

Sources of U.S. Economic Growth in a World of Ideas

Charles I. Jones

U.C. Berkeley and NBER
chad@econ.berkeley.edu
<http://elsa.berkeley.edu/~chad>

American Economic Review, forthcoming
August 17, 2001 – Version 6.02

Abstract

At least since 1950, the U.S. economy has benefited from increases in both educational attainment and research intensity. Such changes suggest, contrary to the conventional view, that the U.S. economy is far from its steady-state balanced growth path. This paper develops a model in which these facts are reconciled with the stability of average U.S. growth rates over the last century. In the model, long-run growth is driven by the worldwide discovery of new ideas, which in turn is tied to world population growth. Nevertheless, a constant growth path can temporarily be maintained at a rate greater than the long-run rate provided research intensity and educational attainment rise steadily over time. Growth accounting with this model reveals that 30 percent of U.S. growth between 1950 and 1993 is attributable to the rise in educational attainment, 50 percent is attributable to the rise in worldwide research intensity, and only about 10 to 20 percent is due to population growth in the idea-producing countries.

JEL Classification: O40, E10

An earlier version of this paper was circulated under the title “The Upcoming Slowdown in U.S. Economic Growth.” I would like to thank Paul David, Zvi Griliches, Michael Horvath, Pete Klenow, Paul Romer, Matthew Shapiro, Robert Solow, Steve Tadelis, John Williams, Alwyn Young, and the participants of various seminars for helpful comments and Jesse Czelusta for excellent research assistance. Three anonymous referees provided especially useful suggestions. Financial support from the National Science Foundation (SBR-9818911), the Olin Foundation, and the Sloan Foundation is gratefully acknowledged.

1 Introduction


Over the last 125 years, the average growth rate of per capita GDP in the U.S. economy has been a steady 1.8 percent per year. Indeed, the stability of U.S. growth rates underlies the conventional view that the U.S. economy is close to its long-run steady-state balanced growth path. This view is supported by a number of stylized facts such as the absence of trends in the U.S. capital-output ratio and U.S. real interest rates, as emphasized by Kaldor (1961).

On the other hand, this conventional view is challenged by two noteworthy changes that have occurred for at least the last 50 years, and probably for much longer. First, time spent accumulating skills through formal education, which we can associate with human capital investment, has increased substantially. As of 1940, less than 25 percent of adults in the United States had completed high school, and only about 5 percent had completed four or more years of college. By 1993, more than 80 percent had completed high school, and more than 20 percent had completed at least four years of college. Second, the search for new ideas has intensified. An increasing fraction of workers in the United States and throughout the OECD consists of scientists and engineers engaged in research and development (R&D). In 1950, for example, the U.S. fraction was about 1/4 of one percent. By 1993, this fraction had risen three-fold to more than 3/4 of a percent.

In virtually any model of economic growth, these changes should lead to long-run increases in income. In neoclassical models, such changes generate transition dynamics in the short run and “level effects” in the long run. The growth rate of the economy rises temporarily and then returns to its original value, but the level of income is permanently higher as a result. In many endogenous growth models, such changes should lead to permanent increases in the growth rate itself.

As shown in Figure 1, however, the growth rate of U.S. per capita GDP

Figure 1: U.S. GDP Per Capita, Log Scale



has been surprisingly stable over the last 125 years: the level of per capita GDP is well-represented by a simple time trend.¹ Jones (1995b) used this evidence to argue against many endogenous growth models. Such models suggest that the long-run growth rate of per capita income should be rising with the increases in R&D intensity or time spent accumulating skills, but the data do not exhibit this phenomenon.

At least on the surface, the evidence also appears puzzling even from the standpoint of a neoclassical growth model. These changes should generate temporarily high growth rates and long-run level effects, but the evidence

¹The data are from Maddison (1995). The growth rate from 1950 to 1994, at an annual rate of 1.95 percent, is slightly higher than the growth rate from 1870 to 1929, at 1.75 percent (see, e.g., Ben-David and Papell (1995) on this increase). At the same time, the growth rate in the 1950s and 1960s at 2.20 percent is slightly higher than the growth rate after 1970 of 1.74 percent, reflecting the well-known productivity slowdown. The main point of the figure is to show that a constant growth trend fits reasonably well to a first approximation, but clearly this is only an approximation. Similar results are obtained with GDP per worker, but there is some difficulty obtaining employment data prior to 1900.

in Figure 1 looks very much like an economy that is fluctuating around its balanced growth path.

How can these facts be reconciled? The explanation proposed in this paper is based on the distinction between a constant growth path and a balanced growth path. Along both paths, growth rates are constant, but the former is driven by transition dynamics while the latter is associated with a steady state. The easiest way to see how this might work is to consider a simple example. Imagine an economy described by a Solow model in which the investment rate, rather than being constant, is growing exponentially. Per capita growth in this economy could settle down to a constant rate that is higher than its long-run rate. Of course, the investment rate cannot grow forever (it is bounded at one), and when the investment rate stops growing, the growth rate of the economy will gradually decline to its long-run rate.

In the long run, the fraction of time that individuals spend accumulating skills and the share of the labor force devoted to research must level off. Over the post-war period, and most likely even before, these variables have been rising steadily. Each increase generates a transition path growth effect and a level effect on income, and the series of increases during the last 50 or 100 years have generated a constant growth path with a growth rate higher than the long-run, sustainable growth rate of the U.S. economy. This appears to be the most plausible way to reconcile the facts that motivate this paper.

A number of authors, most recently Klenow and Rodriguez-Clare (1997) and McGrattan and Schmitz (1999), have observed that while there are a large number of candidate growth models in the literature, there has been surprisingly little attention given to reconciling these models formally with data on economic growth. Progress in this direction has been made with respect to understanding differences in levels of income across countries by Mankiw, Romer and Weil (1992), Eaton and Kortum (1999), and Chari, Kehoe and McGrattan (1997), but almost no research has conducted these “quantitative theory” exercises with a focus on long-run growth. This paper

represents a first step in this direction. A formal growth model — admittedly only one of the many possible candidates — is presented and calibrated in order to understand some of the basic facts of U.S. and world economic growth.

Section 2 of the paper presents a model in which long-run growth is driven by the discovery of new ideas throughout the world. In this respect, the model builds on a large collection of previous research, including Romer (1990), Grossman and Helpman (1991), and Aghion and Howitt (1992) as well as earlier contributions by Phelps (1966), Shell (1966), Nordhaus (1969), and Simon (1986). In the model, growth in the world's stock of useful knowledge is ultimately tied to growth in world research effort. While the model is constructed with an eye toward the quantitative theory exercises that follow, it also yields a number of interesting results in its own right. In particular, the model adds to a growing literature on the way in which “scale effects” matter for economic growth.

Section 3 uses the model to conduct a growth accounting exercise to document the sources of U.S. growth over the period 1950 to 1993. Contrary to the conventional wisdom that the U.S. economy is on a balanced growth path, the accounting suggests that the long-run component of growth was less than 20 percent of the total during these years. More than 80 percent of growth was associated with transition dynamics.

If transition dynamics are so important to recent U.S. growth, why haven't we seen the traditional signature of these dynamics, e.g. a gradual decline in growth rates? Section 4 presents the constant growth path hypothesis and reworks the growth accounting in the presence of this restriction. The results from the two different accounting approaches are similar, suggesting that the constant growth path hypothesis is a reasonable approximation.

Section 5 of the paper discusses the robustness of the results, and Section 6 offers some concluding remarks.

2 Modeling Growth

Consider a world consisting of M separate economies. These economies are similar in that they have the same production possibilities. They differ because of different endowments and allocations. Within an economy, all agents are identical. The economies evolve independently in all respects except one: they share ideas. Until we discuss the creation of ideas, we will focus on a representative economy and omit any subscript to distinguish economies.²

2.1 Production Possibilities

In each economy, individuals can produce a consumption-capital good that we will call output. Total output Y_t produced at time t is given by

$$Y_t = A_t^\sigma K_t^\alpha H_{Yt}^{1-\alpha}, \quad (1)$$

where K_t is physical capital, H_{Yt} is the total quantity of human capital employed to produce output, and A_t is the total stock of ideas available to this economy. We assume $0 < \alpha < 1$ and $\sigma > 0$. Notice that there are constant returns to scale in K and H_Y holding the stock of ideas A constant, and increasing returns to K , H_Y , and A together. This assumption reflects the now-common notion that ideas are nonrivalrous or “infinitely expandable.”³

We now discuss each element of this production function in turn. First, physical capital is accumulated by foregoing consumption:

$$\dot{K}_t = s_{Kt}Y_t - dK_t, \quad K_0 > 0. \quad (2)$$

²In general, Greek letters will be used to denote parameters that are common across countries and constant over time, while Roman letters will denote variables that may differ across countries and may vary over time. The only two exceptions to this rule are the depreciation rate d and the population growth rate n . These Roman letters will denote parameters that are constant over time and identical across countries, as described further below.

³Quah (1996) suggests this latter term, which he attributes to a letter written by Thomas Jefferson in 1813 and discussed by David (1993). See Romer (1990) for a general discussion of this property.

The variable s_{Kt} denotes the fraction of output that is invested ($1 - s_{Kt}$ is the fraction consumed), and $d > 0$ is the exogenous, constant rate of depreciation.

Next, aggregate human capital employed producing output is given by

$$H_{Yt} = h_t L_{Yt}, \quad (3)$$

where h_t is human capital per person and L_{Yt} is the total amount of raw labor employed producing output. An individual's human capital is produced by foregoing time in the labor force. Letting ℓ_h represent the amount of time an individual spends accumulating human capital,

$$h_t = e^{\psi \ell_h t}, \quad \psi > 0. \quad (4)$$

The exponential formulation used here is the most straightforward way of incorporating human capital in a manner that is consistent with the large literature on schooling and wages following Mincer (1974) and with the substantial growth accounting literature that makes adjustments for education. It is a special case of a formulation suggested by Bils and Klenow (2000).

The final factor in the production of output is the stock of ideas, A . In the model, ideas represent the only link between economies; there is no trade in goods, and capital and labor are not mobile. Ideas created anywhere in the world are immediately available to be used in any economy. Therefore, the A used to produce output in equation (1) corresponds to the cumulative stock of ideas created anywhere in the world and is common to all economies.⁴

New ideas are produced by researchers, using a production function like that in Jones (1995a):

$$\dot{A}_t = \delta H_{At}^\lambda A_t^\phi, \quad A_0 > 0, \quad (5)$$

⁴A previous version of this paper considered a model in which the diffusion of ideas was not instantaneous and depended on economic forces. In particular, ideas produced anywhere in the world had to be learned by each person before they could be used in production. This model was more complicated but led to the same ultimate result given in equation (10) below.

where H_A is effective world research effort, given by

$$H_{At} = \sum_{i=1}^M h_{it}^\theta L_{Ait}. \quad (6)$$

In this equation, i indexes the economies of this world, L_{Ai} is the number of researchers in economy i , and $\theta \geq 0$. World research effort is the weighted sum of the number of researchers in each economy, where the weights adjust for human capital.

According to equation (5), the number of new ideas produced at any point in time depends on the number of researchers and the existing stock of ideas. We allow $0 < \lambda \leq 1$ to capture the possibility of duplication in research: if we double the number of researchers looking for ideas at a point in time, we may less than double the number of unique discoveries. We assume $\phi < 1$, which still allows past discoveries to either increase ($\phi > 0$) or decrease ($\phi < 0$) current research productivity.

Finally, there is a resource constraint on labor in this economy. Each economy is populated by N_t identical, infinitely-lived agents. The number of agents in each economy grows over time at the common and constant exogenous rate $n > 0$:

$$N_t = N_0 e^{nt}, \quad N_0 > 0. \quad (7)$$

Each individual is endowed with one unit of time and divides this unit among producing goods, producing ideas, and producing human capital. Because time spent in school is excluded from labor force data, it is helpful to write the resource constraint as

$$L_{At} + L_{Yt} = L_t = (1 - \ell_{ht})N_t. \quad (8)$$

where L_t denotes employment. In addition, we define $\ell_A \equiv L_A/L$ as the fraction of the labor force that works to produce ideas (“research intensity”), and $\ell_Y \equiv L_Y/L$.

2.2 Allocations

It is typical in a paper like this to specify preferences and markets which, given the production possibilities of the model, determine allocations. These equilibrium conditions then provide an additional set of restrictions that can be analyzed and compared to data.

This is not the approach followed here. Instead, we take the allocations as given (ultimately, they will simply be given by the data). We will feed the allocations through the production possibilities just described to see if the “technology” of this model makes any sense. This can be viewed as a precursor to the richer analysis that comes from adding markets to the model and analyzing equilibrium conditions as well as technologies. It is reminiscent of the approach taken originally by Solow (1957) in his growth accounting exercise.

This is not to suggest that explaining the observed allocations is uninteresting. On the contrary — both problems are important, but it is fruitful to consider them one at a time. At the end of the paper, I will suggest a preliminary explanation for some of the observed allocations, but a careful analysis of this question would require another paper.

For the moment, then, we assume that the time paths of s_K , ℓ_A , ℓ_h , and ℓ_Y are given exogenously (and may differ across economies). These variables will be referred to as allocations.

2.3 Key Results from the Model

For the accounting exercises that follow, we need to derive several results from this basic setup. First, notice that the production function in equation (1) can be rewritten in terms of output per worker $y_t \equiv Y_t/L_t$ as

$$y_t = \left(\frac{K_t}{Y_t}\right)^{\frac{\alpha}{1-\alpha}} \ell_{Yt} h_t A_t^{\frac{\sigma}{1-\alpha}} \quad (9)$$

This expression turns out to be quite useful because along a balanced growth path, all terms on the right side except for the last are constant. This

equation will serve as the basis for the first growth accounting exercise in the next section.

A second result is convenient for two reasons: it will be the basis of the second growth accounting exercise and it allows the steady-state growth rate of the economy to be derived easily. Suppose the stocks K and A grow at constant rates (which in turn requires H_A to grow at a constant rate). In this case, as explained in the next paragraph, output per worker in equation (9) can be decomposed as

$$y_t = \left(\frac{s_{Kt}^*}{n + g_k + d} \right)^{\frac{\alpha}{1-\alpha}} \ell_{Yt} h_t \left(\frac{\delta}{g_A} \right)^{\frac{\gamma}{\lambda}} H_{At}^{*\gamma}, \quad (10)$$

where $k \equiv K/L$ and $\gamma \equiv \frac{\sigma}{1-\alpha} \frac{\lambda}{1-\phi}$. Here, and in the remainder of the paper, we use the notation g_x to denote the constant growth rate of some variable x and an asterisk to denote a quantity that is growing at a constant rate.

To see how this equation is derived, it is helpful to keep in mind a key property: when growth rates are constant, stocks can be inferred from flows. Thus, the first term in parentheses in equation (10) is simply the capital-output ratio. This ratio is proportional to the investment rate when the capital stock grows at a constant rate, just as in the standard Solow model. The last term in equation (10) comes from the fact that when the stock of ideas A grows at a constant rate, this stock can be inferred from the flow of research effort H_A .⁵

Finally, this economy can exhibit a stable, balanced growth path, defined as a situation in which all variables grow at constant, exponential rates forever (possibly zero). It is easy to show that along such a path, the allocations must be constant. Then, from equation (10), the growth rate of output per worker is proportional to the growth rate of effective world research H_A . Finally, since h must be constant along a balanced growth path, growth in

⁵More formally, divide both sides of the production function for ideas in (5) by A to get $\dot{A}_t/A_t = \delta H_{At}^\lambda / A_t^{1-\phi}$. When the growth rate of A is constant, this equation can be solved to see that A is proportional to H_{At}^λ . The second-to-last term in (10) is the factor of proportionality, which depends on g_A .

the effective number of world researchers is driven by population growth, and a balanced growth path yields

$$g_y = \frac{\sigma}{1 - \alpha} g_A = \gamma n. \quad (11)$$

2.4 Remarks

At this stage, several remarks about the model are worth noting. First, equation (11) indicates that long-run per capita growth is ultimately tied to world population growth in this model, a result emphasized by Jones (1995a). With allocations given, a larger world population means a larger number of researchers around the world. These researchers produce more ideas, which, being infinitely expansible, raise incomes around the world. This is the intuition behind the scale effect implicit in equation (10). If the level of world population is doubled, keeping all other parameters and allocations constant, then H_A^* is also doubled. This raises the level of income for all countries in the world in the long run by a factor of 2^γ .

The model clarifies the level at which the scale effect associated with the nonrivalry of knowledge operates. The relevant scale variable is the population of the collection of countries that are sufficiently close to the world's technological frontier that they can contribute to the discovery of new ideas. Neither India's large population nor Singapore's small population is particularly relevant to these countries' income levels or growth rates. Rather, it is the scale of world research effort that matters for the economic performance of individual countries.

One final remark on population growth is worth mentioning. It is well-known that cross-country growth regressions typically document a negative correlation between per capita income growth and population growth. But this model appears to predict a positive relationship, as in equation (11). Can these facts be reconciled? The answer is yes. Recall that the standard negative correlation between countries is typically interpreted (e.g. as in

Mankiw et al. (1992)) as reflecting the transition dynamics of the neoclassical growth model: a higher population growth rate reduces the steady-state capital-output ratio because more investment must go simply to maintain the existing capital-output ratio in the growing population. But this effect is also at work here, as indicated by the first term in parentheses in equation (10). In fact, because the H_{At}^* term is the same across countries, this model shares exactly the neoclassical predictions for population growth in a cross-section of countries.

3 Quantitative Analysis

The model developed in the previous section provides a framework for analyzing economic growth in a particular country, recognizing that the engine of growth is the creation of ideas throughout the world. In this section, we apply this model to understand twentieth-century growth in the United States. First, however, we begin by documenting quantitatively the behavior of the key variables emphasized in the model.


3.1 Data

Figure 2 reports data on GDP per hour worked and on the factors of production, corresponding to the quantities in equation (9).⁶ The quantities appear to grow at roughly constant rates, although a slight productivity slowdown in both output per hour and multifactor productivity is apparent. The capital-output ratio is fairly stable, as is commonly accepted, and human capital per worker rises because of the increase in educational attainment discussed below.

Turning to the key allocations from equation (10), Figure 3 plots average educational attainment in the United States for persons age 25 and over, from 1950 to 1993. Educational attainment rises smoothly from a low of

⁶The sources for this and all other data used in the paper are described in detail in the Data Appendix.

Figure 2: Factors of Production in the United States



Note: Multifactor productivity is calculated as the residual in equation (9) and corresponds to the term $A^{1-\alpha}$.

about 8.5 years in 1950 to a high of about 12.5 years by 1993. Ideally, one would also like to measure skills accumulated outside the formal education process, for example through on-the-job training, but this data doesn't seem to be available. To map the educational attainment data into ℓ_h , one needs to divide by some measure of an individual's time endowment or lifetime. Life expectancy has been rising over this period, though clearly at a much slower rate than educational attainment. On the other hand, the length of an individual's working life has actually declined because of the decline in the age of retirement. As a rough compromise between these two trends, we will simply assume that an individual's labor endowment has remained constant.⁷ Also, we will measure ℓ_h directly as educational attainment and

⁷Some rough statistics give an idea of the magnitudes involved. An individual facing a lifetime of the average cross-sectional mortality rates in 1950 would have a life expectancy at birth of 68.1 years. For 1997 the number would be 76.5 years. On the other hand, the

incorporate this constant term in the coefficient ψ .

The last term in equation (10) is the effective number of researchers in the world. Recall from equation (6) that this number is given by a weighted sum of research labor. To provide a rough empirical measure of H_A , we will make two assumptions. First, we assume that only researchers in the G-5 countries (France, West Germany, Japan, the United Kingdom, and the United States) are capable of extending the frontier of knowledge. This is motivated primarily by the lack of data for other countries prior to the 1980s and by the fact that the majority of world research effort is conducted in these countries. Second, we assume that the “quality” of these researchers is the same across the advanced countries and has remained constant over time; this can be implemented artificially by setting $\theta = 0$. This seems like a reasonable assumption if one thinks that to be hired in the first place, a researcher must have a certain level of education. The rise in average educational attainment, then, would not have an important effect on the quality of researchers. Under these assumptions, we can measure world H_A as the sum across the G-5 countries of the number of scientists and engineers engaged in research and development, as reported by the National Science Board (1993, 1996).

Figure 4 displays this series, normalized by the size of G-5 employment. That is, this figure plots a measure of research intensity corresponding to ℓ_A , both for the G-5 as a whole and for the United States. Between 1950 and 1993, research intensity in the G-5 countries increased by more than a factor of four, rising at an average rate of 3.6 percent per year. This rate reflects the very rapid growth in the number of researchers in the G-5, at a rate of 4.8 percent per year, together with the modest increase in G-5 employment at a rate of 1.2 percent per year.⁸

average age of retirement for men fell from 68.7 years in 1950 to 63.7 years in 1989. The life expectancy data are taken from Table 11 in Anderson (1999). The retirement data are taken from Table 1 in Gendell and Siegel (1992).

⁸The lack of a smooth upward trend in U.S. research intensity in Figure 4 is largely

Figure 3: Average U.S. Educational Attainment, Persons Age 25 and Over





Figure 4: Research Intensity in the G-5 Countries



Note: The dashed line indicates data that have been estimated by the author. See the Data Appendix.

The magnitude of research intensity is also worth noting. In the United States and throughout the G-5 countries, less than one percent of the labor force is engaged in research according to the definition employed by the National Science Foundation. This number seems small, and the definition is surely too narrow.⁹ On the other hand, over time R&D has become a more formal activity, possibly suggesting that the measured increase overstates the true increase. Despite these possible problems, measured R&D is the only data we have, and it likely represents a reasonable benchmark provided these caveats are kept in mind.

3.2 Accounting for U.S. Growth

We are now ready to undertake a growth accounting exercise using the model from Section 2. First, we use the form of the production function in equation (9) to decompose the growth rate of output per worker between any two points in time into its components:

$$\hat{y}_t = \frac{\alpha}{1-\alpha}(\hat{K}_t - \hat{Y}_t) + \hat{h}_t + \hat{\ell}_{Yt} + \left(\frac{\sigma}{1-\alpha}\hat{A}_t - \gamma n\right) + \gamma n. \quad (12)$$

where a hat ($\hat{\cdot}$) is used to denote the average change in the log of a variable between two points in time: e.g. $\hat{y}_t = 1/t * (\log y_t - \log y_0)$. By adding and subtracting the steady-state growth rate γn in (12), this equation has a nice interpretation. In particular, all of the terms save the last on the right side of this equation are zero in steady state. If the conventional wisdom that

due to the “bulge” in research intensity between 1955 and 1975 associated with the space program and the defense buildup. Non-defense, non-space research intensity, measured by the spending share of GDP, shows a trend that is closer to monotonic. Because the direct outputs of defense and space spending are measured at cost, these sectors show, by definition, no productivity growth. For this reason, studies of R&D and productivity growth often focus on measures of R&D that exclude these categories. One can, of course, make a case for including these measures, e.g. based on things like the World Wide Web.

⁹In the United States, the definition of R&D focuses on science and engineering. The “research” behind the creation of new consumer products like Odwalla or Jamba juice fruit drinks is not included for this reason. Also, the definition emphasizes research that requires the equivalent of a four-year college degree, meaning that the research undertaken by the young Steve Jobs, Bill Gates, and Marc Andreessen was probably excluded as well.

the U.S. economy is close to its balanced growth path is correct, the last term should account for the bulk of growth.

Several remarks concerning this approach are relevant. First, the decomposition in equation (12) is valid across any two points in time under very weak assumptions; it is not a steady-state relationship. Second, the accounting exercise is in the same spirit as classic work in growth accounting by Solow (1957), Denison (1962), and others, but it differs in some important ways. For example, multifactor productivity growth is made endogenous in this framework by specifying a complete growth model.

In addition, if the economy happens to be growing along a balanced growth path, one hundred percent of growth will be attributed to the last term in equation (12). In traditional accounting for growth in output per worker, even along a balanced growth path there will be a contribution from the capital-labor ratio. This is true even though in some sense the growth in the capital-labor ratio occurs because of total factor productivity growth. Alternatively, if one does the growth accounting in terms of capital intensity (the capital-output ratio), all of growth will be attributed to multifactor productivity growth. This exercise follows the latter approach. In this sense, the exercise is closer to Solow (1956) than to Solow (1957), and follows in the tradition of David (1977), Mankiw et al. (1992), and Klenow and Rodriguez-Clare (1997).

Most of the empirical counterparts of the variables in equation (12) are readily observed. Indeed, the figures already presented contain most of the required data. The key growth rates needed for the accounting decomposition are reported in Table 1.¹⁰

The next step in implementing the growth accounting decomposition

¹⁰One element in Table 1 merits further discussion. The term $\hat{\ell}_Y$ is calculated as the growth rate of the fraction of the labor force working in producing output. Because time spent in school is not considered part of the labor force by the Bureau of Labor Statistics (BLS), $\hat{\ell}_Y$ is computed as the growth rate of $1 - \ell_A$. It shows a slight decline because of the rise in U.S. research intensity, but because less than one percent of the U.S. labor force works as a researcher, the decline is negligible.

Table 1: Average Annual Growth Rates, 1950–1993

Growth Rate of	Variable	Sample Value
Output per hour	\hat{y}	.0200
Capital-output ratio	$\hat{K} - \hat{Y}$	-.0015
Share of labor in goods	$\hat{\ell}_Y$	-.0001
Human capital	\hat{h}	.0063
Multifactor productivity	\hat{A}	.0146
R&D labor	\hat{H}_A	.0483
G-5 labor force	\tilde{n}	.0120
Share of labor in R&D	$\tilde{\ell}_A$.0363
Annual change in ℓ_h	$\Delta\ell_h$.0902

Note: For data sources, see the appendix. A tilde \sim is used to distinguish a “world” aggregate (a G-5 total) from a U.S. value.

suggested by equation (12) is to obtain values for the parameters in that equation. We assume a value of 1/3 for the capital coefficient, α , motivated by the usual data on capital’s share of income.

The parameter ψ is readily inferred from a wealth of microeconomic evidence. Interpreting ℓ_h as years of schooling, the parameter ψ corresponds to the return to schooling estimated by Mincer (1974) and others using log-wage regressions: output per worker, and hence the wage, differs across workers in the same economy with different amounts of schooling with a semielasticity of ψ . The labor market literature suggests that a reasonable value for ψ is .07, which we adopt here. This value implies that an additional year of schooling has a direct effect of raising labor productivity by seven percent.

The parameters σ and γ are the only unknowns that remain in equa-

tion (12). Without observing ideas directly, the parameter σ cannot be identified. We will therefore make the normalization $\sigma = 1 - \alpha$ so that A is measured in units of Harrod-neutral productivity. This normalization is without loss of generality for the purposes of this paper. Empirically, A is recovered from the data in the traditional fashion: it is calculated directly from equation (9).

The parameter γ is a combination of parameters from the idea production function (recall that $\gamma = \frac{\sigma}{1-\alpha} \cdot \frac{\lambda}{1-\phi} = \frac{\lambda}{1-\phi}$) and is more difficult to obtain. Dividing both sides of the production function for ideas in (5) by A_t and rewriting in terms of γ , we have

$$\frac{\dot{A}_t}{A_t} = \delta \left(\frac{H_{At}^\gamma}{A_t} \right)^{1-\phi}. \quad (13)$$

Since A_t is measured as multifactor productivity, this equation states that productivity growth depends on the ratio of the quantity of human capital used in producing ideas to the level of productivity. As an empirical matter, both H_{At} and A_t are trending sharply upward. In contrast, the growth rate of productivity is stationary, or perhaps even declining slightly during the 1950–1993 period because of the productivity slowdown. The parameter γ , then, plays the important role of detrending the ratio H_A^γ/A to deliver a stationary productivity growth rate.

If multifactor productivity growth truly exhibited no trend between 1950 and 1993, the parameter γ would have to equal the ratio of the growth rates of multifactor productivity and H_A . Using the values from Table 1, this would imply a value for γ of $.0146/.0483 = .30$. To the extent that multifactor productivity growth has been declining, the ratio H_A^γ/A would have to decline, which would occur if γ were less than this value of 0.3.

Appendix A provides a more rigorous econometric analysis of the estimation of γ that supports the intuition just given. The estimates in the appendix range from a high value of about 1/3 to a low value of about 0.05. To the extent one believes that productivity growth is mismeasured and that

Table 2: Accounting for U.S. Growth, 1950-1993

— Transition dynamics —						
	Output per Hour	Capital Intensity	Labor Reallo- cation	Educa- tional Attain.	Excess Idea Growth	Steady State Growth
γ	\hat{y}	$\frac{\alpha}{1-\alpha}(\hat{K} - \hat{Y})$	$\hat{\ell}_Y$	\hat{h}	$\hat{A} - \gamma n$	γn
0.050	0.0200 (100.0)	-0.0007 (-3.7)	-0.0001 (-0.6)	0.0063 (31.5)	0.0140 (69.8)	0.0006 (3.0)
0.200	0.0200 (100.0)	-0.0007 (-3.7)	-0.0001 (-0.6)	0.0063 (31.5)	0.0122 (60.9)	0.0024 (12.0)
0.333	0.0200 (100.0)	-0.0007 (-3.7)	-0.0001 (-0.6)	0.0063 (31.5)	0.0106 (52.9)	0.0040 (19.9)

Note: This table reports the growth accounting decomposition corresponding to equation (12). The specifications are sorted according to the value for γ that is used. Numbers in parentheses are percentages of the growth rate of output per hour.

true productivity growth has not declined, one would favor the higher value. Alternatively, if one believes that the productivity slowdown is measured accurately and that the parameter λ is small (for example, about 0.25), then one would favor the smaller value of γ . More generally, this range encompasses the plausible values of γ . Based on the econometric analysis and on the intuition provided above, we will consider values of γ of 0.33 and 0.05, together with an intermediate value of 0.20.

With these parameter values and the data from Table 1, we now turn to the growth decomposition implied by equation (12). This accounting is reported in Table 2. Output per hour grew at an average annual rate of 2.00 percent between 1950 and 1993 in the United States. As mentioned above, the stability of the investment rate translated into a relatively stable capital-output ratio, leading to only a small but negative .07 percentage

point contribution to growth. Similarly, there was a very small shift of labor away from producing goods and in to producing ideas, but the composition effect associated with this change had little effect on output per hour.

The rise in educational attainment contributed 0.63 percentage points to growth in output per hour, accounting for just under 1/3 of growth during this period. Between 1950 and 1993, mean educational attainment in the United States rose by about 4 years. If each year of education leads to a seven percent rise in output per worker, increased educational attainment raised output by about 28 percent over this period, or by an annualized amount of about six-tenths of a percent per year.

The remaining 70 percent of growth is attributed to a rise in the stock of ideas produced by researchers throughout the G-5 countries. This effect is itself the sum of two pieces. First, growth in the stock of ideas in excess of the steady-state rate is the single largest contributor to growth in this decomposition, accounting for between 1.06 and 1.40 percentage points or 53 to 70 percent of growth, depending on the exact value of γ . Finally, the steady-state component, associated with the general rise in G-5 employment, contributed between .06 and .40 percentage points to growth in U.S. output per hour between 1950 and 1993, accounting for only 3 to 20 percent of growth.

In this model, long-run growth arises entirely from world population growth: per capita growth requires growth in the stock of ideas which in turn requires growth in the number of researchers. Nevertheless, the decomposition in Table 2 yields the surprising result that during the period 1950 to 1993, less than 20 percent of growth was attributable to this scale effect. More than 80 percent of growth in the United States during this period is attributed to the transition dynamics associated with educational attainment and the stock of ideas.

4 The Constant Growth Path

A natural question arises at this point. If more than 80 percent of U.S. growth in recent history is associated with transition dynamics, then why do we not see the traditional signature of a transition path, e.g. a gradual decline in growth rates to their steady-state level? Why is it that U.S. growth rates over the last century or more appear so stable?

At some level, it must be that the transition dynamics associated with the various factors of production just happen to offset in such a way as to leave the growth rate of output per worker fairly constant. This could occur if the transition dynamics of the various factors take wildly different paths that, in an amazing coincidence, happen to offset. Alternatively, and perhaps more plausibly, the transition dynamics of the various factors could themselves be well-behaved in a sense that will be made precise, leading to what we will call a constant growth path.

To see this, it is convenient to rewrite equation (10) as

$$y_t = \xi_{Kt}^{\frac{\alpha}{1-\alpha}} \ell_{Yt} e^{\psi \ell_{ht}} \nu \tilde{\ell}_{At}^{\gamma} \tilde{L}_t^{\gamma}. \quad (14)$$

Several new pieces of notation are introduced in this equation. First, $\xi_K \equiv \frac{s_K}{n+g_k+d}$. Second, $\nu \equiv (\delta/g_A)^{\gamma/\lambda}$. Third, a tilde $\tilde{\cdot}$ is used to denote a “world” aggregate (a G-5 total). Finally, we are exploiting the assumption made above about the world’s contribution to research, namely that only G-5 researchers are sufficiently close to the frontier to contribute new ideas, and that researchers have the same, unchanged skill level, which we normalize to one. Therefore, we have $H_A = \tilde{\ell}_A \tilde{L}$, where \tilde{L} is G-5 employment and $\tilde{\ell}_A$ is G-5 research intensity.

It is now appropriate to highlight the distinction between a constant growth path and a balanced growth path. A constant growth path (CGP) is defined as a situation in which all growth rates are constant. It is distinguished from a balanced growth path in that it is not required to be a situation that can continue forever. Notice that nothing in the derivation

of equation (14) requires the allocations to be constant; we only require K and A to grow at constant rates. Based on the evidence in Figure 2, it does not seem implausible that this requirement holds, at least as a first approximation.

The consequence of this observation is that it is possible to observe a constant growth rate of y provided each of the terms in equation (14) is growing at a constant rate. More formally, taking logs and differencing equation (14) to approximate the growth rate gives

$$g_y = \frac{\alpha}{1-\alpha} g_{\xi_K} + g_{\ell_Y} + \psi \Delta \ell_h + \gamma g_{\tilde{\ell}_A} + \gamma \tilde{n}, \quad (15)$$

where \tilde{n} is the growth rate of G-5 employment.

In steady state, every term in equation (15) except the last must be zero, so that this equation reduces back to the condition $g_y = \gamma \tilde{n}$, familiar from equation (11). Out of steady state, however, we see that it is possible for the growth rate of output per worker to be constant and greater than its long-run growth rate. This could occur, for example, because of growth in the human capital investment rate ℓ_h and in research intensity $\tilde{\ell}_A$. Of course, this situation could not exist forever, because these shares are bounded from above at one. But, curiously, the transition dynamics could lead to a temporary, constant growth path.¹¹

Table 3 conducts the growth decomposition under the assumption of a constant growth path. Notice that in this case, all of the terms in equation (15) are observed, with the exception of γ . Instead of using the econometric estimates of γ obtained before, we take this opportunity to provide an independent check on our approach. That is, we calculate the value of γ that makes equation (15) hold exactly. This calculation yields a value of

¹¹There is one problem with this reasoning. Because ℓ_A , ℓ_Y , and ℓ_h are related through the resource constraint, shares cannot grow simultaneously at exponential rates, meaning that a strict constant growth path is not possible. It turns out, however, that because ℓ_Y is close to one, this technicality is not important in practice, as we will see. An alternative would be to focus on Y/L_Y instead of Y/L , in which case an exact CGP is possible.

Table 3: Constant Growth Path Decomposition, 1950-1993

Description	Variable	Sample Value	Percent of g_y
Growth Rate of Y/L	g_y	.0200	100
<i>Equals:</i>			
Capital Intensity Effect	$\frac{\alpha}{1-\alpha}g_{\xi_K}$.0009	4
+ Effect of Labor Reallocation	g_{ℓ_Y}	-.0001	-1
+ Educational Attainment Effect	$\psi\Delta\ell_h$.0063	32
+ G-5 R&D Intensity Effect	$\gamma g_{\tilde{\ell}_A}$.0097	49
+ Scale Effect of G-5 Labor Force	$\gamma\tilde{n}$.0032	16

0.268, at the high end of the range of values for γ used earlier.¹²

The results of the CGP decomposition are roughly in line with the results obtained in the first accounting exercise. Transition dynamics associated with educational attainment and the growth in research intensity account for 80 percent of growth in output per hour. The component of growth associated with rising G-5 employment is approximately 0.3 percentage points, accounting for about 15 percent of growth. This suggests that the CGP interpretation of recent U.S. history is a reasonable approximation.

I have also explored the robustness of the results in Table 3 to assumptions about the mismeasurement of growth or the presence of growth due to other factors left out of the model. If the true growth rate of output per hour is higher than the measured growth rate, a larger proportion of growth is attributed to research, but the basic results are left unchanged. For example, if true growth in output per hour is actually 3 percent, rising

¹²An alternative is to impose the values of γ used before and to include a residual in the growth decomposition. In this case, the percent contribution of the residual to growth is 53 when $\gamma = .05$, 17 when $\gamma = .20$, and -15 when $\gamma = .33$. The CGP approximation, then, is most accurate if γ is between .20 and .33.

research intensity and educational attainment still account for 78 percent of growth and the population growth component is 0.57 percent, accounting for 19 percent of growth. Alternatively, even if 25 percent of growth is actually due to factors outside of the model, transition dynamics associated with human capital and research intensity still account for at least 60 percent of growth. Finally, if the true growth rate of ℓ_A is only 1/2 of the measured growth rate, then the contribution of research intensity and educational attainment remains high at 75 percent while the population component rises to 25 percent of recent growth. However, this increase is due to a rise in γ to .43, which seems high when compared to the earlier econometric results, suggesting that the CGP interpretation is somewhat more strained in this case.

The accounting exercises here begin with the year 1950 because of data limitations. Nevertheless, the available statistics related to research intensity and educational attainment suggest that this interpretation may apply for the preceding half century as well. Romer (2000) documents a steady increase in the share of engineers as a fraction of the labor force going back to 1900; a similar fact is true for chemists. According to Goldin (1999), school enrollment rates increase sharply beginning around 1900, with more modest changes over the previous fifty years.¹³

An implication of these results is that the growth rates experienced in the U.S. economy for the last century or so are not indicative of a steady state. The rise in research intensity and educational attainment that has occurred over this period cannot continue forever. At most the entire labor force can be devoted to producing ideas, and at most individuals can spend their entire lives accumulating human capital. When these variables stabilize, the standard pattern of transition dynamics will presumably set in, and the economy will gradually transit to its long-run rate of growth. This rate is

¹³The evidence that underlies these statements can be found in Bureau of the Census (1975). See especially the following series: D-245, D-255, and H-433.

given by $g_y = \gamma\tilde{n}$ in the model. Between 1950 and 1993, it was approximately equal to 0.3 percent in the U.S. economy, only about 15 percent of the observed growth rate. Obviously, this rate could be even lower in the future if population growth rates decline.

5 Discussion

This section addresses three issues related to the results. First is an exploration of the non-CGP transition dynamics of the model. Second is a comparison to previous growth accounting exercises that measure the contribution of research. Third is a general discussion of the future of economic growth.

5.1 “Traditional” Transition Dynamics

Suppose the economy is growing at 2 percent per year because of increases in educational attainment and increases in research intensity. Then, at some point, ℓ_h and ℓ_A stabilize. What does the transition to the steady state look like?

In the model in this paper, this transition can be analyzed in a straightforward fashion. First, when ℓ_h is constant, there are no transition dynamics associated with human capital. This is an oversimplification which is relatively harmless in the context of the constant growth path analysis, but more generally is probably not a good assumption.¹⁴ Together with the fact that there will be transition dynamics associated with the induced accumulation of physical capital, this oversimplification suggests narrowing our analysis to the transition dynamics for multifactor productivity rather than attempting

¹⁴As just one example, we could instead allow human capital to be accumulated according to $\dot{h}_t = \exp(\beta\ell_h)h_t^\eta A_t^{1-\eta}$ with $0 < \eta < 1$. When h grows at a constant rate, this implies that the level of h is proportional to $\exp(\beta/(1-\eta)\ell_h)$, so that we could just define $\psi = \beta/(1-\eta)$. The constant growth path analysis is then robust to this kind of change. Off a constant growth path, transition dynamics associated with h could be important, and this would potentially affect the first growth accounting exercise in the paper.

to say something about the transition path for output per worker.

In the case considered here in which research intensity has stabilized (or more generally even if research intensity grows at a constant rate), the differential equations governing the growth rate of A and the stock of A can be solved analytically. Let $x_t \equiv \dot{A}_t/A_t$ denote the growth rate of the stock of ideas. With constant research intensity, straightforward analysis reveals that this growth rate satisfies the following differential equation:

$$\frac{\dot{x}_t}{x_t} = \lambda n - \frac{\lambda}{\gamma} x_t \quad (16)$$

This differential equation can be solved to yield

$$\frac{x_t - x^*}{x_t} = \frac{x_0 - x^*}{x_0} e^{-\lambda n t}, \quad (17)$$

where x^* denotes the steady-state growth rate of A .¹⁵

This result is convenient for a couple of reasons. First, it allows us to calculate a half-life for the transition. If research intensity stabilized today, what would the time path of multifactor productivity growth look like? How long would it take the growth rate to fall in half? Table 4 answers these questions for various parameter values, assuming a constant G-5 population growth rate of one percent and starting from an initial multifactor productivity growth rate of 1.46 percent, the average value between 1950 and 1993. To begin, we compute the speed of convergence to steady state using a log-linear approximation. It should not be surprising given the result in equation (17) that this rate is given by λn . Table 4 shows that the associated half-lives from the log-linear approximation are relatively large numbers, into the hundreds of years.

The slow rate of convergence suggested by the log-linear approximation is misleading, however, as the exact calculations in the rest of the table show. A typical value is the half-life of 25.7 years for $\gamma = .20$ and $\lambda = 1/2$. Significantly lower values are possible if λ is larger than $1/2$.

¹⁵ The key integral result used to solve the differential equation is $\int \frac{dx}{x(ax+b)} = \frac{1}{b} \log\left(\frac{x}{ax+b}\right)$.

Table 4: The Half-Life of Multifactor Productivity Growth

λ	λn	Log-Linear	— Exact Half-Life for —		
		Approx.	$\gamma = .33$	$\gamma = .20$	$\gamma = .05$
1	0.010	69.3	20.6	12.8	3.4
1/2	0.005	138.6	41.1	25.7	6.7
1/4	0.003	277.3	82.3	51.4	13.5

Note: Half-lives calculated from equation (17) assuming $x_0 = .0146$ and $n = .01$.

The differential equation in (17) for the growth rate of A can itself be solved.¹⁶ The level of multifactor productivity at time t is given by

$$A_t = A_0 \left(\frac{x_0}{x^*} e^{\lambda n t} + 1 - \frac{x_0}{x^*} \right)^{\gamma/\lambda}. \quad (18)$$

This solution allows us to answer another question of interest. For example, if research intensity had stabilized in 1950 instead of growing so rapidly, how much lower would multifactor productivity be today? Figure 5 plots the time path of A_t on a log scale to answer this question, taking the intermediate value of $\gamma = .20$. For $\lambda = 1$, the level of productivity is 32 percent below trend after 50 years, while for $\lambda = 1/4$, the shortfall is 17 percent. These numbers can be mapped directly to output per worker as well, holding other things equal.¹⁷


5.2 Comparison to Previous Work

A number of studies have employed traditional growth accounting methods to study the effect of R&D on growth; see, for example, Griliches (1988)

¹⁶This solution uses the same integral result from footnote 15.

¹⁷Notice from these results and from those in Table 4 that the convergence to steady state is faster for larger values of λ . Intuitively, recall that this holds γ constant. Therefore a larger value of λ corresponds to a smaller value of ϕ , which speeds up convergence to steady state.

Figure 5: The Transition of Multifactor Productivity to Steady State



Note: Log scale. A_t is calculated using equation (18) assuming $\gamma = .20$, $x_0 = .0146$, and $n = .01$. The dashed line reflects constant growth at a rate of 1.46 percent.

and the Bureau of Labor Statistics (1989) study, as well as the references cited in these papers. Most of these studies report a fairly small accounting contribution of R&D to growth, on the order of 0.2 percentage points per year. In this section, we discuss the relationship between these studies and the findings reported above.

In traditional growth accounting, R&D is treated as a second kind of capital investment: an R&D capital stock is constructed by cumulating past expenditures on R&D. The contribution of R&D to growth is then measured by the factor share of R&D multiplied by the growth rate of the stock of past expenditures.¹⁸ For example, if Z is the total stock of R&D, and if human capital is ignored, the growth accounting equation under such an approach

¹⁸The development below follows Griliches (1988).

looks like

$$\frac{\dot{Y}_t}{Y_t} = \alpha \frac{\dot{K}_t}{K_t} + (1 - \alpha) \frac{\dot{L}_t}{L_t} + \beta \frac{\dot{Z}_t}{Z_t} + \tilde{\mu}, \quad (19)$$

where $\tilde{\mu}$ is the growth rate of exogenous total factor productivity, and β is the elasticity of output with respect to R&D capital. The remaining notation parallels the model presented earlier.

Since $\beta = \frac{\partial Y}{\partial Z} \frac{Z}{Y}$, this equation can be simplified to

$$\frac{\dot{Y}_t}{Y_t} = \alpha \frac{\dot{K}_t}{K_t} + (1 - \alpha) \frac{\dot{L}_t}{L_t} + \rho \frac{R}{Y} + \tilde{\mu}, \quad (20)$$

where $\rho = \partial Y / \partial Z$ is the marginal product of R&D capital (the social return to R&D) and $R = \dot{Z}$ is the flow of R&D expenditures — net investment in R&D capital.

Assessing the impact of R&D on growth in this framework then involves measuring the social return to R&D and the net investment rate. A large number of studies have attempted to estimate these quantities, leading to a wide range of estimates. Griliches (1988) and the BLS (1989) study report social rates of return to R&D of 20 to 50 percent, or even higher; as a benchmark, the BLS chooses a value of 30 percent. The ratio of R&D expenditures to GDP in the United States is measured to be around 2 or 2.5 percent. Assuming no depreciation of R&D capital, this leads to a growth accounting contribution of anywhere between 0.4 and 1.25 percentage points per year.

These back of the envelope calculations yield larger contributions than the results typically reported in this literature. The BLS, for example, obtains estimates in the range of 0.1 to 0.4 percentage points, with a preferred estimated of about 0.2. The difference arises primarily from two sources. First, many of these studies focus on a narrower definition of R&D that excludes federally-funded research. Second, the BLS baseline estimate assumes a depreciation rate of about 8.8 percent; the highest number in the BLS range comes when zero depreciation is assumed, as above.

To compare these contributions to the results reported in Table 2 or 3, one must make one final adjustment. The accounting in Table 2 measures productivity in Harrod-neutral units so that steady-state growth in per capita income is equal to the Harrod-neutral productivity growth rate. To convert the Hicks-neutral units given above, one divides by labor's share of about $2/3$, so that the 0.4 to 1.25 range becomes 0.6 to 1.87. Viewed this way, the contribution reported in Table 2 of 1.46 percentage points or Table 3 of 1.29 percentage points is not inconsistent with the basic growth accounting methodology used in previous studies.¹⁹ Our results are at the upper end of existing estimates, but given the uncertainties surrounding the social return to R&D and the true output share of R&D investment, they are not implausible.²⁰

5.3 The Future of Economic Growth

It is difficult to observe results like those in Tables 2 and 3 without wanting to inquire further about the future of economic growth. A substantial fraction of the growth experienced over at least the last half century, and perhaps before, can be attributed to factors that cannot continue forever. Taken at face value, this suggests that future growth may be only a small fraction of recent growth.

There are a number of considerations and qualifications that need to be taken into account in interpreting this result. First, it turns out that the

¹⁹The 1.46 number is multifactor productivity growth in Table 2. The 1.29 number comes from adding the contributions of research intensity and the scale effect in Table 3.

²⁰One can also ask what social rate of return to R&D is implied by the methodology used to get the results in Table 2 or 3. This turns out to be more complicated. In the model in this paper, the stock of ideas is not simply an R&D capital stock. An additional dollar spent on hiring a researcher yields new ideas tomorrow that increase output, but these ideas may also affect the productivity of research in the future. Jones and Williams (1998) discuss how to measure the social rate of return to R&D in a model related to this one. Applying their methods to the model in this paper, one finds that the social rate of return is not uniquely pinned down by the estimate of γ : it depends on how λ , ϕ , and σ combine to generate the value of γ . However, one can find seemingly plausible values of these parameters that are consistent with a social return to R&D of 30 percent.

U.S. economy may have experienced a similar situation earlier in its history. David (1977) notes that much of nineteenth century U.S. growth was driven by a rising investment rate in physical capital and a corresponding rise in the capital-output ratio, which building on a term used by Hicks, he calls a “grand traverse.” Historically, this particular traverse came to an end, but growth rates did not decline, as other factors took over, namely educational attainment and research intensity. A similar change could occur again.

A second consideration relates to the production function for new ideas. The basic production function considered in this paper is $\dot{A} = \bar{\delta}(\cdot)H_A$, where $\bar{\delta}(\cdot)$ measures the productivity of research effort. Both because it is convenient and because it seems to fit, at least roughly, past experience, we have modeled research productivity as a simple Cobb-Douglas production function that potentially includes both research effort and the existing stock of knowledge as inputs.

It is somewhat natural to imagine that productivity in goods production is monotonically increasing: technologies get better and better over time. In this respect, the productivity of research effort may be very different. We do not know what the “universe” of ideas looks like. It could be that the discovery of past ideas makes future research more and more productive. Or this could be true, but only up to a point: the age of scientific discovery may accelerate right up until the end, and then end. Or perhaps the universe of ideas is laid out such that there are punctuated periods of discovery followed by periods of extremely slow, gradual advance. The point is that it is difficult and perhaps impossible to know the shape of the $\delta(\cdot)$ function, and this imposes sharp limits on our ability to make statements about future growth.

Third, this paper has so far taken population growth to be an exogenous constant. According to the model, the steady-state growth rate of the economy is proportional to the rate of population growth in the idea-producing economies. To the extent that population growth in these countries will

decline in the future, one would expect the long-run growth rate to decline as well.

Finally, consider the rise in educational attainment and research intensity. The latest grand traverse may well come to an end when the increases in these variables cease, but when will this occur? Both measures have been rising at least since 1950 and most probably since before the turn of the century. Measured research intensity is less than one percent of the labor force, so that the upper bound imposed by nature does not seem likely to bind in the near future.

6 Conclusion

This paper presents and calibrates a model of economic growth in a world of ideas. Growth in any particular country is driven in the long run by the implementation of ideas that are discovered throughout the world. In the long run, the stock of ideas is proportional to worldwide research effort, which in turn is proportional to the total population of innovating countries. In this sense, the model points out that the scale effect associated with the nonrivalry of ideas operates at a global level.

The model is employed to conduct two complementary growth accounting exercises and to understand some puzzling facts related to U.S. economic growth. While the per capita growth rate in the United States has been roughly constant on average during the last century, educational attainment and research intensity (both domestically and in the G-5 countries) have increased substantially. These facts are reconciled by highlighting the distinction between a constant growth path on the one hand and a balanced growth path or steady state on the other. It is possible for a sequence of transition dynamics to generate growth at a constant average rate that is higher than the steady-state rate.

The accounting exercises show that U.S. growth between 1950 and 1993

can be decomposed into three key pieces. The first two pieces — associated with transition dynamics — together account for more than 80 percent of growth in output per worker. The rising level of educational attainment accounts for more than one third of growth and increased research intensity in the G-5 countries accounts for about 50 percent of growth. Only about 10 to 20 percent is due to the component of growth associated with the rise in G-5 employment. While long-run growth is ultimately tied to world population growth in the model, more than 80 percent of U.S. growth in recent experience is due to other factors.

This exercise naturally raises questions about the future of U.S. and world economic growth. On the one hand, a plausible conjecture that could explain the rise in research intensity is the increased openness and development of the world economy. This explanation suggests that it is possible for the rise in research intensity to continue for some time into the future, both as the market for ideas continues to expand and as the share of the world's population that is sufficiently skilled to push the technological frontier forward continues to rise.

Still, it is important to recognize that this situation is unsustainable. In the long run, these changes must come to an end, and when this happens, U.S. growth rates can, *ceteris paribus*, be expected to fall considerably.

A Econometric Estimates of γ

This appendix discusses the econometric estimation of γ and justifies the range of 0.05 to 0.33 used in the accounting in Section 3.2. We begin by writing a discrete time version of the idea production function as an approximating model and by distinguishing between observed productivity and the unobserved stock of ideas:

$$\log B_t = \log A_t + \epsilon_t \quad (21)$$

$$\frac{\Delta A_{t+1}}{A_t} = \delta \left(\frac{H_{At}}{A_t^{1/\gamma}} \right)^\lambda. \quad (22)$$

Abusing notation for the purposes of this estimation, we now let B represent measured multifactor productivity and treat A as a latent variable. In (21), A is assumed to be related to multifactor productivity B through a stationary error term ϵ , which could represent measurement error. Equation (22) is simply the discrete time version of the idea production function, rewritten in terms of the growth rate of ideas and the parameter γ .

We will estimate γ econometrically in three different ways. The motivation for each of these approaches was discussed in Section 3.2. In particular, γ effectively detrends the ratio on the right side of equation (22) to produce a stationary productivity growth rate. Intuitively, it is estimated from the time trends in multifactor productivity and H_A , and therefore is (super) consistent.

To see this more formally, it is helpful to log-linearize equation (22) around a path where B_t and H_{At} are growing at constant rates and to write that equation in terms of multifactor productivity using equation (21). This gives a regression specification of the form

$$\Delta \log B_{t+1} \approx \beta_0 + \lambda g_B (\log H_{At} - \frac{1}{\gamma} \log B_t) + \eta_{t+1}, \quad (23)$$

where $\beta_0 \equiv g_B(1 - \log(g_B/\delta))$ is a constant and $\eta_{t+1} \equiv \Delta \epsilon_{t+1} + \frac{\lambda g_B}{\gamma} \epsilon_t$ is an error term.

At first glance, there are a number of problems to worry about in attempting to estimate γ (or $\beta_2 \equiv 1/\gamma$) using ordinary least squares in an equation like (23). First, one might worry about reverse causality using data at an annual frequency: presumably business cycle shocks that produce a boom and raise productivity make it easier for firms in the economy to undertake research. Second, because of measurement error, both $\log B_t$ and η_{t+1} depend on ϵ_t , providing another reason for a biased coefficient. Third, one might worry about the timing of the relationship between research and productivity growth. The approximating model here is written as if research today produces productivity growth next year, but presumably a more complicated distributed lag of research is appropriate (although this is mitigated somewhat by the AR(1) style of the specification, ultimately associated with ϕ).

All of these valid concerns turn out to be addressed by the fact that γ is estimated from the linear time trends in $\log H_{At}$ and $\log B_t$. To see why, notice that equation (23) looks like a standard error-correction model in time series analysis. The coefficient $\beta_2 \equiv 1/\gamma$ is a standard cointegrating coefficient that detrends $\log H_{At} - \beta_2 \log B_t$ to produce stationary productivity growth. Moreover, because there is a linear time trend in $\log H_{At}$, and therefore in $\log B_t$, the coefficient β_2 can be estimated in a straightforward and robust fashion. The results of West (1988), for example, imply that the time trends drive the asymptotic distribution of the OLS estimator of β_2 , so that the estimator is consistent and has a normal distribution. The estimates are surely robust to reverse causality associated with business cycles and to stationary measurement error (for example $\log B_t$ and ϵ_t have a zero population correlation because of the time trend in $\log B_t$). For the same reason, changing the lag structure has very little effect on the coefficients: the time trend in $\log B_t$ is the same as the time trend in $\log B_{t+1}$.

One way to estimate β_2 (and, using the delta method, γ) is simply to run the OLS regression of $\log H_{At}$ on $\log B_t$. This produces an OLS estimate

of γ of 0.323, with a Newey-West robust standard error of 0.019. Running the reverse regression of $\log B_t$ on $\log H_{At}$ produces an alternative estimate of 0.313 with a Newey-West robust standard error of 0.014. These estimates formalize the intuition given in the text in Section 3.2 where we “estimated” γ as the ratio of the growth rates of multifactor productivity and H_A .

In principle, one could stop here. However, because of the productivity slowdown, estimating the error correction model directly proves to be worthwhile. In particular, because $\Delta \log B_{t+1}$ actually has a slight downward trend, γ needs to leave a slightly negative trend in the cointegrating term $\log H_{At} - \beta_2 \log B_t$. The first section of Table 5 reports the results from estimating equation (23) using OLS. The first specification in the Table shows that OLS produces an implausibly large estimate of λ of 4.5. This should not be particularly surprising: the estimate of λ is not driven by time trends and so the OLS estimator is subject to bias because of measurement error and endogeneity. Clearly, the business cycle effects that stimulate both productivity growth and research are dominating here. The remaining three specifications address this bias by imposing the range of reasonable values for λ and then estimating γ subject to this restriction. Recall that λ captures the importance of decreasing returns to research at a point in time: if we double the number of researchers today, the number of new ideas produced today by those researchers rises by 2^λ . It seems reasonable to assume that λ is somewhere between a maximum value of 1.0 and a minimum value of about .25. The estimates of γ then range from 0.178 when $\lambda = 1$ to 0.076 when $\lambda = .25$. As before, the standard errors for these estimates are quite small because the estimates are driven by the time trends in the data.

Finally, one may naturally wonder about the validity of the log-linear approximation. Here, the approximation is convenient primarily because we understand the asymptotic distribution theory of OLS estimators of models like that in equation (23). In contrast, I am not aware of distribution theory for the estimators of the full nonlinear model in equations (21) and (22).

Table 5: Estimating γ , 1950-1993

Parameter	Specification			
	(1)	(2)	(3)	(4)
<i>Log-Linearized Model</i>				
λ	4.535 (2.43)	1.00 —	0.50 —	0.25 —
γ	0.274 (.128)	.178 (.046)	.123 (.044)	0.076 (.033)
R^2	.219	.150	.129	.118
<i>Nonlinear Least Squares</i>				
λ	4.997	1.00	0.50	0.25
γ	0.292	.191	.133	0.083
R^2	.980	.975	.973	.972

Note: Results from estimating equation (23) for the log-linear model and equations (21) and (22) for the nonlinear model. Newey-West robust standard errors are in parentheses. In specifications (2) through (4), specific values of λ are imposed. In the nonlinear model, δ and A_0 are additional parameters that are estimated.

Intuitively, the estimate of γ is still driven by the exponential trends in H_A and B , and therefore one imagines that nonlinear least squares estimation of the model has desirable properties. These estimates are reported in the bottom half of Table 5.²¹ The estimates are quite similar to those in the log-linear model, ranging from a low of 0.083 when $\lambda = .25$ to a high of 0.191 when $\lambda = 1$.

These econometric results suggest a range of estimates for γ that starts at 0.076 at the lower end and reaches 0.323 at the upper end (for our simple regressions of productivity on research). The economic uncertainty regarding the correct specification dominates the sampling uncertainty reflected in the standard errors of any particular estimate. This suggests that the range of 0.05 to 0.33 used in the text very likely includes the true value of γ .

B Data Appendix

The data used in this paper are taken from several different sources. Many of the sources are now available on-line, and the actual data series that I have used are available from the “data sets” section of <http://elsa.berkeley.edu/~chad>.

- *GDP per Hour*. Data on real GDP in chained 1996 dollars are taken from Table 2A, page 130, of Bureau of Economic Analysis (2000), available on the web at <http://www.bea.doc.gov/bea/dn1.htm>. Employment data are from Table B-33 of the *Economic Report of the President, 1997*. Employment is converted to total hours using the Average Weekly Hours of Production Workers (series EEU00500005) for total private industry from the Bureau of Labor Statistics and assuming a constant work year of 50 weeks. The hours data was downloaded from the “National Employment, Hours, and Earnings” link at <http://www.bls.gov/top20.html>. All other employment data used in

²¹Standard errors are omitted because of the lack of a distribution theory. For what it is worth, they were small and looked very much like the errors in the first half of the table.

this paper are in units of bodies rather than hours worked, the implicit assumption being that the hours worked are uniform across categories and countries.

- *Educational Attainment.* Average educational attainment in the population among persons 25 years old and over is calculated from Bureau of the Census (1996), Table 17 (Historical). This source reports the number of persons by “cells” of educational attainment. In computing the average, I assume that each person in a cell has the mean years of schooling for the cell (e.g. persons in the cell corresponding to one to three years of high school are assumed to have 10 years of schooling). Persons in the “four or more years of college” cell, the top cell, are assumed to have four years of college. Missing data are linearly interpolated. This data is available on-line at <http://www.census.gov/population/www/socdemo/education.htm>.
- *Scientists and Engineers Engaged in R&D.* The data for 1965 to 1993 are from National Science Board (1998) and National Science Board (1993). For years prior to 1965 for the U.S., data from the *Historical Statistics of the United States, Colonial Times to 1970* and various editions of the *Statistical Abstract of the United States* are used. Missing data are log-linearly interpolated. National Science Board (1998) is now available on-line at <http://www.nsf.gov/sbe/srs/seind98/start.htm>. For years prior to 1965 for France, Germany, Japan, and the United Kingdom, we assume that the ratio of research intensity between these countries and the United States in 1950 is the same as in 1965. Next, research intensity for intervening years is linearly interpolated for each country and then multiplied by employment (see below) to get an estimate for scientists and engineers engaged in R&D. This data is only used to construct the aggregate G-5 researchers and research intensity; it is not used on a country-by-country basis.

- *Employment.* Data on U.S. employment are from the *Economic Report of the President* (1997), Table B-33, Employed Civilian Labor Force. Employment data for the remaining G-5 countries for 1959 to 1993 are from Bureau of Labor Statistics, U.S. Department of Labor (2000), Table 2, available online at <ftp://ftp.bls.gov/pub/special.requests/ForeignLabor/flslforc.txt>. This data is spliced (using the 1959 observation) onto data on the number of workers from the Penn World Tables Mark 5.6 for 1950 to 1958; see Summers and Heston (1991).
- *Physical Capital.* Data on the real net stock of physical capital in chained 1992 dollars are from “Improved Estimates of Fixed Reproducible Tangible Wealth, 1929–95”, prepared by Katz and Herman (1997), available online at <http://www.bea.doc.gov/bea/an/0597niw/-maintext.htm>.
- *Investment.* Data on real investment in chained 1996 dollars are taken from Table 2A, page 130, of Bureau of Economic Analysis (2000), available on the web at <http://www.bea.doc.gov/bea/dn1.htm>. Note: unlike the capital data, the investment data does not include government investment.

References

- Aghion, Philippe and Peter Howitt**, “A Model of Growth through Creative Destruction,” *Econometrica*, March 1992, 60 (2), 323–351.
- Anderson, Robert N.**, “United States Life Tables, 1997,” *National Vital Statistics Reports*, December 13 1999, 47 (28). National Center for Health Statistics.
- Ben-David, Dan and David H. Papell**, “The Great Wars, the Great Crash, and Steady-State Growth: Some New Evidence about an Old Stylized Fact,” *Journal of Monetary Economics*, December 1995, pp. 453–475.
- Bils, Mark and Peter Klenow**, “Does Schooling Cause Growth?,” *American Economic Review*, December 2000, 90, 1160–1183.
- Bureau of Economic Analysis**, “GDP and Other Major NIPA Series, 1929–99,” *Survey of Current Business*, April 2000, pp. 126–145.
- Bureau of Labor Statistics, U.S. Department of Labor**, “The Impact of Research and Development on Productivity Growth,” September 1989. Bulletin 2331.
- , “Comparative Civilian Labor Force Statistics, 10 Countries, 1959–1999,” 2000. April 17.
- Bureau of the Census**, *Historical Statistics of the United States*, Washington D.C.: U.S. Government Printing Office, 1975.
- , “Current Population Reports: Educational Attainment in the United States: March 1995,” August 1996. P20-489.
- Chari, V.V., Pat Kehoe, and Ellen McGrattan**, “The Poverty of Nations: A Quantitative Investigation,” 1997. Working Paper, Federal Reserve Bank of Minneapolis.
- Council of Economic Advisers**, *Economic Report of the President*, Washington, D.C.: U.S. Government Printing Office, 1997.

- David, Paul A.**, “Invention and Accumulation in America’s Economic Growth: A Nineteenth-Century Parable,” *Journal of Monetary Economics*, 1977, 6 (0), 176–228. Special Supplement.
- , “Knowledge, Property, and the System Dynamics of Technological Change,” *Proceedings of the World Bank Annual Conference on Development Economics, 1992*, 1993, pp. 215–248.
- Denison, Edward F.**, *The Sources of Economic Growth in the United States and the Alternatives Before Us*, New York: Committee for Economic Development, 1962.
- Eaton, Jonathan and Samuel S. Kortum**, “International Technology Diffusion: Theory and Measurement,” *International Economic Review*, 1999, 40, 537–570.
- Gendell, Murray and Jacob S. Siegel**, “Trends in retirement age by sex, 1950–2005,” *Monthly Labor Review*, July 1992, pp. 22–29.
- Goldin, Claudia**, “A Brief History of Education in the United States,” August 1999. NBER Historical Paper No. 119.
- Griliches, Zvi**, “Productivity Puzzles and R&D: Another Nonexplanation,” *Journal of Economic Perspectives*, 1988, 2, 9–21.
- Grossman, Gene M. and Elhanan Helpman**, *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press, 1991.
- Jones, Charles I.**, “R&D-Based Models of Economic Growth,” *Journal of Political Economy*, August 1995, 103 (4), 759–784.
- , “Time Series Tests of Endogenous Growth Models,” *Quarterly Journal of Economics*, May 1995, 110 (441), 495–525.
- and **John C. Williams**, “Measuring the Social Return to R&D,” *Quarterly Journal of Economics*, November 1998, 113 (455), 1119–1135.
- Kaldor, Nicholas**, “Capital Accumulation and Economic Growth,” in F.A. Lutz and D.C. Hague, eds., *The Theory of Capital*, St. Martins Press, 1961, pp. 177–222.

Katz, Arnold J. and Shelby W. Herman, “Improved Estimates of Fixed Reproducible Tangible Wealth, 1929–95,” *Survey of Current Business*, May 1997, p. 69ff.

Klenow, Peter and Andres Rodriguez-Clare, “The Neoclassical Revival in Growth Economics: Has It Gone Too Far?,” in Ben S. Bernanke and Julio J. Rotemberg, eds., *NBER Macroeconomics Annual 1997*, Cambridge, MA: MIT Press, 1997.

Maddison, Angus, *Monitoring the World Economy 1820-1992*, Paris: Organization for Economic Cooperation and Development, 1995.

Mankiw, N. Gregory, David Romer, and David Weil, “A Contribution to the Empirics of Economic Growth,” *Quarterly Journal of Economics*, May 1992, 107 (2), 407–438.

McGrattan, Ellen and James Schmitz, “Explaining Cross-Country Income Differences,” in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, North-Holland Elsevier Science, 1999.

Mincer, Jacob, *Schooling, Experience, and Earnings*, New York: Columbia University Press, 1974.

National Science Board, *Science & Engineering Indicators – 1993*, Washington D.C.: U.S. Government Printing Office, 1993.

—, *Science & Engineering Indicators – 1998*, Washington D.C.: U.S. Government Printing Office, 1998.

Nordhaus, William D., “An Economic Theory of Technological Change,” *American Economic Association Papers and Proceedings*, May 1969, 59, 18–28.

Phelps, Edmund S., “Models of Technical Progress and the Golden Rule of Research,” *Review of Economic Studies*, April 1966, 33 (2), 133–45.

Quah, Danny T., “The Invisible Hand and the Weightless Economy,” April 1996. LSE Economics Department mimeo.

Romer, Paul M., “Endogenous Technological Change,” *Journal of Political Economy*, October 1990, 98 (5), S71–S102.

—, “Should the Government Subsidize Supply or Demand in the Market for Scientists and Engineers?,” June 2000. NBER Working Paper No. 7723.

Shell, Karl, “Toward a Theory of Inventive Activity and Capital Accumulation,” *American Economic Association Papers and Proceedings*, 1966, 56, 62–68.

Simon, Julian L., *Theory of Population and Economic Growth*, New York: Basil Blackwell, 1986.

Solow, Robert M., “A Contribution to the Theory of Economic Growth,” *Quarterly Journal of Economics*, February 1956, 70 (1), 65–94.

—, “Technical Change and the Aggregate Production Function,” *Review of Economics and Statistics*, August 1957, 39 (3), 312–320.

Summers, Robert and Alan Heston, “The Penn World Table (Mark 5): An Expanded Set of International Comparisons: 1950–1988,” *Quarterly Journal of Economics*, May 1991, 106, 327–368.

West, Kenneth, “Asymptotic Normality When Regressors Have a Unit Root,” *Econometrica*, 1988, 56, 1397–1417.