Cross-Country Variations in National Economic Growth Rates: The Role of "Technology"

J. Bradford De Long[1]

University of California at Berkeley
National Bureau of Economic Research

May 1996

Abstract

Technology is both more and less important than the conventional wisdom recognizes as a determinant of differences across national economies in productivity levels. "Technology" in the sense of total factor productivity is *more* important because of the strong endogeneity of population growth and investment rates that magnifies small total factor productivity differentials manyfold in steady state. Thus the apparent paradox of "conditional convergence"—national economies that seem to move toward their steady-state growth paths—coupled with the continuing divergence of relative national GDP per capital levels in the world economy.

By contrast, *technology* proper is less important: much if not most differences in total factor productivity are only tenuously or not at all related to mastery of technology in the sense of the internal combustion engine or the freeze-drying process. Robert Solow (1957) called shifts in total factor productivity "technical change"; his doing so may not have helped economists think clear thoughts over the past forty years.

[1]Prepared for the Federal Reserve Bank of Boston conference on Technology and Economic Growth, June 5-7, 1996. I would like to thank Chad Jones, Lant Pritchett, and Robert Waldmann for helpful conversations.

I. Introduction

I want to praise "technology" as *the* important factor in the relative growth performance of nation-states' economies.

I want to argue that the conventional wisdom substantially understates the role of total factor productivity differences in explaining differentials across nation-state economies in GDP per capita. "Technology" in this sense is *more* important, because of the strong endogeneity of population growth and investment rates. Rich economies are economies in which children are much more "consumption" than "investment" goods, and that have completed their demographic transitions to a régime of low fertility and low population growth. Thus an economy that, initially, finds itself with a small total factor productivity advantage will see that advantage magnified into a larger advantage in output per capita as it converges to a steady-state growth path with lower population growth and a higher capital-output ratio.

Similarly, a rich economy is an economy in which the price of capital goods is relatively low: in a rich economy a given share of national product saved translates into a greater real investment effort than if the economy had the world's average relative price structure. This channel magnifies differences in total factor productivity into larger differences in output per capita working through the steady-state capital output ratio.

Researchers in economic growth have been puzzled by the apparent combination of "conditional convergence" with absolute divergence. Economies appear to be moving toward their individual steady-state growth paths by about two percent per year. Yet the spread of relative output per capita levels across the world continues to increase.

A naive interpretation of this pattern would suggest that at some time in the past nation-states' savings and population growth rates—and thus their output per capita levels—were closer together than they are now, that some shock drove savings and population growth rates apart, and that since then the world's distribution of relative incomes has diverged as economies have traversed toward their steady-state growth paths. But what was this shock that drove savings and population growth rates apart? The evolution of the world's cross-country distribution of income and productivity is much more understandable once one recognizes the endogeneity of factor accumulation, and that relatively poor countries have low investment and high population growth rates in large part because they are relatively poor.

But I also have a caveat: there is also a sense in which I want to bury *technology*. Robert Solow's (1957) article is entitled "Technical Change and the Aggregate Production Function." Certainly since 1957, and perhaps before, economists have used "technical change" and "technology" as shorthand ways of referring to shifts in the aggregate production function. Yet much of difference seen across nations in aggregate total factor productivity has little to do with *technology*—in the sense of knowledge of the internal combustion engine, continuous-casting, the freeze-drying process, or anything that would be recognizable in a model like that of Caballero and Jaffe (1993). *Technology* properly so-called is the ultimate source of our enormous material wealth today relative to our counterparts of a century or so ago: economic growth over the past century in the United States *is* built on our knowledge today of the internal combustion engine, continuous-casting, freeze-drying, and all of our other technologies. Yet differences across nation-states in total factor productivity seem to be related tenuously, or not at all, to *technology*.

‘ Robert Solow may not have done us a big favor when he convinced us to call shifts in the aggregate production function “technical change”; his doing so may not have helped economists to think clear thoughts over the past forty years.

II. Divergence

As best as we can determine from badly flawed data, the economic history of the past century and a quarter is a history not of “convergence” but of “divergence”: the different countries and peoples of the world have not drawn closer together in relative living standards, but have drifted further apart.

Figure 1 below shows the distribution of world real GDP per capita—by percentage of world population, not by nation-state—in 1993 and in 1870, as best as it can be estimated. 1993 estimates of real GDP per capita are purchasing-power-parity concept estimates, measured in the “international dollar” concept that pegs U.S. GDP per capita to its current-dollar value, but that attempts to use the relative price structure not of the advanced industrial economies but of the "world average" economy. They are taken from the 1995 *World Development Report*.

1870 estimates of real GDP per capita are my own extensions and modifications of those found in Angus Maddison's (1995) *Monitoring the World Economy;* by and large they are constructed by "backcasting" individual nation-specific estimates of real GDP per capita growth rates.

Figure 1: World Distributions of Real Income, 1870 and 1993 (PPP Concept)

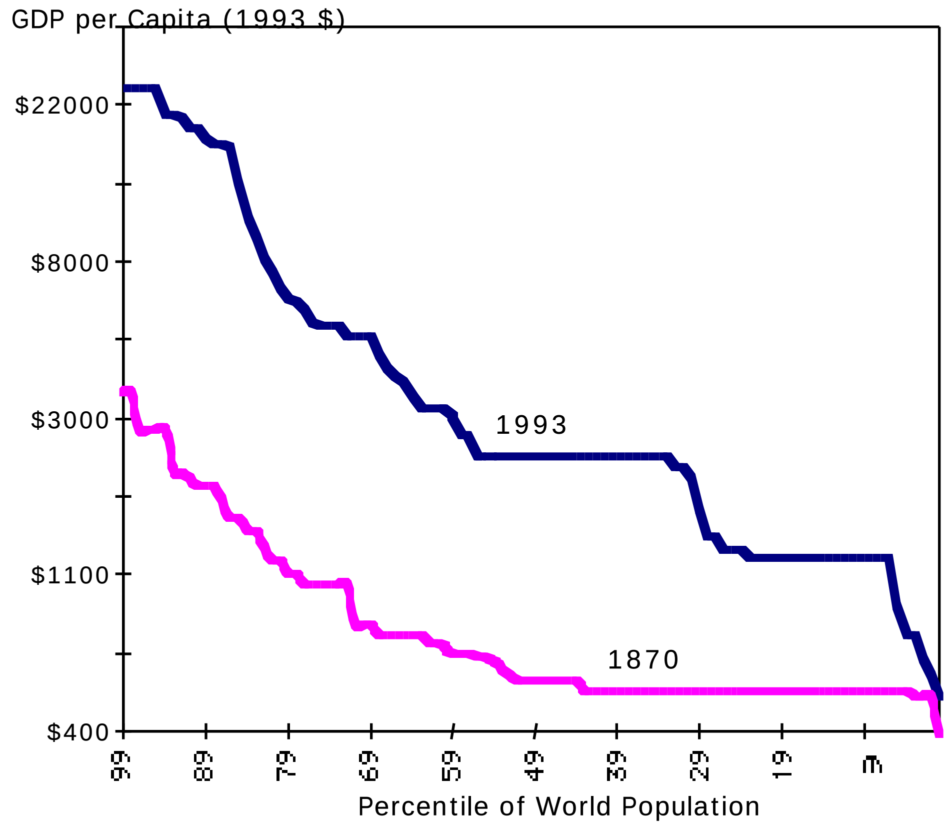

Thus there are a very large number of caveats attached to figure 1:

- Because estimates of 1870 GDP per capita are "backcast," errors in estimating 1993 GDP per capita are necessarily included in estimated 1870 GDP per capita as well.
- The individual nation-specific estimates of growth rates underlying the backcasting are of widely variable quality; they do not use the same methodology.
- Most of the nation-states of today's world did not exist in 1870. Estimates for 1870 cover roughly the same area then that the nation-state occupies now.
- Figure 1 suppresses all variability in productivity and real GDP per capita *inside* of nation-states: everyone in China is assumed to have the 1993 purchasing-power-parity concept real GDP per capita of $2,330.
- Estimates of even 1993 purchasing-power-parity concept real GDP per capita for developing countries are very uncertain. This especially applies to China which, as the World Bank team politely puts it in a footnote, has a GDP per capita estimate that is "subject to more than the usual margin of error."
- The entire enterprise of computing purchasing-power-parity concept real GDP per capita levels may be seriously biased; it may fail to incorporate appropriate allowances for quality differences between products produced in industrialized and developing economies. Certainly purchasing-power-parity concept estimates of relative living standards east and west of the

Iron Curtain made in the 1980s appear, in retrospect, to have wildly exaggerated levels of productivity and material wealth in the former Soviet Union's sphere of influence.[2]

- Estimates of 1870-1993 real GDP per capita growth are unlikely to adequately incorporate changes in quality and in the scope of products that are produced. The thought experiment that underlies constant-dollar cross-time comparisons implicitly involves taking the output produced at a particular date, moving it across time to the base year, and selling it in the base year at the base year's market prices. But suppose you gave me the $2,763 dollars— the estimate of U.S. GDP per capita in 1870—and told me "by the way, you can only spend this sum on products that existed and quality levels that were produced in 1870." Under these stringent restrictions on what I could purchase, I might well value that sum as worth much less than $2,763 of today's dollars.
- Figure 1—plotting approximate GDP per capita by percentile of the world's population—looks significantly different in some respects from figure 2, which plots GDP per capita in 1870 and 1993 by percentile of the world's number of nation-states. Nation-state based calculations show a nearly uniform distribution of log GDP per capita levels over the observed range, especially for 1993. Population-based calculations show a non-uniform distribution with a pronounced upper tail: the difference, of course,

[2]Current exchange rate-based calculations of relative productivity levels and living standards show differences an order of magnitude greater than do purchasing power parity-based

springs from the two very large population nation states of China and India, which are now and were in 1870 relatively poor.

Figure 2: 1870 and 1993 Distributions of PPP GDP per Capita, by Nation-State

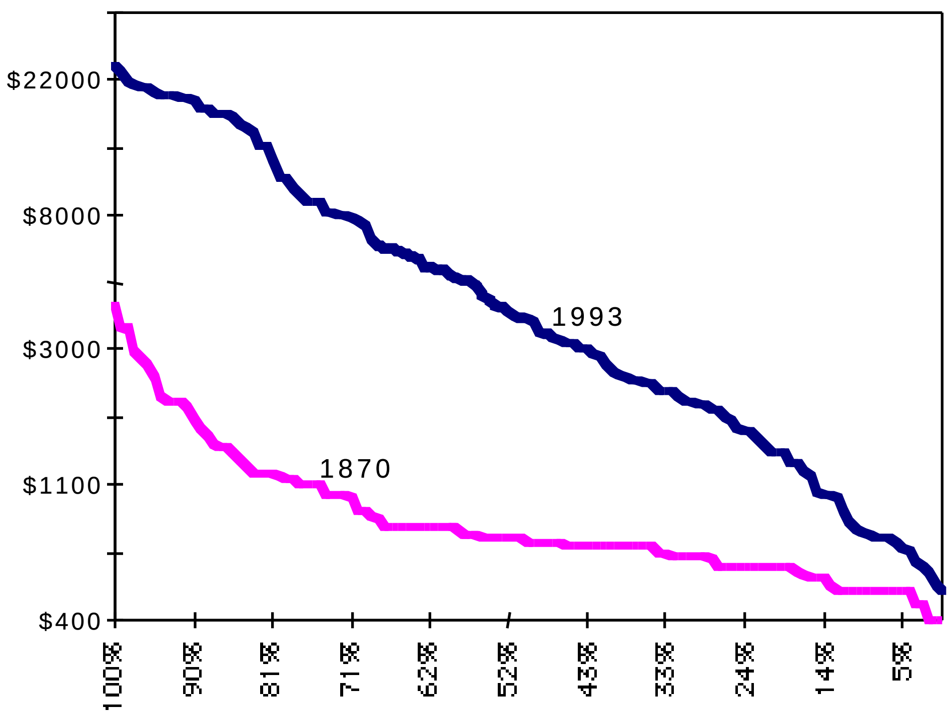

Nevertheless, figure 1 is the best we can do at present.

What are the principal lessons of figure 1? I believe that there are three:

The first is the extraordinary pace of real economic growth over the past century. The highest GDP per capita level attained in 1993 (for the United

---

calculations; it may be that in some senses the exchange rate-based calculations are more informative.

States) was some $24,470 1993-level international dollars; the highest GDP per capita level attained in 1870 (for Australia) was some $4,108 1993-level international dollars. Using this particular metric, the United States today is some six times as wealthy in a material-product real-income sense as was Australia in 1870 (and the United States today is some nine times as well-off as was the United States in 1870).[3]

I stress that this pace of growth is not only very large but also extraordinarily larger than in any previous century that we know of. If 1870-1993 growth were simply a continuation of pre-1870 growth trends, then in 1600 the richest economy in the world would have had a real GDP per capita level of some $110 a year—far too low to support human life.[4]

The twentieth century (extended back to 1870) has seen at least a sixfold multiplication of real GDP per capita at the leading edge of the world's economies; the previous century and a quarter had seen perhaps a doubling during the period of the classical industrial revolution (see Crafts, 1985; Mokyr, 1985). But before that? Perhaps the most prosperous economy of the mid-eighteenth century (probably the Netherlands) held a fifty-percent edge over the most prosperous economy of the mid-fifteenth century (probably the city-states of northern Italy). But perhaps not.

---

[3]This pace of real economic growth would be further magnified if the argument that measured growth in the GDP accounts fails to capture much of the growth in real income that takes the form of improvements in the quality and variety of commodities turned out to be correct. Such factors *might* lead standard estimates to understate "true" economic growth over the past century by a factor of two or three. See, for example, Nordhaus (1994). On the other hand, Simon Kuznets (1963) argued that the constant-dollar current-base-year real GDP calculations that he designed were the most appropriate ones: that we should use the yardstick of the present to assess the past.

And looking more than five hundred years into the past it is hard to see any significant advance in living standards or average productivity levels. Human populations appear to be in a near-Malthusian equilibrium, in which population growth quickly removes the margin for any significant increase in living standards (see Kremer, 1993; Livi-Bacci, 1992; Malthus, 1798). It is not clear that a French peasant of the seventeenth century was any better off than an Athenian peasant of the fourth century B.C.

The second important lesson of figure 1 is the extremely uneven pace of economic growth over the past century. Because the relatively poor economies of the world have not yet completed their demographic transitions to a régime of relatively low fertility, the poorest economies have been the fastest growing over the past century. International migration has not proceeded at a particularly fast pace. Thus the distribution of economic growth appears more uneven and less widely distributed in figure 1, which plots GDP per capita by percentile of the world's population, than in figure 2 which plots GDP per capita by nation-state.

But in both figures the line plotting the world's has rotated clockwise about the bottom right corner: the richest economies today have some six to nine times the GDP per capita of their counterparts in 1870; the economy containing the median today has perhaps four times the GDP per capita of its counterpart in 1870; the poorest economies are little advanced over their counterparts of 1870.

---

[4]A point made by Kuznets (1963), and expanded on in considerable depth by Pritchett (1994).

Figure 3a: 1870 and 1993 GDP per Capita, Top Third of 1870 Distribution

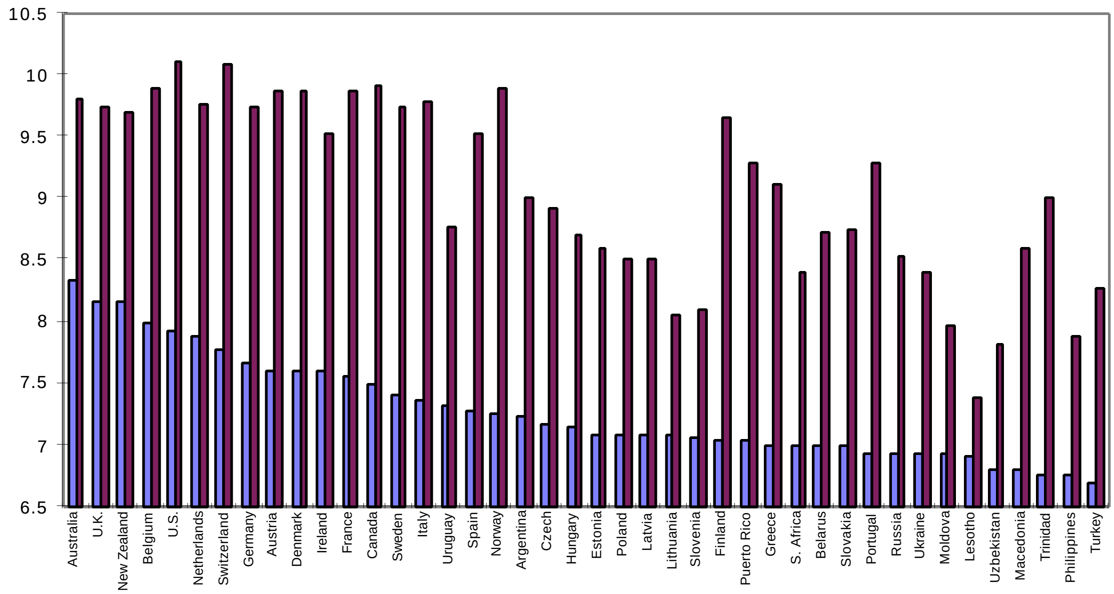

Figure 3b: 1870 and 1993 GDP per Capita, Middle Third of 1870 Distribution

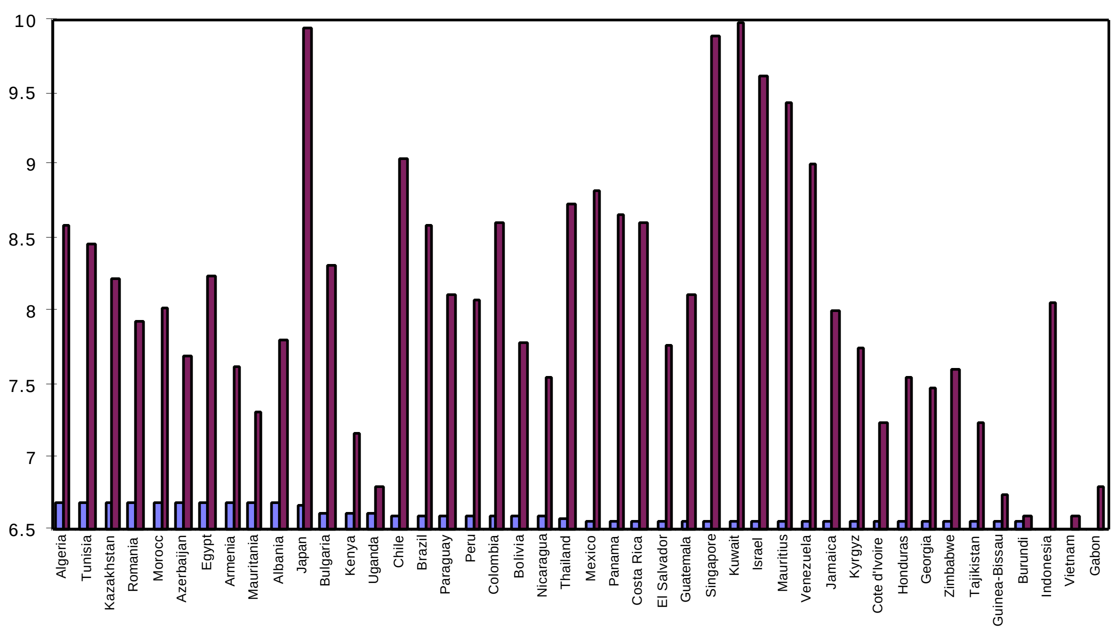

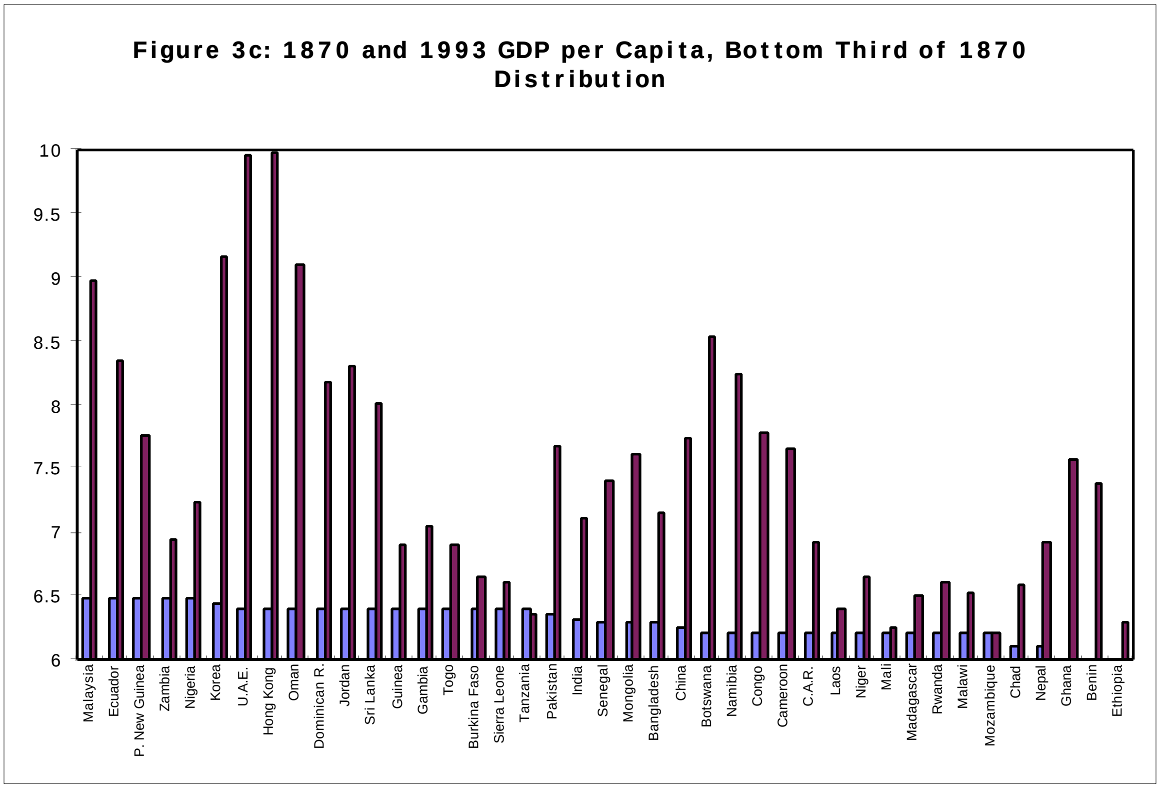

To put this lesson another way, the strong economic growth of the past century—the rise in the geometric average output per capita level in the world from some $760 to some $3150 1993 international dollars per year—has been accompanied by a substantial increase in variance as well. In 1870 the standard deviation of log GDP per capita across the world's population was some 0.53; today it is 1.00. The range from one standard deviation below to one standard deviation above the mean in log GDP per capita took up the interval from $450 to $1310 international dollars in 1870; the same interval runs from $1160 to $8510 international dollars today.

The third lesson is that by and large the economies that were rich in relative terms in 1870 are rich in relative terms today, and that the economies that were poor in relative terms in 1870 are poor in relative terms today.

Barro and Sala-i-Martin (1995) draw a distinction between what they call σ-divergence and β-divergence: they call "σ-divergence" the case where the variance of a distribution grows in spite of a tendency for any given element to revert toward the mean over time; they call "β-divergence" the case where the variance of the distribution would continue to widen even in the absence of all shocks—when there is no systematic regression toward the mean.

The world since 1870 has exhibited not only σ-divergence but also β-divergence: the world's distribution has a greater spread today because there has been a systematic tendency for the relatively rich to grow faster than the relatively poor, and not because shocks to individual nation-states' GDP per capita levels have dominated regression to the mean. Table 1 documents this by reporting simple regressions of nation-states' log GDP per capita levels in 1993 on the level of 1870. If two economies' log GDP per capita levels were separated by an amount X in 1870, they were separated by 1.542(X) in 1993.

The degree of β-divergence is slightly attenuated when continent dummies are added to the right hand side. The continent dummies have the standard pattern: strongly positive for North America, strongly negative for Africa. More interesting, perhaps, is that there is some evidence that GDP per capita levels have tended to converge over the past century and a quarter, if attention is confined to those economies that were in the richer half of the sample in 1870.[5]

Table 1
Simple Convergence Regressions, 1870-1993

	Log 1870 GDP per Capita	Continent Dummies		R^2
Full Sample	1.542			0.689
	(0.145)			
Full Sample	1.316			0.753
	(0.197)			
		0.501	N. America	
		(0.381)		
		0.174	S. America	
		(0.252)		
		0.208	Asia	
		(0.225)		
		-0.592	Africa	
		(0.226)		
Richer Half	0.620			0.533
	(0.126)			
Poorer Half	1.252			0.466
	(0.305)			

The fact that the world's distribution of income and productivity levels across nation states has been diverging goes oddly with a large number of studies (see Cogley and Spiegel, 1996; Mankiw, Romer, and Weil, 1992) that

[5] Williamson (1996) and Taylor and Williamson (1994) point to the factors—largely international migration, increasing trade, and thus converging factor and commodity prices—making for "convergence" among relatively well-off economies before World War I. Dowrick and Nguyen (1989) point to similar factors and document similar "convergence" within the club of relatively rich OECD economies after World War II. W. Arthur Lewis (1978) attempts to account for the failure of relatively poor economies to industrialize before and after World War I.

find evidence for "conditional convergence": gaps between an economy's aggregate income and productivity level and the level corresponding to the steady-state growth path predicted by its investment and population growth rates shrink over time by some two to three percent per year.

A naive interpretation of this pattern would suggest that at some time in the past nation-states' savings and population growth rates must have been more equal. If not, then how did the more concentrated distribution of output per capita levels in the past ever arise? And it would suggest that there was some massive shock that drove savings and population growth rates apart and that since then the world's distribution of relative incomes has diverged, as economies have traversed toward their steady-state growth paths. But there are no candidates for such a shock: the Industrial Revolution initially did not lower but raised population growth rates in the most heavily affected countries.

III. Endogenous Factor Accumulation

Conditional Convergence

Barro (1991) and Mankiw, Romer, and Weil (1992) were among the first to stress the existence of *conditional convergence* in the post-World War II cross-section of the economic growth rates of nation-states' economies. Mankiw (1995) interprets this as indicating that the straightforward Solow growth model is working better and better as time passes: it is becoming more and more the case that differences across nations in relative GDP per capita levels are reflections of the differences in steady-state capital intensity implied by their rates of factor accumulation and population growth.

Yet the appearance of conditional convergence—a coefficient of between -2 and -3 percent per year when the growth rate is regressed on the difference between an economy's initial GDP per capita level and the steady-state level implied by its investment and population growth rates—fits oddly with the fact, documented in the previous section, of unconditional divergence. How can economies be traversing toward their steady states *and* at the same time drawing further and further apart in relative GDP per capita levels?

A naive interpretation of this pattern would suggest that at some time in the past nation-states' savings and population growth rates must have been much more closely bunched together than they are today. This would mean that at that time in the past economies' steady-state and actual output per capita levels were bunched together more closely than they are today. And that some economic shock or series of shocks has since driven their respective savings and population growth rates apart—and thus that the world's distribution of relative incomes has diverged since, as the world's economies have traversed toward their—now distantly separated—steady-state growth paths.

But this naive interpretation has a central problem: what was this shock that drove savings and population growth rates apart? The principal candidate would be the Industrial Revolution. But the Industrial Revolution saw not a fall but a sharp rise in population growth rates in the most heavily affected economies (see Livi-Bacci, 1992). And today very little is left of Rostow's (1957) bold hypothesis that the key to the Industrial Revolution was a sharp rise in investment as a share of national product (see Crafts, 1985; Mokyr, 1985). The shifts in investment and population growth rates brought about by the Industrial Revolution do not go the right way.

Other candidates for a shock to drive economies' investment and population growth rates sharply enough away from one another to generate the observed divergence seen over the past century are simply absent. The overwhelming bulk of divergence in GDP per capita over the past century and a quarter has been due to the uneven spread of the Industrial Revolution, and to differences in relative national rates of total factor productivity growth.

But why, then, the finding of conditional convergence, and the strong positive association of GDP per capita levels with investment rates and the negative association with population growth rates?

Population Growth and the Demographic Transition

One reason is the endogeneity of population growth. Sometime between the fifteenth and the eighteenth centuries the human race passed through what we all hope was its last "Malthusian" episode, in which rising population and limited agricultural resources led to nutritional deficits, higher than average mortality, and population stagnation. Since then the pace of productivity improvement in agriculture has kept ahead of agricultural resource scarcity and population growth (that has carried the world's population from one to six billion so far). Nutrition has been relatively high by historical standards, natural fertility high as well, and natural mortality low.

In the past, the richest human populations appear to have also seen the fastest population growth. But starting perhaps in eighteenth century France a new pattern began to emerge, in which increases in GDP per capita led not to greater fertility and faster population growth but to lower fertility and slower population growth. The number of girls born per potential mother fell, and population growth rates slowed.

Figure 4 shows this pattern at work in the United States over the past two centuries: as GDP per capita has grown, the rate of natural increase of the U.S. population has fallen steadily. Once U.S. GDP per capita grew beyond the $2000 or so 1993 dollars level, fertility began to drop sharply enough to offset the declines in mortality that also accompanied better medicine and rising material prosperity. The rate of population growth, excluding net immigration, is now little over one percent per year—far below the 3.5 percent per year in natural population increase seen in the first half-century of the republic.

The pattern of rising material prosperity and falling natural population increase has had only one significant interruption in the United States in the past two centuries. The Great Depression of the 1930s saw a very sharp fall in childbearing, and a reduction in natural population growth in the 1930s to only 0.7 percent per year. In what Richard Easterlin (1982) sees as a delayed positive response to the Great Depression that balanced out the birth deficit of that decade, births rose to a level not seen since the nineteenth century in the "baby boom" of the 1950s.

Figure 4: U.S. Population Growth and GDP per Capita, 1790-Present

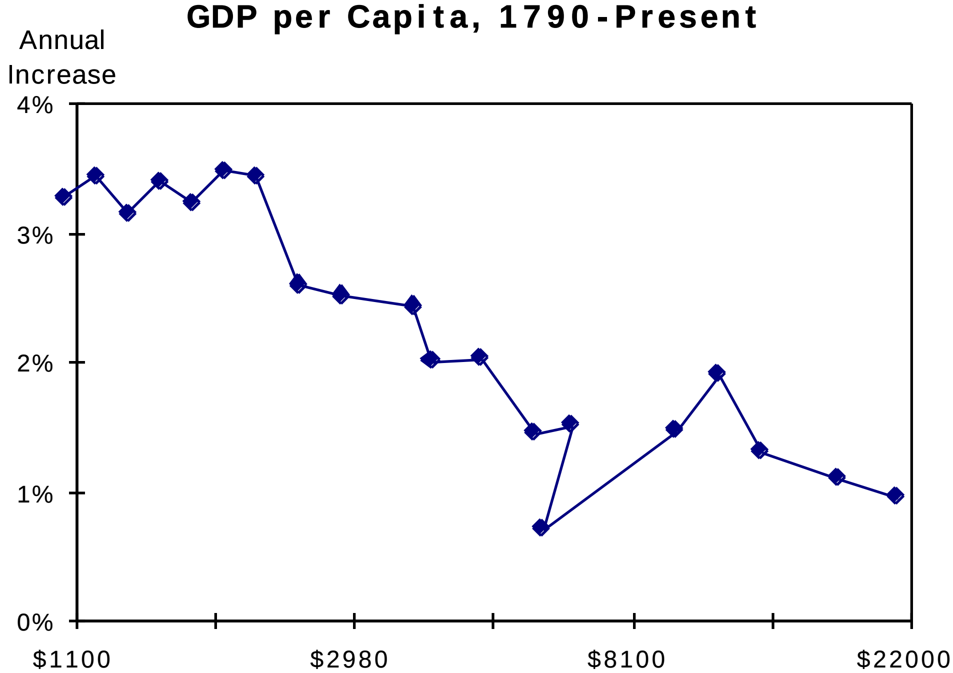

The pattern of increasing material wealth and slowing population growth seen in the United States is completely typical of the pattern that has so far been followed by all nations that have successfully industrialized. Each tripling of GDP per capita is associated with an approximately one percentage point per year fall in the rate of natural population increase.

To my knowledge, no one has ever argued that falling population growth has any sources other than in the increasing material prosperity of the United States, and in the changes in social and economic organization that have followed from the United States's growing material wealth. A richer country has more literate women, and literate women—worldwide—are very interested in effective birth control. In a poorer country the average level of education is low, and children can be put to work at a relatively early age, thus augmenting the production resources of the household. In a richer country the average level of education is high, and children are a major drain on household cash flow for nearly two decades.

Children in relatively poor, low-productivity economies are much more like an "investment" good than are children in rich, high-productivity economies: they are a way to augment the economic resources of the household in a time span of a decade or so. By contrast, children in relatively rich, high-productivity economies are more like a "consumption" good.

Thus we would expect—and we do see—a substantial correlation between high GDP per capita and low population growth arising not so much because low population growth leads to a higher steady-state capital-output ratio but because of the demographic transition: the changes in fertility that have so far been experienced in every single industrialized economy.

The Relative Price of Investment Goods

Begin with the large divergence between purchasing power parity and current exchange rate measures of relative GDP per capita levels. The spread between the highest and lowest GDP per capita levels today, using current exchange rate-based measures, is a factor of 400; the spread between the highest and lowest GDP per capita levels today using purchasing power parity-based measures is a factor of 50. If the purchasing power parity-based measures are correct, real exchange rates vary by a factor of eight between relatively rich and relatively poor economies. And the log GDP per capita level accounts for 80 percent of the cross-country variation in this measure of the real exchange rate, with each one percent rise in GDP per capita associated with an 0.34 percent rise in the real exchange rate.

Real exchange rates are such as to make the prices of traded manufactured goods roughly the same in the different nation-states of the world, putting to one side over- or undervaluations produced by macroeconomic conditions, tariffs and other trade barriers, and desired international investment flows. Thus the eight-fold difference in real exchange rates between relatively rich and relatively poor economies is a reflection of an approximately eight-fold difference in the price of easily-traded manufactured goods: relative to the average basket of goods and prices on which the "international dollar" measure is based, the real price of traded manufactures in relatively rich countries is only one-eighth the real price in relatively poor countries.

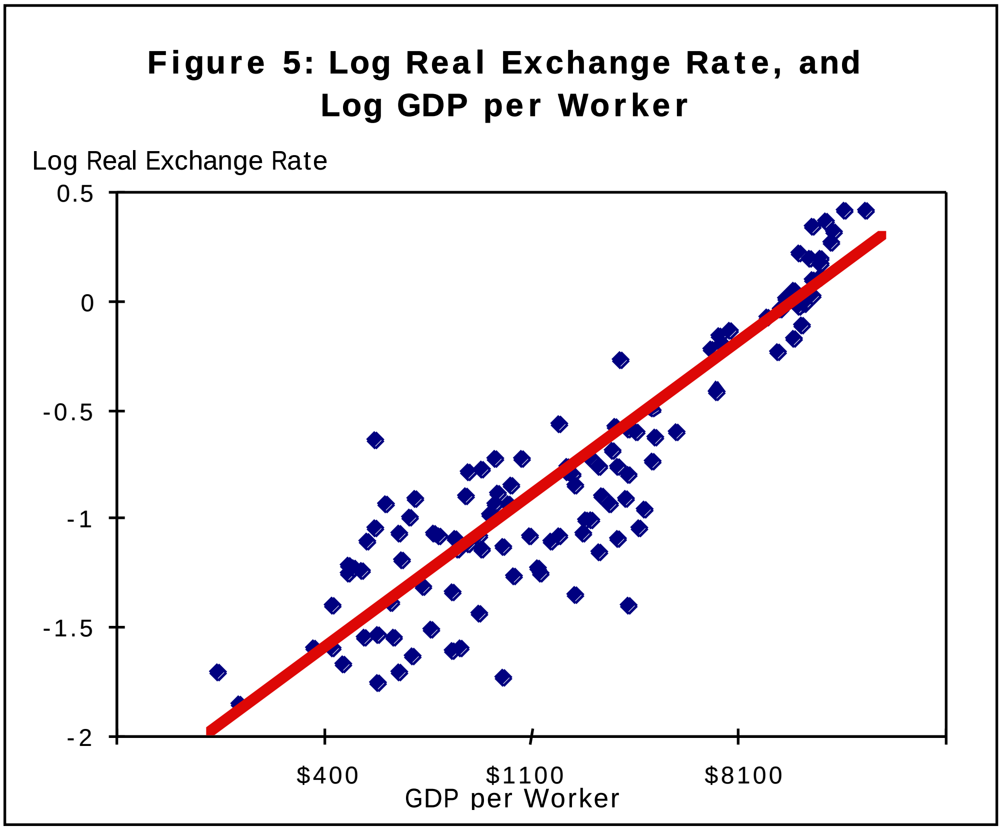

This should come as no surprise. The world's most industrialized and prosperous economies are the most industrialized and prosperous because they have attained very high levels of manufacturing productivity: their productivity advantage in unskilled service industries is much lower than in capital- and technology-intensive manufactured goods.

And a low relative price of technologically-sophisticated manufactured goods has important consequences for nation-states' relative investment rates. In the United States today machinery and equipment account for half of all investment spending; in developing economies—where machinery and equipment, especially imported machinery and equipment is much more expensive—it typically accounts for a much greater share of total investment spending (see Jones, 1994; De Long and Summers, 1991).

Consider the implications of a higher relative price of capital goods for a developing economy attempting to invest in a balanced mix of machinery and structures. There is no consistent trend in the relative price of structures across economies: rich economies can use bulldozers to dig foundations, but poor economies can use large numbers of low-paid unskilled workers to dig foundations. But the higher relative price of machinery capital in developing countries makes it more and more expensive to maintain a balanced mix: the poorer a country, the lower is the real investment share of GDP that corresponds to any given fixed nominal savings share of GDP.

Table 2
Consequences for National Investment of Relative Poverty, and a High Price of Capital Goods

RER GDP per Capita Level	Price of Machinery	Savings Share of GDP	Investment Share
$24,000	100	20.0%	20.0%
$6,000	160	20.0%	15.4%
$1,500	257	20.0%	11.2%
$375	411	20.0%	7.8%
$95	659	20.0%	5.3%

Table 2 shows the consequences—the gap between nominal savings and real investment shares of GDP—that follow from the high relative price of machinery and equipment in poor countries that wish to maintain a balanced mix of investment in structures and equipment. For a country at the level of the world's poorest today—with a real exchange rate-based GDP per capita level of some $95 a year—saving 20% of national product produces a real investment share (measured using the "international dollar" measure) of only some 5% of national product.

In actual fact poor economies do *not* maintain balanced mixes of structures and equipment capital: they cannot afford to do so, and so economize substantially on machinery and equipment. Thus here are two additional channels by which relative poverty is a cause slow growth: first, relative poverty is the source of a high real price of capital, a low rate of real investment corresponding to any given nominal savings effort, and a low steady-state capital-output ratio; second, to the extent that machinery and equipment are investments with social products that significantly exceed the profits earned by investors (see De Long and Summers, 1991), the price structures in relatively poor developing economies lead them to economize on exactly the wrong kinds of capital investment.

Implications

The standard Solow (1956) and Swan (1956) growth model, written in per worker terms and expressed in logs, contains the production function:

(1) $$\ln(y) = \alpha \ln(k) + \tau$$

where y is output per worker, k is capital per worker, α is the capital share in the production function, and τ is the log of total factor productivity. If the economy has a constant investment rate I, a constant population growth rate n, and has labor efficiency growth and depreciation rates g and d, then in steady state at any point in time output per worker will be given by:

$$\text{(2)} \qquad \ln(y) = \frac{\alpha}{1-\alpha}\left(\ln(I) - \ln(n+g+\delta)\right) + \frac{\tau}{1-\alpha}$$

Suppose, however, that we take account of the feedback from GDP per capita levels on population growth rates:

$$\text{(3)} \qquad \ln(n+g+\delta) = -\phi \ln(y) + \nu$$

where n is that portion of ln(n+g+δ) that is not accounted for by the combination of the dependence of population growth on output and the background rates of labor efficiency growth and depreciation. The pattern of the demographic revolution from the United States's historical experience suggests that the parameter φ is, over the relevant range, approximately equal to 0.2.

And suppose we take account of the feedback from GDP per capita levels to the real investment share:

(4) $\ln(I) = \ln(s) - \ln(p_k) = \ln(s) + \theta \ln(y) - \eta$

where s is the economy's nominal savings share, p_k is the real price of capital goods, η is the deviation of the price of capital goods from what would have been predicted given the level of real output, and θ—the elasticity of capital goods prices with respect to output—is roughly equal to 0.3 over the range relevant for developing economies.

Combining (2), (3), and (4) produces an expression for the steady-state level of output allowing for the endogeneity of population growth rates as a result of the demographic transition, and for the dependence of the relative price of investment on output per worker:

(5) $$\ln(y) = \frac{\alpha \ln(s) - \alpha\eta - \alpha\nu + \tau}{1 - \alpha - \alpha\theta - \alpha\phi}$$

Equation (5) allows us to calculate, for various possible values for the share α of produced capital goods in the production function and for the chosen values of φ and θ, the impact on the level of the steady-state growth path of a shift in the exogenous component of savings, capital goods prices, population growth, or total factor productivity. Because they enter symmetrically into equation (5) the effects of the first three are the same.

Table 3
Consequences for Steady State of Endogenous Population Growth and Capital Goods Prices

Capital Share α	Denominator of (5)	Effect of s, ν, h	Effect of Total Factor Pdty.
0.20	0.70	0.29	1.43
0.40	0.40	1.00	2.50
0.60	0.10	6.00	10.00
0.67	0.00	∞	∞

Table 3 reports that—with a produced factor inputs share in the production function of 0.4—a one percent increase in the savings rate (or a one percent fall in the exogenous component of capital goods prices) carries with it a one percent increase in the steady-state level of output. But a one percent increase in total factor productivity raises the steady-state level of output by fully 2.5 percent. Growth-accounting decompositions would, if applied to such an economy, attribute only one percent of the higher level of output to higher total factor productivity—less than one-sixth of the total effect. The growth accounting decomposition is not wrong, but incomplete: to the extent that the higher capital stock is a result of higher total factor productivity reducing the relative price of capital, and to the extent that higher total factor productivity pushes an economy further along its demographic transition to low population growth, exogenous shifts in total factor productivity have effects orders of magnitude greater than growth accounting procedures suggest, even without any powerful externalities in the production function.

As interesting, perhaps, is the case in which there *are* externalities to investment—whether in infrastructure, in research and development, in human capital, or in machinery and equipment—and in which the true capital share α in the production function is substantially greater than the 0.4 found in the usual specifications of the Solow model. The true capital share cannot get as high as 0.67 without triggering explosive paths for output per capita, in which very small boosts to total factor productivity set in motion patterns of population growth reduction and investment increase that converge to no steady state at all, but simply grow until the log-linear approximations in equations (3) and (4) break down.

It is difficult to look at the cross-country pattern of growth over the past century without thinking that the determinants of the steady-state growth paths toward which countries converge must be nearly singular. What difference between Canada and Argentina in 1870 would have led anyone to forecast their—now more than two and a half-fold—difference in GDP per capita? Or the twenty-fold gap between Taiwan and India? Recognizing the endogeneity of the demographic transition and of investment has the potential to help us understand why the economic history of the past century and a quarter has proceeded as it did, without requiring assumptions of external effects that seem—perhaps—implausibly large.

The endogeneity of the demographic transition, and of investment, also helps make sense of the odd combination of global divergence together with “conditional convergence.” To the extent that relatively low productivity today is a cause of an economy’s attraction to a low steady-state growth path, it is less necessary to look for shocks in the past that both (a) pushed economies away from their long-run growth paths, and (b) pushed economies’ GDP per capita levels together, if we want to account for the evolution of the world’s distribution of income.

IV. Conclusion

Caveat

But I do have one important caveat: I want to praise “technology,” as it appears in the aggregate economic growth models; but there is a sense in which I also want to bury it.

Robert Solow’s (1957) article is entitled “Technical Change and the Aggregate Production Function.” Certainly since 1957 economists have used “technical change” and “technology” as shorthand ways of referring to shifts in the aggregate production function.

Yet do we really want to do this? Much of the difference seen across nations in aggregate total factor productivity seems to have little to do with *technology*—in the sense of knowledge of the internal combustion engine, continuous-casting, or the freeze-drying process.

Consider Greg Clark's (1987) excellent study of productivity in the cotton textile industry *circa* 1910. Table 4 reports some of Clark's calculations, most strikingly the seven-fold difference in labor productivity found between mills in the United States and cotton mills in the region of China near Shanghai.