pp. 1169–213.
- Banerjee, Abhijit and Andrew F. Newman. 1991. "Risk-Bearing and the Theory of Income Distribution," *Rev. Econ. Stud.*, 58:2, pp. 211–35.
- . 1993. "Occupational Choice and the Process of Development," *J. Polit. Econ.*, 101:2, pp. 274–98.
- Barro, Robert. 1999. "Inequality, Growth and Investment," mimeo, Harvard U.
- Barro, Robert and Xavier Sala-i-Martin. 1995. *Economic Growth*. NY: McGraw-Hill.
- Benabou, Roland. 1996. "Inequality and Growth," *NBER Macroeconomics Annual*, 11, pp. 11–74.
- Berman, Eli; John Bound, and Zvi Griliches. 1994. "Changes in Demand for Skilled Labor within US Manufacturing: Evidence from the Annual Survey of Manufactures," *Quart. J. Econ.*, 109:2, pp. 367–97.
- Blundell, Richard and Stephen Bond. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models," *J. Econometrics*, 87:1, pp. 115–43.
- Bolton, Patrick and Mathias Dewatripont. 1994. "The Firm as a Communication Network," *Quart. J. Econ.*, 109:4, pp. 809–39.
- Borjas, George J.; Richard B. Freeman, and Lawrence F. Katz. 1992. "On the Labor Market Effects of Immigration and Trade," in *Immigration and The Workforce: Economic Consequences for the United States and Source Areas*. Borjas and Freeman, eds. Chicago: Chicago U. Press, pp. 213–44.
- Bound, John and George Johnson. 1992. "Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations," *Amer. Econ. Rev.*, 82:3, pp. 371–92.
- Bourguignon, François. 1981. "Pareto-Superiority of Unegalitarian Equilibria in Stiglitz' Model of Wealth Distribution with Convex Savings Function," *Econometrica* 49:6, pp. 1469–75.
- . 1998. "The Measurement of Multidimensional Poverty," presented at the 1998 Congress of the European Economic Association.
- Breen, Richard and Cecilia García-Peñalosa. 1998. "Income Inequality and Macroeconomic Volatility: An Empirical Investigation," mimeo, Nuffield College, Oxford.
- Bresnahan, Timothy; Erik Brynjolfsson, and Lorin Hitt. 1998. "How Do Information Technology and Workplace Organization Affect Labor Demand? Firm-Level Evidence," mimeo, MIT.
- Brynjolfsson, Erik and Lorin Hitt. 1998. "Information Technology and Organizational Design: Evidence from Micro Data," mimeo, MIT.

- Burtless, Gary. 1999. "Effects of Growing Wage Disparities and Changing Family Composition on the US Income Distribution," *Europ. Econ. Rev.*, 43:4, pp. 853-65.
- Cappelli, Peter and Kermit Daniel. 1995. "Technology, Work Organization and the Structure of Wages," Working Paper, U. Pennsylvania, Wharton.
- Caroli Eve. 1998. "New Technologies, Organizational Change and the Skill Bias: A Go into the Black Triangle," forthcoming in *Technology and the Future Employment of Europe*. Petit Pascal and Luc Soete.
- Caroli, Eve; Nathalie Greenan, and Dominique Guellec. 1997. "Organizational Change and Human Capital Accumulation," Working Paper CEPREMAP 9719.
- Caroli, Eve and John Van Reenen. 1999. "Skill-Biased Organizational Change? Evidence from a Panel of British and French Establishments," CEPREMAP Working Paper 9917.
- Chennells, Lucy and John Van Reenen. 2000. "The Effects of Technical Change on Skills, Wages, and Employment: A Survey of the Micro-Economic Evidence," in *Information and Communication Technologies: Productivity, Employment and Earnings*. Jacques Mairesse and Nathalie Greenan, eds. Cambridge: MIT Press, forthcoming.
- Davis, Steven J. 1992. "Cross-Country Patterns of Change in Relative Wages," *NBER Macroeconomics Annual*, 7, pp. 239-92
- David, Paul A. 1990. "The Dynamo and the Computer: A Historical Perspective on the Productivity Paradox," *Amer. Econ. Rev.*, 80:2, pp. 355-61.
- Deininger, Klaus and Lyn Squire. 1996. "A New Data Set Measuring Income Inequality," *World Bank Econ. Rev.*, 10:3, pp. 565-91.
- DiNardo, John E.; Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages 1973-1992: A Semi-Parametric Approach." *Econometrica*, 64:5, pp. 1001-44.
- Doms, Mark; Timothy Dunne, and Kenneth R. Troske. 1997. "Workers, Wages and Technology," *Quart. J. Econ.*, 112: 1, pp. 253-89.
- Duguet, Emmanuel and Nathalie Greenan. 1997. "Skill Biased Technical Change: an Econometric Study at the Firm Level," *Revue Economique*, 48:5, pp. 1061-89.
- Dunne, Timothy; John Haltiwanger, and Kenneth R. Troske. 1997. "Technology and Jobs: Secular Changes and Cyclical Dynamics," *Carnegie-Rochester Conf. Ser. Public Policy*, 46:0, pp. 107-78.
- Easterly, William and Sergio Rebelo. 1993. "Fiscal Policy and Economic Growth: An Empirical investigation," *J. Monetary Econ.*, 32:3, pp. 417-58.
- Eicher, Theo S. 1996. "Interaction between Endogenous Human Capital and Technological Change," *Rev. Econ. Stud.*, 63:1 pp. 127-44.
- Evans, David S. and Boyan Jovanovic. 1989. "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints" *J. Polit. Econ.*, 97:4, pp. 808-27.
- Falk, Martin and Bertrand Koebel. 1997. "The Demand for Heterogeneous labor in Germany," Zentrum für Europäische Wirtschaftsforschung GmbH Discuss. Paper 97-28, Mannheim.
- Feenstra, Robert C. and Gordon H. Hanson. 1999. "The Impact of Outsourcing and High-Technology Capital on Wages: Estimates for the United States, 1979-90," *Quart. J. Econ.*, 114:3, pp. 907-40.
- Forbes, Kristin. 1998. "A Reassessment of the Relationship Between Inequality and Growth," mimeo, MIT.
- Fortin, Nicole M. and Thomas Lemieux. 1997. "Institutional Changes and Rising Wage Inequality: Is there a Linkage?" *J. Econ. Perspect.*, 11:2, pp. 75-96.
- Freeman, Richard B. 1993. "How Much Has Deunionization Contributed to the Rise in Male Earnings Inequality?" in *Uneven Tides: Rising Inequality in America*. Sheldon Danziger and Peter Gottschalk, eds. NY: Russell Sage Foundation, pp. 133-63.
- Galor, Oded and Omer Moav. 1998. "Ability Biased Technological Transition, Wage Inequality and Economic Growth," mimeo, Brown U.
- Galor, Oded and Daniel Tsiddon. 1997. "Technological Progress, Mobility, and Economic Growth," *Amer. Econ. Rev.*, 87:3, 363-82.
- Galor, Oded and Joseph Zeira. 1993. "Income Distribution and Macroeconomics," *Rev. Econ. Stud.*, 60:1, pp. 35-52.
- García-Peñalosa, Cecilia. 1995. "The Paradox of Education or the Good Side of Inequality," *Oxford Econ. Papers*, 47:2, pp. 265-85.
- Glomm, Gerhard and B. Ravikumar. 1992. "Public versus Private investment in Human Capital: Endogenous Growth and Income Inequality," *J. Polit. Econ.*, 100:4, pp. 818-34.
- Gosling, Amanda; Stephen Machin, and Costas Meghir. 1998. "The Changing Distribution of Male Earnings in the UK," IFS Working Paper 98/9.
- Gottschalk, Peter and Timothy M. Smeeding. 1997. "Cross-National Comparisons of Earnings and Income Inequality," *J. Econ. Lit.*, 35:2, pp. 633-87.
- Greenan, Nathalie. 1996. "Innovation Technologique, Changements Organisationnels et Evolution des Compétences," *Economie et Statistique*, 298, pp. 35-44.
- Griliches, Zvi. 1986. "Productivity, R&D and Basic Research at the Firm-Level in the 1970s," *Amer. Econ. Rev.*, 76:1, pp. 141-55.
- . 1994. "Productivity, R&D and the Data Constraints," *Amer. Econ. Rev.*, 84:1, pp. 1-23.
- Griliches, Zvi and Jacques Mairesse. 1984. "Productivity and R&D at the Firm-Level," in *R&D Patents and Productivity*. Zvi Griliches, ed. Chicago: U. Chicago Press, pp. 339-74.
- Groschen, Erika and David Levine. 1998. "The Rise and Decline (?) of US Internal Labor

- Markets," Mimeo, Hass School, U. California Berkeley.
- Hausmann, Ricardo and Michael Gavin. 1996. "Securing Stability and Growth in a Shock-Prone Region: The Policy Challenges for Latin America," in *Securing Stability and Growth in Latin America*. R. Hausmann and Helmut Reisen, eds. Paris: OECD, pp. 23–64.
- Ichniowski, Casey and Kathryn Shaw. 1995. "Old Dogs and New Tricks: Determinants of the Adoption of Productivity-Enhancing Work Practices," *Brookings Pap. Econ. Act., Microeconomics*, 0:0, pp. 1–55.
- Inter-American Development Bank. 1995. "Overcoming Volatility in Latin America," in *Report on Economic and Social Progress in Latin America: 1995*, Inter-American Development Bank, Washington DC.
- Johnson, George E. 1997. "Changes in Earnings Inequality: The Role of Demand Shifts," *J. Econ. Perspect.*, 11:2, pp. 41–54.
- Juhn, Chinhui; Kevin Murphy, and Brooks Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill," *J. Polit. Econ.*, 101:3, pp. 410–42.
- Koebel, Bertrand. 1997. "Aggregation in Production Analysis: Heterogeneity and Representativity," PhD Thesis, Louis Pasteur U., Strasbourg.
- Koike, Kazuo. 1988. *Understanding Industrial Relations in Modern Japan*. NY: Saint-Martin's Press.
- Kremer, Michael and Eric Maskin. 1996. "Wage Inequality and Segregation by Skill," NBER Working Paper 5718.
- Krugman, Paul R. 1995. "Growing World Trade: Causes and Consequences," *Brookings Pap. Econ. Act.*, 0:1, pp. 327–62.
- Kuznets, Simon. 1955. "Economic Growth and Income Inequality," *Amer. Econ. Rev.*, 45:1, pp. 1–28.
- . 1963. "Quantitative Aspects of the Economic Growth of Nations," *Econ. Devel. and Cultural Change*, 11:2, pp. 1–80.
- La Porta, Rafael; Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert W. Vishny. 1997. "Legal Determinants of External Finance," *J. Finance*, 52:3, pp. 1131–1150.
- Lawrence, Robert Z. 1996. *Single World, Divided Nations?* Paris: OECD Development Centre.
- Levy, Frank and Richard J. Murnane. 1992. "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations," *J. Econ. Lit.*, 30:3, pp. 1333–82.
- Lindbeck, Assar and Dennis J. Snower. 1997. "Reorganization of Firms and Labor-Market Inequality," *Amer. Econ. Rev. (Papers and Proceedings)*, 86:2, pp. 315–21.
- Lindert, Peter H. 1996. "What Limits Social Spending?" *Explorations Econ. Hist.*, 33:1, pp. 1–34.
- Lucas, Robert E. 1988. "On the Mechanics of Economic Development," *J. Monetary Econ.*, 22:1, pp. 3–42.
- . 1993. "Making a Miracle," *Econometrica*, 61:2, pp. 251–72.
- Lynch, Lisa M. 1993. "The Economics of Youth Training in the United States," *Econ. J.*, 103:420, pp. 1292–302.
- Machin, Stephen. 1996a. "Wage Inequality in the UK," *Oxford Rev. Econ. Policy*, 12:1, pp. 47–64.
- . 1996b. "Changes in the Relative Demand for Skills," in *Acquiring Skills. Market Failures, their Symptoms and Policy Responses*. Alison L. Booth and Dennis J. Snower, eds. Cambridge: Cambridge U. Press, pp. 129–46.
- Machin, Stephen. 1997. "The Decline of Labor Market Institutions and the Rise in Wage Inequality in Britain," *Europ. Econ. Rev.*, 41:3–5, pp. 647–57.
- Machin, Stephen and John Van Reenen. 1998. "Technology and Changes in Skill Structure: Evidence from Seven OECD Countries," *Quart. J. Econ.*, 113:4, pp. 1215–44.
- Marsden, David. 1990. "Institutions and Labor Mobility: Occupational and Internal Labor Markets in Britain, France, Italy and West Germany," in *Labor Relations and Economic Performance*. Renato Brunetta and Carlo Dell'Aringa, eds. London: MacMillan, pp. 414–38.
- Mirrlees, James A. 1971. "An Exploration in the Theory of Optimum Income Taxation," *Rev. Econ. Stud.*, 38:114, pp. 175–208.
- Murphy, Kevin M. and Finis Welch. 1992. "The Structure of Wages," *Quart. J. Econ.*, 107:1, pp. 285–326.
- Neven, Damien J. and Charles Wyplosz. 1998. "Relative Prices, Trade and Restructuring in European Industry," in *Trade and Jobs in Europe: Much Ado about Nothing?* Mathias Dewatripont, André Sapir, and Khalid Sekkat, eds. Oxford: Clarendon Press.
- Nickell, Stephen J. 1997. "Unemployment and Labour Market Rigidities: Europe versus North America," *J. Econ. Perspect.*, 11:3, pp. 55–74.
- Nickell, Stephen J. and Brian Bell. 1996. "Changes in the Distributions of Wages and Unemployment in OECD Countries," *Amer. Econ. Rev.*, 86:2, pp. 302–07.
- OECD. 1993. *Employment Outlook*, Paris: OECD.
- . 1996. *Employment Outlook*, Paris: OECD.
- Okun, Arthur M. 1975. *Equality and Efficiency: The Big Tradeoff*, Washington: The Brookings Institution.
- Osterman, Paul. 1994. "How Common is Workplace Transformation and Who Adopts It?" *Ind. Lab. Rel. Rev.*, 47:2, pp. 173–88.
- Owens, Trudy and Adrian Wood. 1997. "Export-Oriented Industrialization Through Primary Processing?" *World Devel.*, 25:9, pp. 1453–70.
- Perotti, Roberto. 1992. "Fiscal Policy, Income Distribution, and Growth," Columbia U. Working Paper 636.
- . 1993. "Political Equilibrium, Income Distribution, and Growth," *Rev. Econ. Stud.*, 60:4, pp. 755–76.
- . 1996. "Growth, Income Distribution, and Democracy: What the Data Say," *J. Econ. Growth*, 1:2, pp. 149–87.

- Persson, Torsten and Guido Tabellini. 1994. "Is Inequality Harmful for Growth?" *Amer. Econ. Rev.*, 84:3, pp. 600–21.
- Piketty, Thomas. 1996. "Inégalités et Redistribution," *Revue d'Economie Politique*, 104:6, pp. 769–800.
- . 1997. "The Dynamics of the Wealth Distribution and the Interest Rate with Credit Rationing," *Rev. Econ. Stud.*, 64:2, pp. 173–90.
- Piore, Michael. 1994. "Corporate Reform in American Manufacturing and the Challenge to Economic Theory," in *Information Technology and the Corporation of the 1990s*. Thomas J. Allen and Michael S. Scott-Morton, eds. Oxford: Oxford U. Press, pp. 43–60.
- Prendergast, Canice. 1999. "The Provision of Incentives in Firms," *J. Econ. Lit.*, 37:1, pp. 7–63.
- Psacharopoulos, George. 1986. *Financing Education in Developing Countries*. Washington: World Bank.
- Radner, Roy. 1993. "The Organization of Decentralized Information Processing," *Econometrica*, 61:5, pp. 1109–46.
- Ramey, Garey and Valerie A. Ramey. 1995. "Cross-Country Evidence on the Link Between Volatility and Growth," *Amer. Econ. Rev.*, 85:5, pp. 1138–51.
- Rebelo, Sergio. 1991. "Long-Run Policy Analysis and Long-Run Growth," *J. Polit. Econ.*, 99:3, pp. 500–21.
- Revenga, Ana L. 1992. "Exporting Jobs?: The Impact of Competition on Employment and Wages in US Manufacturing," *Quart. J. Econ.*, 107:1, pp. 255–84.
- Rubinstein, Yona and Daniel Tsiddon. 1998. "Coping with Technological Progress: The Role of Ability in Making Inequality So Persistent," mimeo, Tel Aviv U.
- Scott, Elisabeth D.; K. C. O'Shaughnessy, and Peter Cappelli. 1996. "Management Jobs in the Insurance Industry: Organization, Deskilling and Rising Pay Inequality," in *Broken Ladders: Managerial Careers in the New Economy*. Paul Osterman, ed. NY: Oxford U. Press, pp. 126–54.
- Shleifer, Andrei. 1986. "Implementation Cycles," *J. Polit. Econ.*, 94:6, pp. 1163–90.
- Slaughter, Matthew J. 1998. "What Are the Results of Product-Price Studies and What Can We Learn from Their Differences?" mimeo, Dartmouth College.
- Soskice, David. 1993. "Social Skills from Mass Higher Education: Rethinking the Company-Based Initial Training Paradigm," *Oxford Rev. Econ. Policy*, 9:3, pp. 101–13.
- Stiglitz, Joseph E. 1969. "The Distribution of Income and Wealth Among Individuals," *Econometrica*, 37:3, pp. 382–97.
- Stole, Lars A. and Jeffrey Zwiebel. 1996. "Intra-Firm Bargaining under Non-Binding Contracts," *Rev. Econ. Stud.*, 63:3, pp. 375–410.
- Tamura, Robert. 1991. "Income Convergence in an Endogenous Growth Model," *J. Polit. Econ.*, 99:3, pp. 522–40.
- Temple, Jonathan. 1999. "The New Growth Evidence," *J. Econ. Lit.*, 37:1, pp. 112–56.
- Thesmar, David and Mathias Thoenig. 1998. "The End of Fordism: Macroeconomic Consequences of Endogenous Organisational Responses," mimeo, Delta, Paris.
- Tinbergen, Jan. 1975. *Income Distribution: Analysis and Policies*. Amsterdam: North-Holland.
- Todaro, Michael P. 1997. *Economic Development*. London: Longman.
- Van Zandt, Timothy. 1997. "Decentralized Information Processing in the Theory of Organizations," in *Contemporary Economic Development Reviewed, Vol. 4: The Enterprise and its Environment*. Murat Sertel, ed. London: MacMillan.
- Violante, Gianluca. 1996. "Equipment Investment and Skill Dynamics: A Solution to the Wage Dispersion Puzzle?" mimeo, University College, London.
- Wolff, Edward N. 1996. "The Productivity Slowdown: The Culprit at Last? Follow-Up on Hulten and Wolff," *Amer. Econ. Rev.*, 86:5, pp. 1239–52.
- Wood, Adrian. 1994. *North-South Trade, Employment and Inequality: Changing Fortunes in a Skill Driven World*. Oxford: Clarendon Press.
- . 1998. "Globalisation and the Rise in Labour Market Inequalities," *Econ. J.*, 108:450, pp. 1463–82.
- Wood, Adrian and Cristobal Ridao-Cano. 1996. "Skill, Trade and International Inequality," Institute of Development Studies, Working Paper 46.
- World Bank. 1990. *World Development Report*. Washington DC: World Bank.
- . 1993. *The East Asian Miracle*. Washington DC: World Bank.