EXPLORING THE PROCESS OF ECONOMIC GROWTH

Arnold C. Harberger

University of California, Los Angeles

June 2004

This paper is intended as an introduction to the study of economic growth. It attempts to convey an understanding of the various constituent elements of the process of growth, and to show how they work. It also attempts to give readers an idea of the quantitative magnitudes of these various growth components -- how important they have been historically, and what sort of orders of magnitude we can reasonably expect (or hope) them to attain under various future scenarios.

My personal objective, as author, is to leave readers with a realistic, feet-on-the-ground vision of the growth process, with a proper regard for the achievements of the past as well as an appreciation of the reasons for its failures, and with a good sense of the order of magnitude of future benefits that can reasonably be expected to flow from different types of action and/or levels of effort.

This paper is addressed to a broad audience -- one that includes professional economists but is not limited to them. I thus try to keep the exposition at a reasonably nontechnical level, but not so nontechnical as to obscure important analytical issues. Readers are not just supposed to look at the engines of growth, but to gain a pretty clear picture of how they work. I also hope that they gain an appreciation of some of the issues involved in our measurement of a country's growth rate and its various components.

The Breakdown of Growth

For more than half a century, economists have been using a breakdown of the growth rate that distinguishes between:

a) a contribution explained by increments to the country's work force,

b) a contribution coming from increments to the country's stock of physical capital, and

c) a contribution that is variously labeled as coming from "technical advance", "increased productivity", or "real cost reduction"!

This breakdown is extremely useful, because with its aid we can think more clearly not only about the sources of growth but also about the orders of magnitude of the various components.

The Labor Contribution

Consider a simple example of a labor force broken down into three components -- unskilled, skilled and professional. Let their relative numbers, wages and total earnings be as follows:

Category	Number	Annual Earnings Per Worker	Total Earnings (Millions)
Unskilled	8 million	$1,000	$8,000
Skilled	1.5 million	2,000	3,000
Professional	0.5 million	4,000	2,000
Total or Average	10 million	$1,300	$13,000

Suppose now that over a decade the labor force grows by 10%. This could lead, at the end of the decade to the following distribution:

Unskilled	8.8 million

Skilled	1.65 million
Professional	0.55 million

In this case the analysis is easy, each category has gone up by 10%, so we would on this account consider that the labor contribution to growth was 1.3 billion.

But suppose the labor force increments were not in the same proportion. Say the professionals added 0.1 million, the skilled added 0.2 million and the unskilled added 0.7 million.

Category	Increment In Number	Initial Earnings Per Worker	Increment to GDP (million)
Unskilled	0.7 million	$1,000	$700
Skilled	0.2 million	2,000	400
Professional	0.1 million	4,000	400
Total or Average	1 million	$1,500	$1,500

Faced with this pattern of growth, the preferred position is to attribute $1300 million to the growth of the labor force and $200 million to the improvement in its quality

Actual Increment In Number	Increment If Proportional Growth	Difference In Numbers	Avg. Wage	Difference In Value
0.7 million	0.8 million	-0.1 million	$1000	$-100,000
0.2 million	0.15 million	+0.05 million	2000	+100,000
0.1 million	0.05 million	+0.05 million	4000	+200,000
TOTAL	1 million	0		$+200,000

Sometimes the increment of output due to the change in quality of the labor force is measured as part of the labor contribution to growth. When this is not the case it is counted as part of the third term -- growth due to improvement in productivity.

The most common way of measuring labor's contribution is to simply multiply the percentage increment in the labor force (here 10%) by the fraction of GDP represented by the earnings of labor. Thus, if initial GDP were 26 billion, labor's contribution would be measured as 1/2 (= labor's initial share of GDP) times 10% (% increase in labor force). This calculation yields an increment to GDP of 5% over a decade, or 1/2 percent per year over this period.

When this method is used, the contribution to the increment in average quality would appear as part of the third "productivity improvement" term in the breakdown of growth. Note that in this case all the information that one needs is the initial labor force (10 million), the increment in numbers (1 million) and the initial average wage ($1,300). For an explicit estimation of the contribution of labor, including the improvement in average quality, one either needs detailed information on the breakdown of the increase in labor by skill category, or must use other, more sophisticated methods. But readers should not be alarmed. Even when the easiest, least demanding methods are used, the contribution of improved labor quality is not "left out" -- it is simply shifted to the third category, where it is counted as a part of the improvement in productivity.

The Capital Contribution

The easiest way to visualize the capital contribution to growth is to think of individuals who put their savings into bonds carrying interest. If they save 12% of their income, and put it in bonds yielding 6%, this will generate an increment to income equal to .72 (= .12 × .06) of one percent. In the case of an economy, we want to know the fraction of that economy's income that was added to the capital stock, together with the relevant rate of return to be attributed to the incremental investment.

The data on a country's gross investment is typically readily available, but unfortunately, it is not sufficient for our purposes. We need also to get information on the amount by which the economy's existing capital stock depreciates during the period in question. This information is hard to come by, so one or another estimation procedure typically has to be employed. One such procedure is to build up a time series of the country's capital stock, accumulating past investments and applying assumed depreciation rates to the existing stock at the beginning of each period. Another procedure posits an average capital/output ratio for the country, and takes annual depreciation (also at an assumed depreciation rate) based on this ratio times each individual year's real GDP.

The second element in the capital contribution is the rate of return. Ideally, one would like to have, for this purpose, the true economic marginal productivity of the net increment of capital that is being added to the capital stock. The procedures typically used by economists approximate this by the average rate of return earned on the existing capital stock.[1]

To measure the rate of return to capital, one simply adds up all of the income accruing to capital in the economy during the time period in question. Some of the components of this income -- particularly those generated by corporations, are relatively easy to measure. Others,

---

[1]Most studies define the capital contribution symmetrically with that of labor. That is, they measure this contribution as the share of capital in GDP times the percentage increment of the capital stock. Writing the share of capital as $[(\rho+\delta)K/Y]$ and the percentage increment as $[\Delta K/K]$, it is easily seen that their product simplifies to $(\rho+\delta)(\Delta K/Y)$, which is the expression we have used. This procedure makes it clear that the traditional procedure assigns as the marginal product of capital the average gross-of-depreciation rate of return $(\rho+\delta)$ observed in the economy in question.

while not measured directly, are regularly estimated in the official national income statistics. The best example here is the rental income attributed to dwelling units occupied by their owners. Here, quite obviously, no actual rent is paid, nor does a monthly or annual rental rate exist. So the national income statisticians impute a rental rate, usually based on an assumed (e.g., 6% per annum) rate of return to capital. Finally, we have items which the national income accountants simply do not estimate. These fit broadly under the umbrella "income of unincorporated enterprises" and basically comprise a whole range of incomes which nobody has taken the trouble to separate into a part attributable to labor and a part attributable to capital. Family-owned retail stores, farms and other enterprises are examples. Professional activities with significant capital investments are another. For modern industrial countries, the portion of GDP that falls into this "limbo" category is relatively small, so no great error is made by making an arbitrary split (e.g., half due to labor, half to capital) on limited evidence and sensible judgments. But for many developing countries the problem is much bigger, owing to the proliferation of such enterprises (particularly in the informal economy) combined with a paucity of hard statistical data, even on the more formal parts of the economy.

Building a Coherent Picture

One way to try to surmount the problems caused by insufficient data is to build a picture of an economy in such a fashion that one has a number of points at which one can check the internal consistency and also the general plausibility of the picture that emerges. To see how one can quickly run into implausible results, consider that somebody suggests that in country X the ratio of capital stock to GDP is around 5, and that the net rate of return is around 20% per annum.

This is implausible on its face, for it would mean the capital's net return would take up the whole GDP. Nothing would be left for labor, or even for the depreciation of the capital stock itself. Drop the rate of return to 15% and one is still in big trouble, because with a plausible rate of a depreciation (say 5%), capital would again be eating up 100% of GDP [(.15 + .05) × 5].

So if one is going to consider that a net rate of return of 15% is plausible, one must link it to a much more modest capital-output ratio. Thus if the capital-output ratio is 2.5 and the depreciation rate is 5%, then capital would get half of the country's GDP, and labor the other half. If one raises the capital-output ratio to, say, 3, it behooves one to think of a lower rate of return than 15%. If one links a K/Y ratio of 3.0 with a net rate of return of 10%, together with the depreciation rate of 5%, one gets a share of capital equal to 45%, with labor getting 55%.

Our plausibility check is not yet complete, however. For we should also explore the shares of capital and labor in the net product generated in the economy. This is obtained by deduction depreciation both from the aggregate GDP and from capital's share of it. So with a 15% rate of return and a capital-output ratio of 2.5 we would obtain a share of capital in net domestic product of 37.5/87.5, or about 43%. Taking a 10% rate of return and a capital/output ratio of 3 gives us a capital share in net domestic product of 30/85 or about 35%.

Alternative Growth Scenarios

Following the reasoning of the previous section, we constructed three alternative scenarios on which to base the breakdown of growth rates for a number of countries. Using these scenarios we will be able to use each country's actual data, year by year, on the ratio of investment to GDP, and on the growth rates of GDP and of the labor force. These data are available for a whole range of countries for which the data required by other methods are lacking.

The last row of Table 1 is solely intended as a plausibility check. The investment that is measured in the basic national accounts data typically includes (as it should from an economic point of view) investment in roads, ports, public buildings, and the machinery and equipment connected thereto. It also includes investment in housing, whether it be public housing, housing to be occupied by its owners, or housing that will actually be rented on the commercial rental market. Yet when the returns to capital are presented in the same account, they will not incorporate an implicit rate of return on roads or public buildings, for example. And they will typically contain only a relatively modest rate of return for the housing sector. The rate of return actually collected on rental housing is typically lower than that on most business capital. That on owner-occupied housing is (as previously noted) typically imputed at a modest rate of return, and that on public housing typically appears at subsidized (often heavily subsidized) rental rates.

In order to get an idea of what rates of return on business capital are implied by our various scenarios, I have assumed that one quarter of the capital stock yields a zero return (as recorded in the national income accounts) that another quarter represents housing of all three types and yields an average real return (as counted) of 6% per annum. The figures in row g of Table 1 are then found so as to produce the assumed average overall return. Thus if business capital yields 27% and accounts for half the total, we have for column (1), $(.5 \times .27) + (.25 \times .06) + (.25 \times .00) = .15$, giving us the 15% overall rate of return assumed in the scenario. For scenario (2) we have $(.5 \times .17) + (.25 \times .06) + (.25 \times .00) = .10$. Finally the scenario (3) we have $(.5 \times .12) + (.25 \times .06) + (.25 \times .00) = 0.75$.

Plausibility checks and Sensitivity Analysis

As can be seen from Table 1, the three scenarios cover a quite wide gamut of possibilities. An educated guess would be that scenario 1 looks reasonably good for a high-growth developing country, that scenario 3 looks reasonably good for an advanced industrial country and that scenario 2 seems quite plausible as a central tendency for most of the rest.

Our main purpose, however, is not to assign particular scenarios to particular countries, but to show that important conclusions can be drawn, no matter which of these quite distinct but internally consistent scenarios we apply.

Why should we expect that the results will be robust, as we move from scenario to scenario? The main reason is that we are using these scenarios to explain the same observed growth, on the basis of the same measured gross investment and of the same measured increase in the labor force. Thus if we assign a higher rate of return to capital, this will add to capital's contribution but at the same time reduce the amount of growth that is explained by productivity improvement.

In most of what we do with the scenarios we will work with the middle one. And, roughly speaking, one can estimate what would happen quite easily. To move from scenario 2 to scenario 1, add something up to about three quarters of a percentage point per year to the capital contribution yielding an R1 that is less than R2 by a similar amount. To move from scenario 2 to scenario 3, subtract up to perhaps one fourth of a percentage point for the capital contribution, yielding an R3 that is greater than R2 by a similar amount.

The case of France (1976-2001) is typical. If we use scenario 1 and assign a 15% real net rate of return (20% gross), we get a capital contribution to growth equal to 1.9% per year. If we go to the other extreme (scenario 3) and assign only a 7 1/2% net real rate of return, the capital contribution falls to 0.7% per year. In the first case the residual to be explained by productivity improvement is 0.4% per year; in the latter case it is 1.5%. The main result is shifting the explanation of growth from one cell to the other, and it is rarely more than about one percentage point per annum that is thus shifted. Obviously, scenario 2 lies in the middle, yielding a capital contribution of 1.0 percent per year together with a contribution of 1.2% per year from improved productivity.

Looking At The Data

Later we will have the opportunity to describe our data in more detail. For the moment let me simply list the country groupings and periods covered.

OECD: Australia (69-00), Belgium (70-00), Canada (66-01), Denmark (66-01), Greece (60-01), Germany (72-01), France (70-01), Finland (60-01) Japan (61-01), Ireland (71-00), Netherlands (71-01), New Zealand (60-00), Norway (60-00), Portugal (71-01), Spain (71-01), Sweden (65-01), Switzerland (65-01), Turkey (68-01), UK (60-01), U.S. (60-00).

Asian Tigers: China (65-01), Hong-Kong, China (65-01), Indonesia (79-01), Korea (61-01), Malaysia (60-00), Singapore (65-00), Thailand (60-01).

Other Asia: Bangladesh (60-01), India (60-01), Iran (65-01), Israel (61-01), Pakistan (60-01), Philippines (60-01).

Africa: Cameroon (75-01), Egypt (60-01), Ghana (67-01), Kenya (64-01), Morocco (60-01), Nigeria (75-01), South Africa (60-01).

Latin America/Caribbean: Argentina (60-01), Brazil (70-01), Chile (60-01), Colombia (60-01), Costa Rica (60-01), Ecuador (65-01), Honduras (60-01), Guatemala (60-01), E Salvador (65-01), Jamaica (65-01), Mexico (60-00), Nicaragua (60-98), Panama (60-01), Paraguay (60-01), Peru (60-01), Uruguay (60-01), Venezuela (60-01).

One can see in the above listing that we have an ample representation of important countries in each of the listed categories. The missing category is the transition countries, which were intentionally left out, in part for lack of data, in part because our methodology is difficult to apply in these countries.

The distribution of R's (rates of productivity [TFP] increase under the three scenarios) are shown in Figures 1-3. Figure 4 shows the results under a composite scenario in which instances of high growth are deemed to imply high rates of return to capital (scenario 1) and instances of low growth are deemed to imply a low rate of return to capital (scenario 3).

Visual inspection confirms a great similarity among all four distributions of R, the most notable difference being that the distribution of R* -- the implied rate of productivity increase under the composite scenario -- has, as expected, somewhat less variance than the others. The means and medians are all between 1 and 2 percent per year, and the fraction of negative observations (productivity decline) is always between a fifth and a third.

Table 2 summarizes this information. There are indeed significant differences among the means, medians, and fractions of negative observations, as one moves from one definition to another. But the differences are not great, and particularly not so when it is realized that where R is greater, the capital contribution to growth that we measure or estimate is correspondingly less.

<u>Differences Among Regions</u>

To construct Table 3, we took data that were averaged for each country over the entire period that we had data for it. Then, regional medians were taken for all the relevant variables.

As is well known, medians and other fractiles do not necessarily "add up". In this case we show this via two different figures for R2. The first of these is obtained by taking the median of the measured full-period average R2 for the different countries in the region. This is called the "measured R2". The second figure, called the "implied R2" is obtained by taking "Median Growth Rate minus Median Labor Contribution minus Median Capital Contribution" for each region. As one can see, the differences between measured and implied rates of productivity improvement are relatively modest, reaching as high as 0.5 percent per year in only 2 cases. Note also that in one of these cases -- that of the Asian tigers, the difference between measured and implied productivity improvements eliminated entirely, when we use the calculations based on scenario #1 (high rate of return to capital). There are many reasons for thinking this to be the most appropriate scenario for such high-growth episodes.

What conclusions can we draw from Table 3? Labor contributions are sizeable for all regions except the OECD. The median capital contribution exceeds 1% per year only for the Asian Tigers. The median contribution of productivity improvement is 2% or less, except for the two Asian areas where it is higher. But interestingly, if we are prepared to accept that scenario #1 is factually more plausible than scenario #2 for the Asian Tigers, the median productivity improvement in no case gets as high as 3% per year. I believe on the basis of this evidence, that rates of real cost reduction in the range of 2% per annum have to be considered to represent a quite reasonable performance. Certainly they cannot plausibly be considered as failures.

<u>The Phenomenon of Negative Productivity Changes</u>

One of the notable characteristics of most exercises in the breakdown of economic growth into its components is the substantial fraction of the time in which productivity changes in the wrong direction. Obviously, nobody is aiming for this result but it happens nonetheless. It can come from a considerable number of causes -- a world recession, a serious fall in the price of a principle export good (coffee, copper, oil, etc.), a crop failure or natural disaster at home, a balance of payments or banking crisis, or a general shift from good to bad economic policies at home. Here I want to emphasize that some, but by no means all such incidents are the consequence of bad policies. And this means that we should look further into each case before "convicting" a country's policymakers or the international entities that influence and advise them.

It is somewhat surprising that economists have paid so little attention to the phenomenon of declining productivity (TFP) or its equivalent, increasing real costs. This might be more easily understandable if the phenomenon were rare, but in actual fact it appears in nearly every data set in which economic growth is presented in a disaggregated fashion.

In the present case, our basic data are in the form of annual changes of a country's aggregate GDP. Yet nonetheless something like a quarter of all our observations show negative productivity change. Some of this undoubtedly reflects cyclical changes in the economy. It has long been noted by both labor economists and business cycle analysts that productivity tends to decline in recessions. Some of this can be ascribed to the phenomena of "labor hoarding": employers are reluctant to lay off experienced workers during recessions, because they fear that those workers may not be available when business recovers. Thus, employers reluctantly but willingly sustain some potentially avoidable costs, feeling they would lose more in the long run by laying off experienced workers in recession times.

But instances of declining productivity are far more frequent than can be explained by recessions -- for it is a rare country in which a quarter of all years are recession years. Hence we must seek other explanations for a goodly share of these observations.

We have been interested in this problem for several years and have explored it using a number of different data sets. Empirically, the phenomenon of declining productivity is found even more frequently as one deals with more and more disaggregated data. Most often, declining productivity is associated with a) declining output and b) a reduced rate of profit, which very often means last year's profits turning into this year's losses. These, in turn, are typically the consequence of declining demand for a firm's (or industry's) product. The declining demand sometimes derives from a recession in the economy, but it can also result from competition. A firm's competitors may have encountered innovations that lower costs or improve quality, thus drawing demand away from this firm's products. Foreign competitors may do the same, offering equal or better products at lower costs. An appreciation of a country's real exchange rate may reduce the cost of imports of all kinds, causing widespread shifts of demand away from domestic producers.

The above list should leave readers at ease with the idea that real cost increases can be prompted by many causes, of which cyclical recessions cover only a part. The key "umbrella" that covers most of the non-recessionary cases is the notion of competition, stemming from innovations at home and abroad, and from real exchange rate adjustments that reflect a changing pattern of a country's comparative advantage. One must realize that these forces are on the whole very positive from the point of view of the economy as a whole. The essence of progress in a dynamic market economy is what Schumpeter called "creative destruction" -- the process by which new products and new methods displace older ones, leading to a constant churning of economic activities, with only the fittest surviving. This apparently cruel story is what leads to constantly rising standards of living and real wages in a successful, growing economy. It is hard to imagine an extended process of successful economic growth that does not contain an important element of creative destruction. This, then, is our best explanation of that other part of observed real cost increases. Those that do not reflect recessionary trends in the economy are probably suffering the effects of increased competition and "creative destruction".

In spite of the fact that these explanations help us understanding most episodes of declining productivity, it remains true that this important part of the process of economic growth has received very little study. More work is needed before we have a complete picture of it. And until we understand it well enough to prove otherwise, there always remains the possibility and the hope that greater understanding will enable us to take actions that can help promote greater economic growth by shrinking the size of this important negative component.

The Importance of Real Cost Reductions

In this section I will show how important is the third term -- productivity improvement -- in accounting for differences in growth rates. In Figures 5 through 10, GDP growth rates are measured on the horizontal axes, and their middle measure of productivity increase -- R2 -- appears in the vertical axis. The correlations shown in these diagrams all speak for themselves. In every case, more than 90% of the variance is explained. Moreover, the slope of the regression line is such that in each case more than three quarters of each percentage point difference in GDP growth rates is accounted for by the difference (along the regression line) in the productivity change component of growth.

This result confirms the findings of many previous studies. The standard finding is that yes, productivity change is always an important component of each period's measured growth rate. The productivity-change factor is even more important when we try to explain the differences in one country's growth rate between high-growth and low-growth periods, or the differences among countries between high-growth and low-growth episodes. Figures 5 through 10 do precisely this. Each observation in these graphs represents an episode of growth. And an episode was defined as a period in which a country's growth proceeded at a more-or-less steady rate. An upward or downward shift in the trend of growth would signify the onset of a new episode. All cases of significant declines in GDP are registered as separate episodes.

Figure 5 shows the scatter diagram for all countries and all episodes. Figures 6 through 10 do the same, but region by region. It is easy to see that the same story is being told, over and over again, as one passes from region to region.

Figure 11 modifies the procedure in two respects. first, the observations are year-to-year changes rather than average annual rates of change for a particular episode. And second, productivity change is measured by the composite index R* rather than by R2. Recall that R* is based on a net rate of return of 15% for high-growth periods, 10% for middle-growth periods, and 7.5% for low-growth periods. It is comforting that even with both these alterations, the basic message remains the same. Moreover, examination of all the figures reveals that the points representing positive movements in both GDP and productivity change lie essentially on and around the regression lines as shown. This shows that our inferences could equally well have been drawn from this evidence, even if we excluded the observations derived from cyclical recessions, economic crises and other sources of declining GDP. That is, our conclusion as to the great importance of the rate of productivity improvement as an explainer of differences on observed GDP growth rates -- that conclusion remains firm even if we base our evidence only on instances of “normal” positive GDP growth rates.

Introducing “The Growth Syndrome”

Having established the importance of productivity changes in influencing GDP growth rates -- particularly in accounting for the difference between high-growth and low-growth episodes -- I now want to try to convince you that productivity changes are even more important than that. The proposition is that this year’s productivity growth not only contributes to this year’s growth but also to that of future periods. It does so by causing investment to be higher for a period of years after the initial real cost reduction, and by causing the rate of return to capital to be higher, also for a period of years after the initial stimulus.

To pursue this topic we shall turn to some other data sets, to which a more elaborate methodology has been applied. Recall that in the exercises presented up to now, the capital contribution to growth was estimated by incorporating the country’s actual gross investment, together with an assumed real rate of return. Period-by-period depreciation was estimated using “scenarios” each of which was built to be consistent with “its” built-in rate of return.

In the data sets to which we now turn, the real rates of return to the existing capital stock are estimated period-by-period, using actual measured returns and a capital stock built up on the basis of observed data on real investment. (These methods were not used up to now, because we wanted to get as wide a coverage as was reasonably possible.)

With data sets that separately specify investment and the rate of return, we can directly inquire whether these elements are in any sense “predictable” on the basis of earlier data. In particular, we seek to predict both the investment rate and the real rate of return of future years on the basis of the current year's rate of real cost reduction (= productivity increase). When predicting the change in the rate of return from the "present" to the future, the "present" is taken as the midpoint of the year for which the productivity increase is measured. Thus our base for the rate of return would be $(\rho_t+\rho_{t-1})/2$, and our base for the rate of investment (NI/Y) is $[NI_t/Y_t) + (NI_{t-1}/Y_{t-1})]/2$.

Table 4 presents the results of exercises using the productivity increase between t-1 and t, in order to predict the change in the real rate of return and in the rate of net investment for a number of periods into the future. In all cases what is predicted is the change from the "base" value to the "future' value of each variable.

The data of Table 4 refer to 23 countries for which the U.N. National Accounts present data on "compensation to workers" (Necessary to get labor's share in GDP). Where those accounts also give "income of unincorporated enterprises", half of this rubric was assigned as compensation to labor, the other half as return to capital.[2]

To interpret the results of Table 4, one should think of an exercise in which a coin is flipped 23 times. How likely is it that one would get 20, 21, 22, or 23 heads in such an exercise -- or even 15 or 17? It turns out that one would get 20 or more less than 1% of the time, and 15 or more less than 5% of the time. In short, these results are highly significant in a statistical sense.

[2]Otherwise the following formula was used: total return to labor = "compensation to workers" plus "0.3 (GDP minus indirect taxes minus compensation to workers)". This gives labor 30% of the net income not already counted under "compensation to workers."

They reveal a significant predictive value for productivity changes. Before elaborating further, I turn now to present more evidence.

The next evidence comes from a huge data set -- the National Bureau of Economic Research's Survey of 458 Manufacturing Industries. Our survey covers the years 1958-92. Again we are predicting the changes in the rate of return (ρ) and the rate of net investment (ni) on the basis of the measured change in productivity taking place at a specific time (t). Again, the change we are predicting is the change between the "base" and the specified future years. The "base" is $(\rho_{t-1}+\rho_t)/2$ for the rate of return and $(ni_{t-1}+ni_t)/2$ for the rate of investment.

For each question we asked, we ran 458 regressions -- one for each of the 458 manufacturing industries covered in the NBER survey. The results of these regressions were very similar to those reported in Table 4, with the overwhelming majority of correlations being positive as we use this year's productivity increase to predict the future change in the rate or return or the future change in the rate of investment.

But the most interesting evidence from this data set concerned our investigation of causality. The question here is whether the connections between productivity and the rate of return and between productivity and the rate of investment might not be two-way streets -- in the sense that yes, productivity changes might predict changes in the rate of return, but might not changes in the rate of return also predict changes in productivity?

To answer this question we used huge panel regressions in which the data for all 458 industries were pooled, making for around 15,000 total observations of annual changes. A separate regression intercept was allowed for each industry to allow for so-called "fixed effects". And since in our individual industry studies we had found predictive power for up to four years, we allowed our explanatory variables to appear with four separate lags. Thus, to predict the rate of return at time t, we used the change in productivity at times t-1, t-2, t-3, t-4. and when we turned the question around we used the rate of return at times t-1, t-2, t-3 and t-4 to predict the change in productivity at time t. A similar procedure was followed for the investment rate, first using four lags of the investment rate to predict each year's change in productivity, and then using four lags of the change in productivity to predict each year's investment rate. The results of this exercise are summarized in Table 5.

The message from Table 5 is clear. Changes in productivity work much better as predictors of the rate of return than the rate of return works as a predictor of the change in productivity. The same is true for the rate of net investment. That is, changes in productivity predict the rate of investment much better than the rate of investment predicts changes in productivity.

Finally, we turn to yet another data set -- the United Nations Industrial Statistics, broken down into 3-digit manufacturing sectors. Here the calculations were made for four industrialized countries, four rapid growers from East Asia and four Latin American countries. Table 6 shows the percentage of correlations that were positive, as one used the current rate of productivity improvement as a predictor of the change in the rate of return between the "base" and each given future period. In the entire table, there are only four cases in which less than half of the correlations were positive, and three of these were predictions that extended to period t+4.

Figure 12 shows histograms representing the frequency distribution of correlation coefficients, using the change in productivity at time t to predict the change in rate of return from the "base" first to period t+1, then to periods t+2, t+3, and t+4.

Figure 13 does the same, predicting the percentage change in investment for four future periods. In all cases there is a large preponderance of positive correlations, confirming the results of our earlier exercises.

Conclusions on the "Growth Syndrome"

I believe the evidence is overwhelming that todays productivity improvement sets in motion a chain of consequences that carries for several future periods. The most significant of these consequences have to do with the rate of return to capital and the rate of investment. Consider the following scenario: In a given firm or industry a way is found at time t to reduce real costs. This adds immediately to the profitability of the operation. But the increase in profitability also carries forward to future years. The extra profits are not likely to last forever, because competition will likely enter to erode them away, in the process transferring the benefits of the innovation from the innovating firm or industry to the consumers or other demanders of the product in question. But the extra profitability does tend to last a significant period of time, carrying forward for at least four years in all of the different data sets that we have examined.

Innovations and other improvements in productivity do not necessarily call for more investment. Some may have to do with management style or advertising. Others may simply entail abandoning unprofitable operations. So it is not at all inevitable that productivity improvements should predict increases in investment, but in fact they seem to do so, once again in all our data sets. So we conclude that not necessarily but quite regularly, productivity improvements lead to increases in the rate of investment. This is true at the aggregate level (for countries), at the 3-digit industry level (from the U.N. Industrial Statistics covering 3-digit industries in 12 countries) and for industries very narrowly classified (from the NBER survey of 458 manufacturing industries in the United States).

The impacts of productivity change on future investment and on the future real rate of return signify future growth dividends from today's productivity improvement. Thus, when economists have assigned a very important role to productivity improvements as key elements in the growth process, they have typically understated that role by not counting the impact of today's productivity change on the future "capital contribution" to the growth rate.

The big message of the "growth syndrome", then, is that real cost reductions (= productivity improvements) are even more important than we previously thought. They not only have a direct effect on this year's growth, but also an important indirect effect, via higher future investment and higher future rates of return.

A Further Important Lesson: Real Cost Reduction Occurs Within Firms

From the exercises reported in the previous sections, readers should quickly gather that productivity improvement is not something that we only measure at the level of the national economy. No, in fact it can be broken down into the finest of subdivisions. The NBER survey of 458 manufacturing industries is just illustrative. Each of these can be broken down into still narrower categories until finally we come down to the productive enterprise itself. This is really the basic unit for the study of productivity, and also for the study of economic growth.

The components of growth that we have mentioned all work at the level of the individual firm or other productive operation. Increments to the work force, improvements in the quality of labor, increments to the capital stock, increases in the real economic rate of return to capital, and finally increases in total factor productivity (real cost reduction) itself -- all these happen at the level of the firm.

What lesson or lessons can we extract from this observation? The most important, in my opinion, is that it helps us to understand the infinite diversity of the growth process. Growth happens in one firm because of computerization, in another because a new insecticide has been found, in yet another because a new variety of hybrid corn has been invented, and yet another because internet marketing has lowered its costs.

Basically, the GDP of a national economy is simply the sum total of the value added of all the productive entities that are located there. The increment to the work force of the nation is the sum of the increments in the constituent entities. If the work force improves in quality, those quality improvements are distributed among the entities where the work force is employed. The capital investment that takes place is also distributed across the economy, including individual households and government agencies. The returns to capital that we measure are those that are counted in the GDP. These too take place at the level of the productive enterprise where they are generated.

Economic Growth as the Result of Hard Work

The vision of economic growth taking place in thousands, even millions of productive entities in an economy is very useful, because it helps us to realize what a huge amount of effort goes into the process. Vast numbers of business people -- owners, managers, engineers, accountants, supervisors, foremen, skilled and unskilled workers -- all contribute their energies to the process. Finding new investment opportunities, meeting competitive challenges, making current and future operations more efficient -- all these are part of the daily lives of those engaged in production throughout the economy.

It is hard to generalize about so complex a process, but I would like to underline how much of it flows directly from the economic motivation of individuals. In general, workers look for the jobs that provide the best wages for them -- these are the jobs to which their talents are the best suited. Those who embark on further education do so, at least in part, in recognition of the increase in their earning power that will ensue. Business people look for investments with high real returns. Likewise, they are constantly on the lookout for additional ways to reduce real costs. There are, as I have said elsewhere, many more than 1001 ways in which real costs are in fact reduced. It could be foolhardy to try to list them all. But the underlying motivation to reduce real costs is completely natural -- just as natural as the urge of consumers to improve their levels of satisfaction.

TABLE 1

Alternative Growth Scenarios

		(1) High Rate Of Return	(2) Middle Rate Of Return	(3) Low rate Of Return
a.	Net Rate of Return (ρ)	.15	.10	.075
b.	Rate of Depreciation	.05	.05	.05
c.	Capital-Output Ratio	2.5	3.0	3.2
d.	Share of Capital in GDP	.50	.45	.40
e.	Share of Labor in GDP	.50	.55	.60
f.	Share of Capital in NDP	.43	.35	.29
g.	Approximate Net Rate of Return on Business Capital	.27	.17	.12

TABLE 2

Average Rates of Productivity Increase

All Countries

	Mean	Median	Fraction < 0
R1	.0115	.0131	32.8%
R2	.0186	.0201	24.7%
R3	.0205	.0218	23.1%
R*[a]	.0150	.0155	25.6%

[a]R* is obtained by "assigning" the high rate-of-return scenario (#1) to countries with high GDP growth rates (≥ 5%), and assigning the low rate-of-return scenario (#3) to those with low GDP growth rates (≤ 1%). Those with growth rates between 1% and 5% were "assigned" scenario #2.

TABLE 3

Regional differences In Breakdown of Growth

(Medians of full-period averages for each country)

Region	Median GDP Growth	Median Capital Contributed	Median Labor Contributed	Measured R2	Implied R2
OECD	3.2	1.0	0.6	1.4	1.6
Asian Tigers	6.8	1.7	1.5	4.1	3.6
Other Asia	4.7	0.7	1.5	2.6	2.5
Africa	3.7	0.7	1.5	2.0	1.5
Latin America/Caribbean	4.0	0.7	1.6	1.9	1.7
Asian tigers Scenario #1	6.8	2.8	1.3	2.7	2.7

TABLE 4

Predicting Rates of Return and Rates of Investment

(23 Countries: Data From U.N. National Accounts and

World Bank Economic Indicators)

Correlations Between	Number of Correlations Positive	Negative
R_t and $\rho_{t+1} - [(\rho_{t-1}+\rho_t)/2]$	22	1
R_t and $\rho_{t+2} - [(\rho_{t-1}+\rho_t)/2]$	20	3
R_t and $\rho_{t+3} - [(\rho_{t-1}+\rho_t)/2]$	17	6
R_t and $\rho_{t+4} - [(\rho_{t-1}+\rho_t)/2]$	15	8
R_t and $ni_{t+1} - [(ni_{t-1}+ni_t)/2]$	21	2
R_t and $ni_{t+2} - [(ni_{t-1}+ni_t)/2]$	23	0
R_t and $ni_{t+3} - [(ni_{t-1}+ni_t)/2]$	22	1
R_t and $n_{t+4} - [(ni_{t-1}+ni_t)/2]$	21	2

R_t = rate of increase of productivity (TFP) between t-1 and t

ρ_t = real rate of return in capital in year t

ni_t = net investment rate (NI/Y) in year t.

Countries: Canada, Colombia, Costa Rica, Ecuador, Finland, France, Greece, Italy, Jamaica, Japan, Korea, Mexico, Myanmar, Paraguay, South Africa, Sri Lanka, Thailand, United Kingdom, United States, Uruguay, Venezuela, Zambia, Zimbabwe.

Period: 1960-1994 (some early years missing due to lack of data).

TABLE 5

Causality Experiment

Connection	R^2	R^2	Connection
a) ΔTFP to ρ	**.80**	**.06**	**c) ρ to ΔTFP**
b) ΔTFP to (ΔK/K)	**.21**	**.02**	**d) (ΔK/K) to ΔTFP**

a) $\rho_t = a_0 + a_1(\Delta TFP_{t-1}) + a_2(\Delta TFP_{t-2}) + a_3(\Delta TFP_{t-3}) + a_4(\Delta TFP_{t-4})$

b) $\Delta TFP_t = b_0 + b_1(\rho_{t-1}) + b_2(\rho_{t-2}) + b_3(\rho_{t-3}) + b_4(\rho_{t-4})$

c) $(\Delta K/K)t = c_0 + c_1(\Delta TFP_{t-1}) + c_2(\Delta TFP_{t-2}) + c_3(\Delta TFP_{t-3}) + c_4(\Delta TFP_{t-4})$

d) $\Delta TFP_t = d_0 + d_1(\Delta K/K)_{t-1} + d_2(\Delta K/K)_{t-2} + d_3(\Delta K/K)_{t-3} + d_4(\Delta K/K)_{t-4}$

Source: NBER Survey: 458 Manufacturing Industries from 1958 to 1992.

TABLE 6

Correlations Between Contemporaneous TFP Growth Rate And
Rates of Return for three Digit Manufacturing Sectors From U.N. Industrial Statistics

% ΔTFP_t Predicting $\Delta\rho$ From "Base" $[(\rho_{t-1}+\rho_t)/2]$

(percentage of r > 0)

	$\rho_{t+1}-(\rho_t+\rho_{t-1})/2$	$\rho_{t+2}-(\rho_t+\rho_{t-1})/2$	$\rho_{t+3}-(\rho_t+\rho_{t-1})/2$	$\rho_{t+4}-(\rho_t+\rho_{t-1})/2$
USA	92.9%	67.9%	60.7%	89.3%
Japan	100.0%	88.9%	100.0%	96.3%
U.K.	89.3%	85.7%	78.6%	89.3%
Canada	89.3%	82.1%	42.9%	25.0%
Indonesia	88.0%	64.0%	60.0%	64.0%
Korea	92.9%	82.1%	67.9%	78.6%
Malaysia	82.1%	78.6%	64.3%	50.0%
Philippines	89.3%	89.3%	89.3%	57.1%
Chile	92.9%	85.7%	75.0%	67.9%
Colombia	64.3%	60.7%	53.6%	46.4%
Mexico	53.8%	73.1%	65.4%	42.3%
Panama	84.0%	84.0%	72.0%	72.0%
Average	84.9%	78.5%	69.1%	64.8%