

Demographics and recent productivity performance: insights from cross-country comparisons

Paul Beaudry *Department of Economics, University of British Columbia*

Fabrice Collard *CNRS-GREMAQ and IDEI, University of Toulouse*

David A. Green *Department of Economics, University of British Columbia*

Abstract. Over the last few decades, countries have experienced quite different patterns of productivity growth. In this paper, we emphasize the role of country level demographics in explaining these differences. In particular, looking over the period 1960–2002, we show that cross-country data support the notion that, starting in the late 1970s, countries went through a period of technological transition that lasted at least until the mid-1990s for the fastest adjusting countries and is still proceeding for the slower adjusting countries. The main claim of the paper is that the country-level rate of labour growth was a key factor driving the speed of adjustment to the new technological paradigm, implying that much of the cross-country difference in economic performance over recent decades can be explained by demographic differences across countries as opposed to the many other factors emphasized in the literature. JEL classification: O33, O41

Démographie et productivité: éclaircissements et conjectures dérivés de comparaisons entre pays. Au cours des quelques dernières décennies, divers pays ont fait l'expérience de patterns différents dans la croissance de la productivité. Dans ce mémoire, les auteurs présentent une interprétation des patterns récents de productivité du travail qui met l'accent sur la démographie du pays comme facteur explicatif de ces différences. En particulier, pour la période 1960–2002, on montre que les données comparatives entre pays supportent l'hypothèse que, à partir de la fin des années 1970, les pays ont vécu une période de transition technologique qui a duré jusqu'au milieu des années 1990 pour les pays qui se sont ajustés le plus vite, et qui est encore en cours pour ceux qui se sont ajustés plus lentement. La proposition principale mise de l'avant dans le mémoire est que la croissance de la main d'oeuvre d'un pays a été un facteur clé dans la définition de

The authors would like to thank Angelo Melino and David Andolfatto for comments. This paper was written while the second author was visiting the Department of Economics, University of British Columbia. Email: paulbe@interchange.ubc.ca; fabrice.collard@univ-tlse1.fr; green@econ.ubc.ca

la vitesse d'ajustement au nouveau paradigme technologique. Ceci suggère qu'une bonne partie des différences de performance économique entre pays au cours des dernières décennies est attribuable aux différences dans la démographie de ces pays plutôt qu'aux multiples facteurs mentionnés habituellement dans la littérature. Compte tenu du fait que la transition technologique met du temps à s'accomplir (20 à 30 ans), cette perspective suggère que qu'il faut être extrêmement prudent quand on se fonde sur des comparaisons de croissance de productivité entre pays et sur des périodes courtes (de 5 à 10 ans) pour tirer des conclusions quant aux politiques économiques à mettre en place.

1. Introduction

Over the last few decades, there has been a wide variety of patterns of economic growth across countries. For example, in the 1970s and through much of the 1980s countries such as Japan, Germany, and several other continental European countries experienced substantial growth in living standards, while productivity growth was sluggish in North America. In contrast, by the mid-1990s the situation was almost entirely reversed, with countries such as the United States, Canada, and Australia exhibiting strong growth, while Japan and many European countries were performing rather poorly. Given such a change in growth patterns, it is natural to ask how best to explain it.

One common explanation for these observations is that there was a major technological change in the world economy somewhere in the early 1990s, the effect being to change the institutional structure most favourable to growth. In particular, according to this view, the improvements in communication technology observed in the early 1990s favoured countries with a more flexible and incentive driven work environment – such as that of the United States – and were detrimental to more corporatist countries, such as Japan and many continental European countries. The implication of this view is that strong growth is now most likely achieved by adopting U.S.-style institutions, while the institutions of Japan and Germany – which were very often promoted and admired in the 1980s – are now viewed as being less favourable to growth.

In this paper, we present evidence in support of an alternative viewpoint regarding the differential growth patterns observed across countries over recent history. Although our explanation is also based on the effects of a major change in technology, it downplays the role of institutions in explaining the phenomena and instead emphasizes the interaction between technological change and demographic factors. In short, in this paper we make two main claims. First, we claim that much of the difference in growth paths observed across countries in recent history reflects the adjustment process associated with a major technological and organizational change that began in the mid to late 1970s, with the fastest adjusting countries completing the transition by the mid-1990s. Second, and most important, we claim that differences in the rate of labour growth explain why countries adjusted at different speeds and hence explain the different observed growth paths across countries.

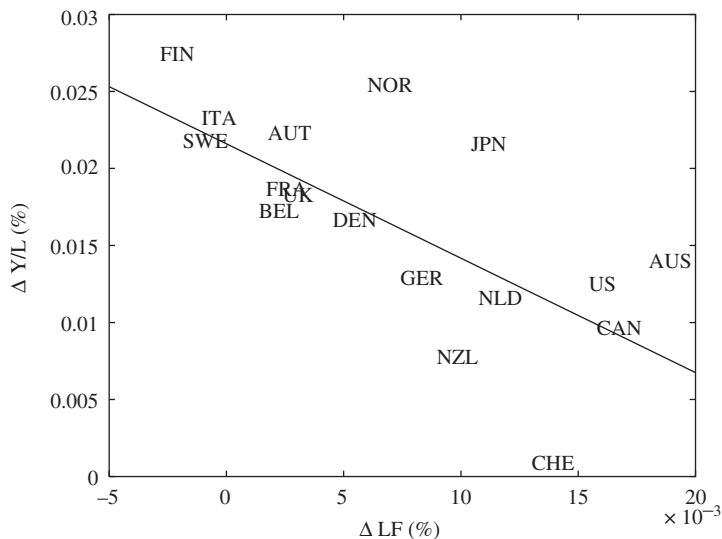


FIGURE 1 Labour productivity growth and labour force growth
NOTE: The plain line corresponds to the regression model:

$$\Delta \log(Y/L) = 0.0216 - 0.743 \Delta LF, \bar{R}^2 = 0.55$$

(0.002) (0.190)

At first pass, it may appear extremely unlikely that differences in rates of labour growth are key factors in explaining differences in economic performance over the last 30 years. It is easy to think of many alternative factors that seem more plausible. However, our reading of the data suggests that such demographic factors are an important and generally unrecognized factor in explaining differences in growth performance over this period. In order to provide initial motivation for this claim, in figure 1 we plot annualized labour productivity growth against the annualized growth of the labour force for major industrial economies over the period 1978 to 1995. The figure reveals a very strong negative correlation between these variables over the period. This is confirmed by the examination of the coefficient from a regression of the change in output-per-worker on labour force growth, which is -0.743 with a standard error of 0.19.

As we will make clear shortly, the pattern presented in figure 1 is robust to including additional variables, appears causal, and is not present in earlier years such as over the period 1960–78. Since our conjecture regarding the role of labour growth is clearly controversial, we will take care to provide both theory and data in support of it. In particular, in the first section of the paper, we present a very simple model of structural change that illustrates the type of empirical patterns that one should expect in the presence of technological paradigm shifts. The model is based on the adjustment process associated

with the arrival and dissemination of a new means of production,¹ and it highlights why countries with different demographic profiles should exhibit systematic differences in growth patterns during the adjustment process induced by a structural change in technology. As we will discuss, our simple theory implies that the role of labour growth in affecting economic performance changes radically in times of structural change in comparison with more normal periods. As we shall show, the empirical patterns implied by the model regarding the relationship between productivity growth and population growth are strongly borne out in the data.

This paper builds on our previous work aimed at understanding recent changes in the economy. We have been working on this research project, together or in pairs, for several years. During this time, our view of how cross-country comparisons can be used to gain insight regarding economic change has evolved. Initially, in Beaudry and Green (1998, 2000, 2003), we questioned whether changes in the wage structure observed in Canada and the United States over the 1980s and 1990s were simply the reflection of an ongoing, exogenous technological change, as emphasized by a large segment of the literature, or if instead it was a reflection of endogenous technological adoption driven by certain country-specific factors. In these papers, we used and refined models of endogenous technological adoption to articulate how one could evaluate such claims. Our results, based on wage data for Canada, the United States, Germany, and the United Kingdom, supported the notion that this period most likely reflected a major change in paradigm, where the speed of change was endogenous to country-specific factors. In particular, we emphasized the role of capital deepening in this process, especially the effect of demographics on physical capital deepening. It is at this point that we began to recognize that demographic differences across countries may be important for understanding recent economic history. In subsequent work, we wanted to look at a broader set of countries and over a longer time span in order to better assess the robustness of our previous findings. This objective led us to focus on patterns of output growth across countries (i.e., output per worker) instead of wage patterns, since such data are more widely available. We began by examining the extent to which the patterns of output growth across OECD countries were suggestive of a major change in technological paradigm (Beaudry and Green 2002; Beaudry and Collard 2003). Once again, we found that the data patterns conformed rather well to the predictions of a model of an endogenous switch in techniques of production driven by differences in population growth. In more recent work, we have extended our exploration to include both developed and developing countries (Beaudry, Collard, and Green 2004), and we have explored the potential role of globalization (Beaudry and Collard 2004) as an alternative explanation for the observed transformations.

1 Similar insights can be drawn from a model that emphasizes a major change in the trade opportunities as shown in Beaudry and Collard (2004).

The current paper brings together some methodological and empirical insights we have drawn from our previous work, as well as providing a more up-to-date treatment of the data. As will become clear, we believe that the data strongly suggest that, over the period 1960–2002, we witnessed a period of technologically driven structural change, as believed by many, but that the nature of the change appears at odds with many preconceptions. Moreover, we will show that demographics appear to have played a central role in explaining how different countries have adjusted to the change. As for interpreting the nature of the observed structural change, we present our preferred interpretation of the observations at the end of the paper. However, at this time, we recognize that our preferred interpretation involves a substantial amount of conjecture and therefore still calls for more research.

The remaining sections of the paper are structured as follows. In section 2, we present a model of technologically induced structural change, which motivates our approach to the data. The main observation that we will derive from the model is that during a period of technological transition, the relationship between growth and demographics is likely to change radically and in a predictable fashion. In section 3, we present a series of empirical results based on the co-movement of output growth and population growth aimed at examining whether we have likely witnessed a structural change somewhere over the period 1960–2002. In this section, we report results based on comparing economic performance across major OECD countries as well as results based on a much wider set of countries. In section 4, we reassess the value of the model of technological revolution presented in section 1 and offer a modified version, which appears more consistent with a larger set of observations. In the conclusion, we discuss the implications of our research for a set of policy issues.

2. A model of technological transition

In this section we present a multi-country model consisting of a set of small open economies that all produce the same final good, which can be either consumed or invested. Our objective with this model is to illustrate how the emergence of a new technological paradigm causes a structural break in the process determining growth. The model is similar in spirit to that presented, for example, in Greenwood and Yorukoglu (1997) (also see Aghion and Howitt 1998; Caselli 1999). In particular, we will use the model to show how the arrival of a new mode of production (i.e., a General Purpose Technology) (see Bresnahan and Trajtenberg 1995, or Helpman and Trajtenberg 1998) generates a drastic change in the relationship between labour productivity growth and population growth. In section 3, we exploit this implication of the model to motivate our exploration of the time varying relationship between labour productivity growth and labour growth as a means of searching for a possible technology-induced structural change. Hence this theoretical section

provides guidance on how to approach data to evaluate the occurrence of a technological revolution.

2.1. *The economy in a pre-technological revolution era*

Consider an economy comprising a large number of identical households, each of which derives utility from the consumption, $c(t)$, of an homogeneous good. The household's utility function takes the form

$$U = \int_0^{\infty} e^{-\rho t} u(c(t)) L(t) dt, \quad (1)$$

where $\rho > 0$ is a constant discount factor, $u(c(t))$ is the per person utility in period t , and $L(t)$ is the size of the dynasty in year t . The latter is assumed to move in the same manner as the population of the economy as a whole, according to

$$\dot{L}(t) = nL(t),$$

where $n > 0$ is the constant population rate of growth, and $L(0) > 0$ is given.

In every period, each household is assumed to supply inelastically $L(t)$ units of labour at the real wage rate $w^o(t)$.² The household owns the capital stock, $K^o(t)$, and rents it to firms at a real rate, $q^o(t)$. Each household has access to international capital markets and holds foreign bonds, $B(t)$, which yield a rate of return equal to the world real interest rate r . These revenues are used to purchase new bonds, investment, $I^o(t)$, and consumption goods. The household therefore faces the following budget constraint.

$$\dot{B}(t) + c(t)L(t) + I^o(t) = w^o(t)L(t) + rB(t) + q^o(t)K^o(t). \quad (2)$$

Investment purchases are then used to form the capital stock according to

$$\dot{K}^o(t) = I^o(t) - \delta K^o(t), \quad (3)$$

where $\delta \in (0, 1)$ is the rate of depreciation of capital and $K^o(t) \geq 0$ is given. Investment is specific to the technology and assumed to be irreversible and, thus, $I^o(t) \geq 0$. Each household determines its optimal consumption/savings plans by maximizing (1) subject to (2) and (3).

The homogeneous good is produced using capital, $K^o(t)$, and labour, $L(t)$, according to a constant returns to scale technology represented by the following production function

$$Y^o(t) = K^o(t)^\alpha L(t)^{1-\alpha}, \quad (4)$$

2 Note that the 'o' superscript corresponds to "old", to distinguish prices and quantities associated with the initial, or old, technology from those that will be associated with the new technology, the introduction of which is the focus of our study.

where $\alpha \in (0, 1)$. In order to keep notation to a minimum, we abstract from labour augmenting technological progress, but all results survive when we take it into account. Optimal production plans are then determined by maximizing profits.

Since households have access to international capital markets and since we are considering a small open economy, the equilibrium capital-labour ratio, $k^o(t) = K^o(t)/L(t)$ is given by³

$$k^o(t) = \left(\frac{\alpha}{r + \delta} \right)^{\frac{1}{1-\alpha}}.$$

Therefore, output per worker is given by

$$y(t) = \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}}.$$

From the last equation, we immediately derive the following proposition.

PROPOSITION 1. *In the pre-technological revolution era, population growth exerts no effect on the equilibrium path of labour productivity.*

This is a result common to many simple growth models.

2.2. The economy during a technological revolution

We now assume that in period $t = t^*$, the economy experiences an unexpected technological revolution, in the sense that the economy has access to a new means of production that is more productive than the old technology. This technology produces an amount, $Y^N(t)$, of the same homogeneous good as was produced in the pre-technological revolution era using technology specific capital, $K^N(t)$, and skilled workers, $S(t)$, according to a constant returns to scale technology represented by the production function

$$Y^N(t) = K^N(t)^\alpha ((1 + \gamma)S(t))^{1-\alpha}, \quad (5)$$

where $\gamma > 0$ fully captures the technological improvement from using skilled labour in the new technology relative to using unskilled labour, $L(t)$, in the old technology. The representative firm now has the choice between using the earlier technology or the new technology. This arbitrage is made by maximizing current profits,

$$Y^o(t) + Y^N(t) - q^o(t)K^o(t) - w^o(t)L^o(t) - q^N(t)K^N(t) - w^N(t)S(t).$$

3 A technical appendix, downloadable from <http://fabcol.free.fr>, gives a more formal statement of all the results.

The introduction of the new technology leaves households' preferences unchanged. It does, however, affect households' decisions. In particular, households must now decide whether to continue working in the old technology or to opt for a training program that will allow them to obtain the skills necessary to take part in the new technology. Workers are assumed to be unproductive when they are in a firm's training program, with individuals in training having an instantaneous probability $\Omega > 0$ of becoming a skilled worker at any point during training. We also assume that new entrants in the labour market come into the market at a skill level proportional to the other individuals in the economy. This assumption eliminates any steady-state effect of population growth and can be justified on the grounds of social learning. We assume that this spillover effect of skill to new entrants plays out at the aggregate level and is therefore not internalized by the household. This implies that the stock of skilled workers, $S(t)$, evolves according to the relationship:

$$\dot{S}(t) = \Omega H(t) + n\bar{s}(t)L(t), \quad (6)$$

where $H(t)$ is the number of individuals in training and $\bar{s}(t)$ is the fraction of skilled individuals in the population, that is, $\bar{s}(t) = S(t)/L(t)$. The remaining $L(t) - H(t) - S(t)$ workers work in the old technology.

The main assumption we make about the new technology and the associated training process is that $\Omega\gamma > r$. This condition guarantees that the expected discounted gain associated with training is sufficiently high to ensure a complete switch to the new technology in the long run.⁴

When the new technology has become available, the old capital still evolves as in (3), and the law of motion for the new capital stock is given by

$$\dot{K}^N(t) = I^N(t) - \delta K^N(t), \quad (7)$$

where, just as in the old technology, investment $I^N(t)$ must be greater than zero because of irreversibility.

The household's budget constraint now becomes

$$\begin{aligned} \dot{B}(t) + c(t)L(t) + I^N(t) + I^o(t) &= rB(t) + w^o(t)(L(t) - S(t) - E(t)) \\ &+ w^N(t)S(t) + q^o(t)K^o(t) + q^N(t)K^N(t). \end{aligned} \quad (8)$$

Each household determines its optimal consumption/savings plans by maximizing (1) subject to (3), (6), (7) and (8).

Given that the household faces a fixed interest rate on the international capital market, the household's optimal allocation of time between training

4 Since the effect of training translates into a marginal increase in productivity of magnitude γ with probability Ω , this condition states that the expected discounted gain ($\Omega\gamma/r$) associated with training is greater than the long run reservation cost which is 1.

and working can be reduced to the following optimization problem (stated in intensive form),

$$\max_{\{h(t), s(t)\}_{t=t^*}^{\infty}} \int_{t^*}^{\infty} e^{-(r-n)(t-t^*)} (w^o(t)(1-h(t)-s(t)) + w^N(t)s(t)) dt$$

subject to

$$\dot{s}(t) = \Omega h(t) + n\bar{s}(t) - ns(t),$$

where, $\bar{s}(t)$ is taken as given by the household (see the technical appendix for details).

The first-order conditions for this problem are

$$\begin{aligned} w^o(t) &= \Omega \mu(t) \\ \dot{\mu}(t) &= -(w^N(t) - w^o(t)) + r\mu(t) \end{aligned}$$

$$\lim_{t \rightarrow \infty} \mu(t)s(t)e^{-(r-n)(t-t^*)} = 0,$$

where $\mu(t)$ is the shadow price associated with being a skilled worker. Since the new capital stock can be financed on the international capital market at rate $r + \delta$, it implies that the capital labour ratio in the new technology, and therefore the wage rate, are constant. In particular, it implies that the wage paid to skilled workers is

$$w^N(t) = (1 - \alpha)(1 + \gamma) \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}}.$$

Then, from the optimality conditions for the allocation of labour, it can be shown that the co-state variable doesn't change over time, so it follows that the wage paid in the traditional technology is given by

$$w^o(t) = \frac{\Omega}{r + \Omega} w^N(t).$$

Since $w^o(t)$ will also be equal to the marginal productivity of labour in the old technology, this implies that the fraction of workers in training is given by

$$h(t) = 1 - s(t) - \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \left(\frac{r + \delta}{\alpha} \right)^{\frac{1}{1-\alpha}} k^o(t).$$

It can be shown that firms will choose $I_t^o = 0$ after the introduction of the new technology.⁵ Therefore, the evolution of skilled labour for $(t \geq t^*)$ is equal to

⁵ This results stems from the fact that the assumption $\gamma\Omega > r$ implies that, in equilibrium, the marginal efficiency of capital in the old technology is lower than in the new technology.

$$s(t) = 1 - e^{-\Omega(t-t^*)} - \frac{\Omega}{\Omega - \delta - n} \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \left(\frac{r + \delta}{\alpha} \right)^{\frac{1}{1-\alpha}} \cdot k^o(t^*) (e^{-(\delta+n)(t-t^*)} - e^{-\Omega(t-t^*)}). \quad (10)$$

Given the dynamics of skill and of training, the dynamics of output per worker during the technological transition ($t \geq t^*$) can be directly derived from the production functions and is given by

$$y(t) = (1 + \gamma) \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}} + \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \frac{(1 + \gamma)(r + \delta)}{\alpha} \left(\frac{\Omega}{r + \Omega} - \frac{\Omega}{\Omega - \delta - n} \right) \cdot k^o(t^*) e^{-(\delta+n)(t-t^*)} + (1 + \gamma) \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}} \cdot \left(\frac{\Omega}{\Omega - \delta - n} \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \left(\frac{r + \delta}{\alpha} \right)^{\frac{1}{1-\alpha}} k^o(t^*) - 1 \right) e^{-\Omega(t-t^*)},$$

where $k^o(t^*)$ denotes the capital-labour ratio in the old technology at time t^* . We are then in a position to characterize the effects of introducing the new technology.

PROPOSITION 2. *Starting from the steady state of the economy where only the old technology is available, the introduction of the new technology triggers a temporary drop in labour productivity.*

This proposition is illustrated in panel (a) of figure 2, which depicts the evolution of labour productivity before and after the introduction of the new technological paradigm. In this figure, we report typical time paths for output per worker for both a country with a zero rate of population growth and a country with a positive rate of population growth. In both economies, labour productivity initially drops after the introduction of the new technology and then recovers to a higher level. In the pre-transition period, $t \in (0, t^*)$, only the old technology is available, and labour productivity is essentially determined by the world real interest rate, as indicated in proposition 1. Therefore, during this period, there are no differences between economies with high and those with low population growth. As soon as the new technology is introduced, labour productivity drops in each economy, as indicated in panel (a). The explanation for this phenomenon is simple and is illustrated in figure 3. This figure plots the evolution of the number of individuals choosing to be trained (panel (a)) and the evolution of skilled labour (panel (b)) during the technological transition. As figure 3 illustrates, as soon as the new technology is introduced, households move from working with the old technology to training. This is so because training is the only way for workers to become skilled and eventually take advantage of a higher real wage in the new technology.

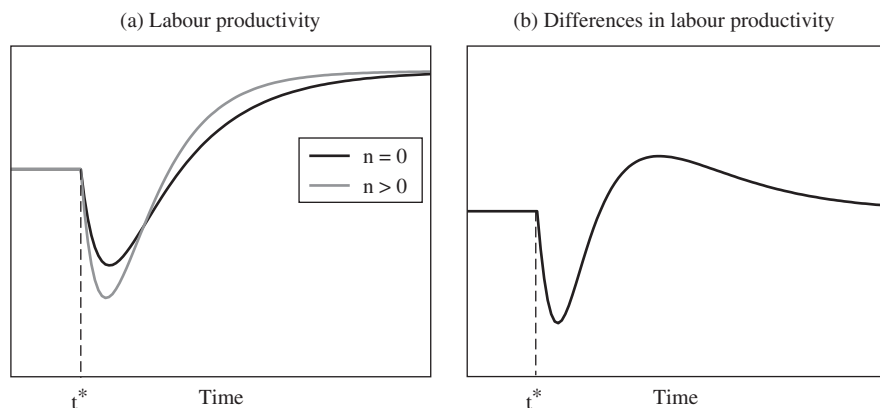


FIGURE 2 Output during a technological transition

The result is that fewer workers are operative in the old technology and, in consequence, output drops. Another phenomenon is also taking place during this transition. The rental rate of capital served by the old technology is lower than in the new technology. There is therefore no incentive for the household to keep investing in the old capital stock, which simply depreciates. This effect reinforces the pattern generated owing to skill upgrading and explains why labour productivity keeps on decreasing for so long. When the new technology is widely adopted, the process reverses and labour productivity regains strength.

It is worth noting that this process does not take place at the same speed in the two economies, as the next proposition makes clear.

PROPOSITION 3. *After the introduction of the new technology, the higher population growth country will first experience less growth in labour productivity than a lower population growth country, followed by a period of higher growth.*

In panel (b) of figure 2 we plot the (log) labour productivity differential between economies with high and those with low population rates of growth after the introduction of the new technology. The pattern displayed in this figure indicates that the economy with the highest population growth initially performs less well in terms of labour productivity than the other economy. Underlying this, as is evident from panel (a) of figure 3, is the fact that the economy with the highest rate of population growth allocates more workers to training in the earlier periods of adoption of the new technology. This occurs because that economy is less tempted to remain in the old technology, since the effective depreciation rate ($\delta + n$) of the old capital-labour ratio is greater in the economy with high population growth. In other words, the relative importance of the old capital stock in the economy is decreasing faster in the high population growth economy. Therefore, this economy forms skilled

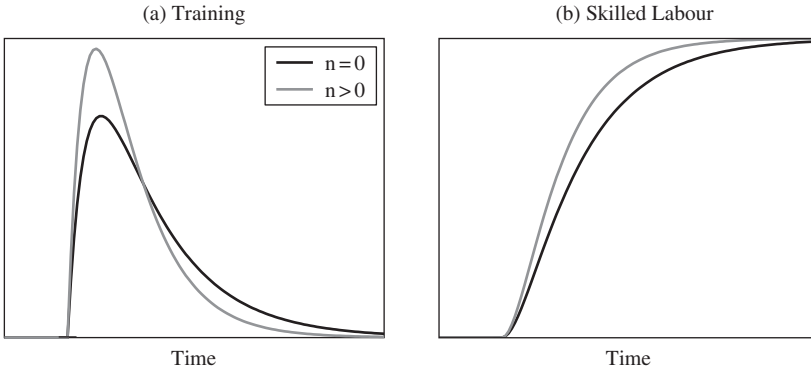


FIGURE 3 Training and skilled labour during a technological transition

workers at a faster pace, implying that it has fewer productive workers in the early stages of the transition. Hence, in the earlier periods of the adoption process, the high population growth economy loses more in terms of productivity than the low population growth economy. But this eventually reverses, since both economies are converging towards the same steady state. Indeed, over time, the high population growth economy takes advantage of its earlier training efforts and implements the new technology at a faster pace as more skilled labour becomes available. It therefore eventually experiences a period of faster growth.⁶

3. Examining growth patterns

The model of the previous section suggests that if there is a structural change due to the arrival and dissemination of a new technological paradigm, then the growth patterns observed in a country during such a transition period will be affected by its demographic characteristics. In particular, the model indicates that, in comparison with a low population growth country, a country with a higher rate of population growth will first experience worse growth outcomes after the arrival of the new technological paradigm followed, later, by better growth outcomes. In contrast, before the technological transition (i.e., in normal times), the model has the property that growth outcomes are unrelated to country demographics. The model therefore suggests a way of exploring whether a set of countries has been subject to a major technological paradigm shift by looking at whether there has been a major change in the relationship between output growth and population growth. Based on these predictions, we

⁶ Throughout this process, the accumulation of capital specific to the new technology is growing faster in the high population growth economy since this capital stock is proportional to the number of skilled workers.

can test the key implications of the model in a regression framework. In particular, consider a simple structural change model given by

$$\forall t < t^*, \frac{(\log y(t) - \log y(0))}{t} = \alpha_{0t} + \alpha_1 \frac{(\log L(t) - \log L(0))}{t} + \alpha_2 X(t) + \varepsilon(t)$$

and

$$\forall t \geq t^*, \frac{(\log y(t) - \log y(t^*))}{(t - t^*)} = \beta_{0t} + \beta_{1t} \frac{(\log L(t) - \log L(t^*))}{(t - t^*)} + \beta_2 X(t) + \varepsilon(t),$$

where X is a set of other variables that may affect growth. Stated in the context of these regressions, the model has three key implications: (i) that a well defined t^* (or small possible range for t^*) exists and is evident in the data; (ii) $\alpha_1 = 0$; and (iii) the β_{1t} s are significantly different from zero and imply a time varying pattern similar to that depicted in figure 2.

The main difficulty in examining these model implications is determining an appropriate value for t^* . Once we assign that value, testing the implications for the α_1 and β_{1t} parameters is straightforward. Indeed, testing for whether there is a structural break of the kind described here is also straightforward, given a particular hypothesized value for t^* . However, we are partly interested in uncovering where any such structural break might lie in the data. To accomplish this, we estimate the preceding system of equations for every possible value of t^* from 1960 to the end of the sample. In each case, we construct a set of fitted values for $\log(y)$ from our estimates and then calculate the sum of squared prediction errors. We are interested in whether a plot of these sums of squared prediction errors from the whole set of possible structural break points indicates that there is a clear candidate (or range of candidates) for a structural break. We then proceed by taking that candidate as given and test the model implications for the α_1 and β_{1t} parameters.⁷

We begin our examination by focusing on the experience of the major industrialized countries over the period 1960–2002.⁸ We choose to focus on only the richest countries, since it is a set for which assuming common access to frontier technological opportunities appears most plausible. The 17 countries forming our

7 Our approach is in the spirit of, though not exactly the same as, using a series of rolling F-tests for structural breaks. In our case, the dependent variable is redefined at each potential value for t^* because the growth rate after the structural break is normalized relative to the value of $\log(y)$ at t^* . This means the F-statistics corresponding to testing for the existence of a structural break are not directly comparable across different values for t^* . Our response to this is to examine the implications of differing potential structural break points for fitting the common underlying series: $\log(y)$.

8 We choose countries with more than 1 million people and with GDP per capita in 1980 greater than 50% of the U.S. level. We found it natural to cut the sample at this point, since it is where there was a rather large break in the data. For example, the next richest industrialized countries had per capita incomes below one-third of the U.S. level in 1980.

sample are Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Italy, Japan, Netherlands, New Zealand, Norway, Sweden, Switzerland, the United Kingdom, and the United States. The data are taken from the 6.0 version of the World Penn Tables and the OECD statistical compendium. In particular, the output data are taken from the World Penn Tables to assure international comparability,⁹ while employment and labour force data are taken from the OECD statistical compendium. Details of data construction are presented in appendix A.¹⁰ The measure we use for $L(t)$ is the population aged 15 to 64. An alternative measure for $L(t)$ would be the country's labour force. However, since labour market participation decisions may respond to variations in economic growth, we chose to focus on the adult population measure of $L(t)$, since it is more likely exogenous to the phenomena examined. Nevertheless, we get similar results using the size of the labour force as an alternative measure for $L(t)$, and we will show some of those results below.

In figure 4, we present the sum of squared prediction errors described above. The horizontal axis corresponds to the sequence of values for t^* we examine, while the vertical axis corresponds to the value of the sum of squared prediction errors. The only additional variable we include in the underlying regressions (i.e., variables in X) is the initial level of output per worker in each sub-sample. However, the results depicted in this figure are robust to including many other variables that may affect growth. The figure indicates a clear minimum of the model selection criterion in the 1975 to 1978 range. Thus, the first implication of the model (that there is a well defined year or a small range of years in which a structural break occurred) is confirmed in the data. The other implications are also met in the data. In particular, setting the structural break at any year in the 1975 to 1978 range, we cannot reject the restriction that α_1 equals zero at any conventional significance level (p-value = 0.49 when $t^* = 1975$ and 0.88 when $t^* = 1978$), while the set of β_{1t} , $t = (t^*, 2002)$, are jointly significantly different from zero (p-value < 0.01 for both $t^* = 1975$ and 1978).¹¹

Of course, the model predictions are stronger than just a zero impact of population growth before the structural change and some non-zero impact after the change: it implies a specific post-change relationship depicted in figure 2. To examine those implications and to further illustrate the relationship between output growth and population growth before and after our estimated t^* (i.e., between 1975 and 1978), we plot the estimated cumulative effect of a 1% difference in population growth on

9 The data in version 6.0 of the WPT end in 1999. We extended the data to 2002 using data from the OECD.

10 Note that in the World Penn data, employment is proxied by the number of people between the ages of 15 and 64. Since this is a poor proxy for employment, we do not use this proxy in constructing output per worker but instead use the employment data available in the OECD statistical compendium.

11 The corresponding p-values when we use growth in the labour force rather than growth in population as the key right-hand-side regressor are 0.79, when $t^* = 1975$; and 0.38, when $t^* = 1978$ for α_1 ; and 0.00 for both $t^* = 1975$ and 1978 for the β_{1s} .




FIGURE 4 Model selection (population 15–64)

labour-productivity in figure 5. Specifically, in the left panel of figure 5, we report the sequence of $\alpha_{1,t}$ s obtained by estimating the regression

$$\frac{\log y(t) - \log y(1960)}{(t - 1960)} = \alpha_0(t - 1960) + \alpha_{1,t}(\log L(t) - \log L(1960)) \\ + \alpha_2(t - 1960) \times \log y(1960) + \varepsilon(t)$$

for $t = 1961$ to 1978. The solid dark line in the figure corresponds to the estimated set of $\alpha_{1,t}$ s. The grey line corresponds to a fourth-order polynomial in time fitted to the $\alpha_{1,t}$ series in order to ease interpretation of potential trends, and the dashed lines correspond to the 90% confidence interval around the estimated effects. The $\alpha_{1,t}$ s capture the effects of annualized population growth rates on annualized growth rates in output as we move forward, year by year, from 1960. It is clear from the figure that population growth had essentially zero impact on output growth throughout this period. In contrast, in the right panel of the figure, we plot the same the type of sequence of estimates, but now starting in 1975 instead and letting t go from 1976 to 2002. The estimated effects of population growth on output performance were substantially different over this period. In particular, the point estimates indicate that a country with a 1% rate of population growth over the period 1975 to 1995 experienced 15% less growth in output per worker over the same period relative to a country with zero population growth. The declines from the mid-1970s to the mid-1990s are statistically significant at any conventional significance level. Then, starting in 1995, we see that the process starts to reverse as the negative effect of population growth decreases,




FIGURE 5 The cumulative effect of 1% population growth on output per worker
 NOTES: The gray line corresponds to a fourth-order polynomial of time fitted on the series of time varying coefficients.

indicating faster growth for the high-population-growth countries after 1995.¹² What is perhaps most striking about this figure is its concurrence with the pattern generated from our theoretical model after the introduction of a new technology.

The key implication from this initial set of regressions is that the rate of adult population growth may have played an important role in shaping medium-run growth in output per worker over the 1978–95 period. It is interesting to verify whether this pattern is evident even in the simplest plots of the data. Thus, figure 6 reports the relationship between annualized labour productivity growth and the annualized rate of growth of the population aged between 15 and 64 for the period 1978 to 1995. The figure reveals a very strong negative correlation between these variables over this period. Recall that figure 1 showed a very strong negative correlation between annualized productivity growth and labour force growth for the same period. As discussed earlier, we carry out our main investigations by examining the impacts of the rates of growth of the non-dependent age (15–64) population as our measure of labour growth to avoid potential endogeneity issues that would arise from using actual labour force growth. Nonetheless, the model is actually written in terms of labour force growth and it is interesting to see that the same patterns arise regardless of the precise measure of labour growth, that is, labour force growth or adult population growth.

The combination of figure 4 and figure 5 suggests that the relationship between output and population growth over the last 40 years differs within three sub-periods: 1960 to the late 1970s, the late 1970s to the mid-1990s, and the mid-1990s to the present. To provide a more concise description of the relationship within these periods, we report the estimates from regressions of output growth on population growth and the initial level of output for the three periods, 1960–78,

¹² Computing the same test as the one reported in figure 4 but starting in 1978, we find another break in 1995.




FIGURE 6 Labour productivity growth and population (15–64) growth
NOTE: The plain line corresponds to the regression model:

$$\Delta \log(Y/L) = 0.026 - 0.937 \Delta \text{POP}_{1564}, \bar{R}^2 = 0.53.$$

(0.002) (0.223)

1978–95, and 1995–2002 in table 1. At the bottom of the table we report the p-values associated with tests of the null hypotheses that the slope parameters in the regressions do not change between the 1960–78 and 1978–95 periods (reported at the bottom of the second column of estimates) and that they do not change between the 1978–95 and 1995–2002 periods (at the bottom of the third column). We report both on tests that all slope parameters do not change and that just the population growth rate parameter did not change.

The results in the table accord closely with what we saw in figure 5. In particular, the impact of population growth on output growth is small and statistically insignificant in the 1960–78 period but becomes large and highly significant in the second period. The apparent turnaround in this relationship after the mid-1990s that is discernable in figure 5 is confirmed by the population growth coefficient in the 1995–2002 regression becoming positive and relatively substantial in size. However, owing to the shortness of this last sub-period, this effect is not well defined, and the conclusion that the process has reversed remains more a tantalizing suggestion than something in which we can place strong faith. The stability tests at the bottom of the table indicate statistically significant (at the 5% significance level) changes in the population parameter between both the first and second and the second and third periods. Thus, we can make a strong case that the relationship between output and population growth changed radically after about 1978 and appears to have

TABLE 1
Baseline regression (population 15–64)

	1960–78	1978–95	1995–2002
Cst	0.132 (0.018)	0.047 (0.026)	0.002 (0.044)
$(Y/L)_0$	-0.033 (0.006)	-0.006 (0.008)	0.002 (0.011)
ΔN	-0.038 (0.249)	-0.846 (0.246)	0.402 (0.390)
R^2	0.74	0.54	0.07
$Q(\text{Total})$	-	0.016	0.002
$Q(\Delta N)$	-	0.016	0.001

NOTE: Standard errors are in parenthesis. The stability test lines report the p-value of the test.

changed again after 1995, though, given the imprecision of estimates in the last period, it is hard to be sure whether the relationship has reversed sign or just been reduced in magnitude. Because of this uncertainty, reflected in a tendency for regression coefficients for the 1995–2002 period to be unstable across specifications, we focus on the two earlier periods (1960–78 and 1978–95) in the remainder of this section.¹³ Finally, the coefficient on initial-period output per worker, which is intended to capture convergence effects, is of about the order of magnitude seen in earlier studies in the pre-1978 period but becomes much smaller, as well as statistically insignificant, after 1978. Thus, we move from a period in which standard convergence processes play an important role in determining output growth and population growth plays essentially no role to one in which the reverse is true.

The results in figure 5 and table 1 represent the basic patterns in the data. However, it is possible that the population growth variable is really capturing the effects of other determinants of growth, and so we want to check the robustness of the conclusions we have presented so far to the inclusion in our growth regressions of other sets of variables that are often hypothesized to affect growth. As a first step, in table 2 we introduce variables capturing physical and human capital effects. In particular, in the first two columns, we present estimates from a regression specification that extends that in table 1 by bringing in the investment to GDP ratio – a standard variable in a Solow growth model. The investment rate variable does not enter significantly in either period and the conclusions from the previous table are unchanged. In particular, the growth rate of population moves from being small and statistically insignificant in the first period to being large, negative, and highly

13 Alternatively, we could examine the whole post-1978 period as one unit. Doing this does not have a substantial effect on our conclusions. However, we believe there is a strong possibility that something has changed after 1995, and lumping this period in with the 1978–95 period would then lead to misrepresentations of the key patterns.

TABLE 2
Robustness of timing

	1960–78	1978–95	1960–78	1978–95
Cst	0.128 (0.048)	0.017 (0.043)	0.225 (0.059)	0.011 (0.055)
(Y/L) ₀	-0.033 (0.008)	-0.003 (0.008)	-0.035 (0.007)	-0.003 (0.008)
ΔN	-0.043 (0.254)	-0.849 (0.240)	0.268 (0.260)	-0.883 (0.312)
I/Y	0.001 (0.009)	0.006 (0.007)	-0.006 (0.009)	0.006 (0.007)
H	–	–	-0.034 (0.015)	0.002 (0.013)
R ²	0.74	0.56	0.81	0.57
Q(Total)	–	0.034	–	0.012
Q(ΔN)	–	0.018	–	0.005

NOTES: Standard errors are in parenthesis. The stability test lines report the p-value of the test.

significant in the second. In a second pass, we control for human capital types of effects. Indeed, human capital is thought by most economists to be one of the most prominent factors contributing to growth (see Uzawa 1965 or Lucas 1988). Furthermore, in their seminal paper, Mankiw, Romer, and Weil (1992) argue that the data fitting properties of the neoclassical growth model are greatly enhanced by including both types of capital – physical and human. Although this finding has been challenged (see Durlauf and Quah 1999 or Klenow and Rodriguez-Clare 1997 for discussions), education clearly remains an important factor, which deserves attention. Therefore, in the third and fourth columns, we introduce a human capital variable (average years of schooling) in the specification. Our main findings regarding the importance of population growth over the second period remain unchanged. The only real difference is that the human capital variable itself enters significantly (though with what might be viewed as a perverse sign) in the pre-1978 period, and the pre-1978 population growth coefficient is of greater magnitude, though still highly statistically insignificant. Nonetheless, the key conclusion that the impact of population growth went from being insignificant before 1978 to strong and negative after 1978 does not change with the introduction of controls for investment in physical and human capital.

We continue with our robustness investigation by seeing whether the large population effect in the 1978–95 period stands up to the inclusion of other variables. Earlier papers have suggested that institutions play an important role in economic growth and that, in particular, Anglo-Saxon institutions are strongly related to good growth outcomes. An inspection of figure 6 might lead a reader to wonder whether that is what we are really picking up in our estimates, since there is a cluster of Anglo-Saxon countries in the bottom

TABLE 3
Robustness (1978–95)

	OLS	OLS	OLS	OLS	OLS	IV
Cst	0.047 (0.026)	0.041 (0.025)	0.041 (0.025)	0.068 (0.025)	0.048 (0.026)	0.038 (0.028)
$(Y/L)_0$	-0.006 (0.008)	-0.005 (0.007)	-0.005 (0.007)	-0.011 (0.007)	-0.007 (0.007)	-0.003 (0.008)
ΔN	-0.854 (0.286)	-0.734 (0.248)	-0.772 (0.278)	-0.758 (0.218)	-0.820 (0.251)	-1.048 (0.290)
AS	0.000 (0.003)	-	0.001 (0.003)	-	-	-
SC	-	0.004 (0.003)	0.004 (0.003)	-	-	-
ΔU	-	-	-	-0.070 (0.031)	-	-
$\Delta C/T$	-	-	-	-	0.300 (0.369)	-
$\Delta E/T$	-	-	-	-	0.137 (0.161)	-
R^2	0.54	0.59	0.59	0.65	0.58	-
F	-	-	-	-	-	0.00

NOTES: Standard errors are in parenthesis. F is the p-value associated with the excluded instrumental variables test.

right (high population growth-low economic growth) corner of the figure. To assess this, we re-estimated our initial specification, including a dummy variable for predominantly Anglo-Saxon countries (Australia, Canada, New Zealand, the United Kingdom, and the United States), denoted *AS*. The estimates, presented in the first column of table 3, reveal a coefficient of zero on the Anglo-Saxon dummy variable and no effect on the large estimated population growth effect. In columns 2 and 3, we also try introducing a dummy variable for Scandinavian countries (Denmark, Norway, and Sweden), denoted *sc*, first on its own and then in combination with the Anglo-Saxon dummy. The inclusion of the Scandinavian dummy does reduce the population growth effect somewhat, even though it is not statistically significant in its own right. Nonetheless, the basic conclusion that the population growth rate is a large, negative, and statistically significant determinant of economic growth in this period does not change. As an alternative approach to controlling for institutional impacts, we next include the percentage change in each country's unemployment rate as an additional regressor. This experiment addresses whether the high correlation between productivity growth and population growth could be simply the result of certain labour market policies that may have favoured labour productivity at the cost of increasing unemployment.¹⁴

¹⁴ See Blanchard (1997) for a discussion along these lines.

In such a case, technological adoption, as discussed in section 2, would not be a relevant candidate explanation for the type of phenomenon we are documenting. The results from this specification, reported in column 4, indicate that the introduction of the change in unemployment does not affect the observed negative correlation between labour force growth and labour productivity growth, although the effect of changes in unemployment is itself significant.

Yet another potential interpretation of our results is that we are really picking up changes in the age structure of populations, something that has recently been argued is a determinant of growth (see, e.g., Feyrer 2002). To investigate this possibility, we estimate a specification, including two age structure variables. The first variable is the percentage change in the ratio of the number of children under age 15 to the population as a whole, denoted $\Delta(C/T)$. The second variable corresponds to the percentage change in the ratio of the elderly (people over age 65) to the total population, denoted $\Delta(E/T)$. The results from this specification are reported in column 5. The age structure variables themselves do not enter statistically significantly and, more important for our discussion, also do not change the result that population growth is strongly negatively related to output growth in this period. Finally, there may be some concern that over extended time periods, population growth may be endogenously related to economic growth, particularly to the extent that immigration and emigration flows respond to changes in the economic fortunes of an economy. To address this, we instrument for population growth using the average annual population growth for the given country in the beginning of the previous period (1960–70). Essentially, this means we are checking whether persistently high population growth countries fared particularly poorly in the 1978–95 period. This instrument is valid if we assume that demographic decisions in the previous 20 years were not based on accurate predictions of patterns of economic growth after 1978. Given our evidence that there was a structural change in the growth process around 1978, this seems like a reasonable assumption. Note, also, that one might be concerned that there might be a permanent country-specific effect that relates to both economic performance and population growth. Such an effect, if not controlled for, would render this instrumental variables strategy ineffective, since the instrument would also reflect that effect. However, we have already conditioned on such effects by including the start of period output-per-worker variable in our specifications.

The results from the instrumental variable specification, given in column 6, are stronger than what we observed without instrumenting. In particular, the coefficient on the population growth rate variable increases in magnitude from -0.846 to -1.048 . The implication is that, if anything, any endogeneity problems in the simple OLS estimation were causing us to underestimate the full impact of population growth on economic growth.




FIGURE 7 The cumulative effect of 1% population growth on output per worker

3.1. Results for the world

To this point, we have confined our analysis to a set of rich countries. We did this, in part, as a means of controlling for other factors, such as political instability, that might confound our attempts to estimate basic underlying patterns of economic growth. However, we are also interested in whether the patterns we have identified are only a developed world phenomenon or show up in a broader sample of countries. To this end, we repeat our main investigation for the full set of countries for which we have data throughout our data period. In the Penn World tables, this amounts to 106 countries, for which there are data spanning the years 1960–98. Figure 7 reproduces figure 5 using the whole world sample. The results from the two samples show strong similarities. In particular, in the pre-1978 period, we again see only limited effects of population growth on economic growth. Though the estimates from the full world sample show a dip in the population effect in the early 1970s and a slight negative trend in the effect over the whole period, the implication from comparing the pre and post-1978 periods remains the same: limited pre-1978 population effects are converted into much more substantial effects after 1978. In the post-1978 period, we see the same evidence of a limited relationship up until the early to mid-1980s followed by a sharp increase (in absolute value) in the estimated effects. The drop is much more substantial than what was observed for our restricted sample and there is no evidence of a slow down or turnaround in the trend after 1995. It is worth noting, though, that this sample runs only to 1998 and thus may simply miss any potential turnaround. Thus, the overall pattern of differences before and after 1978 is, if anything, even stronger in the whole-world sample.

These rough impressions from the figure are confirmed in regression analysis. Table 4 contains estimates of our basic specification and the physical and human capital robustness checks. We again divide the sample in 1978 but allow the second period to run to 1998, which is the end of the sample for this data.

TABLE 4
Regression analysis

	1960–78	1978–98	1960–78	1978–98	1960–78	1978–98
Cst	0.024 (0.019)	0.030 (0.020)	0.109 (0.024)	0.147 (0.024)	0.152 (0.023)	0.140 (0.029)
$(Y/L)_0$	0.001 (0.002)	-0.000 (0.002)	-0.006 (0.002)	-0.008 (0.002)	-0.015 (0.003)	-0.010 (0.003)
ΔN	-0.230 (0.214)	-0.976 (0.185)	-0.155 (0.192)	-0.806 (0.157)	-0.092 (0.166)	-0.808 (0.179)
I/Y	-	-	0.014 (0.003)	0.023 (0.003)	0.011 (0.003)	0.019 (0.004)
H	-	-	-	-	0.017 (0.003)	0.008 (0.006)
R^2	0.02	0.24	0.22	0.47	0.43	0.45
$Q(\text{Total})$	-	0.000	-	0.000	-	0.000
$Q(\Delta N)$	-	0.004	-	0.006	-	0.002

NOTES: Standard errors are in parenthesis. The stability test lines report the p-value of the stability test. There are 106 observations. When education is introduced as regressor, the number of observations is 86.

The results from the basic specification in the first two columns support our discussion of the figures. The pre-1978 population growth effect is larger in absolute value than what was observed for the richer country sample but is still nowhere near statistical significance at any conventional significance level. In contrast, the post-1978 population growth effect is larger than what was observed in the smaller sample. It is also, again, statistically significantly different from zero at any conventional significance level and negative. The test-statistics at the bottom of the second column also indicate that there is no question that the impact of population growth changed between the two periods. The other difference relative to the earlier sample is that the initial output level has small and insignificant effects in both periods for the whole-world sample, while for the smaller sample of countries it had a significant coefficient in the pre-1978 period. As we did in the results from the richer country sample, we check the robustness of these results by introducing physical and human capital variables.¹⁵ Introducing the investment rate variable reduces the size of the estimated population growth effects in both periods but does not alter the overall pattern of an insignificant and relatively small effect before 1978 being followed by a much larger and significant second-period effect. In contrast to the results with the smaller sample, the investment rate itself enters significantly and becomes significantly larger in the second period than in the first. This is in line with the results in Beaudry, Collard, and Green (2004), where we showed that, for a broad sample of countries, both population

¹⁵ Note that we have only 86 observations in the last specification reported in the table because of limitations on education data.

and investment effects changed after 1978 and that this could explain the “hollowing-out” of the middle of the world income distribution that occurred over this period. Finally, the results are similar when we also introduce the human capital variable, with the only exception being that the convergence parameter is now statistically significant in both periods. The education effect itself is significant in the first period but insignificant and much smaller in the second period.

Overall, the results of these empirical investigations strongly support the notion of a structural break that took place in the late 1970s. That break is reflected in a switch from population growth’s having no impact on economic growth patterns to its having large, statistically significant and negative effects after 1978. Moreover, the pattern of the population effects varies over time in a manner that is in close concordance with the pattern our theoretical model predicts for a period following the introduction of a major new technology. There is also some evidence of a turnaround in the relationship between population and economic growth in the late 1990s. However, while the basic pattern of no effect from population growth followed by a strong negative effect is evident both in samples of countries covering the whole world and covering only rich countries, the potential turnaround is evident only in the rich countries. Finally, these conclusions stand up to substantial robustness checks, suggesting that this is a strong basic pattern in the data.

4. Interpretation

The concordance between the theory presented in section 2 and the empirical patterns described in section 3 provide support for the view that (i) we have witnessed a major technological revolution starting somewhere in the mid- to late 1970s, and (ii) demographic differences across countries have been an important, and generally unrecognized, force driving the differential growth performances across countries during this period. An important element of these results is the duration of the patterns being described. The patterns play out in a medium-term time frame: 20 to 30 years. This implies caution in trying to interpret patterns in the type of time frames often examined in public discourse, the most recent being 5 to 10 years. For example, some observers have used comparisons of relative growth performance across countries since the mid 1990s to argue that the relatively successful countries in that time frame have superior economic and social policies. The ‘American Miracle’ of the late 1990s, in particular, has been used in this way. But the model and data presented here suggest that recent differences in growth may, instead, simply reflect that demographic differences across countries have led them to adjust to a technological revolution at different speeds.

As emphasized above, the data patterns presented in section 3 are surprisingly consistent with the very simple model of a technological revolution presented in section 2. Based on this, we could decide to conclude the paper

here, since the data and model provide answers to most of the questions posed in the introduction. However, we would like to take advantage of this forum to discuss in some detail what we see as the way forward from here. That necessarily involves the somewhat unconventional step of criticizing what we ourselves have done and suggesting alternative solutions. Those potential solutions build on what we view as the important insights gained from the model presented to this point but require modifications to that model to match other patterns in the data.

In this spirit, we want to point out two shortcomings of the model of section 2. First, a key pattern in the data is the very long duration of the productivity slowdown that began in the mid-1970s and lasted at least until the mid-1990s. Recall that in the model of section 2, a productivity slowdown arises because agents are investing to learn how to use the new technology. Thus, a prolonged productivity slowdown implies particularly patient agents who are ready to forgo current income in order to invest in learning new skills. In our quantitative exploration of the model, we found that extremely low discount rates were needed to reproduce the pattern observed in the data. Since such low discount rates appeared unreasonable, it suggests to us a need to adjust the model in a manner that relies less on the forward looking behaviour of workers. Second, the time paths of the level of productivity for specific countries do not appear to match the model well. For example, consider the time path of labour productivity observed in the United States since 1947 as depicted in figure 8.

In this figure, we plot U.S. labour productivity¹⁶ as well as two trend lines with identical slopes. The first trend line passes through the points 1947 and 1973, while the second trend line passes through the points 1995 and 2002. This graphically illustrates that U.S. labour productivity growth over the period 1995–2002 was essentially identical to that observed over the period 1947 to the mid-1970s (note that the first trend line goes almost through the year 1978). The only difference between the two trend lines is that the second trend line has an intercept that is almost 15% less than that of the first trend line. The lower trend line for the 1995–2002 period suggests that the economy witnessed a downward level shift starting somewhere in the 1970s and ending around 1995. In contrast, the model implies that by the end of the process countries should experience an upward level shift. In fact, the pattern in figure 8 suggests that the U.S. economy may have gone through a transition between 1975 and 1995, as suggested by the model, but this transition appears to have led to a lower level of productivity instead of a higher level as suggested by the model of section 2. The obvious question, then, is whether one can explain why rational agents may adopt a new technology that would lead to a lower level of

16 This series is taken for the Bureau of Labour statistics and represents output per hour of work.




FIGURE 8 U.S. labour productivity

productivity, while simultaneously having the speed of adoption being driven by rates of population growth.

In the light of these difficulties, in the next subsection we present a modified model of a technological revolution that (i) does not rely on low discount rates to explain the relatively poor growth performance of high population growth countries during the period 1978–95, (ii) explains a downward level shift in labour productivity by the end of the process. We chose to present this alternative model at the end of the paper, as opposed to presenting in section 2, since it is a model built on quite controversial assumptions. Nonetheless, we believe it is a useful model, since it illustrates how the puzzles can be reconciled.¹⁷ We leave the work of developing and testing refutable implications from this model (beyond the data patterns it was built to explain) to future work. Our goal here is to stimulate discussion over whether this, or an alternative, model provides an appropriate explanation for the data patterns we have emphasized.

4.1. A modified interpretation

The model we present here builds on the theoretical and empirical insights we gained from the model in section 2.¹⁸ In particular, it focuses on how economies adapt to a new technological paradigm as represented by a discrete change in the

¹⁷ In Beaudry and Collard (2004) we present an alternative explanation building on the process of globalization instead of technological change.

¹⁸ The interested reader is referred to the technical appendix to this paper for a more formal derivation of the model.

production process. Furthermore, it again emphasizes the central role of population growth in explaining cross-country differences in growth performance during the transition. The main difference in the model is that instead of introducing a new technological paradigm that requires new skills to operate, the current model introduces a new technology that renders the motivation or monitoring of work effort easier. In order to capture such a possibility, we exploit ideas developed in the efficiency wage literature (especially the gift exchange literature).¹⁹ To this end, let us consider an initial situation where a good X can be produced using the following production function:

$$X(t) = K_x^o(t)^\alpha (L_x^o(t)e(t))^{1-\alpha},$$

where $0 \leq \alpha \leq 1$. K_x^o is physical capital, L_x^o is labour time and e is effort (or initiative). The main assumption we draw from the efficiency wage literature is that e reacts to the degree to which the wage paid by a producer of X exceeds the competitive wage. That is, if we let $w_z(t)$ represent the competitive wage for labour and we denote by $w_x^o(t)$ the wage paid to a worker producing X , our main assumption is that effort is an increasing function of $w_x^o(t)/w_z(t)$ as stated below²⁰

$$e(t) = e\left(\frac{w_x^o(t)}{w_z(t)}\right), \quad e'\left(\frac{w_x^o(t)}{w_z(t)}\right) > 0.$$

In this case, the optimal choice of w_x^o satisfies the condition:

$$\frac{e'\left(\frac{w_x^o(t)}{w_z(t)}\right) w_x^o(t)}{e\left(\frac{w_x^o(t)}{w_z(t)}\right) w_z(t)} = 1.$$

Let us denote the optimal relative wage that solves the above decision by γ^* and let us denote $e(\gamma^*)$ by e^* . Thus, e^* is the optimal level of effort to implement.

Now consider the introduction of an alternative means of producing X given by $F(K_x^N(t), L_x^N(t))$, where $K_x^N(t)$ is a new type of capital. The main feature of this new technology is that it no longer depends explicitly on effort. This implies that a firm using this technology will find it optimal to pay workers the reservation wage $w_z(t)$ instead of paying an above-market wage. The technological change that would allow a transformation of this type could be associated with making tasks easier to monitor or less demanding on motivation.

19 Note that our results do not hinge on the adoption of the efficiency wage hypothesis. The results would hold under alternative modellings of the labour market providing they generate rents.

20 We are assuming that the effort function is such that it is optimal for the firm to pay a wage premium.

However, there may be a drawback with this new form of work organization in that it may be less productive than the conventional technology when the latter is operated at the optimal level of effort. To capture this idea, let $F(K_x^N(t), L_x^N(t))$ be given by

$$F(K_x^N(t), L_x^N(t)) = K_x^N(t)^\alpha (\tilde{e} L_x^N(t))^{1-\alpha},$$

where $\tilde{e} < e^*$ is the level of effort implemented in the new technology without a need to pay a wage premium. The new technology therefore saves on wage costs at the price of inducing lower productivity. The question then is how does an economy adjust to the introduction of such a technology? To answer this question, we need to pose the problem in an equilibrium framework that permits an endogenous determination of $w_z(t)$. To this end, consider the situation where there is a tradeable final good Y that is produced using to non-traded intermediate goods X and Z according to the function,

$$Y(t) = X(t)^\varphi Z(t)^{1-\varphi}, \quad 0 < \varphi \leq 1,$$

and assume that the good Z is produced according to the production function

$$Z(t) = A_z K_z(t)^\alpha L_z(t)^{1-\alpha},$$

which implies that the production of the good Z is not subject to motivation issues.²¹ We assume, as in section 2, that there is international trade in capital and that there is an irreversibility constraint on domestic capital accumulation for all three types of capital (K_x^o, K_x^N, K_z). The accumulation equations for capital are standard, with physical depreciation rate given by δ . The population in the economy grows at rate n and households maximize the discounted utility obtained by the consumption of the final good. In the absence of the routinized technology for the production of good X , it is easy to verify that the equilibrium level of output per worker, $y(t) = Y(t)/L(t)$, in the economy is given by

$$y(t) = A_z^{\frac{1-\varphi}{1-\alpha}} \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}} (\varphi^\varphi (1 - \varphi)^{1-\varphi})^{\frac{1}{1-\alpha}} \frac{e^{*\varphi} \gamma^{*1-\varphi}}{\varphi + (1 - \varphi)\gamma^{*}},$$

which is independent of n . Hence, before the introduction of the new way of organizing work, demographics exert no effect on labour productivity.

Now consider introducing into such an economy the new method of production for X where \tilde{e} satisfies the condition $\gamma^* \tilde{e} / e^* > 1$. Under this condition, it is profitable at initial prices to switch the production of X from the older technology towards the new routinized technology. However, the switch is not immediate, since the new

21 Note that we assume that the elasticity of output with respect to each factor is the same in each sector and each technology. This assumption is made to keep analytical tractability. This assumption can be easily relaxed without altering our results at the price of losing a simple analytical solution.

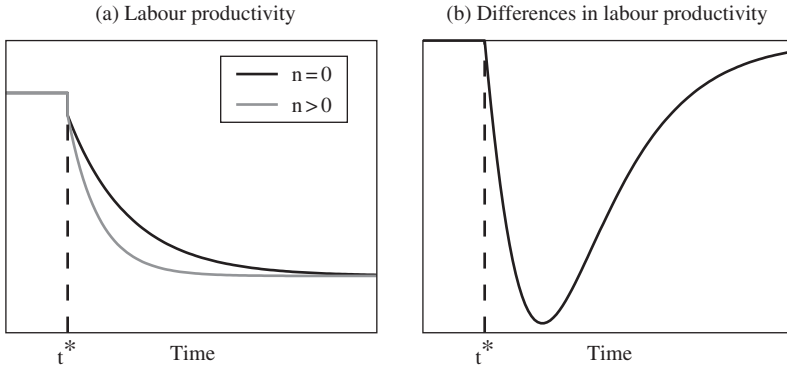


FIGURE 9 Labour productivity in the transition

technology requires another form of capital, which is built up gradually through investment. In the meantime, the capital in the old technology depreciates slowly. In this case, the time path of output per worker in the economy is given by

$$\begin{aligned}
 y(t) = & (\varphi^\varphi(1 - \varphi)^{1-\varphi} A_z^{1-\alpha})^{\frac{1}{1-\alpha}} \left(\frac{\alpha}{r + \delta}\right)^{\frac{\alpha}{1-\alpha}} \tilde{e}^\varphi \\
 & + \frac{\gamma^* - 1}{\gamma^*} \frac{r + \delta}{\alpha} \left(\frac{e^*}{\gamma^* \tilde{e}}\right)^{\frac{1-\alpha}{\alpha}} k_x^o(t^*) e^{-(\delta+n)(t-t^*)}.
 \end{aligned}$$

One important aspect to note from the above equation is that the speed of transition to the new steady state does not depend on the discount rate of workers, as was the case in the model of section 2. Instead, the speed of transition now depends only on the depreciation rate of capital and the rate of population growth, since workers are not making an investment decision with respect to the new technology. Hence, the transition process may be very extended even if individuals have high discount rates.

In figure 9, we report the evolution of labour productivity and the (log) labour productivity differential between an economy with a high and a low population rate of growth after the introduction of the new technology. Inspection of this figure indicates that the model possesses many of the same dynamic properties as the model developed in section 2. First, introducing the new method of production triggers a drop in labour productivity, as depicted in panel (a) of the figure.²² Second, after the introduction of the new technology, the higher population growth country initially experiences a lower rate of growth in labour productivity than the smaller population growth country.

22 To see this, compute the difference between $y(t)$, $t < t^*$, as produced with the non-routinized technology, and $y(t)$ at the time of introduction of the routinized technology ($t = t^*$).

Later, as it catches up to the same steady state, the labour productivity in the high population growth economy will grow at a faster pace.

The main difference relative to the model of section 2 is that the new technology, once introduced, induces a long-run a negative effect on the level of labour productivity. As we argued earlier, this is more consistent with that observed in the United States over the period 1960–2002 (see figure 8) than the pattern predicted by the model of section 2. Hence, the modified model presented in this section suggests that since the late 1970s we may have witnessed the introduction of a new technology that has rendered the labour market more competitive and less dependent on efficiency wages, at the cost of a permanent negative level effect on productivity.²³ Furthermore, the speed of transition of this model has been shown to be independent of the workers' rate of time preference. It is our belief that such a model provides an example of the type of building blocks needed to explain this period.

5. Conclusion

The object of this paper has been to show how cross-country comparisons of medium-run economic performance can help refine our understanding of recent economic history. The main empirical pattern we documented was a drastic change in the co-movement between labour productivity growth and population growth over the period 1960–2002. In particular, we showed that prior to the mid seventies, this co-movement was close to zero. From the late seventies to the mid nineties, this co-movement was strongly negative. There is also some evidence that the correlation may have become positive after 1995. To interpret such a pattern, we have presented two models of major technological change, in the spirit of the General Purpose Technology literature. In both models, population growth interacts with the adoption speed of a new technological paradigm to produce the type of U-shaped pattern of correlations between labour productivity growth and population growth observed in the data. Thus, patterns of productivity and population growth across countries provide support for the view that the world economy has witnessed a major technological revolution starting around 1978, and that the speed of adjustment to that structural change for a particular country has been, to a large extent, driven by its demographics.

23 The model presented here can easily be extended to be consistent with an increase in the returns to skill. For example, consider a situation where the good X requires only unskilled labour, while the good Z can be produced by both skilled and unskilled labour in the production of good Z , where (i) skilled labour is substitutable to unskilled labour but is more productive and (ii) the productivity differential between skilled and unskilled labour is greater than the rent paid to unskilled labour in sector X . Under these conditions, the average wage paid to unskilled workers will drop thereby leading to an increase in the average differential between skilled and unskilled workers.

These observations, in and of themselves, are important in policy discussions, since they suggest that relative economic performances across countries since 1995 may in large part reflect a reversal of patterns observed pre-1995, as opposed to the effects of recent policies. In contrast, the relatively strong growth performance of several countries, most notably the United States, since the mid-1990s has been used by some observers as evidence of the superiority of U.S. institutions. Given that the recent patterns can be explained by a medium-term economic model driven by demographics, we believe that such policy conclusion should be treated with caution.²⁴

At a more speculative level, we have suggested that productivity patterns over the last 30 years may best be explained by a model of a technological revolution where the new technological paradigm favours the development of a more competitive labour market at the cost of a negative level shift in productivity. We argued in favour of this interpretation based on both the pattern of labour productivity observed in the United States over the last 50 years, as well as the view that it takes excessive forward-looking behaviour by agents to explain a process that involves lost productivity for 20 years before seeing any gains. However, only more research on the issue will allow us to provide a more definite answer.

Appendix A. Data

Three datasets were used in this study. Data on output, population and investment shares are taken from the Penn World Table 6.0, downloadable from <http://webhost.bridgew.edu/baten/>. Education data are taken from Barro and Lee's (2000) dataset, which is downloadable from <http://www.cid.harvard.edu/ciddata/ciddata.html>. Finally, data regarding employment, unemployment, and labour force are taken from the OECD Economic Outlook 75.

Our measure of income, y , is the logarithm of real GDP chain per worker (RGDPW in PWT 6.0), where the definition of a worker is based on economically active population.

Population is POP in PWT 6.0. Workers are computed as the population from 15 to 64 obtained from

$$\text{POPW}_t = \frac{\text{real GDP chain per capita}}{\text{real GDP chain per worker}} \times \text{population} = \frac{\text{RGDPL}}{\text{RGDPW}} \times \text{POP}.$$

n then denotes the rate of growth of the 15–64 population.

The PWT 6.0 stopped in 1998. We therefore used the rate of growth of investment, output and population obtained from OECD sources in order to extend the dataset to 2002.

24 Indeed, one might plausibly build a political economy model in which the higher population growth economies going through the initial productivity slowdown turn to radically altering their social programs as a policy response. Later, when the inevitable turnaround occurs, it would appear that the cuts in social policy caused the improved growth.

In the version of the OECD Economic Outlook we used, employment, unemployment, and labour force data started later than 1960: Australia (1964), Austria (1965), and Netherlands (1969). We therefore used the rate of growth of each of these variable (obtained from an earlier version of the Compendium) to retruncate the series.

Appendix B. Solving the model

The interested reader is referred to the technical appendix of this paper for detailed proofs of the propositions (available from <http://fabcol.free.fr>). In this section, we aim only at solving the labour allocation problem and use it to derive output per worker.

As indicated in the body text, the labour allocation problem, stated in intensive form, is

$$\begin{aligned} \max_{\{s(t), h(t)\}_{t=t^*}^{\infty}} & \int_{t^*}^{\infty} e^{-(r-n)(t-t^*)} (w^o(t)(1-h(t)-s(t)) + w^N(t)s(t)) dt \\ \text{s.t. } & \dot{S}(t) = \Omega h(t) + n\bar{s}(t) - ns(t). \end{aligned}$$

Letting $\mu(t)$ denote the shadow price of skilled labour, we have the following set of optimality conditions.

$$w^o(t) = \Omega \mu(t) \tag{B1}$$

$$\dot{\mu}(t) = -(w^N(t) - w^o(t)) + r\mu(t) \tag{B2}$$

As soon as the economy adopts the new technology, households stop investing in the old capital, which just depreciates over time, such that, using the demand for old capital, we have²⁵

$$k^o(t) = k^o(t^*)e^{-(\delta+n)(t-t^*)}, \forall t > t^*. \tag{B3}$$

On the contrary, it is always worthwhile to invest in the new technology and the rental price of capital is given by $q^N(t) = r + \delta$. Then, the capital stock is given by

$$k^N(t) = (1 + \gamma) \left(\frac{\alpha}{r + \delta} \right)^{\frac{1}{1-\alpha}} s(t). \tag{B4}$$

Using this result in the demand for skilled labour, we get

$$w^N(t) = (1 - \alpha)(1 + \gamma) \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}}. \tag{B5}$$

25 See the technical appendix for a formal proof of this statement.

Since $w^N(t)$ actually jumps to its steady-state value, the system (B1)–(B2) can be simply solved to give

$$w^o(t) = \frac{\Omega}{r + \Omega} w^N(t) = (1 - \alpha)(1 + \gamma) \frac{\Omega}{r + \Omega} \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}}. \quad (\text{B6})$$

Then, using the demand for unskilled labour, we obtain an expression for $h(t)$:

$$h(t) = 1 - s(t) - \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \left(\frac{r + \delta}{\alpha} \right)^{\frac{1}{1-\alpha}} k^o(t^*) e^{-(\delta+n)(t-t^*)}. \quad (\text{B7})$$

Plugging the latter expression in the law of motion of skilled labour and solving the implied differential equation, we get

$$s(t) = 1 - e^{-\Omega(t-t^*)} - \frac{\Omega}{\Omega - \delta - n} \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \left(\frac{r + \delta}{\alpha} \right)^{\frac{1}{1-\alpha}} \cdot k^o(t^*) \left(e^{-(\delta+n)(t-t^*)} - e^{-\Omega(t-t^*)} \right).$$

Plugging this result in the definition of the old and the new technologies, we obtain the following expression for output per worker

$$\begin{aligned} y(t) = & (1 + \gamma) \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}} + \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \frac{(1 + \gamma)(r + \delta)}{\alpha} \\ & \cdot \left(\frac{\Omega}{r + \Omega} - \frac{\Omega}{\Omega - \delta - n} \right) k^o(t^*) e^{-(\delta+n)(t-t^*)} + (1 + \gamma) \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}} \\ & \cdot \left(\frac{\Omega}{\Omega - \delta - n} \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \left(\frac{r + \delta}{\alpha} \right)^{\frac{1}{1-\alpha}} k^o(t^*) - 1 \right) e^{-\Omega(t-t^*)}, \end{aligned}$$

which corresponds to the law of motion of output per worker reported in the main text.

Appendix C. Proof of propositions

Proof of proposition 2. The steady-state output in the economy where only the old technology is available is given by

$$\bar{y}^o = \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}}.$$

The immediate effect of introducing the new technology is that output jumps to the new level

$$y(t^*) = \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1-\alpha}{\alpha}} \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}},$$

where we set $k^o(t^*)$ to the steady-state value of the economy with the old technology. Noting that this rewrites

$$y(t^*) = \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1-\alpha}{\alpha}} \bar{y}^o,$$

and that by assumption $\Omega\gamma/r > 1$, we have $y(t^*) < \bar{y}^o$. □

Proof of proposition 3. In order to prove the first part of the proposition, it is convenient to rewrite the dynamics of output per worker as

$$y(t) = \psi_0 + \left(\psi_1 - \frac{\psi_2}{\Omega - \delta - n} \right) e^{-(n+\delta)(t-t^*)} + \left(\frac{\psi_2}{\Omega - \delta - n} - \psi_0 \right) e^{-\Omega(t-t^*)},$$

where

$$\psi_0 = (1 + \gamma) \left(\frac{\alpha}{r + \delta} \right)^{\frac{\alpha}{1-\alpha}}, \psi_1 = \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \frac{(1 + \gamma)(r + \delta)}{\alpha} \frac{\Omega}{r + \Omega} k^o(t^*)$$

and

$$\psi_2 = \left(\frac{r + \Omega}{\Omega(1 + \gamma)} \right)^{\frac{1}{\alpha}} \frac{(1 + \gamma)(r + \delta)}{\alpha} \Omega k^o(t^*).$$

What we need to show is that the larger the rate of population growth, the lower the growth in output per worker at the beginning of the transition, which amounts to proving that the first-order derivative of the rate of growth of output per worker with respect to n is negative when t is close to t^* .

The rate of growth of output per worker writes as

$$\frac{\dot{y}(t)}{y(t)} = - \frac{(\delta + n) \left(\psi_1 - \frac{\psi_2}{\Omega - \delta - n} \right) e^{-(n+\delta)(t-t^*)} + \Omega \left(\frac{\psi_2}{\Omega - \delta - n} - \psi_0 \right) e^{-\Omega(t-t^*)}}{\psi_0 + \left(\psi_1 - \frac{\psi_2}{\Omega - \delta - n} \right) e^{-(n+\delta)t} + \left(\frac{\psi_2}{\Omega - \delta - n} - \psi_0 \right) e^{-\Omega(t-t^*)}} = - \frac{u(t, n)}{v(t, n)}.$$

We therefore have that

$$\frac{\partial \dot{y}(t)/y(t)}{\partial n} = - \frac{\frac{\partial u(t, n)}{\partial n} v(t, n) - \frac{\partial v(t, n)}{\partial n} u(t, n)}{v(t, n)^2}.$$

Straightforward calculation shows that at the time of introduction of the new technology, $t = t^*$, we have $u(t, n) = (\delta + n)\psi_1 - \Omega\psi_0 + \psi_2$, $v(t, n) = \psi_1$, $\partial u(t, n)/$

$\partial n = \psi_1$, and $\partial v(t, n)/\partial n = 0$. Hence, plugging these results into the derivative of the rate of growth of output per worker evaluated at time t^* , we get

$$\left. \frac{\partial \dot{y}(t)/y(t)}{\partial n} \right|_{t=t^*} = -1,$$

which proves the first part of the proposition.

The second part of the proposition is trivial, since the steady state of the economy does not depend on the rate of population growth. So, no matter n , all economies will tend to the same limit. Since the rate of growth in the high-population-growth economy is lower at the beginning of the transition than it is in the low-population-growth economy, it has to be greater at some point to converge to the same steady state. \square

References

- Aghion, P., and P.W. Howitt (1998) 'On the macroeconomic effects of major technological change,' in *General Purpose Technologies and Economic Growth*, ed. E. Helpman (Cambridge MA: MIT Press)
- Barro, R.J., and J.W. Lee (2000) 'International Data on Educational Attainment Updates and Implications,' NBER Working Paper 7911
- Beaudry, P., and D. Green (1998) 'What is driving US and Canadian wages: endogenous technical change or endogenous choice of technique?' NBER Working Paper 7697
- (2000) 'Population growth, technological adoption and economic outcomes in the information era,' *Review of Economic Dynamics* 5, 749–74
- (2003) 'The changing structure of wages in the US and Germany: what explains the differences?' *American Economic Review* 93, 573–602
- Beaudry, P., and F. Collard (2002) 'Why has the employment-productivity tradeoff among industrialized countries been so strong?' NBER Working Paper 8754
- (2003) 'Recent technological and economic change among industrialized countries: insights from population growth,' *Scandinavian Journal of Economics* 105, 441–63
- (2004) 'Globalization, returns to accumulation and the world distribution of output,' Working Paper 10565, NBER 2004
- Beaudry, P., F. Collard, and D. Green (2004) 'Decomposing the twin peaks: a study of the changing world distribution of output per worker,' NBER Working Paper 9240
- Blanchard, O.J. (1997) 'The medium run,' in *Brookings Papers on Economic Activity*, Vol. 16 (Cambridge MA: MIT Press)
- Bresnahan, T., and M. Trajtenberg (1997) 'General purpose technologies: engines of growth?' *Journal of Econometrics* 65, 83–108
- Caselli, F. (1999) 'Technological Revolutions,' *American Economic Review* 89, 78–102
- Durlauf, S.N., and D.T. Quah (1999) 'The new empirics of economic growth,' in *Handbook of Macroeconomics*, ed. J.B. Taylor and M. Woodford Vol. 1 (Amsterdam: Elsevier Science)
- Feyrer, J. (2002) 'Demographics and productivity,' mimeo, Dartmouth College
- Greenwood, J., and M. Yorukoglu (1997) '1974,' *Carnegie-Rochester Conference Series* 46, 49–95
- Helpman, E., and T. Trajtenberg (1998) 'A time to sow and a time to reap: growth based on general purpose technologies,' in *General Purpose Technologies and Economic Growth*, ed. E. Helpman (Cambridge MA: MIT Press)

- Klenow, P.J., and A. Rodriguez-Clare (1997) 'The neoclassical revival in growth economics: has it gone too far?' in *NBER Macroeconomics Annual*, ed. B.S. Bernanke and J. Rotemberg, Vol. 12 (Cambridge MA: MIT Press)
- Lucas, R. (1988) 'On the mechanisms of economic development,' *Journal of Monetary Economics* 22, 3–42
- Mankiw, N.G., D. Romer, and D.N. Weil (1992) 'A contribution to the empirics of economic growth,' *Quarterly Journal of Economics* 107, 407–37
- Uzawa, H. (1965) 'Optimum technical change in an aggregative model of economic growth,' *International Economic Review* 6, 18–31