Explaining Productivity Growth: The Role of Demographics

Paul Beaudry
University of British Columbia
Fabrice Collard
University of Toulouse
David A. Green[1]
University of British Columbia

PRODUCTIVITY GROWTH IN THE United States has been higher than in most other industrial countries since the mid 1990s, and has picked up since 2000. Based on this observation, it has become quite common to hear that flexible, U.S.-style institutions are the key to economic growth in the current technological era. Growth and technological change, though, are longer term phenomena by their nature, and much can be learned by extending our examination of growth rates back several decades. Doing so, we will show, puts cross-country differences in economic growth in a different light and suggests that demographic differences across countries may have played an important role in shaping growth patterns in the period from the mid to late 1970s to the present.

It is our claim that the post-1995 period of economic growth should not be seen in isolation, but instead should be viewed in the context of a major economic transformation that started somewhere in the mid to late 1970s and is still affecting us now. The main fact emerging from our research is that, within this period, differences in the rate of growth of the labour force appear to have played a key role in determining the different timing across countries of the productivity slowdown and subsequent pick-up. For example, we will show that countries with high rates of labour force growth experienced particularly slow productivity growth over the period 1978-1995, before experiencing a turn-around after 1995. In contrast, countries with low rates of labour force growth appear to have experienced a prolonged period of slow productivity growth that extended into the post-1995 era. We will also show that the effects of labour force growth evident in the post-1975 data were not present in earlier years, suggesting that the post-1975 period may well have been special in terms of economic transformation.

In this article, we will begin by briefly reviewing the research path that led us to hypothesize a substantial role for labour force growth in explaining cross-country differences in economic growth since the mid 1970s. We then present several pieces of evidence that illustrate our claim regarding the role of demographics in shaping labour productivity profiles over this period. Finally, in the last section, we compare the patterns observed in the data with those

1 Paul Beaudry is Professor of Economics, Canada Research Chair and Fellow of the Bank of Canada. Fabrice Collard is Researcher and was Visiting Professor of Economics at the University of British Columbia in 2002. David A. Green is Professor of Economics. Email: green@econ.ubc.ca.

predicted by theories of major technological change. In particular, we will discuss the extent to which cross-country patterns in economic growth fit with a model in which countries began to make a transition to a new technological and organizational mode of production beginning in the mid 1970s. This, in itself, is not a particularly novel claim since many observers have argued that the advent of widespread use of computers has been associated with substantial changes in the organization of production. However, we argue that the empirical evidence fits with a model in which countries with rapid labour force growth have been induced to undergo the transition first, with the fastest adjusting countries completing the transition by the mid 1990s. This is a story which takes a longer term view of technological change and highlights the potential pitfalls of using shorter time frames for making inferences regarding the forces affecting growth patterns. Whether or not our interpretation of recent economic history ultimately stands up to more extended testing, it serves as a caution to those who use five or ten years' worth of economic data to make pronouncements on the relationship between policy, institutions and economic growth.

Recent Trends in Economic Growth

In our ongoing research on the effects of technological change on the labour market, and on the economy more generally since the mid 1970s, we have come to the conclusion that labour force growth has played a much larger role than is usually recognized. We first formulated this conjecture in the context of examinations of changes in the wage structure in Canada, the United States and Europe over the past thirty years. In a series of papers (Beaudry and Green, 1998, 2003 and forthcoming) we argued that relative movements in skilled and unskilled wages, the skilled and unskilled labour force, and physical capital fit with the notion that we have been witnessing a period of technological transition in which the rate of transition is endogenously determined by relative supplies of inputs in a given country. Empirical examinations revealed that capital deepening (i.e., movements in amounts of physical capital per worker) was the major driving force behind differences in rates of technological change and that this deepening was in turn driven primarily by demographic differences. That is, the countries we studied (Canada, the United States, the United Kingdom and Germany) experienced differences in rates of capital deepening not because of differences in rates of growth of the physical capital stock but because of differences in the rates of growth of the labour force. Building on those findings, we extended our investigation to patterns of growth in output per worker both across developed and developing countries (Beaudry and Green, 2002, Beaudry and Collard, 2003 and Beaudry, Collard and Green, 2004 and 2005). It was in these investigations that we uncovered both the importance of labour force growth as a determinant of differences in economic growth rates across countries in recent decades and the fact that this pattern represented a substantial shift from earlier decades.

Our first goal in this article is to illustrate that there has been a major change in the process determining economic growth since the mid to late 1970s, and that this change has involved an increased role of labour force growth. More specifically, we will review the main pieces of evidence presented in Beaudry, Collard and Green (2005). Our investigation focuses primarily on the experiences of the set of major industrialized countries: Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Italy, Japan, the Netherlands, New Zealand, Norway, Sweden, Switzerland, the United Kingdom and

the United States,[2] which we follow over the period 1960-2002. The key issue we examine is the growth patterns in output per worker and their relationship with the rate of labour force growth.[3] Our main measure of the labour force is the population aged 15 to 64. We have also used estimates of the actual labour force, but we are concerned that these may respond contemporaneously to changes in economic growth. In contrast, movements in the size of the population at labour force-relevant ages are most likely driven by demographic forces that can reasonably be seen as exogenous in the medium run. However, let us stress that our results do not change if we use estimates of the actual labour force.

Econometric Evidence for Selected OECD Countries

The first step in our investigation relies on the tests presented in Beaudry, Collard and Green (2005) which establish that there was a structural break in the relationship between economic growth and labour force growth over the period 1960-2002. We refer the reader to that paper for a complete exposition of the tests we perform. Here, we will simply state the main conclusion: using regressions of growth in output per worker on initial levels of output per worker and labour force growth, we find strong evidence of a structural break in approximately 1978.[4] There is also some evidence of a further change in the process in approximately 1995. To illustrate the size and nature of this structural break, in Chart 1 we plot the cumulative effect of a one per cent difference in annual labour force growth on the growth of output per worker. In the first panel, we report the estimated effects for the period 1960-1978, with each point on the line representing the cumulative effect on output observed between 1960 and the given year. Essentially, one can think of the figure showing the difference in the growth that would have been achieved by a given year if a country had an ongoing labour force growth rate of one per cent per year versus what would have been achieved if there had been no labour force growth. The results in the first panel indicate that countries with higher rates of labour force growth did not experience any significant difference in productivity growth relative to lower labour force growth countries. For example, over the entire period 1960-1978, our point estimates indicate that countries with one per cent per year higher average annual labour force growth grew approximately 0.5 per cent more over the entire period (i.e., 0.5/18 per cent per year more), which is extremely small and not statistically significant. In panel 2, we report the same type of estimates for the period 1975-2002. In contrast to the results in the first panel, we see that the effect of labour force growth on labour productivity growth was substantial and time varying over this period. If we focus on the estimated effect as of 1995, we see that a country with a one per cent per year higher rate of labour force growth experienced approximately 15 per cent less growth in output per worker over the period 1975-1995, which implies a slower growth rate of about 0.75 per cent per year for 20 years. Then after 1995, we see that the pattern slightly reverses itself, indicating that countries with high rates of labour force growth started doing better than their low labour force growth counterparts

2 The specific countries were chosen as those with more than a million people and with GDP per capita in 1980 greater than 50 per cent of the U.S. level. We chose to focus on the richest countries since they seemed to be the strongest candidates for being at the technological frontier and therefore for being affected by the type of forces we wish to discuss.

3 Our output data come from the 6.0 version of the Penn World Tables while our labour force data come from the OECD Statistical Compendium.

4 The initial level of output per worker is a standard regressor in the economic growth literature, intended to capture convergence in growth over time. The tests indicate that a set of years from 1975 to 1980 are all strong candidates for the break point, with 1978 being the preferred year in that set. Use of any of the other years in the set as the break point does not alter our results or conclusions.

Chart 1

The Cumulative Effect of One per cent per year Growth in the Working Age Population on Output per Worker Growth

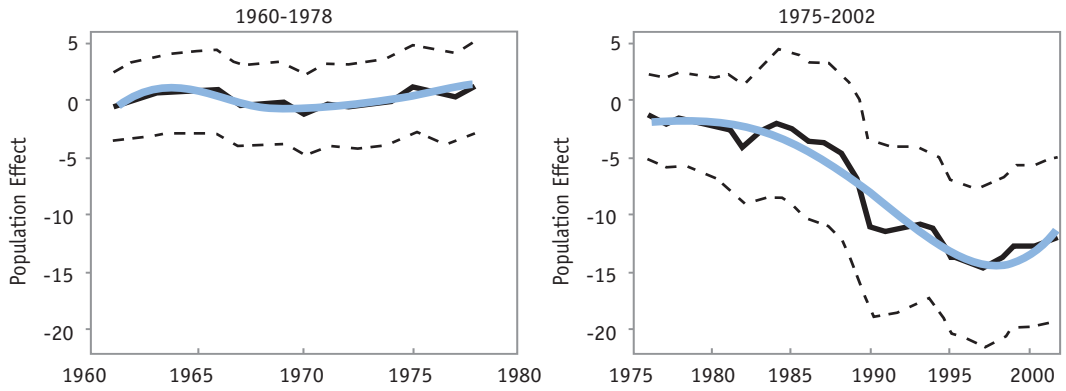

Notes: The dotted line represents a 95 per cent confidence band. The grey line corresponds to a fourth order polynomial of time fitted on the series of time varying coefficients.

Table 1

Regression Results for Output per Worker Growth in Selected OECD Countries

Baseline Regressions

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

Notes: Standard errors in parentheses. The stability test lines report the p-value of the test.

in the post-1995 era. However, note that even by 2002, countries with higher rates of labour force growth had still not entirely reversed the drop in labour productivity they experienced over the period 1975-95 relative to lower labour force growth countries.

To summarize and give precision to the observations in Chart 1, in Table 1 we present the results of separate regressions of the growth in output per worker on a constant, initial output per worker, and labour force growth for each of the periods 1960-1978, 1978-1995 and 1995-2002. The estimated coefficients in this table effectively capture the patterns observed in the chart. In particular, in the pre-1978 period, the impact of labour force growth on economic growth is quite small in economic terms and is not statistically significantly different from zero at any conventional significance level. In addition, the coefficient on the initial (in this case 1960) level of output per worker is negative and highly significant, fitting with convergence models: those countries that start farthest behind have the strongest growth rates. The size of this latter coefficient is similar to that found in investigations by numerous other authors.

After 1978, the process changes dramatically. First, the initial (in this case, 1978) level of output per worker no longer has an economically substantial or statistically significant effect. More strikingly, though, the coefficient on labour force growth has become very large and highly statistically significant. The numbers in the last two rows of the table are p-values corresponding to tests of the hypotheses that there is no change in

either of the slope parameters in the regression between the first and second periods (the first of the two values) and that there is no change in the effect of the change in labour force growth (the second of the two values). Both reflect a strong rejection of the restriction of no change. That is, there is strong statistical evidence of a change in the process determining growth after versus before 1978, with the main feature of that change being a large increase in the importance of labour force growth in determining productivity growth over the period 1978-95.

The last column in the table indicates that there may have been a further change in the post-1995 period. The labour force growth effect becomes relatively large and positive. However, neither convergence forces nor labour force growth have statistically significant effects. Thus, there is some indication that high labour force growth countries have moved from being laggards in growth in the 1978-1995 period to leaders in the post-1995 period. However, the period is too short to draw firm conclusions from it, as reflected in the large standard errors on the estimated effects.

The patterns evident in Table 1 are also present even in the simplest plots of the data. In Chart 2 we plot the relationship between average annual output per worker growth and the average annual rate of growth of the population aged 15 to 64 for the period 1978 to 1995. The figure shows the strong negative relationship between the two variables for this period, supporting the results in the middle column of Table 1. The figure includes a simple regression line fit to the plotted points. That line has a slope of -0.937 with a standard error of 0.223. The interesting implication (which we do not pursue here) is that this set of developed countries experienced almost identical growth rates in output in this period, resulting in differences in growth rates in output per worker, on average, being entirely associated with differences in labour force growth rates.

Chart 2
Labour Productivity Growth and Population (15-64) Growth in Selected OECD Countries, 1978-1995

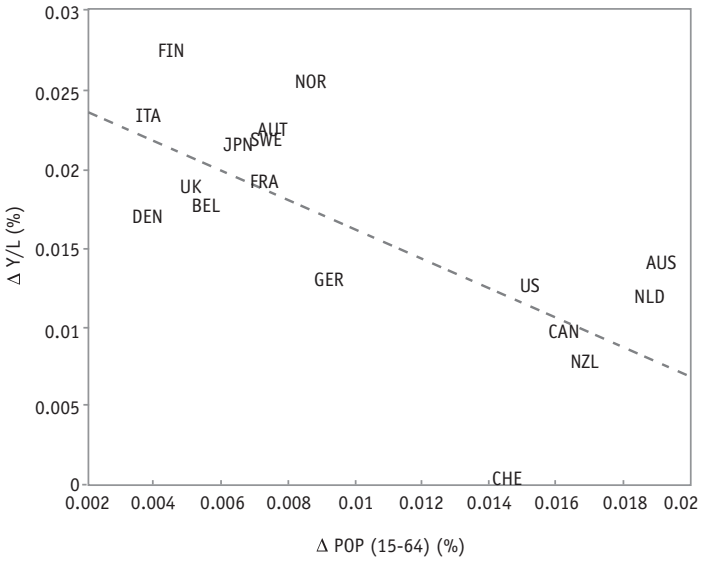

Notes: The plain line corresponds to the regression model:
$\Delta \log(Y/L) = 0.026 - 0.937 \Delta POP1564$, $\bar{R}^2 = 0.53$
(0.002) (0.223)

Checking the Robustness of the Results

To this point, we have documented a surprising relationship between labour force growth and productivity growth since 1978. However, it is possible that this link is not causal and that we are really capturing other, so far omitted, factors. Two obvious candidates are the investment rate, which is a standard variable of interest in neo-classical growth theory, and measures of human capital, which have received considerable attention in recent discussions of growth (see, for example, Uzawa, 1965, Lucas, 1988, Mankiw, Romer and Weil, 1992, Durlauf and Quah, 1999 and Klenow and Rodrìguez-Clare, 1997). In Table 2, we present results from extending our specification by introducing the investment rate (measured as the ratio of investment to GDP) and a human capital measure (average years of schooling). Because we already viewed the post-1995 period as too short to draw firm conclusions, we focus on the first two periods when we bring in the additional

Table 2
Regression Results for Output per Worker Growth in Selected OECD Countries
Robustness of Timing

	1960-1978	1978-1995
Cst.	0.225	0.011
	(0.059)	(0.055)
$(Y/L)_0$	-0.035	-0.003
	(0.007)	(0.008)
ΔN	0.268	-0.883
	(0.260)	(0.312)
I/Y	-0.006	0.006
	(0.009)	(0.007)
H	-0.034	0.002
	(0.015)	(0.013)
R^2	0.81	0.57
Q(Total)	—	0.012
$Q(\Delta N)$	—	0.005

Notes: Standard errors in parentheses. The stability test lines report the p-value of the test.

variables. (In the post-1995 period, almost all variables are found to be insignificant, which is not very surprising over such a short period.) An inspection of Table 2 reveals that the new regressors have effects that are, by and large, economically insubstantial and statistically insignificant. But what is most important from our perspective is that introducing these variables does not alter the main patterns from Table 1. The labour force growth effect is now positive in the pre-1978 period but is still far from being statistically significantly different from zero, and it again becomes negative, large and statistically significant in the second period.

In Beaudry, Collard and Green (2005) we provide further robustness checks for our main result. Based on recent arguments that institutions, and particularly Anglo-Saxon institutions, play an important role in economic growth, we introduce dummy variables corresponding to both Anglo-Saxon and Nordic countries. We also introduce, as a separate regressor, the percentage change in the country's unemployment rate within each period in an attempt to capture the notion that some countries may have institutions that favour higher labour productivity at the cost of rising unemployment. Finally, we also introduce measures of changes in the age structure of the population, in an attempt to control for the possibility that differences in experiences with baby booms had effects on age structures that ultimately had impacts on growth. In all of these cases, the results in terms of the labour force growth effects are the same. Specifically, those effects remain large, negative and statistically significant in the 1978-1995 period across all specifications.

One final concern is with the potential endogeneity of the labour force growth rate: we may be picking up the impact of improved economic growth on fertility decisions rather than the causal impact of changes in labour force growth on economic growth. In Beaudry, Collard and Green (2005), we address this by essentially examining the relationship of economic growth in the 1978-1995 period to historic population growth rates from before that period. The idea is that these historic growth rates still allow us to identify perennially high labour force growth economies but that population growth rates from before 1978 should not be affected by changes in economic growth after 1978.[5] The results of this exercise imply, if anything, a larger negative impact of population growth on economic growth in the 1978-1995 period.

Extending the Investigation to Developing Countries

Our investigation so far has focused on a set of rich countries. We did this, in part, as a means of controlling for other factors, such as political

5 Formally, we use an instrumental variables estimator employing each country's age 15-64 population growth from 1960 to 1970 as an instrument for the 1978-1995 age 15-64 population growth rate. This instrument is valid if families did not accurately predict the post-1978 growth patterns and alter their fertility decisions in a way that affected the adult population growth over the 1960 to 1970 period.

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 1960-1998. The results from re-estimating both our basic specification and the physical and human capital robustness checks using the full world sample are presented in Table 3. The results from the basic specification, shown in the first two columns, lead to the same conclusions as those drawn from the rich country sample. The pre-1978 labour force 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 labour force 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 labour force growth changed between the two periods. The other difference relative to the earlier sample is that the initial output per worker 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 in the results from the richer country sample, we check the robustness of these results by introducing physical and human capital variables.[6] Introducing the investment rate variable reduces the size of the estimated labour force 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 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 labour force growth 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.

Table 3
Regression Results for Output per Worker Growth in Selected Developed and Developing Countries

	1960-1978	1978-1998	1960-1978	1978-1998
Cst.	0.024	0.030	0.152	0.140
	(0.019)	(0.020)	(0.023)	(1.029)
$(Y/L)_0$	0.001	-0.000	-0.015	-0.010
	(0.002)	(0.002)	(0.003)	(0.003)
ΔN	-0.230	-0.976	-0.092	-0.808
	(0.214)	(0.185)	(0.166)	(0.179)
I/Y	—	—	0.011	0.019
			(0.003)	(0.004)
H	—	—	0.017	0.008
			(0.003)	(0.006)
R^2	0.02	0.24	0.43	0.45
Q(Total)	—	0.000	—	0.000
$Q(\Delta N)$	—	0.004	—	0.002

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

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

Chart 3
Index of Output per Hour in the U.S. Business Sector, 1960-2004

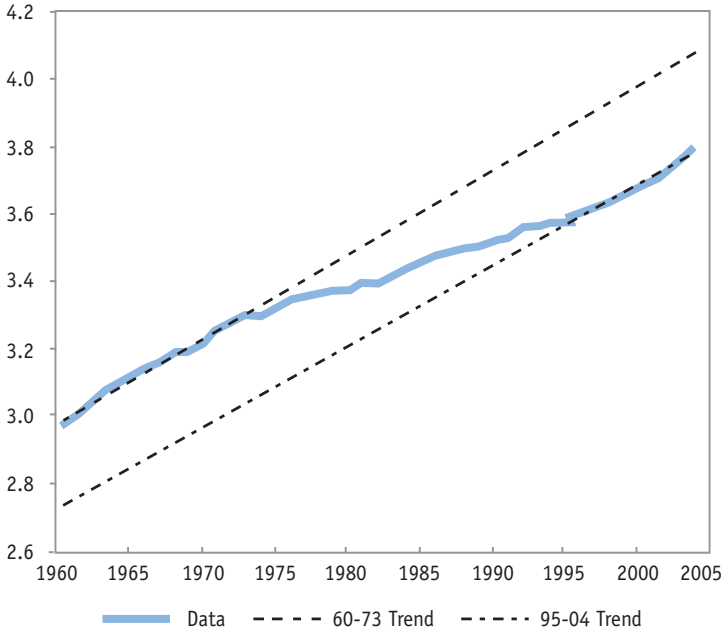

Chart 4
Index of Output per Hour in the Canadian Business Sector, 1960-2004

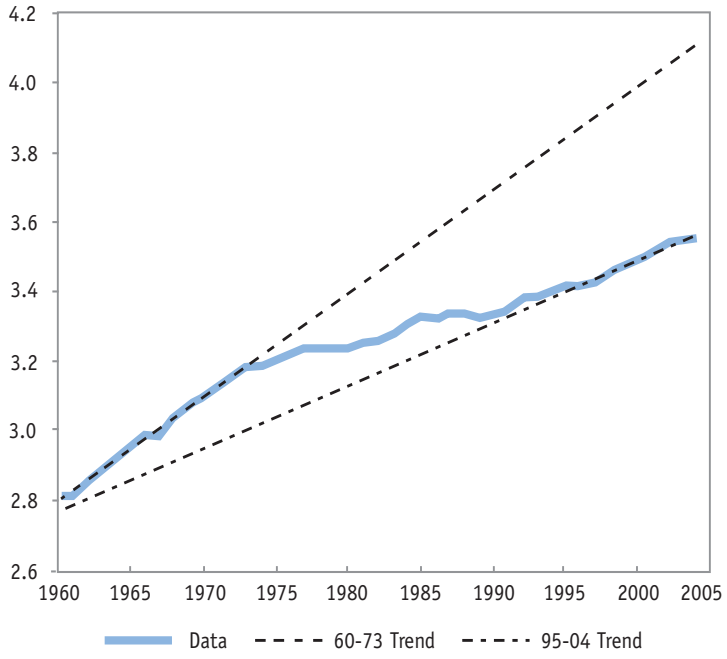

Productivity Growth Within Selected OECD Countries Over Time

To this point, we have focused on cross-country patterns in economic growth. It is informative, also, to present plots of growth patterns within specific countries over time. Chart 3 provides a plot of annual output per hour worked for the United States over the period 1960 to 2004. In the figure, we have also plotted the trend line fitted to the pre-1973 data, extending the line beyond 1973. We chose 1973 only because an inspection of the plotted data suggested a break in the series at that point. Examining the plotted series relative to the pre-1973 trend line makes it apparent that the period from 1973 to the mid 1990s was a period of dramatically lower growth than the pre-1973 period. After 1995, output per hour begins to grow at a better rate again. We also fit a trend line to the post-1995 data. Interestingly, the recent period of growth, which is often depicted as something of a "U.S. miracle", actually reflects a growth rate very similar to the pre-1973 period.

Chart 4 replicates the plot in Chart 3 but with Canadian data. The general pattern is similar to that for the United States: strong pre-1973 productivity growth is followed by a prolonged slow-growth period extending to the mid 1990s and then a more recent period of renewed growth. In contrast to the United States, Canada's post-1995 growth rate is still lower than its pre-1973 rate. While we do not show it here for the sake of brevity, a similar pattern is evident in other developed countries with high labour force growth rates, such as Australia and the Netherlands. There is some variation in whether, as in the United States, their most recent growth experiences match their pre-1973 growth or, as in Canada, their recent growth is slower than their pre-1973 rates. However, the common pattern is the shifting down of the growth trend over time, with recent strong growth outcomes actually only reflecting at most a return to earlier growth rates. If these countries could have

continued on their pre-1973 growth paths, their levels of output per hour, and consequently of output per capita, would be dramatically higher today.

Explaining the Patterns

We turn now to considering potential explanations for the growth patterns presented in the previous section. In Beaudry, Collard and Green (2005), we present formal models of two main possibilities. Here, we will provide heuristic descriptions of those models and refer readers to our earlier paper for details.

The models we wish to consider need to capture the fact that we observed a structural break in the relationship between productivity growth and labour force growth starting in the mid to late 1970s, with the break being followed by a 18 to 20 year period during which high labour force growth countries under-performed relative to low labour force growth countries. In particular, our objective is to show how the emergence of a new technological paradigm can cause such a structural break in the process determining productivity growth, thereby causing a temporary negative relationship between productivity growth and labour force growth.

Modeling the Productivity Effects Over Time of a Technological Revolution

The first model we present is, at heart, a standard neo-classical growth model populated by individuals who consume goods produced using capital and labour. The consumers (who are also the workers) maximize their lifetime utility by allocating spending in each period of their lives. For this discussion, labour force growth is assumed to be exogenous and to occur at a constant rate. On the other side of the market, the consumption good is produced using capital and labour in a constant returns to scale technology.

In the initial era, before the major technological change which we introduce later, the model yields the implication that the rate of labour force growth has no effect on the equilibrium growth of output per worker. This is a standard result from neo-classical growth theory.

We next consider a technological revolution which takes the form of the sudden emergence of a new, alternative means (technology) for producing the consumption good. We first consider a new technology which is unambiguously more productive than the old technology. This new technology uses technology-specific physical capital and a new type of skilled worker, where these new skills can only be learned on the job by using the new technology. In contrast, the old technology uses the old form of physical capital and does not require the new skills. We will refer to workers that have the new skills as skilled workers, while those who do not have such skills will be referred to as unskilled workers. The capital previously created for use with the old technology depreciates but cannot be shifted to use with the new technology. Workers and firms must therefore coordinate themselves to retrain and acquire skills needed to work with the new technology. The cost of that investment in skills is the foregone production during the learning period. Thus, shifting fully to the new technology takes time, as investment in the new physical capital and skills unfolds. As a result, not all firms will immediately shift to the new technology nor will all workers be allocated to learning how to use the new technology. Instead, there will be a transition period during which both technologies are in use.

Results of the Model versus Observed Productivity Trends

The results that emerge from an examination of this model are most easily explained using charts. To that end, in Chart 5 we show simulated time paths of output per worker from a model economy with zero labour force

Chart 5
Productivity During a Technological Transition

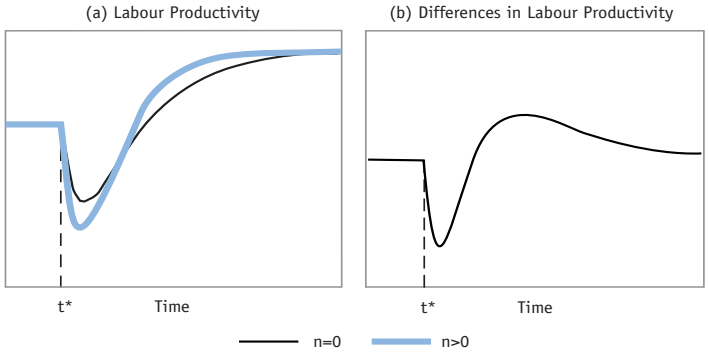

Chart 6
Training and Skilled Labour During a Technological Transition

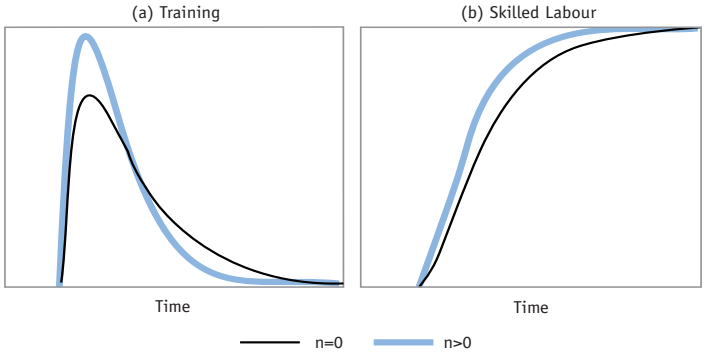

growth and in an otherwise identical model economy with a positive labour force growth rate. In this and subsequent charts, the new technology is depicted as becoming available at time t^*. As can be seen in the chart, in both economies labour productivity initially drops after the introduction of the new technology and then recovers to a higher level.

In the pre-transition period (before t^*), only the old technology is available and, as we stated earlier, there are no differences between economies with high and low labour force 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 Chart 6. This figure plots the evolution of the number of individuals allocated to learning the new technology (panel (a)) and the resulting number of skilled workers (panel (b)). As Chart 6 illustrates, as soon as the new technology is introduced, some workers transfer from using the old technology and start learning how to use the new technology. The result is that fewer workers are operative in the old technology and, in consequence, output drops since workers are not yet productive in the new technology. When there are a sufficient number of skilled workers, the process reverses and labour productivity regains strength. Another phenomenon is also taking place during this transition. The return to capital obtained in the old technology is lower than in the new technology and therefore there is no incentive for the household to keep investing in the old capital stock. Hence, the old technology slowly disappears from use in the economy as its capital stock depreciates.

It is worth noting that this process does not take place at the same speed in the two model economies. In particular, after the new technology is introduced, higher labour force growth countries will initially experience slower growth in labour productivity than their lower labour force growth counterparts. Later, though, the high labour force growth economy has the higher productivity growth rate. This pattern is depicted in the second panel of Chart 5, which shows the time path of the productivity differential between the zero and positive labour force growth economies. Underlying this pattern, as shown in Chart 6, is the fact that there is greater training/learning undertaken in the high labour force growth economy in the early part of the transition to the new technology. This arises from the fact that what is ultimately important for production is the amount of capital available to each worker. In both high and low labour force growth economies, investment in capital

for the old technology stops as soon as the new technology becomes available. The combination of this depreciating old-technology physical capital stock with a faster growing labour force means a more rapid decline of the old-technology capital per worker and, thus, a more rapid decline in the relative importance of the old technology in the higher labour force growth economy. This, in turn, implies a greater incentive to learn the new technology in the higher labour force growth economy, which then means that such an economy has more of its workers engaged in learning and fewer engaged in production in the early stages of the transition. Hence, in the earlier periods of the adoption process, the high labour force growth economy loses more in terms of productivity than the low labour force growth economy. But this eventually reverses itself since both economies are converging toward the same steady state. Indeed, over time, the high labour force 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 productivity growth.

The labour productivity profiles in Chart 5 bear a reasonable resemblance to the American and Canadian productivity paths presented in Charts 3 and 4. In particular, the actual productivity paths follow a pattern of a strong slowdown in growth followed by more rapid growth that is also evident in the predictions from the model. This, in fact, is one of the key points we would like to emphasize. While discussions of productivity movements often focus only on the directly preceding 5 to 10 years, it is often more appropriate to think in terms of longer time frames. Here, we have constructed a model in which a rapid productivity turn-around is not due to new policies but is, instead, a reflection of a longer term technological transition. That such a turn-around may arise only after an earlier period of poor performance appears relevant when assessing the exemplary performance of the U.S. economy since the mid 1990s. In the model, all economies will ultimately pass through this same transition, but high labour force growth economies go through it faster. This, in fact, is the second point we wish to emphasize. The model presents a plausible explanation for the fact that labour force growth has emerged as a strong determinant of productivity growth since the mid to late 1970s. Before the arrival of the new technology (i.e., the new mix of computers and organizational form), both the model and the data indicate that differences in labour force growth across countries are uncorrelated with economic growth rates. Once the new technology is introduced, all economies pay a productivity price for switching both labour and capital resources to the new technology. Because faster labour force growth economies make the transition faster, they first appear to be relatively sluggish and then relatively fast growing economies as the transition proceeds. This is what is observed in the data.

Possible Shortcomings

While this model provides a reasonably good match to the data, it also has some shortcomings which suggest a need to continue exploring alternative explanations. One potential difficulty with the model involves the duration of the technological transition. Based on Charts 3 and 4, one would say that the productivity slowdown phase of the transition lasted from the mid 1970s to the mid 1990s. Recall that in the model this slowdown is primarily caused by workers being reallocated from productive work in the old technology to engaging in learning activities using the new technology. However, it requires an extremely high level of patience for workers and firms to choose to invest in something that does not pay off for approximately 20 years. Indeed, in our quantitative explorations of this

model we found that extremely low discount rates were needed to reproduce the pattern observed in the data. Since such low discount rates seem to us to be implausible, it suggests a need to adjust the model in a manner that relies less on the forward looking behaviour of workers and firms.

The within-country over-time productivity patterns in Charts 3 and 4 also suggest a need to modify the theoretical explanation. In both countries, there is clear evidence of the slowdown in growth followed by improved growth rates that our model predicts. However, the model is built around the assumption that we are witnessing a transition to a higher productivity technology. This would imply that the recent growth trajectories should be steep enough to take these economies to levels of output per capita above what they would have reached had they stayed with the growth path allowed by the old technology. Instead, Charts 3 and 4 indicate that recent growth paths are, at best, on the same trend as the pre-transition paths. While the process may be far from complete, the data so far point to the controversial conclusion that we are undergoing a transition to a new technology that may be less productive.

An Alternative Model: The Possibility of Productivity-Decreasing Technological Change

Given these two concerns with the previous theoretical explanations, we want to discuss briefly an alternative model which still has the feature of explaining why labour force growth suddenly became so important for output per worker growth over the period 1978-1995 but requires less heroic assumptions about individual patience and allows for the possibility that we are making a transition towards a less productive technology. Thus, consider a situation with most of the same characteristics as the environment described earlier. The main difference will be in the nature of the technological change. In particular, instead of introducing a new technological paradigm that requires new skills to operate, we introduce a new technology which renders the motivation or monitoring of work effort easier. Thus, suppose that there is an old technology in which production is accomplished using physical capital and effective hours of labour. The latter is the product of what we might call "face hours" (the time the worker shows his or her face in the workplace) times the effort actually expended per face hour. This implies that we are considering a model that is standard in the efficiency wage literature. In this case, some firms will offer a wage premium over the standard competitive wage in order to elicit more effort from their workers. This, in turn, implies that there are workers who are receiving a wage above the competitive market wage. In the standard language of economics, their wage includes an economic rent component and therefore marginal productivity is not equated to the dis-utility of work time.

Now consider the introduction of an alternative means of production that once again requires a new type of capital. The main feature of this new technology is that it no longer depends explicitly on effort. This might be the case, for example, if the new physical capital included computers that allowed for better monitoring of effort. One implication of this new technology is that firms can pay workers the competitive wage: it effectively eliminates the rents some workers would have been receiving. 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. Thus, the new technology may save on wage costs at the price of inducing lower productivity. If the wage savings are large enough relative to the productivity decline, it is profitable for firms to switch to the new technology

after its introduction. However, by the time all firms follow that profit option and switch to the new technology, the economy is less productive than it would have been if it had stayed with the old technology.

The implied time path for productivity in this modified model is quite similar to the one derived from the earlier model. Chart 7 contains the implied productivity profiles for a positive and a zero labour force growth economy for the case where the new technology is less productive but saves on workers' rents. As in the earlier model, there is a transition period during which both the old and the new technologies are in use. The length of the transition period no longer depends on worker patience – since there is no skills investment in this version of the model – but is determined by the rate of depreciation of the pre-existing capital stock built up for use with the old technology. As in the earlier model, and for the same reasons, the higher labour force growth economy adopts the new technology faster and, as a result, experiences relatively slower productivity growth directly after the new technology is introduced and faster productivity growth later. The key difference, as emphasized by a comparison of Charts 5 and 7, is that after the transition the productivity path is now lower than that achieved with the old technology. Individual firm decisions, in which firms attempt to increase profits by reducing rents going to workers, ultimately lead to a less productive economy, though no firm would likely recognize this as a side effect of its decisions. This is an intriguing possibility which, to our view, is consistent with the productivity data we presented earlier. We are not claiming to have proved that this model completely captures the situation we either are currently facing or have faced in the recent past, but it is a plausible candidate explanation with important ramifications for policy discussions.

Chart 7
Labour Productivity in the Transition

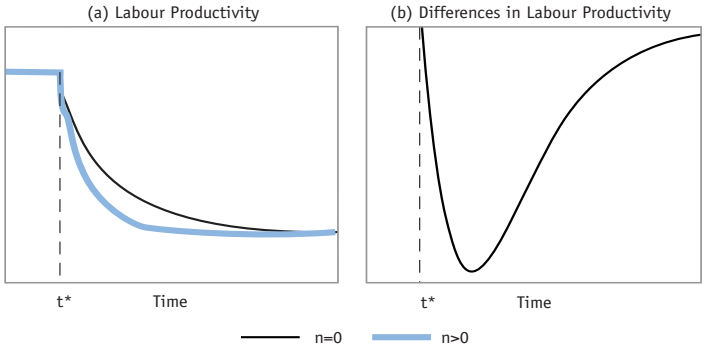

Concluding Remarks

Our main goal in this paper has been to demonstrate that one can gain considerable insight into differences in productivity within and across countries by taking a medium-run perspective. By doing so, we are able to document a dramatic change in the co-movement between productivity growth and labour force growth over the period 1960-2002. In particular, we showed that prior to the mid 1970s, this co-movement was close to zero. From the late 1970s to the mid 1990s, 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 discussed two models of major technological change, in the spirit of the General Purpose Technology literature. In both models, labour force growth interacts with the adoption speed of a new technological paradigm to produce the type of U-shaped pattern of correlations between productivity growth and labour force growth observed in the data. Thus, patterns of productivity and labour force growth across countries provide support for the view that the world economy 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.

These empirical patterns are important when considering economic policy since they suggest that relative economic performances across countries since 1995 may in large part reflect a reversal of patterns observed before 1995, as opposed to the effects of recent policies. In contrast, some observers have argued that the relatively strong productivity growth performance of the United States since the mid 1990s should be viewed as evidence that the United States has found the optimal institutional mix for dealing with the current era of technological change and globalization. Given that the recent patterns can be explained by a medium-term economic model driven by demographics, we believe that such policy conclusions should be treated with caution. In fact, it would not be difficult to build a model in which recent U.S. institutional changes are an effect rather than a cause of the productivity changes we document. In particular, higher labour force growth economies might respond to their relative productivity slowdown in the initial phase of the transition by casting about for new policies, including cutting back on the generosity of their social programs. Later, when the inevitable turn-around in productivity occurred, it would appear that the cuts in social policy caused the improved growth.

At a more speculative level, we have suggested that productivity patterns over the last 30 years may be explained by a model of a technological revolution where the new technological paradigm favours the development of a more competitive labour market without increasing the time path of productivity. We argued in favour of this interpretation primarily based on the labour productivity patterns observed in both Canada and the United States over the last 50 years. However, more research is required before this, undoubtedly controversial, hypothesis can be either accepted or rejected.

References

Beaudry, P. and F. Collard (2003) "Recent Technological and Economic Change among Industrialized Countries: Insights from Population Growth," *Scandinavian Journal of Economics*, vol. 105 no. 3, pp. 441-463.

Beaudry, P., F. Collard and D. Green (2004) "Changes in the World Distribution of Output-per-Worker 1960-98: How a Standard Decomposition Tells an Unorthodox Story," Institute for Fiscal Studies Working Paper W04/15 and forthcoming in the *Review of Economics and Statistics*.

Beaudry, P., F. Collard and D. Green (2005) "Demographics and Recent Productivity Performance: Insights from Cross-Country Comparisons," *Canadian Journal of Economics*, vol. 38, no. 2, pp. 309-344.

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.

Beaudry, P. and D. Green (2002) "Population Growth, Technological Adoption and Economic Outcomes in the Information Era," *Review of Economic Dynamics*, vol. 5 no. 4, pp. 749-774.

Beaudry, P. and D. Green (2003) "The Changing Structure of Wages in the US and Germany: What Explains the Differences?" *American Economic Review*, vol. 93 no. 3, pp. 573-602.

Beaudry, P. and D. Green (forthcoming) "Changes in US Wages 1976-2000: Ongoing Skill Bias or Major Technological Change?" *Journal of Labor Economics*.

Durlauf, S. N. and D. T. Quah (1999) "The New Empirics of Economic Growth," in J. B. Taylor and M. Woodford (eds.) *Handbook of Macroeconomics*, vol. 1, (Amsterdam:Elsevier Science).

Klenow, P. J. and A. Rodrìguez-Clare (1997) "The Neoclassical Revival in Growth Economics: Has it Gone too Far?" in B. S. Bernanke and J. Rotemberg (eds.) *NBER Macroeconomics Annual*, vol. 12 (Cambridge:MIT Press), chapter 2.

Lucas, R. (1988) "On the Mechanisms of Economic Development," *Journal of Monetary Economics* vol. 22, pp. 3-42.

Mankiw, N. G., D. Romer and D. N. Weil (1992) "A Contribution to the Empirics of Economic Growth," *Quarterly Journal of Economics* vol. 107, no. 2, pp. 407-437.

Uzawa, H. (1965) "Optimum Technical Change in an Aggregative Model of Economic Growth," *International Economic Review*, vol. 6, no. 1, pp. 18-31.