The Facts of Economic Growth

Charles I. Jones*

Stanford GSB and NBER

April 6, 2015 – Version 0.4

Preliminary and Incomplete

Abstract

Why are people in the richest countries of the world so much richer today than 100 years ago? And why are some countries so much richer than others? Questions such as these define the field of economic growth. This paper documents the facts that underlie these questions. How much richer are we than 100 years ago, and how large are the income gaps between countries? The purpose of the paper is to provide an encyclopedia of the fundamental facts of economic growth upon which our theories are built, gathering them together in one place and updating the facts with the latest available data.

*If you know of a fact that would fit well in this paper, I'd love to hear about it. The paper is definitely a work in progress...

> "[T]he errors which arise from the absence of facts are far more numerous and more durable than those which result from unsound reasoning respecting true data." — Charles Babbage, quoted in Rosenberg (1994), p. 27.

> "[I]t is quite wrong to try founding a theory on observable magnitudes alone... It is the theory which decides what we can observe." — Albert Einstein, quoted in Heisenberg (1971), p. 63.

Why are people in the United States, Germany, and Japan so much richer today than 100 or 1000 years ago? Why are people in France and the Netherlands today so much richer than people in Haiti and Kenya? Questions like these are at the heart of the study of economic growth.

Economics seeks to answer these questions by building quantitative models — models that can be compared with empirical data. That is, we'd like our models to tell us not only that one country will be richer than another, but by how much. Or to explain not only that we should be richer today than a century ago, but that the growth rate should be 2 percent per year rather than 10 percent. Growth economics has only partially achieved these goals, but a critical input into our analysis is knowing where the goalposts lie — that is, knowing the facts of economic growth.

The goal of this paper is to lay out as many of these facts as possible. Kaldor (1961) was content with documenting a few key stylized facts that basic growth theory should hope to explain. Jones and Romer (2010) updated this list to reflect what we've learned over the last 50 years. The purpose of this paper is different. Rather than highlighting a handful of stylized facts, we draw on the last thirty years of the renaissance of growth economics to lay out what is known empirically about the subject. These facts are updated with the latest data and gathered together in a single place. The hope is that they will be useful to newcomers to the field as well as to experts. The result, I hope, is a fascinating tour of the growth literature, from the perspective of the basic data.

The paper is divided broadly into two parts. First, I present the facts related to the growth of the "frontier" over time: what are the growth patterns exhibited by the richest countries in the world? Second, I focus on the spread of economic growth throughout

Figure 1: GDP per person in the United States

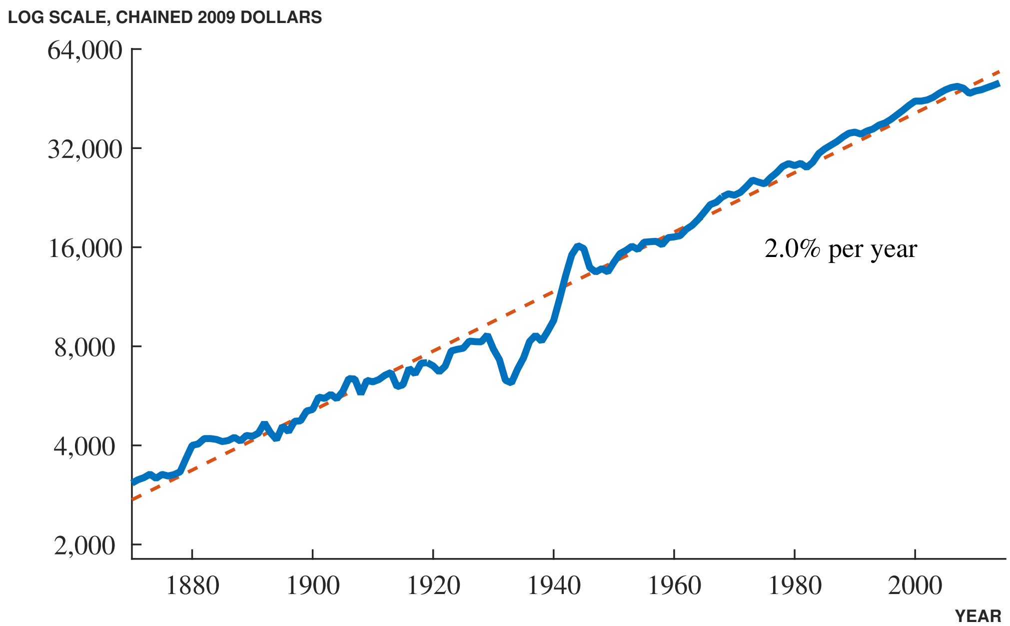

Note: Data for 1929–2014 are from the U.S. Bureau of Economic Analysis, NIPA Table 7.1. Data before 1929 are spliced from Maddison (2008).

the world. To what extent are countries behind the frontier catching up, falling behind, or staying in place? And what characteristics do countries in these various groups share?

1. Growth at the Frontier

We begin by discussing economic growth at the "frontier." By this I mean growth among the richest set of countries in any given time period. During the last century or so, the United States often serves as a stand in for the frontier, and we will follow this tradition.

1.1. Modern Economic Growth

Figure 1 shows one of the key stylized facts of frontier growth: For nearly 150 years, GDP per person in the U.S. economy has grown at a remarkably steady average rate of around 2 percent per year. Starting at around $3,000 in 1870, per capita GDP rose to more than $50,000 by 2014, a nearly 17-fold increase.

Beyond the large, sustained growth in living standards, several other features of this graph stand out. One is the significant decline in income associated with the Great Depression. However, to me this decline stands out most for how anomalous it is. Many of the other recessions barely make an impression on the eye: over long periods of time, economic growth swamps economic fluctuations. Moreover, despite the singular severity of the Great Depression — GDP per person fell by nearly 20 percent in just four years — it is equally remarkable that the Great Depression was *temporary*. By 1939, the economy is already passing its previous peak and the macroeconomic story a decade later is once again one of sustained, almost relentless, economic growth.

The stability of U.S. growth also merits some discussion. With the aid of the trend line in Figure 1, one can see that growth was slightly slower pre-1929 than post. Table 1 makes this point more precisely. Between 1870 and 1929, growth averaged 1.76 percent, versus 2.23 percent between 1929 and 2007 (using "peak to peak" dates to avoid business cycle problems). Alternatively, between 1900 and 1950, growth averaged 2.06 percent versus 2.16 percent since 1950. Before one is too quick to conclude that growth rates are increasing, however, notice that the period since 1950 shows a more mixed pattern, with rapid growth between 1950 and 1973, slower growth between 1973 and 1995, and then rapid growth during the late 1990s that gives way to slower growth more recently.

The interesting "trees" that one sees in Table 1 serves to support the main point one gets from looking at the "forest" in Figure 1: steady, sustained exponential growth for the last 150 years or so is a key characteristic of the frontier. All modern theories of economic growth — Solow (1956), Lucas (1988), Romer (1990), Aghion and Howitt (1992), for example — are designed with this fact in mind.

The sustained growth in Figure 1 also naturally raises the question of whether such growth can and will continue for the next century. On the one hand, this fact more than any other helps justify the focus of many growth models on the balanced growth path, a situation in which all economic variables grow at constant exponential rates forever. And the logic of the balanced growth path suggests that the growth can continue indefinitely. On the other hand, as we will see, there are reasons from other facts and theories to question this logic.

Table 1: The Stability of U.S. Growth

Period	Growth Rate	Period	Growth Rate
1870–2007	2.03	1973–1995	1.82
1870–1929	1.76	1995–2007	2.13
1929–2007	2.23		
1900–1950	2.06	1995–2001	2.55
1950–2007	2.16	2001–2007	1.72
1950–1973	2.50		
1973–2007	1.93		

Note: Annualized growth rates for the data shown in Figure 1.

1.2. Growth over the Very Long Run

While the future of frontier growth is surely hard to know, the stability of frontier growth suggested by Figure 1 is most certainly misleading when look back further in history. As shown in Figure 2, sustained exponential growth in living standards is an incredibly recent phenomenon. For thousands and thousands of years, life was, in the evocative language of Thomas Hobbes, "nasty, brutish, and short." Only in the last two centuries has this changed, but in this relatively brief time, the change has been dramatic.[1]

Between the year 1 C.E. and the year 1820, living standards in the "West" (measured with data from Western Europe and the United States) essentially doubled, from around $600 per person to around $1200 per person, as shown in Table 2. Over the next 200 years, however, GDP per person rose by more than a factor of twenty, reaching $26,000.

The era of modern economic growth is in fact even more special than this. Evidence suggests that living standards were comparatively stagnant for thousands and thousands of years before. For example, for much of pre-history, humans lived as simple

[1]Papers that played a key role in documenting and elaborating upon this fact include Maddison (1979), Kremer (1993), Maddison (1995), Diamond (1997), Pritchett (1997), and Clark (2001). This list neglects a long, important literature in economic history; see Clark (2014) for a more complete list of references.

Figure 2: Economic Growth over the Very Long Run

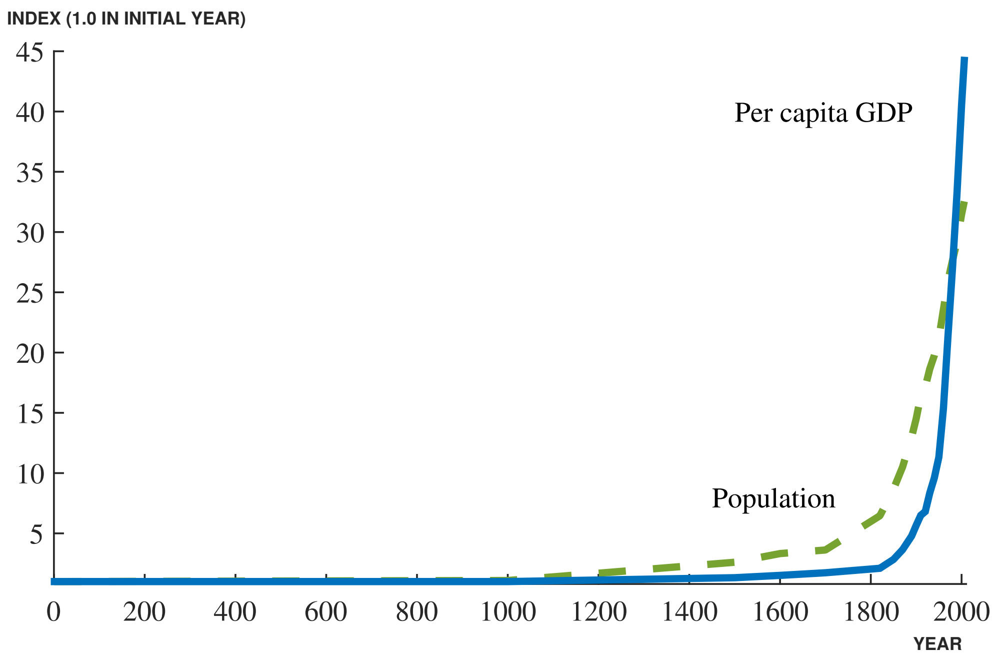

Note: Data are from Maddison (2008) for the "West," i.e. Western Europe plus the United States. A similar pattern holds using the "world" numbers from Maddison.

Table 2: The Acceleration of World Growth

Year	GDP per person	Growth rate	Population (millions)	Growth rate
1	590	...	19	...
1000	420	-0.03	21	0.01
1500	780	0.12	50	0.17
1820	1,240	0.15	125	0.28
1900	3,350	1.24	280	1.01
2006	26,200	1.94	627	0.76

Note: Data are from Maddison (2008) for the "West," i.e. Western Europe plus the United States. Growth rates are average annual growth rates in percent, and GDP per person is measured in real 1990 dollars.

hunters and gatherers, on the edge of subsistence. From this perspective — say for the last 200,000 years or more — the era of modern growth is spectacularly brief. It is the economic equivalent of Carl Sagan's famous "pale blue dot" image of the earth viewed from the outer edge of the solar system.

Table 2 reveals several other interesting facts. First and foremost, over the very long run, economic growth at the frontier has accelerated — that is, the rates of economic growth are themselves increasing over time. Romer (1986) emphasized this fact for living standards as part of his early motivation for endogenous growth models. Kremer (1993) highlighted the acceleration in population growth rates, dating as far back as a million years ago, and his evidence serves as a very useful reminder. Between 1 million B.C.E. and 10,000 B.C.E., the average population growth rate in Kremer's data was 0.00035 percent per year. Yet despite this tiny growth rate, world population increased by a factor of 32, from around 125,000 people to 4 million. As an interesting comparison, that's similar to the proportionate increase in the population in Western Europe and the United States during the past 2000 years, shown in Table 2.

Various growth models have been developed to explain the transition from stagnant living standards for thousands of years to the modern era of economic growth. A key ingredient in nearly all of these models is Malthusian diminishing returns. In particular, there is assumed to be a fixed supply of land which is a necessary input in production. Adding more people to the land reduces the marginal product of labor (holding technology constant) and therefore reduces living standards. Combined with some subsistence level of consumption below which people cannot survive, this ties the size of the population to the level of technology in the economy: a better technology can support a larger population.

Various models then combine the Malthusian channel with different mechanisms for generating growth. Lee (1988), Kremer (1993), and Jones (2001) emphasize the positive feedback loop between "people produce ideas" as in the Romer models of growth with the Malthusian "ideas produce people" channel. Provided the increasing returns associated with ideas is sufficiently strong to counter the Malthusian diminishing returns, this mechanism can give rise to dynamics like those shown in Figure 2. Lucas (2002) emphasizes the role of human capital accumulation, while Hansen and Prescott (2002) focus on a neoclassical model that features a structural transformation from agriculture to manufacturing. Oded Galor, with his coauthors, has been one of the most significant contributors, labeling this literature "unified growth theory." See Galor and Weil (2000) and Galor (2005).

2. Sources of Frontier Growth

The next collection of facts related to economic growth are best presented in the context of the famous growth accounting decomposition developed by Solow (1957) and others. This exercise studies the sources of growth in the economy through the lens of a single aggregate production function. It is well-known that the conditions for an aggregate production function to exist in an environment with a rich underlying microstructure are very stringent. The point is not that anyone believes those conditions hold. Instead, one often wishes to look at the data "through the lens of" some growth model that is much simpler than the world that generates the observed data. A long list of famous papers supports the claim that this is a productive approach to gaining knowledge, Solow (1957) itself being an obvious example.

While not necessary, it is convenient to explain this accounting using a Cobb-Douglas specification. More specifically, suppose final output Y_t is produced using stocks of physical capital K_t and human capital H_t:

$$Y_t = \underbrace{A_t M_t}_{\text{TFP}} K_t^{\alpha} H_t^{1-\alpha} \tag{1}$$

where α is between zero and one, A_t denotes the economy's stock of knowledge, and M_t is anything else that influences total factor productivity (the letter "M" is reminiscent of the "measure of our ignorance" label applied to the residual by Abramovitz (1956) and also is suggestive of "misallocation," as will be discussed in more detail later). The next subsection provides a general overview of growth accounting for the United States based on this equation, and then the remainder of this section looks more closely at each individual term in equation (1).

2.1. Growth accounting

It is traditional to perform the growth accounting exercise with a production function like (1). However, that approach creates some confusion in that some of the accumulation of physical capital is caused by growth in total factor productivity (e.g. as in a standard Solow model). If one wishes to credit such growth to total factor productivity, it is helpful to do the accounting in a slightly different way.[2] In particular, divide both sides of the production function by Y_t^{α} and solve for Y_t to get

$$Y_t = \left(\frac{K_t}{Y_t}\right)^{\frac{\alpha}{1-\alpha}} H_t Z_t \tag{2}$$

where $Z_t \equiv (A_t M_t)^{\frac{1}{1-\alpha}}$ is total factor productivity measured in labor-augmenting units. Finally, dividing both sizes by the aggregate amount of time worked, L_t, gives

$$\frac{Y_t}{L_t} = \left(\frac{K_t}{Y_t}\right)^{\frac{\alpha}{1-\alpha}} \frac{H_t}{L_t} \cdot Z_t \tag{3}$$

In this form, growth in output per hour Y_t/L_t comes from growth in the capital-output ratio K_t/Y_t, growth in human capital per hour H_t/L_t, and growth in labor-augmenting TFP, Z_t. This can be seen explicitly by taking logs and differencing equation (3). Also, notice that in a neoclassical growth model, the capital-output ratio is proportional the

[2] Klenow and Rodriguez-Clare (1997), for example, takes this approach.

the investment rate in the long-run and does not depend on total factor productivity. Hence the contributions from productivity and capital deepening are separated in this version, in a way that they were not in equation (1).

The only term we have yet to comment on is H_t/L_t, the aggregate amount of human capital divided by total hours worked. In a simple model with one type of labor, one can think of $H_t = h_t L_t$, where h_t is human capital per worker which increases because of education. In a richer setting with different types of labor that are perfect substitutes when measured in efficiency units, H_t/L_t also captures composition effects. The Bureau of Labor Statistics, from which I've obtained the accounting numbers discussed next, therefore refers to this term as "labor composition."

Table 3 contains the growth accounting decomposition for the United States since 1948, corresponding to equation (3). Several well-known facts emerge from this accounting. First, growth in output per hour at 2.5 percent is slightly faster than the growth in GDP per person that we saw earlier. One reason is that the BLS data measure growth for the private business sector, excluding the government sector (in which there is zero productivity growth more or less by assumption). Second, the capital-output ratio is relatively stable over this period, contributing almost nothing to growth. Third, labor composition (a rise in educational attainment, a shift from manufacturing to services, and the increased labor force participation of women) contributes 0.3 percentage points to growth. Finally, as documented by Abramovitz, Solow, and others, the "residual" of total factor productivity accounts for the bulk of growth, coming in at 2.0 percentage points, or 80 percent of growth since 1948.

The remainder of Table 3 shows the evolution of growth and its decomposition over various periods since 1948. We see the rapid growth and rapid TFP growth of the 1948–1973 period, followed by the well-known "productivity slowdown" from 1973 to 1995. The causes of this slowdown are much debated but not convincingly pinned down, as suggested by the fact that the entirety of the slowdown comes from the TFP residual rather than from physical or human capital; Griliches (1988) contains a discussion of the slowdown.

Remarkably, the period 1995–2007 sees a substantial recovery of growth, not quite to the rates seen in the 1950s and 1960s, but impressive nonetheless, coinciding with the dot-com boom and the rise in the importance of information technology. Byrne, Oliner

Table 3: Growth Accounting for the United States

Period	Output per hour	Contributions from: K/Y	Labor Composition	Labor-Aug. TFP
1948–2013	**2.5**	**0.1**	**0.3**	**2.0**
1948–1973	3.3	-0.2	0.3	3.2
1973–1990	1.6	0.5	0.3	0.8
1990–1995	1.6	0.2	0.7	0.7
1995–2000	3.0	0.3	0.3	2.3
2000–2007	2.7	0.2	0.3	2.2
2007–2013	1.7	0.1	0.5	1.1

Note: Average annual growth rates (in percent) for output per hour and its components for the private business sector, following equation (3). Source: Authors calculations using Bureau of Labor Statistics, *Multifactor Productivity Trends*, August 21, 2014.

and Sichel (2013) provide a recent analysis of the importance of information technology to growth over this period and going forward. Lackluster growth in output per hour since 2007 is surely in large part attributable to the Great Recession, but the slowdown in TFP growth (which some such as Fernald (2014) date back to 2003) is troubling.

2.2. Physical capital

The fact that the contribution of the capital-output ratio was modest in the growth accounting decomposition suggests that the capital-output ratio is relatively constant over time. This suggestion is confirmed in Figure 3. The broadest concept of physical capital ("Total"), including both public and private capital as well as both residential and non-residential capital, has a ratio of 3 to real GDP. Focusing on non-residential capital brings this ratio down to 2, and further restricting to private non-residential capital leads a ratio of just over 1.

The capital stock is itself the cumulation of investment, adjusted for depreciation. Figure 4 shows nominal spending on investment as a share of GDP back to 1929. The share is relatively stable for much of the period, with a notable decline during the last

Figure 3: The Ratio of Physical Capital to GDP

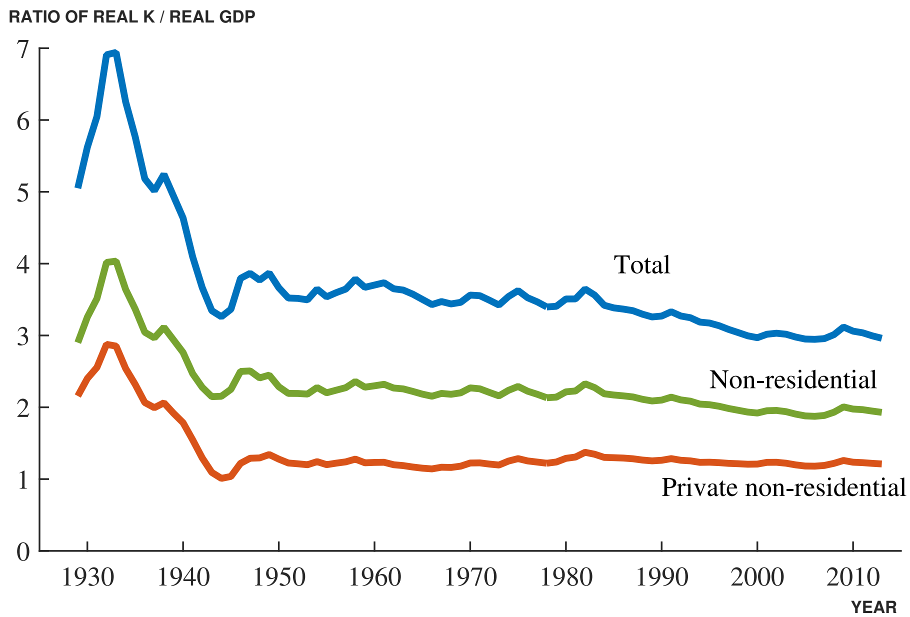

Source: Burea of Economic Analysis Fixed Assets Tables 1.1 and 1.2. The numerator in each case is a different measure of the real stock of physical capital, while the denominator is real GDP.

two decades.

In addition to cumulating investment, however, another step in going from the (nominal) investment rate series to the (real) capital-output ratio involves adjusting for relative prices. Figure 5 shows the price of various categories of investment, relative to the GDP deflator. Two facts stand out: the relative price of equipment has fallen sharply since 1960 by more than a factor of 3 and the relative price of structures has risen since 1929 by a factor of 2 (for residential) or 3 (for non-residential).

A fascinating observation comes from comparing the trends in the relative prices shown in Figure 5 to the investment shares in Figure 4: the nominal investment shares are relatively stable when compared to the huge trends in relative prices. For example, even though the relative price of equipment has fallen by more than a factor of 3 since 1960, the nominal share of GDP spent on equipment has remained steady.

The decline in the fall of equipment prices has featured prominently in parts of the growth literature; for example, see Greenwood, Hercowitz and Krusell (1997) and Whelan (2003). These papers make the point that one way to reconcile the facts is

Figure 4: Investment in Physical Capital (Private and Public), United States

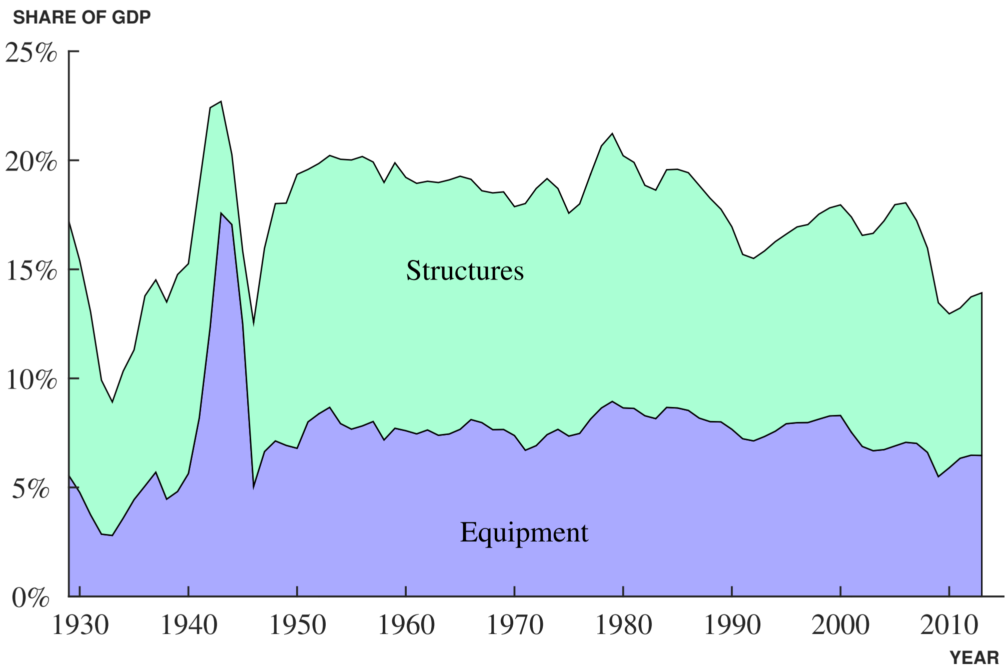

Source: National Income and Product Accounts, U.S. Bureau of Economic Analysis, Table 5.2.5. Intellectual property products and inventories are excluded. Government and private investment are combined. Structures includes both residential and nonresidential investment. Ratios of nominal investment to GDP are shown.

Figure 5: Relative Price of Investment, United States

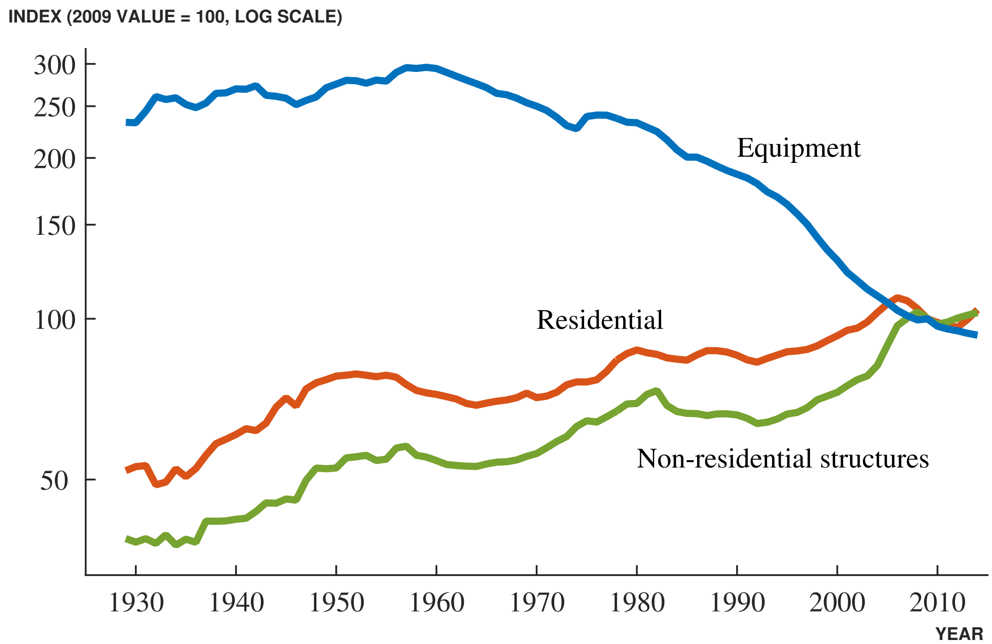

Note: The chained price index for various categories of private investment is divided by the chained price index for GDP. Source: National Income and Product Accounts, U.S. Bureau of Economic Analysis Table 1.1.4.

with a two sector model in which technological progress in the equipment sector is substantially factor that technological progress in the rest of the economy — an assumption that rings true in light of Moore's Law and the tremendous decline in the price of a semiconductors. Combining this assumption with Cobb-Douglas production functions leads to a two-sector model that is broadly consistent with the facts we've laid out. A key assumption in this approach is that better computers are equivalent to having more of the old computers, so that technological change is, in a rough sense, capital (equipment)-augmenting. The Cobb-Douglas assumption ensures that this non-labor augmenting technological change can coexist with a balanced growth path and delivers a stable nominal investment rate.[3]

2.3. Factor Shares

One of the original Kaldor (1961) stylized facts of growth was the stability of the shares of GDP paid to capital and labor. Figure 6 shows these shares using two different data sets, but the patterns are quite similar. First, between 1948 and 2000, the factor shares were indeed quite stable. Second, since 2000 or so, there has been a marked decline in the labor share and a corresponding rise in the capital share. According to the data from the Bureau of Labor Statistics, the capital share rose from an average value of 34.2% between 1948 and 2000 to a value of 38.7% by 2012. Or in terms of the complement, the labor share declined from an average value of 65.8% to 61.3%.

It is hard to know what to make of the recent movements in factor shares. Is this a temporary phenomenon, perhaps amplified by the Great Recession? Or are some more deeper structural factors at work? Karabarbounis and Neiman (2014) document that the fact extends to many countries around the world and perhaps on average starts even before 2000. Other recent papers looking at this question include Elsby, Hobijn and Şahin (2013), and Bridgman (2014).

2.4. Human capital

The other major neoclassical input in production is human capital. Figure 7 shows a time series for one of the key forms of human capital in the economy, education. More

[3] This discussion is related to the famous Uzawa theorem about the restrictions on technical change required to obtain balanced growth; see Schlicht (2006) and Jones and Scrimgeour (2008).

Figure 6: Capital and Labor Shares of Factor Payments, United States

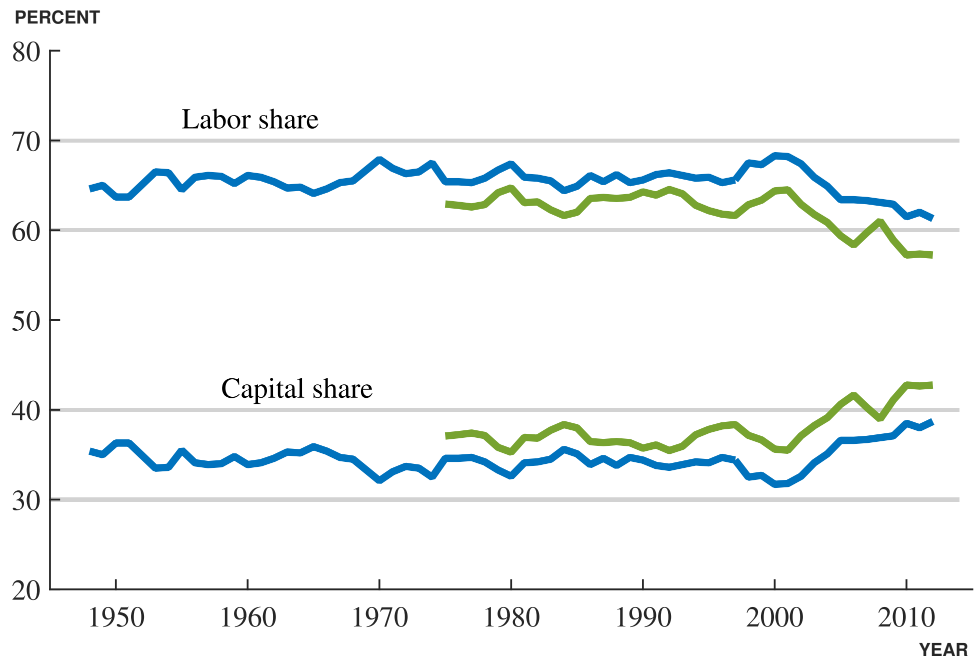

Note: The series starting in 1975 are from Karabarbounis and Neiman (2014) and measure the factor shares for the corporate sector, which the authors argue is helpful in eliminating issues related to self-employment. The series starting in 1948 is from the Bureau of Labor Statistics *Multifactor Productivity Trends*, August 21, 2014, for the private business sector. The factor shares add to 100 percent.

specifically, the graph shows educational attainment by birth cohort, starting with the cohort born in 1875.

Two facts emerge. First, for 75 years, educational attainment rose steadily, at a rate of slightly less than one year per decade. For example, the cohort born in 1880 got just over 7 years of education, while the cohort born in 1950 received 13 years of education on average. As shown in the second (green) line in the figure, this translated into steadily rising educational attainment in the adult labor force. Between 1940 and 1980, for example, educational attainment rose from 9 years to 12 years, or about 3/4 of a year per decade. With a Mincerian return to education of 7 percent, this corresponds to a contribution of about 0.5 percentage points per year to growth in output per worker.

The other fact that stands out prominently, however, is the levelling-off of educational attainment. For cohorts born after 1950, educational attainment rose more slowly than before, and for the latest cohorts, educational attainment has essentially

Figure 7: Educational Attainment, United States

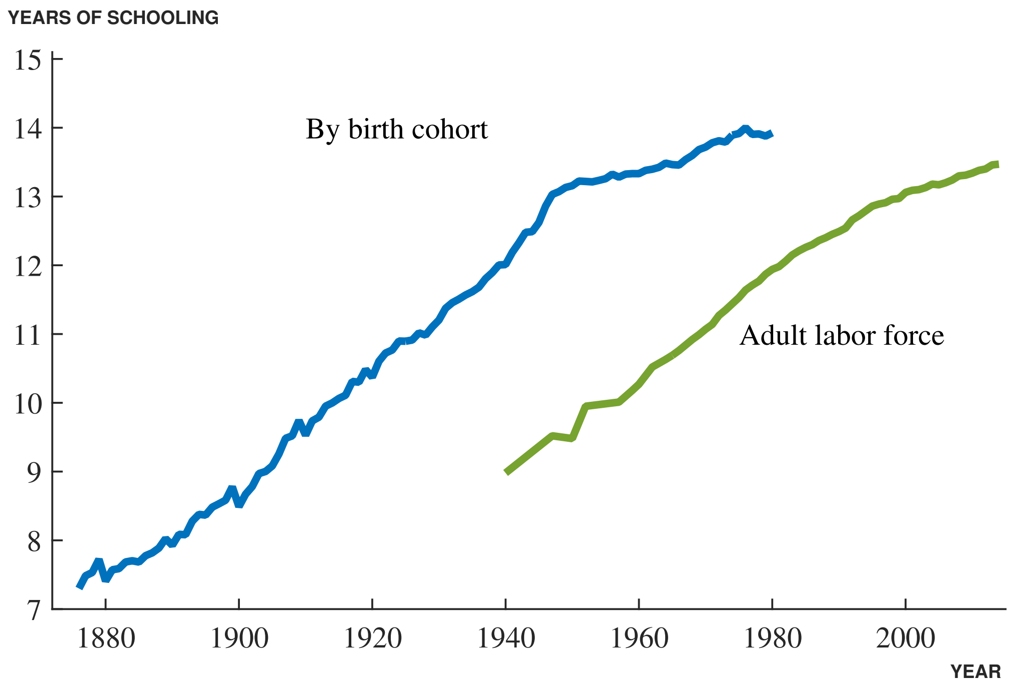

Note: The blue line shows educational attainment by birth cohort from Goldin and Katz (2007). The green line shows average educational attainment for the labor force aged 25 and over from the Current Population Survey.

Figure 8: The Supply of College Graduates and the College Wage Premium, 1963–2012

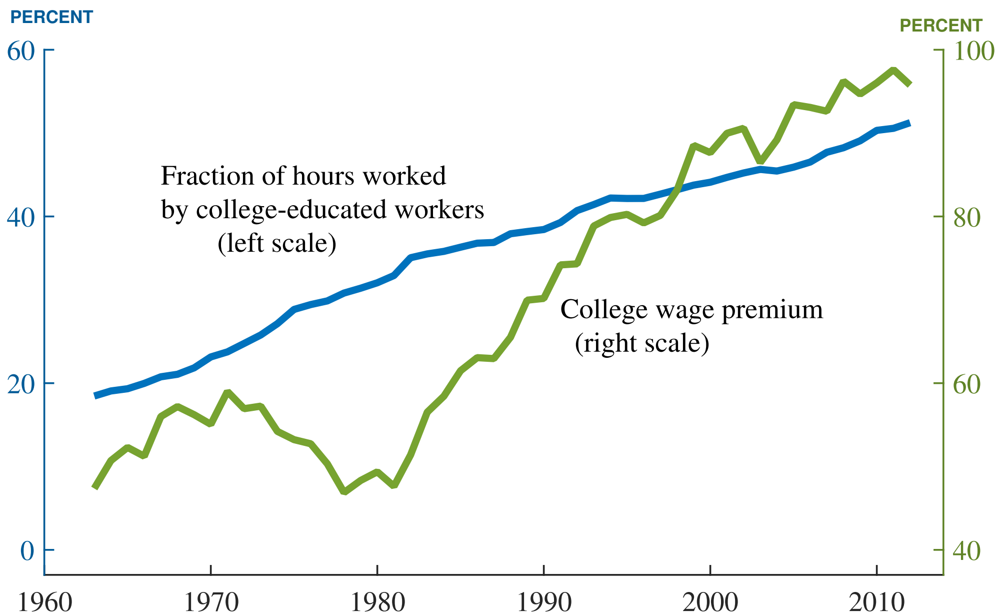

Note: The supply of U.S. college graduates, measured by their share of total hours worked, has risen from below 20 percent to more than 50 percent by 2012. The U.S. college wage premium is calculated as the average excess amount earned by college graduates relative to non-graduates, controlling for experience and gender composition within each educational group. Source: Autor (2014), Figure 3.

flattened out. Over time, one expects this to translate into a slowdown in the increase of educational attainment for the labor force as a whole, and some of this can perhaps be seen in the last decade of the graph.

Figure 8 shows another collection of stylized facts related to human capital made famous by Katz and Murphy (1992). The blue line in the graph shows the fraction of hours worked in the U.S. economy accounted for by college-educated workers. This fraction rose from less than 20 percent in 1963 to more than 50 percent by 2012. The figure also shows the college wage premium, that is the excess amount earned by college graduates over non-graduates after controlling for experience and gender. This wage premium averaged around 50 percent between 1963 and the early 1980s but then rose sharply through 2012 to peak at nearly 100 percent. Thus, even though the supply of college graduates was growing rapidly, the wage premium for college graduates was increasing sharply as well.

Katz and Murphy (1992) provide an elegant way to understand the dynamics of the college wage premium. Letting "coll" and "hs" denote two kinds of labor ("college graduates" and "high school graduates"), the human capital aggregate that enters production is given by a CES specification:

$$H = ((A_{coll}L_{coll})^{\rho} + (A_{hs}L_{hs})^{\rho})^{1/\rho} \tag{4}$$

An increase in the supply of college graduates lowers their marginal product, while an increase in the technology parameter A_{coll} raises their marginal product. Katz and Murphy (1992) show that with an elasticity of substitution of around 1.4, a constant growth rate of A_{coll}/A_{hs}, which Katz and Murphy call "skill-biased technical change," together with the observed movements in L_{coll}/L_{hs} can explain the time series for the college wage premium.

Human capital includes more than just education, of course. Workers continue to accumulate skills on the job. This human capital shows up as higher wages for workers, but separating this into a quantity of human capital and a price per unit of human capital requires work. One simple approach is to assume each year of work experience leads to a constant increase in human capital, and this approach is commonly pursued in growth accounting. Examples of richer efforts to measure human capital in a growth setting include Lucas (2009), Lucas and Moll (2014), and Manuelli and Seshadri (2014).

2.5. Ideas

Our next set of facts relate to the economy's stock of knowledge or ideas, the A in the production function that we began with back in equation (1). It has long been recognized that the "idea production function" is hard to measure. Where do ideas come from? Part of the difficulty is that the answer is surely multidimensional. Ideas are themselves very heterogeneous, some clearly arise through intentional research, but others seem to arrive by chance out of seemingly nowhere. Confronted with these difficulties, Solow (1956) modeled technological change as purely exogenous, but this surely goes too far. The more people there are searching for new ideas, the more likely it is that discoveries will be made. This is true if the searching is intentional, as in research, but even if it is a byproduct of the production process itself as in models

Figure 9: Research and Development Spending, United States

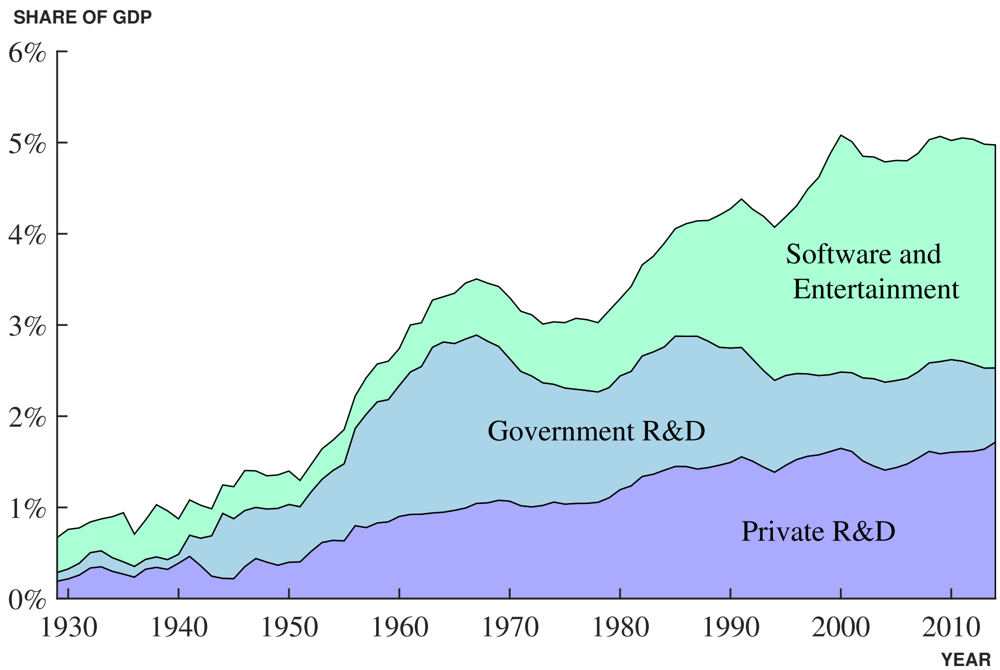

Source: National Income and Product Accounts, U.S. Bureau of Economic Analysis via FRED database. "Software and Entertainment" combines both private and public spending. "Entertainment" includes movies, TV shows, books, and music.

of learning by doing. The production of new ideas plays a fundamental role in the modern understanding of growth; see Romer (1990), Grossman and Helpman (1991), and Aghion and Howitt (1992).[4]

With this in mind, Figure 9 shows spending on research and development, as a share of GDP, for the United States. These data can now be obtained directly from the National Income and Product Accounts, thanks to the latest revisions by the Bureau of Economic Analysis. The broadest measure of investment in ideas recorded by the NIPA is investment in "intellectual property products." This category includes traditional research and development, spending on computer software, and finally spending on "entertainment," which itself includes movies, TV shows, books, and music.

Definition of R&D — WalMart does none!!

Several facts stand out in Figure 9. First, total spending on investment in intellectual

[4]Various perspectives on the idea production function are presented by Mokyr (1990), Griliches (1994), Weitzman (1998), and Fernald and Jones (2014).

Figure 10: Research Employment Share

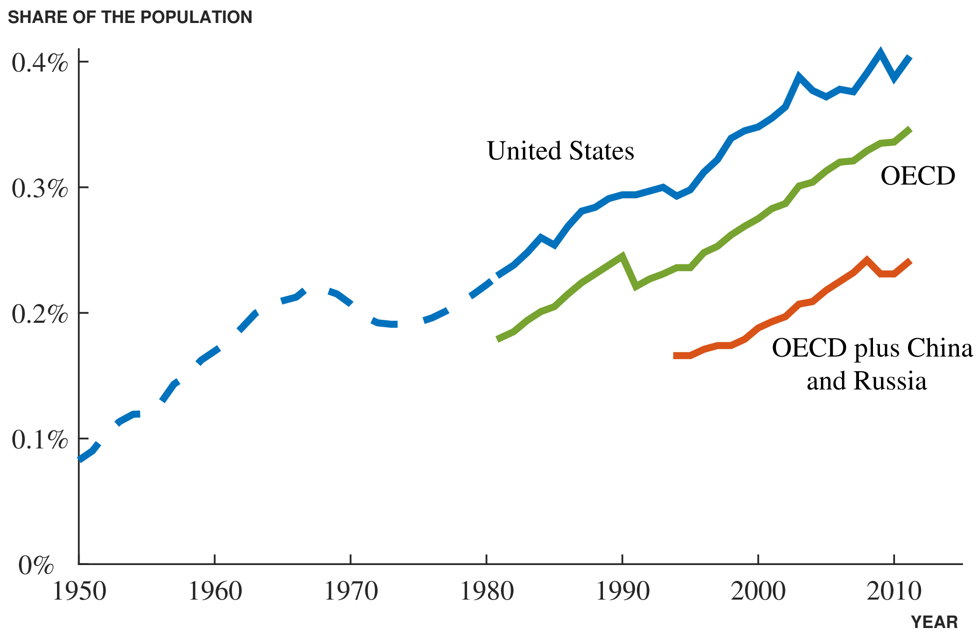

Source: Data for 1981–2001 are from *OECD Main Science and Technology Indicators*, http://stats.oecd.org/Index.aspx?DataSetCode=MSTI_PUB. Data prior to 1981 for the United States are spliced from Jones (2002), which uses the NSF's definition of "scientists and engineers engaged in R&D."

property products has rise from less than 1 percent of GDP in 1929 to nearly 5 percent of GDP in recent years. This overall increase reflects a large rise in private research and development and a large rise in software and entertainment investment, especially during the last 25 years. Finally, government spending on research and development has been shrinking as a share of GDP since peaking in the 1960s with the space program.

Figure 10 provides an alternative perspective on R&D in two dimensions. First, it focuses on employment rather than dollars spent, and second it brings in an international perspective. The figure shows the number of researchers in the economy as a share of the population.[5]

Each of the three measures in the figure tells the same story: the fraction of the population engaged in R&D has been rising in recent decades. This is true within the United States, within the OECD, and even if we incorporate China and Russia as well.

[5]According to the OECD's Frascati Manual 2002, p. 93, researchers are defined as "professionals engaged in the conception or creation of new knowledge, products, processes, methods and systems and also in the management of the projects concerned."

Figure 11: Patents Granted by the U.S. Patent and Trademark Office

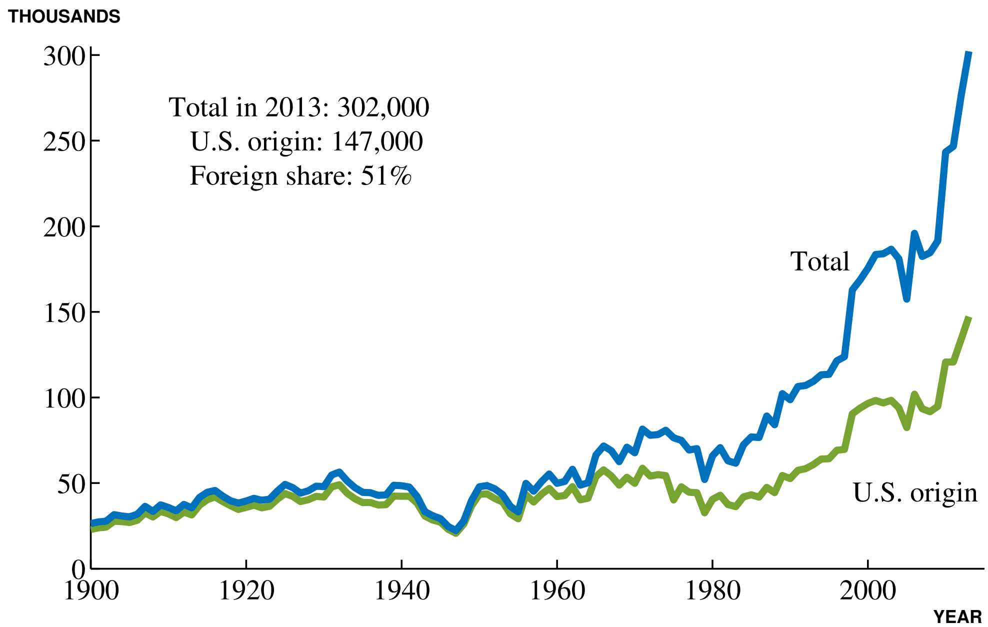

Source: http://www.uspto.gov/web/offices/ac/ido/oeip/taf/h_counts.htm

So far, we have considered the input side of the idea production function. We now turn to the output side. Unfortunately, the output of ideas is even harder to measure than the inputs. One of the more commonly-used measures is patents, and this measure is shown in Figure 11.

On first glance, it appears that patents, like many other variables reviewed in this essay, have grown exponentially. Indeed, at least since 1980 one sees a very dramatic rise in the number of patents granted in the United States, both in total and to U.S. inventors. The difference between these two lines — foreign patenting in the U.S. — is also interesting, and one testimony to the global nature of ideas is that 56 percent of patents granted by the U.S. patent office in 2013 were to foreigners.

A closer look at Figure 11, though, reveals something equally interesting: the number of patents granted to U.S. inventors in 1915, 1950, and 1985 was approximately the same. Put another way, during the first 85 years of the 20th century, the number of patents granted to U.S. residents appears to be stationary, in sharp contrast to the dramatic increase since 1985 or so.

Griliches (1994) combined these two basic facts related to ideas (rapid growth in the inputs, stable production of patents) to generate a key implication: the productivity of research at producing patents fell sharply for most of the 20th century. Kortum (1997) developed a growth model designed to match these facts in which he emphasized that patents can be thought of as *proportional* improvements in productivity. If each patent raises GDP by a constant percent, then a constant flow of new patents can generate a constant rate of economic growth. The problem with this approach (or perhaps the problem with the patent data) is that it breaks down after 1980 or so. Since 1980, the number of patents has risen by more than a factor of four, while growth rates are more or less stable. The bottom line is that the idea production function remains something of a black box perhaps precisely because we do not have great measures of ideas or the inputs used to produce them.[6]

2.6. Misallocation

The organizing principle for this section of the paper is the production function given back in equation (1). In specifying that production function, I broke total factor productivity into two pieces: the stock of ideas, A, and everything else, which I labeled "M" either for the "measure of our ignorance" or for "misallocation." It is this latter interpretation that I wish to take up now.

One of the great insights of the growth literature in the last 15 years is that misallocation at the micro level can show up as a reduction in total factor productivity at a more aggregated level. This insight appears in various places, including Banerjee and Duflo (2005), Chari, Kehoe and McGrattan (2007), Restuccia and Rogerson (2008) and Hsieh and Klenow (2009).

The essence of the insight is quite straightforward: when resources are allocated optimally, the economy will operate on its production possibilities frontier. When resources are misallocated, the economy will operate inside this frontier. But that is just another way of saying that TFP will be lower: a given quantity of inputs will produce less output.

As we explain in detail in the second part of this paper (in Section 4.6.), there is a

[6] Examples of progress include Caballero and Jaffe (1993), Acemoglu, Akcigit, Bloom and Kerr (2013), and Akcigit, Celik and Greenwood (2014b).

large literature on misallocation and development — this is our best candidate answer to the question of why are some countries so much richer than others. There is much less discussion of the extent to which misallocation is related to frontier growth, the subject at hand.

While it is clear conceptually that even the country or countries at the frontier of growth can suffer from misallocation and that changes in misallocation can contribute to growth, there has been little work quantifying this channel. Indeed, my own working hypothesis for many years was that this effect was likely small in the United States in the last 50 years. I now believe this is wrong.

Hsieh, Hurst, Jones and Klenow (2013) highlight a striking fact that illustrates this point: in 1960, 94 percent of doctors and lawyers were white men; by 2008, this fraction was just 62 percent. Given that innate talent for these and other highly-skilled professions is unlikely to differ across groups, the occupational distribution in 1960 suggests that a large number of innately talented African Americans and white women were not working in the occupations dictated by comparative advantage. The paper quantifies the macroeconomic consequences of the remarkable convergence in the occupational distribution between 1960 and 2008 and finds that 15 to 20 percent of growth in aggregate output per worker is explained by the improved allocation of talent. In other words, declines in misallocation may explain a significant part of U.S. economic growth during the last 50 years.

Examples to drive home these statistics are also striking. Sandra Day O'Connor — future Supreme Court Justice — graduated third in her class from Stanford Law School in 1952. But the only private sector job she could get upon graduation was as a legal secretary (Biskupic, 2006). Closer to our own profession, David Blackwell, of contraction mapping fame, was the first African American inducted into the National Academy of Sciences and the first tenured at U.C. Berkeley. Yet despite getting his Ph.D. at age 22 and obtaining a post-doc at the Institute for Advanced Studies in 1941, he was not permitted to attend lectures at Princeton and was denied employment at U.C. Berkeley for racial reasons. He worked at Howard University until 1954, when he was finally hired as a full professor in the newly-created statistics department at Berkeley.[7]

Another potential source of misallocation is related to the economics of ideas. It has

---

[7] See http://en.wikipedia.org/wiki/David_Blackwell. I'm grateful to Ed Prescott for this example.

long been suggested that knowledge spillovers are quite significant, both within and across countries. To the extent that these spillovers are increasingly being internalized or addressed by policy — or to the extent that the opposite is true — the changing misallocation of knowledge resources may be impacting economic growth.[8] Quantifying these and other types of misallocation affecting frontier growth is a fruitful direction for future research.

Hsieh and Moretti on land use...

2.7. Explaining the Facts of Frontier Growth

While this essay is primarily about the facts of economic growth, it is helpful to step back and comment briefly on how multiple facts have been incorporated into our models of growth.

The basic neoclassical growth framework of Solow (1956) and Ramsey (1928) / Cass (1965) / Koopmans (1965) has long served as a benchmark organizing framework for understanding the facts of growth. The nonrivalry of ideas, emphasized by Romer (1990), helps us understand how sustained exponential growth occurs endogenously. I review this contribution and some of the extensive research it sparked in Jones (2005).[9]

The decline in the relative price of equipment and the rise in the college wage premium are looked at together in Krusell, Ohanian, Rios-Rull and Violante (2000). That paper considers a setting in which equipment capital is complementary to skilled labor, so that the (technologically-driven) decline in the price of equipment is the force of skill-biased technological change. That paper uses a general CES structure. One of the potential issues in that paper was that it could lead to movements in the labor share. But perhaps we are starting to see those in the data.

The presence of trends in educational attainment and research investment opens up interesting opportunities for future research. Why are educational attainment and the share of labor devoted to research rising over time? What are the implications of these trends for future growth? Restuccia and Vandenbroucke (2013) suggest that

[8]For evidence on knowledge spillovers, see Griliches (1992), Coe and Helpman (1995), Jones and Williams (1998), Klenow and Rodriguez-Clare (2005), and Bloom, Schankerman and Reenen (2013).

[9]Romer's insights are expanded upon in various directions. Aghion and Howitt (1992) and Grossman and Helpman (1991) emphasize the important role of creative destruction. Jones (1995), Kortum (1997), and Segerstrom (1998) clarify the way in which nonrivalry interacts with population growth to explain sustained growth in living standards.

skill-biased technological change is itself responsible for driving the rise in educational attainment. Acemoglu (1998) considers the further interaction when the direction of technological change is endogenous. Jones (2002) considers the implication of the trends in education and research intensity for future growth, suggesting that these trends have substantially raised growth during the last 50 years above the economy's long-run growth rate.

3. Frontier Growth: Beyond GDP

The next collection of facts related to frontier growth look beyond the aggregate of GDP. These facts are related to structural change (the decline of agriculture and the rise of services, especially health), changes in leisure and fertility, rising inequality, and falling commodity prices.

3.1. Structural Change

Figure 12 shows one of the most dramatic structural changes affecting frontier economies over the last two hundred years and beyond: the decline of agriculture. In 1840, about two out of every three workers in the U.S. economy worked in agriculture. By 2000, this share had fallen to just 2.4 percent. Similar changes can be seen in value-added in agriculture as a share of GDP as well as in other countries. For example, the chart also shows agriculture's share of employment in Japan, declining from 85 percent around 1870.[10]

The structural transformation has several other dimensions and connections in the growth literature. For example, the decline in agriculture is first associated with a rise in manufacturing, which is ultimately replaced by a rise in services, including health and education; more on this below.

Another form of structural transformation that has seen renewed interest is the possibility that machines (capital) may substitute for labor. Autor, Levy and Murnane (2003) look at detailed occupational classifications to study the impact of computerization on labor demand. They emphasize a polarization, with computerization being

[10]The literature on structural transformation and economic growth is surveyed by Herrendorf, Rogerson and Valentinyi (2014). More recent contributions include Boppart (2014) and Comin, Lashkari and Mestieri (2015), who emphasize demand systems with heterogeneous income effects.

Figure 12: Employment in Agriculture as a Share of Total Employment

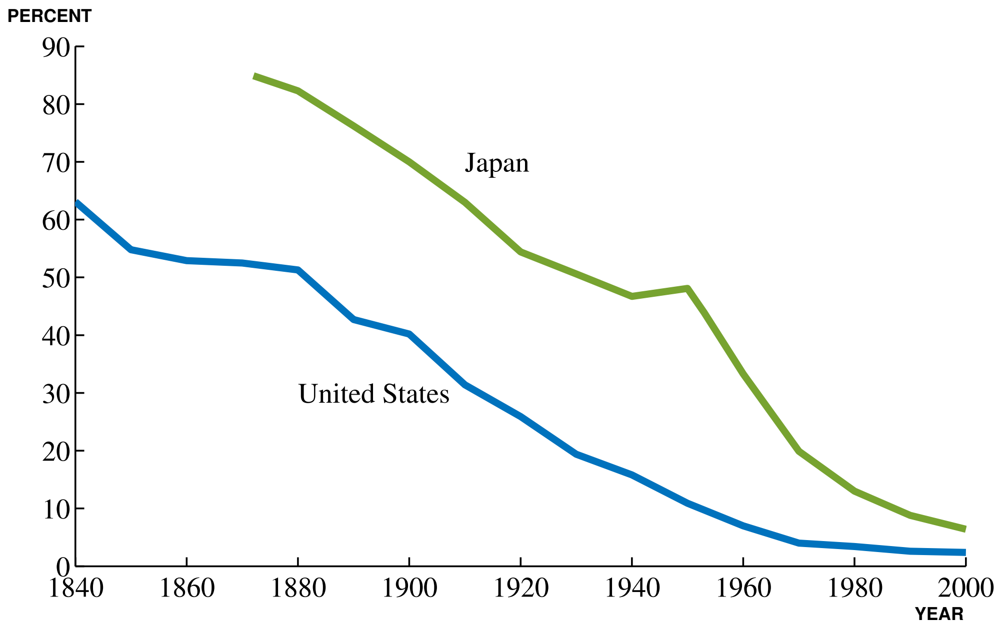

Source: Herrendorf, Rogerson and Valentinyi (2014).

particularly substitutable for routine cognitive tasks that can be broken into specific rules but complementary to nonroutine, cognitive tasks. That is, computers substitute for bank tellers and low-level secretaries, while increasing the demand for computer programmers and leaving untouched manual jobs like janitorial work. Brynjolfsson and McAfee (2012) highlight broader ramifications of artificial intelligence, whereby computers might start driving cars, reading medical tests, and combing through troves of legal documents. That is, even many tasks thought to be cognitive and not easily routinized may be subject to computerization. What impact will such changes have on the labor market?

The answer to this question is obviously complicated and the subject of ongoing research.[11] One useful reference point is the enormous transformation that occurred as the agricultural share of the U.S. labor force went from 2/3 to only 2 percent, largely because of mechanization and technological change. There is no doubt that this had a transformative affect on the labor market, but by and large this transformation was

[11] For some examples, see Acemoglu (1998), Zeira (1998), Caselli (1999), Hemous and Olsen (2014).

Figure 13: Health Spending as a Share of GDP

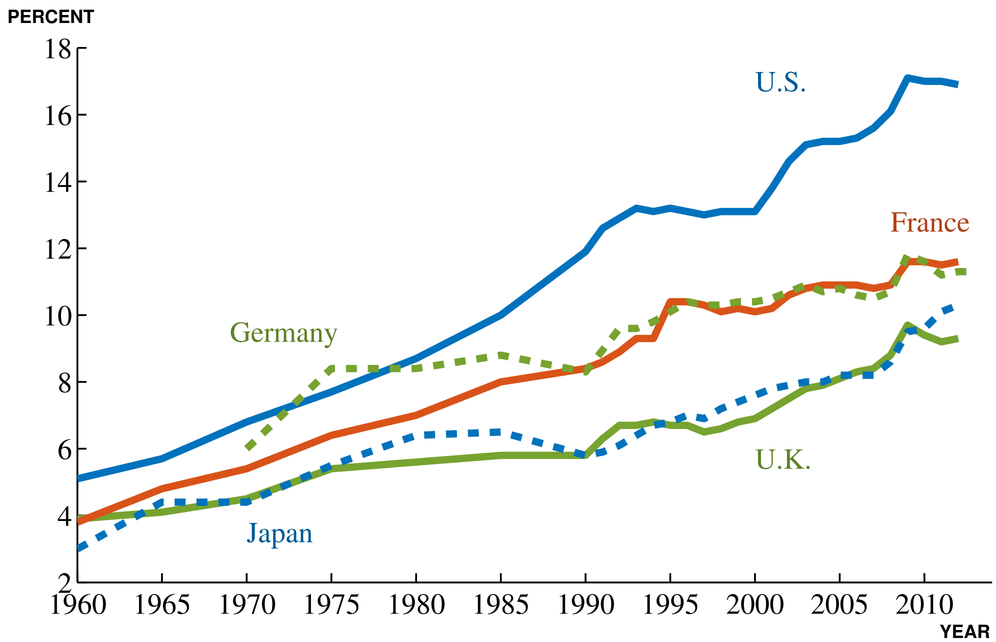

Source: *OECD Health Statistics 2014*.

overwhelmingly beneficial. That's not to say that it must be that way in the future, but the example is surely worth bearing in mind.

3.2. The Rise of Health

A different structural transformation has been predominant during the last 50 years: the rise of health spending as a share of GDP. Figure 13 shows this fact for the United States and for several other OECD countries. In the United States, the health share more than tripled since 1960, rising from 5 percent in 1960 to 17 percent in recent years. Large trends are present in other countries as well, with the share in France, for example rising from under 4 percent to nearly 12 percent.

Hall and Jones (2007) propose that the widespread rise in the prominence of health care is a byproduct of economic growth. With standard preferences, the marginal utility of consumption declines rapidly. This is most easily seen for CRRA preferences in which the intertemporal elasticity of substitution is below one, in which case flow utility is bounded. As we get richer and richer, the marginal utility of consumption on any given

Figure 14: Life Expectancy at Birth and at Age 65, United States

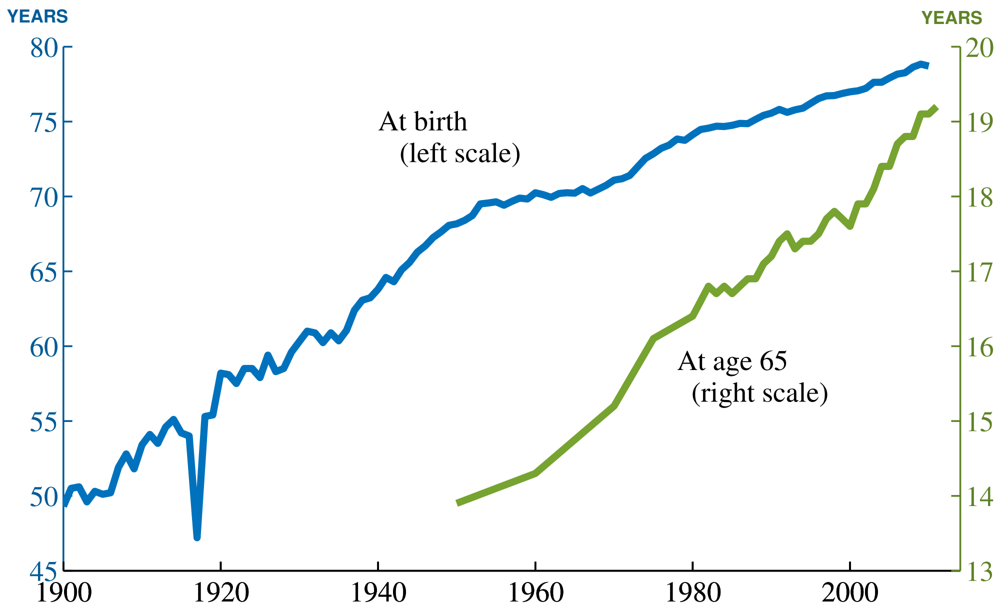

Source: *Health, United States 2013* and https://www.clio-infra.eu.

day declines rapidly; what people really need are more days of life to enjoy their high level of consumption. Hence there is an income effect tilting spending toward life-saving categories.

One of the few time series related to economic growth that does not grow exponentially is life expectancy, where the increases tend to be arithmetic rather than exponential. Figure 14 shows life expectancy at birth and at age 65 in the United States. Life expectancy at birth increased rapidly in the first half of the 20th century, thanks to improvements in public health and large declines in infant mortality. Since 1950, the rate of improvement has been more modest, around 1.8 years per decade. The figure also shows that the rise in life expectancy occurs at old ages. Life expectancy conditional on reaching age 65 has risen by just under one year per decade since 1950.[12] Interestingly, the mortality rate itself seems to grow exponentially with age, a phenomenon known as

[12] Nordhaus (2003) and Murphy and Topel (2006) discuss the rise in life expectancy and the economic returns to reducing mortality in more detail. Oeppen and Vaupel (2002) suggest that "record life expectancy" (i.e. the maximum life expectancy across countries) has grown linearly at 2.5 years per decade for more than 150 years.

Figure 15: Average Annual Hours Worked, Select Countries

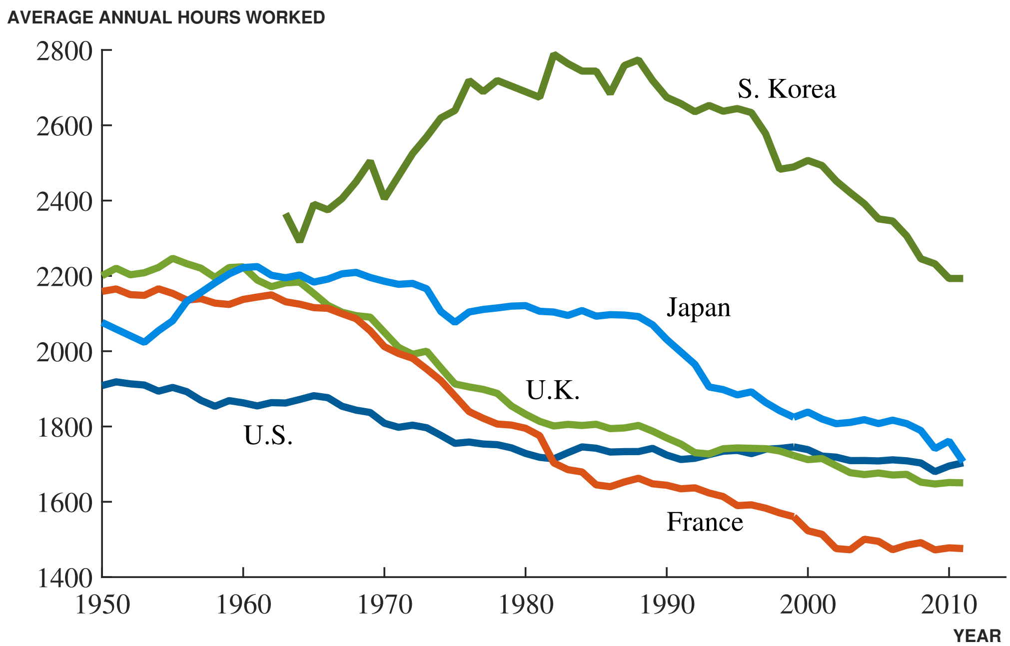

Source: Average annual hours worked per person employed, from the Penn World Tables 8.0.

the Gompertz-Makeham Law; see Dalgaard and Strulik (2014).

3.3. Hours Worked and Leisure

A standard stylized fact in macroeconomics is that the fraction of the time spent working shows no trend despite the large upward trend in wages. The next two figures show that this stylized fact is not really true over the longer term, although the evidence is somewhat nuanced.

Figure 16 shows average annual hours worked per person engaged in work from the Penn World Tables, which takes its data in turn from the Total Economy Database of the Conference Board. Among advanced countries, annual hours worked has fallen significantly since 1950. Average hours worked in the United States, for example, fell from 1909 in 1950 to 1704 in 2011. In France, the decline is even more dramatic, from 2159 to 1476. The decline starts slightly later in Japan after their recovery from World War II, with hours falling from 2222 in 1960 to 1706 in 2011.

Figure 16 breaks the U.S. evidence down into more detail, courtesy of Ramey and

Figure 16: Average Weekly Hours Worked, United States

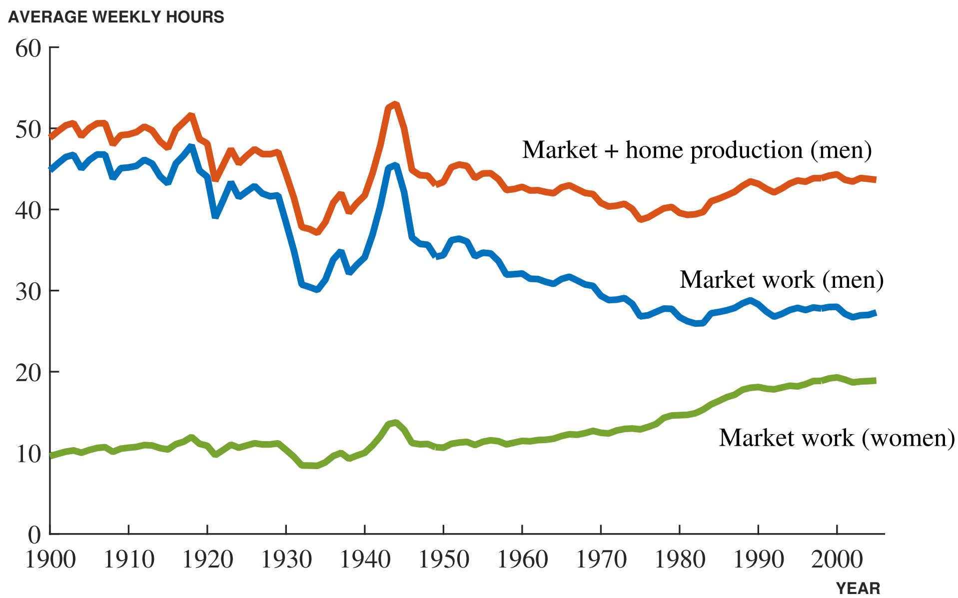

Source: Average weekly hours per worker, from Ramey and Francis (2009).

Francis (2009). First, the figure shows the split between men and women. Average weekly hours of market work by men fell sharply between 1900 and 1980, before leveling off. In contrast, market work by women has been on an upward trend. Ramey and Francis (2009) also use time diaries to estimate home production, and this is where the story gets more complicated. As men are substituting away from market work, they are also substituting into home production. Home production by men rose from just 4 hours per week in 1900 to more than 16 hours per week in 2005. The increase in leisure, then, was much smaller than the decline in market hours suggests.

3.4. Fertility

The facts we have presented thus far in this section — the decline in agriculture and the rise in services like health, the rise in life expectancy, the decline in hours worked — are all consistent with a particular kind of income effect. As people get richer, the marginal utility of consumption falls and people substitute away from consumption and toward actions that conserve on the precious time endowment. Time is the one thing that technological progress cannot create!

Figure 17: Fertility in the United States and France

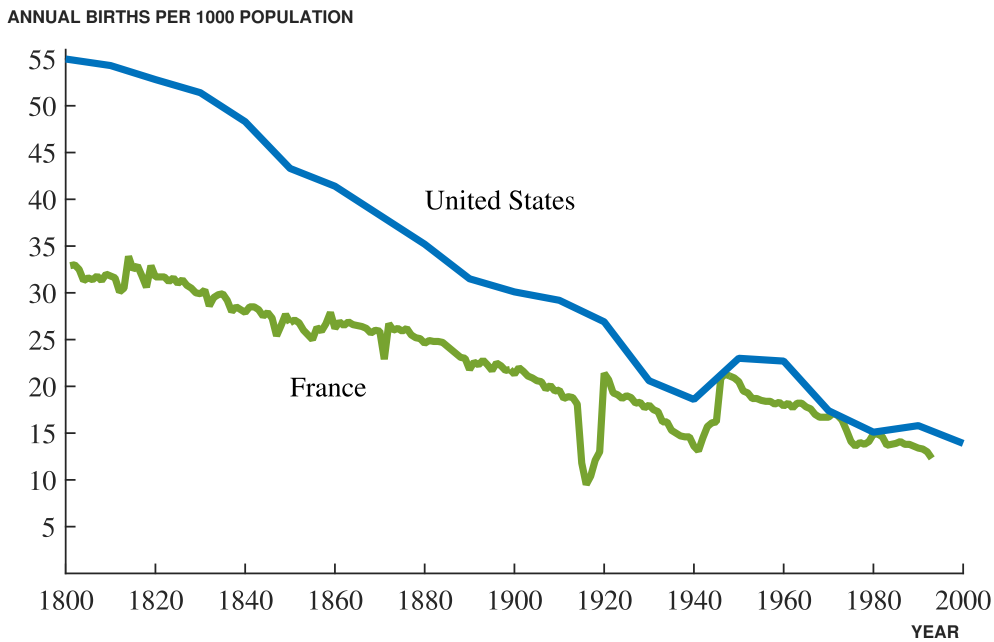

Source: Data for the United States are from Haines (2008) and data for France are from Greenwood and Vandenbroucke (2004).

The next fact on fertility raises interesting questions about this hypothesis. In particular, Figure 17 shows the large decline in fertility dating back at least to 1800, known as the demographic transition. Since 1800, the birth rate has fallen from 5.5 percent in the United States and 3.3 percent in France down to less than 1.5 percent in recent years.

In dynastic models like Barro and Becker (1989), in which having more children is essentially a way of increasing one's own effective lifetime or time endowment, there is a force that tends to raise fertility, at least if income effects dominate substitution effects. But instead, we see strong declines in fertility in the data. A large literature seeks to understand the declines in fertility and the hump-shape in population growth that are together known as the demographic transition. A key part of the standard explanation is that children are themselves time intensive, in which case conserving on children also conserves on time as people get richer.[13]

[13]For example, see Galor and Weil (1996), Doepke (2005), Greenwood, Seshadri and Vandenbroucke (2005), Jones, Schoonbroodt and Tertilt (2010), Cordoba and Ripoll (2014), and Jones and Tertilt (forthcoming).

Figure 18: Top Income Inequality in the United States and France

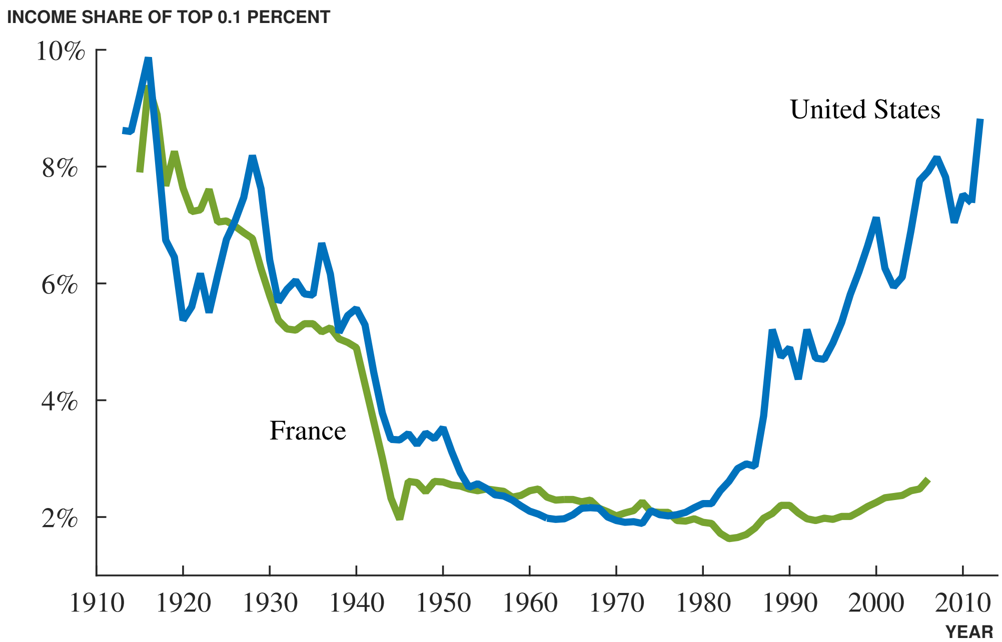

Source: World Top Incomes Database, Alvaredo, Atkinson, Piketty and Saez (2013).

3.5. Top Inequality

One of the more famous facts documented during the last decade is shown in Figure 18. This is the top income inequality graph of Piketty and Saez (2003). In both the United States and France, the share of income earned by the top 0.1 percent of households was around 9 percent in 1920, and in both countries the share declined sharply until the 1950s to around 2 percent. It stayed at this low level until around 1980. But then a very large difference emerged, with top income shares rising in the U.S. to essentially the same level as in 1920, while the share in France remains relatively low. Much of the decline in the first part of the century is associated with capital income, and much of the rise in U.S. inequality since 1980 is associated with labor (and business) income.[14]

It is also worth stepping back to appreciate the macroeconomic consequences of this inequality. Figure 19 merges the Piketty-Saez top inequality data with the long-

[14] Possible explanations for this pattern are discussed by Castaneda, Diaz-Gimenez and Rios-Rull (2003), Cagetti and Nardi (2006), Atkinson, Piketty and Saez (2011), Benhabib, Bisin and Zhu (2011), Aoki and Nirei (2013), Jones and Kim (2014), Piketty (2014), Piketty, Saez and Stantcheva (2014), and Saez and Zucman (2014).

Figure 19: GDP per person, Top 0.1% and Bottom 99.9%

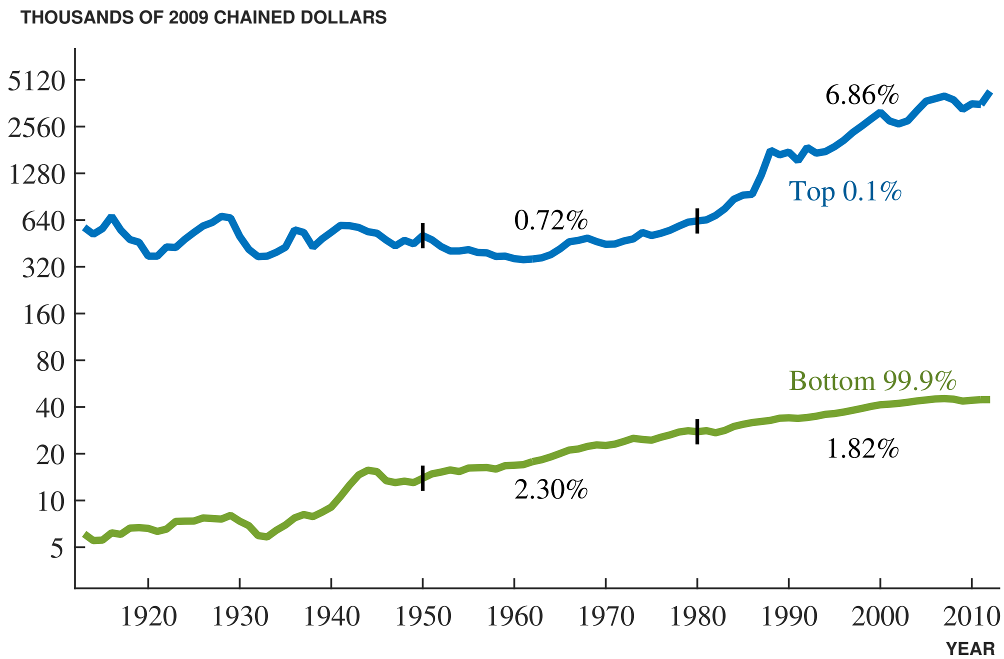

Note: This figure displays an estimate of average GDP per person for the top 0.1% and the bottom 99.9%. Average annual growth rates for the periods 1950–1980 and 1980–2007 are also reported. Source: Aggregate GDP per person data are from Figure 1. The top income share used to divide the GDP is from the October 2013 version of the world top incomes database, from http://g-mond.parisschoolofeconomics.eu/topincomes/.

run data on GDP per person for the United States shown at the start of this paper in Figure 1. In particular, the figure applies the Piketty-Saez inequality shares to average GDP per person to produce an estimate of GDP per person for the top 0.1% and the bottom 99.9%.[15]

Two key results stand out. First, until recently, there is surprisingly little growth in average GDP per person at the top: the value in 1913 is actually *lower* than the value in 1977. Instead, all the growth until around 1960 occurs in the bottom 99.9%. The second point is that this pattern changed in recent decades. For example, average growth in GDP per person for the bottom 99.9% declined by around half a percentage point, from 2.3% between 1950 and 1980 to only 1.8% between 1980 and 2007. In contrast, after being virtually absent for 50 years, growth at the top accelerated sharply: GDP

[15]It is important to note that this estimate is surely imperfect. GDP likely does not follow precisely the same distribution as the Adjusted Gross Income data that forms the basis of the Piketty-Saez calculations: health benefits are more equally distributed, for example.

per person for the top 0.1% exhibited growth more akin to China's economy, averaging 6.86% since 1980. Changes like this clearly have the potential to matter for economic welfare and merit the attention they've received.

3.6. The Price of Natural Resources

This next fact is very different from what we've been discussing, but it is one of the more surprising facts related to frontier growth. Figure 20 shows the real price of industrial commodities, consisting of an equally-weighted basket of aluminum, coal, copper, lead, iron ore, and zinc, deflated by the consumer price index. During the 20th century, world demand for these industrial commodities exploded with the rise of the automobile, electrification, urbanization, and the general industrialization that occurred in the United States and around the world. The surprise shown in the figure is that the real price of these commodities *declined* over the 20th century. Moreover, the magnitude of the decline was large — a factor of 5 between the year 1900 and 2000. Evidently, some combination of increased discoveries and technological changes led the effective supply to grow even faster than enormous rise in demand.[16]

Also striking, though, is the large increase in the real price of these commodities since 2000. Part of the explanation could be the rapid growth of China and India over this period and the large increase in demand for commodities that their growth entailed. Interestingly, we will see later that many developing countries performed quite well in the 2000s. Some of that growth contributed to the rise in demand for commodities, but some of that success may also reflect commodity-driven growth resulting from the rise in demand from China and India.

4. The Spread of Economic Growth

Up until now, we've been primarily concerned with the growth of the frontier: what are the facts about how the frontier is moving over time? Now, we turn to how growth is spreading across countries: how are different countries moving relative to the frontier?

[16]This fact has been noted before, for example by Simon (1981).

Figure 20: The Real Price of Industrial Commodities

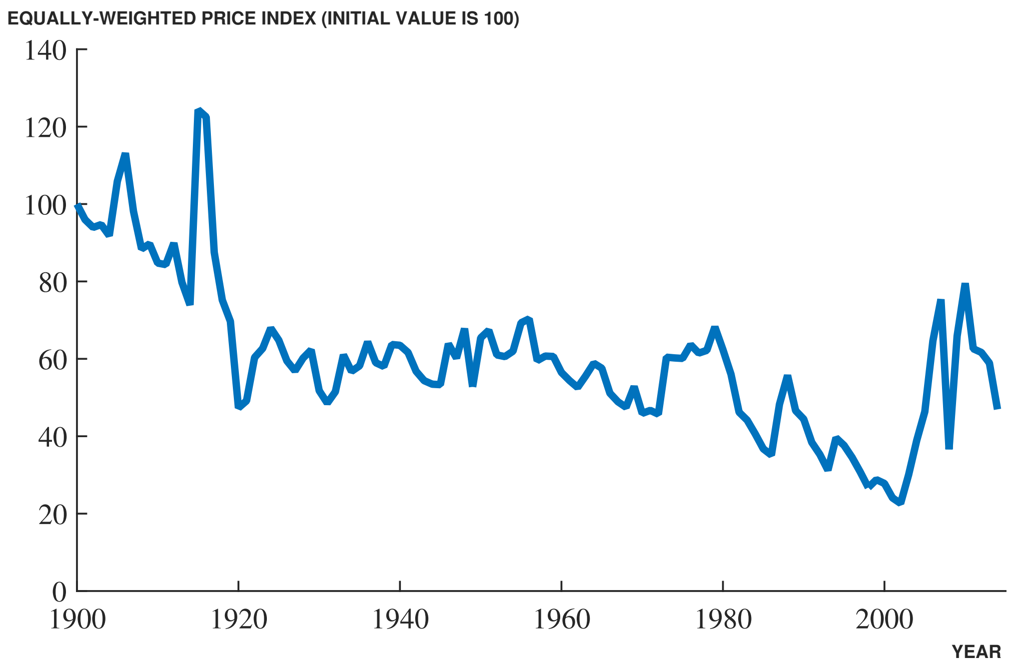

Note: The price of an equally-weighted basket of aluminum, coal, copper, lead, iron ore, and zinc, deflated by the consumer price index. Commodity prices are from www.globalfinancialdata.com and the CPI is from www.measuringworth.com.

Figure 21: The Great Divergence

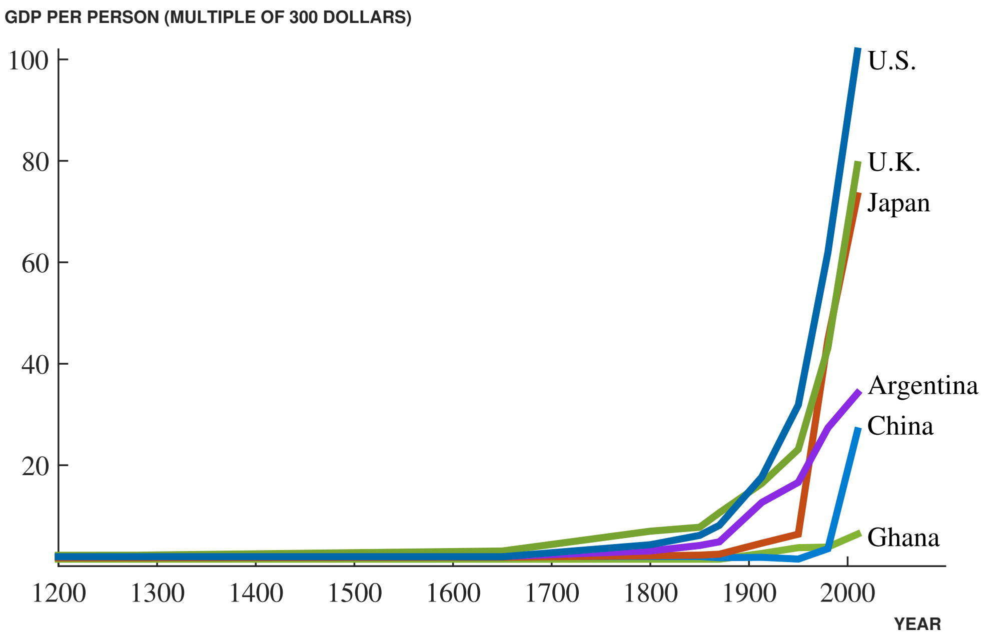

Note: The graph shows GDP per person for various countries, normalized by the value in the United Kingdom in the initial year. Source: The Maddison Project, Bolt and van Zanden (2014).

4.1. The Long Run

One of the key facts about the spread of growth over the very long run is that it occurred at different points in time, resulting in what is commonly referred to as "The Great Divergence."[17] Figure 21 illustrates this point. GDP per person differs modestly prior to the year 1600 according to The Maddison Project data. For example, GDP per person in the year 1300 ranges from a high of $1620 in the Netherlands (in 1990 dollars) to a low of $610 in Egypt. But Egypt was surely not the poorest country in the world at the time. Following an insight by Pritchett (1997), notice that the poorest countries in the world in 1950 had an income around $300, and this level — less than one dollar per day — seems very close to the minimum average income likely to prevail in any economy at any point in time. Therefore in 1300, the ratio of the richest country to the poorest was on the order of $\$1620/\$300 \approx 5$. Even smaller ratios are observed in Maddison's data prior to the year 1300.

[17] See Maddison (1995), Pritchett (1997), Lucas (2000), and Pomeranz (2009).

Figure 21 shows how this ratio evolved over time for a small sample of countries, and one sees the "Great Divergence" in incomes that occurs after the year 1600. The ratio of richest to poorest rises to more than 10 by 1870 (for the United Kingdom) and then to more than 100 by 2010 for the United States. Across the range of countries, rapid growth takes hold at different points in time. Argentina is relatively rich by 1870 and growth takes off in Japan after World War II. In 1950, China was substantially poorer than Ghana — by more than a factor of two according to Maddison. Rapid growth since 1978 raises China's living standards to more than a factor of 25 over the benchmark level of $300 per year.

Figure 22 shows the spread of growth since 1870 in an alternative way, by plotting incomes relative to the U.S. level. A key fact that stands out when the data are viewed this way is the heterogeneity of experiences. Some countries like the U.K., Argentina, and South Africa experience significant declines in their incomes relative to the United States, revealing the fact that their growth rates over long periods of time fell short of the 2% growth rate of the frontier. Other countries like Japan and China see large increases in relative incomes.

4.2. The Spread of Growth in Recent Decades

Figure 23 focuses in on the last 30 years using the Penn World Table 8.0 data, again showing GDP per person relative to the U.S. Several facts then stand out. First, incomes in the countries of Western Europe have been roughly stable, around 75 percent of the U.S. level. It is perhaps surprising that countries like France, Germany, and the U.K. are this far behind the United States. Prescott (2004) observes that a large part of the difference is in hours worked: GDP per hour is much more similar in these countries, and it is the fact that work hours per adult are substantially lower in Western Europe that explains their lower GDP per person. Jones and Klenow (2015) note that in addition to the higher leisure, Western Europeans tend to have higher life expectancy and lower consumption inequality. Taking all of these factors into account in constructing a consumption-equivalent welfare measure, the Western European countries look much closer to U.S. levels than the simple GDP per person numbers imply.

Figure 23 also illustrates the "lost decades" that Japan has experienced. After rapid growth in the 1980s (and before), Japan peaked at an income relative to the U.S. of 85

Figure 22: The Spread of Economic Growth since 1870

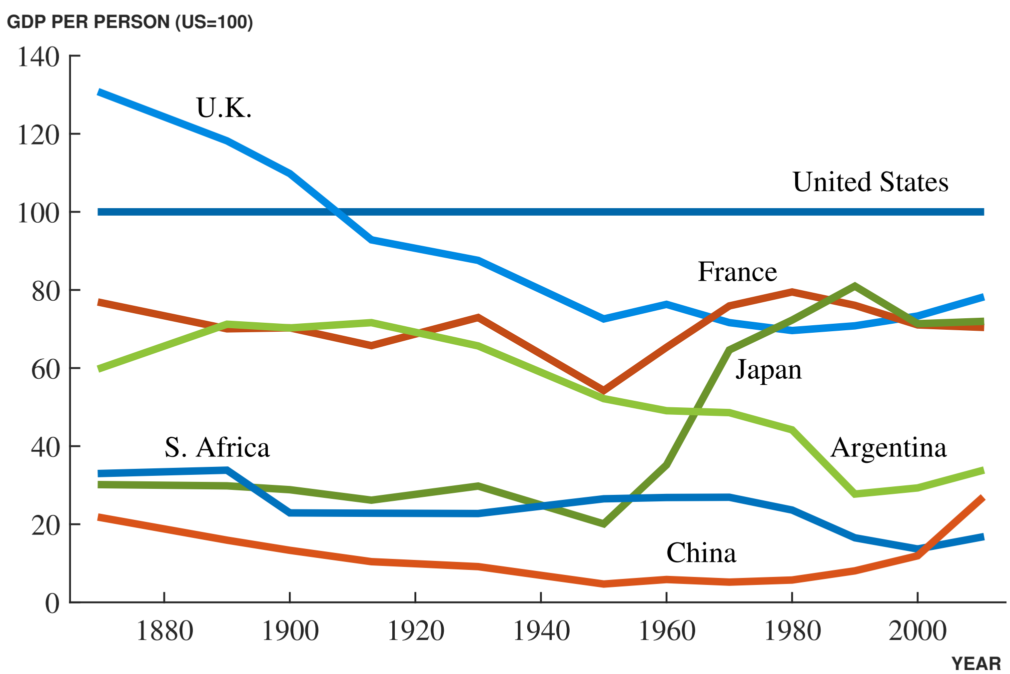

Source: The Maddison Project, Bolt and van Zanden (2014).

Figure 23: The Spread of Economic Growth since 1980

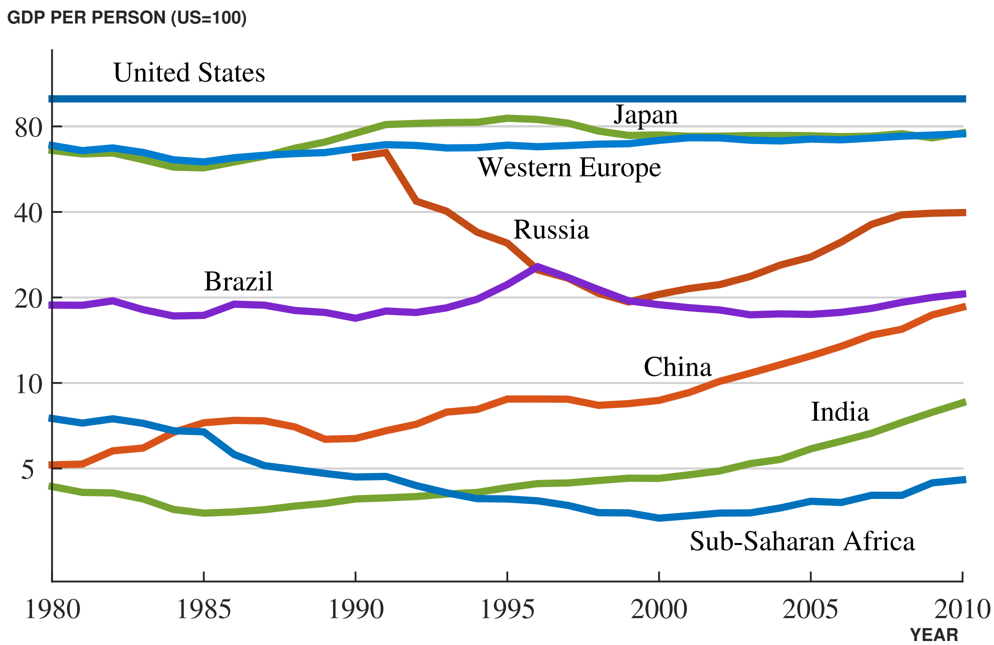

Source: The Penn World Tables 8.0.

percent in 1995. Since 1995, though, Japan has fallen back to around 75 percent of the U.S. level. The rapid growth of China since 1980 and India since around 1990 are also evident in this figure. The contrast with Sub-Saharan Africa is particularly striking, as that region as a whole falls from 7.5 percent of U.S. income in 1980 to just 3.3 percent by 2000. Since 2000, many of the countries and regions shown in Figure 23 exhibit catch-up to the United States.

Figure 24 shows GDP per person relative to the United States in 1960 and 2011 for 107 countries. Countries scatter widely around the 45-degree line, and the first impression is that there is no systematic pattern to this scattering. Some countries are moving up relative to the U.S. and some countries are falling further behind, and the movements can be large, as represented by the deviations from the 45-degree line.

Looking more closely at the graph, there is some suggestion that there are more middle-income countries above the 45-degree line than below. At least between 1960 and 2011, countries in the middle of the distribution seemed more likely to move closer to the U.S. than to fall further behind. In contrast, for low income countries, the oppo-

Figure 24: GDP per Person, 1960 and 2011

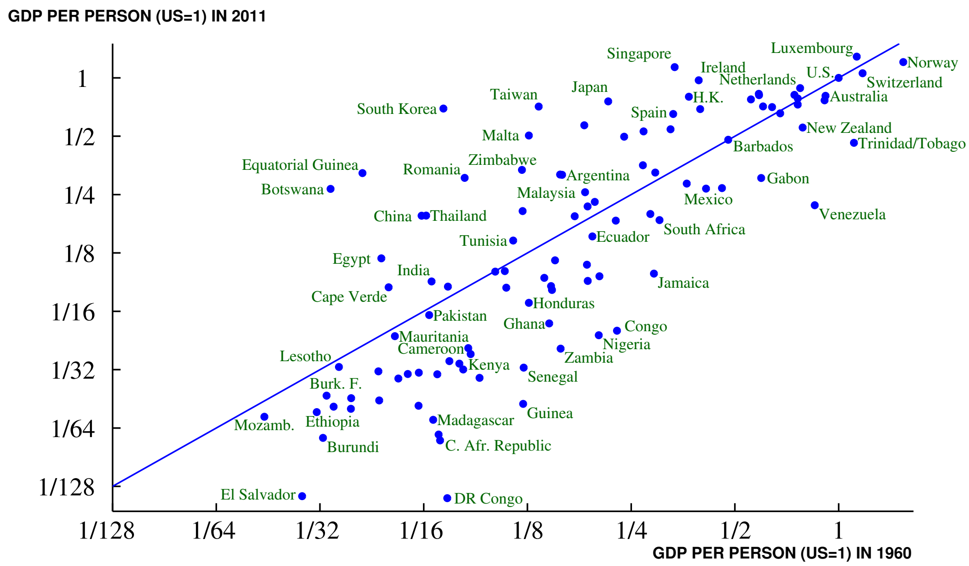

Source: The Penn World Tables 8.0.

site pattern appears in the data: poor countries are on average more systematically below the 45-degree line rather than above.

Figure 25 shows one of the more famous graphs from the empirical growth literature, illustrating the "catch-up" behavior of OECD countries since 1960. Among OECD countries, those that were relatively poor in 1960 — like Japan, Portugal, and Greece — grew rapidly, while those that were relatively rich in 1960 — like Switzerland, Norway, and the United States — grew more slowly. The pattern is quite strong in the data; a simple regression line leads to an R-squared of 75%.[18]

Figure 26 shows that a simplistic view of convergence does not hold for the world as a whole. There is no tendence for poor countries around the world to grow either faster or slower than rich countries. For every Botswana and South Korea, there is an El Salvador and Niger.

Barro (1991), Barro and Sala-i-Martin (1992), and Mankiw, Romer and Weil (1992) provide a key insight into why the convergence pattern appears in Figure 25 but not in Figure 26. In particular, they show that the basic predictions of neoclassical growth theory hold for the world as a whole. Countries around the world are converging —

[18]See also Baumol (1986) and DeLong (1988).

Figure 25: Convergence in the OECD

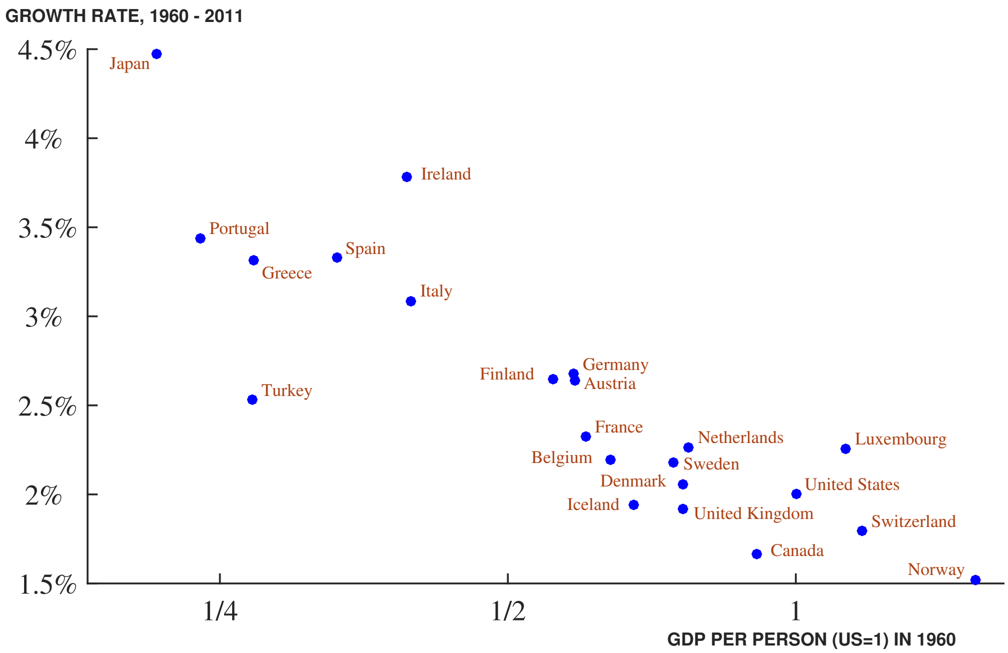

Source: The Penn World Tables 8.0. Countries in the OECD as of 1970 are shown.

Figure 26: The Lack of Convergence Worldwide

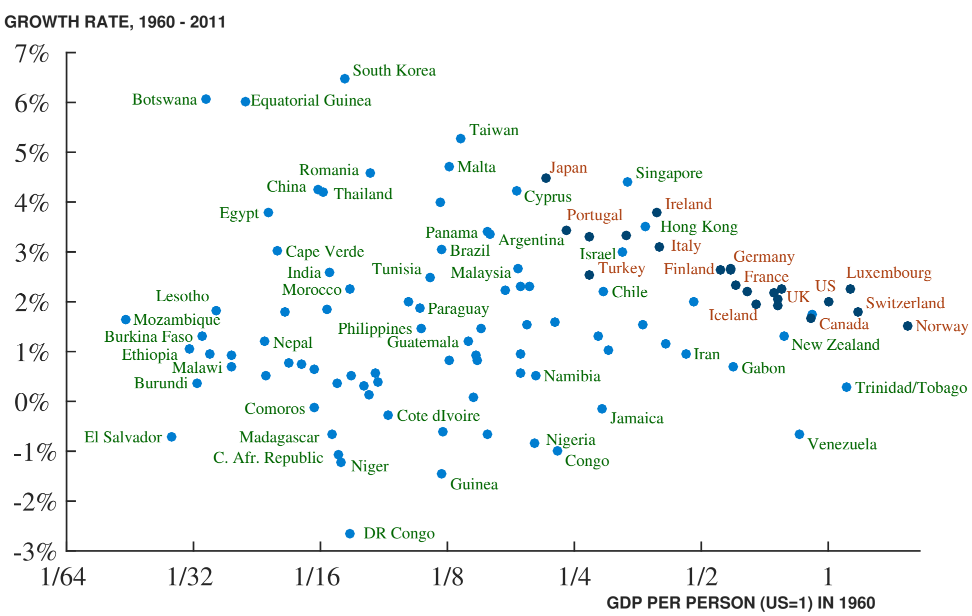

Source: The Penn World Tables 8.0.

but to their own steady-states, rather than to the frontier. If one conditions on determinants of a country's steady state (such as the investment rates in physical and human capital), then one sees that countries below their steady states grow rapidly and those above their steady states grow slowly (or even decline). The rate at which countries converge to their own steady state — often called the "speed of convergence" — seems to be around 2% per year, a fact sometimes known as "Barro's iron law of convergence." The interpretation of the OECD countries in Figure 25, then, is that these countries have relatively similar steady state positions, so that even if we do not condition on these determinants formally, the convergence phenomenon appears. Confirming this logic, the implied speed of convergence for the OECD countries estimated from the slope of the best-fit line for Figure 25 is 1.8% per year.[19]

These general patterns are examined in more detail in the following graphs and tables. Figure 27 shows the standard deviation of log GDP per person over time for this stable 107-country sample. As an alternative measure of dispersion, it also shows the ratio of GDP per person between the 5th richest and 5th poorest countries in the sample. Both measures reveal the same thing: between 1960 and the late-1990s, there was a widening of the world income distribution, at least when each country is a unit of observation. In the last decade or so, this pattern seems to have stabilized. In fact, some of this pattern was already evident back in Figure 24. The poorest countries in 1960 such as Mozambique were only about 50 times poorer than the United States. By 2011, there are many countries with relative incomes below this level, and both El Salvador and the Democratic Republic of the Congo are more than 128 times poorer than the United States.

Table 4 examines the dynamics of the distribution of incomes across countries in a more systematic fashion, following Quah (1993). First, we sort the 107 countries for which we have data in both 1960 and 2010 into bins based on their income relative to the world frontier, represented by the United States. Then, using decadal growth rates for the 5 decades between 1960 and 2010, we calculate the sample probabilities that countries move from one bin to another. Finally, we compute the stationary distribution of countries across the bins that is implied by assuming these sample probabilities are constant forever.

[19] See Barro (2012) for a recent discussion of convergence.

Figure 27: Divergence since 1960

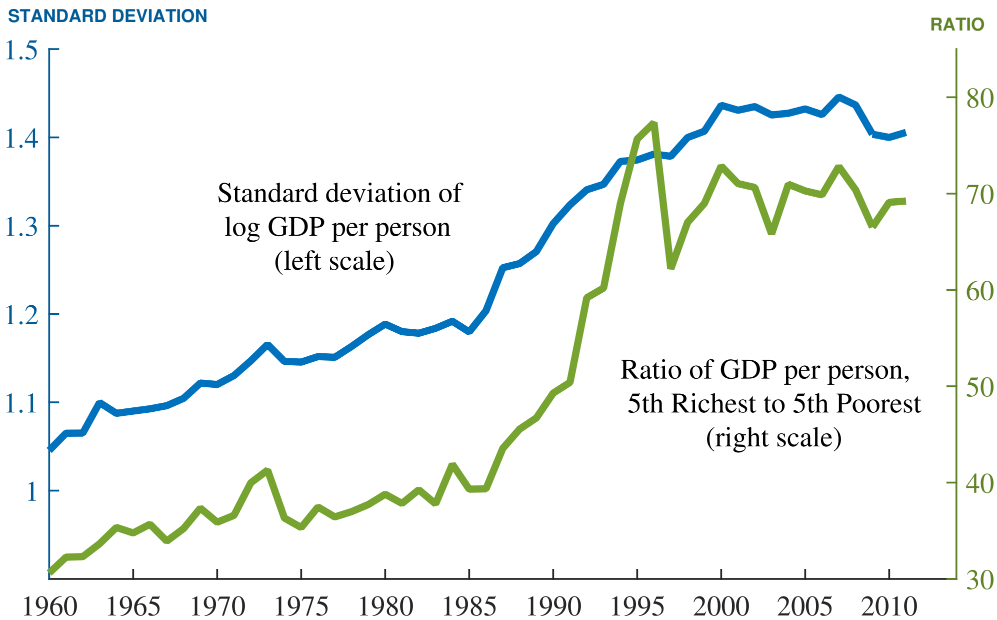

Source: The Penn World Tables 8.0, calculated across a stable sample of 107 countries.

Table 4: The Very Long-Run Distribution

	—— Distribution ——			Years to
"Bin"	1960	2010	Long-Run	"Shuffle"
Less than 5 percent	14	29	27	1470
Between 5 and 10 percent	21	12	9	1360
Between 10 and 20 percent	25	13	8	1040
Between 20 and 40 percent	18	16	8	1120
Between 40 and 80 percent	14	18	28	1450
More than 80 percent	7	12	20	1500

Entries under "Distribution" reflect the percentage of countries with relative (to the U.S.) GDP per person in each bin. "Years to Shuffle" indicates the number of years after which the best guess as to a country's location is given by the long-run distribution, provided that the country begins in a particular bin. Computed following Jones (1997) using the Penn World Tables 8.0.

First, note that some of the patterns we have remarked upon already are present in the 1960 and 2010 distributions shown in Table 4. For example, there is an increase in the fraction of countries in the highest two bins in 2010 relative to 1960. There is also a decrease in the fraction of countries between 5 and 40 percent of the U.S. level. On the other hand, there is a (surprisingly?) large increase in the fraction of countries with less than 5 percent of the U.S. income level.

Iterating over the dynamics implied by the sample transition probabilities leads to the stationary distribution shown in the third main column of the table.[20] Many more countries are projected to move out of the middle and into the top of the distribution. But there is a large mass of countries — 27 percent here — that inhabit the bin of countries with less than 5 percent of the U.S. income.

Overall, the picture that emerges from this kind of analysis is that there is a basic dynamic in the data for the last 50 years or more that says that once countries get on the "growth escalator," good things tend to happen and they grow rapidly to move closer to the frontier. Where they end up depends, as we will discuss, on the extent to which their institutions improve. And this process is itself captured in the Markov transition dynamics estimated in Table 4. It is interesting that only 20 percent of countries end up in the top bin, which may say something about how hard it is to adopt and maintain the best institutions.

Equally striking, however, is the fact that more than a quarter of countries are in the bottom bin, below 5 percent of the frontier, even in the long-run distribution. This surprised me a bit, as it didn't happen with earlier versions of this exercise, for example in Jones (1997). Nevertheless, a dynamic that appears in the latest version of the Penn World Tables is that the very poorest countries seem to be falling further behind.

One thing that may be misleading about this kind of exercise is that it implies that the stationary distribution is ergodic. That is, countries move around this distribution over time, so that, given enough time, even the U.K. can end up in the poorest bin. (The last column of the table suggests that these dynamics are very slow.) Lucas (2000), in his "Macroeconomics for the 21st Century," suggests that from the standpoint of the year 2100, the most striking fact of macroeconomics may end up being how many countries

[20]Mathematically, the computation is easily illustrated. We estimate the Markov transition probabilities of countries across the bins. Multiplying this matrix by the initial distribution yields an estimate of the distribution of income for the next decade. Doing this many times yields an estimate of the long-run distribution.

have moved close to the frontier. In other words, the Great Divergence of the last 200 years may turn into a Great Convergence over the next century. Perhaps the diffusion of good rules and good institutions leads to a more or less permanent improvement in the distribution, which is only partially captured by the kind of calculation done here.[21]

4.3. The Distribution of Income by Person, not by Country

Up until now, we've taken the country as the unit of observation. This is appropriate for some purposes, but there is also a good case to be made for taking the person as the unit of observation: why should the 500,000 people in Luxembourg get the same weight as the 1.4 billion people in China?

Figure 28 approaches the data from the standpoint of the individual. We assume each person in a country gets that country's GDP per person and then compute the world income distribution by person. More detailed calculations along these lines incorporate the distribution of income *within* each country as well and provide further support for the basic point in Figure 28; see Bourguignon and Morrisson (2002) and especially Sala-i-Martin (2006).

In 1960, 50 percent of the world's population lived on less than 3 dollars per day (measured in 2005 U.S. dollars). By 2011, only 5 percent of the world's population lived below this level. The big difference, of course, is China and India, which between them contain more than one third of the world's population. In 1960, China and India were very poor, with incomes below $2.75 per day, while by 2011 average incomes were $10 per day in India and $22 per day in China. From the standpoint of the individual most outstanding fact of economic growth over the last 50 years is how many people have been elevated out of poverty.

4.4. Development Accounting

Development accounting applies the logic of Solow's growth accounting to explain differences in the levels of GDP per worker across countries. Countries can be rich because they have large quantities of inputs or because they use these inputs efficiently. Quantitatively, how important are each of these components?

[21] Buera, Monge-Naranjo and Primiceri (2011) study the diffuson of market-oriented institutions in a setting where countries learn from the growth experiences of their neighbors.

Figure 28: The Distribution of World Income by Population

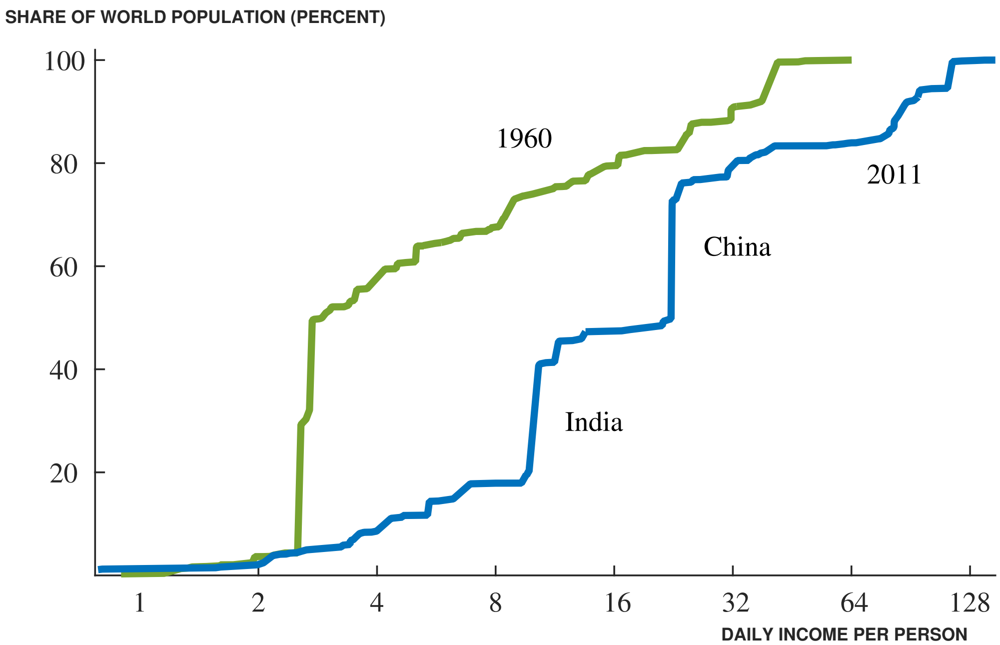

Source: The Penn World Tables 8.0, calculated across a stable sample of 107 countries.

An early version of development accounting is Denison (1967), who compared 8 European economies to the United States in 1960. Christensen, Cummings and Jorgenson (1981) work with a similar set of countries and extend the analysis to include human capital. King and Levine (1994) focus on the role of physical capital versus TFP in a broad set of countries, and provide the first use of the phrase "development accounting" that I have found.[22] Klenow and Rodriguez-Clare (1997) and Hall and Jones (1999) incorporate human capital differences and consider a broad range of countries. Caselli (2005) provides a detailed overview and analysis of this literature.[23]

The basics of development accounting follow closely upon the analysis of growth accounting that we conducted back in Section 2. To see this link, recall the production

[22] Bob Hall and I (Hall and Jones, 1996) proposed the phrase "levels accounting" which doesn't have nearly the same ring!

[23] The papers cited to this point assume a known production function — typically Cobb-Douglas with an exponent on capital around 1/3. A related set of papers including Mankiw, Romer and Weil (1992), Islam (1995), and Caselli, Esquivel and Lefort (1996) conduct a similar exercise in a regression framework. The limitation of the regression framework in its simplest form is that it imposes an orthogonality between inputs and total factor productivity which seems inappropriate. Estimating production functions is notoriously difficult.

function we considered there:

$$Y_t = \underbrace{A_t M_t}_{\text{TFP}} K_t^{\alpha} H_t^{1-\alpha}. \tag{5}$$

Some versions of development accounting work directly with this production function. The advantage is that it is the most straightforward approach. The disadvantage is familiar from growth accounting and the standard neoclassical growth model: differences in TFP induce capital accumulation that leads to differences in K across countries. Hence some of what is attributed to K in this approach might more more naturally be attributed to TFP.

An alternative approach — pursued by Klenow and Rodriguez-Clare (1997) and Hall and Jones (1999) — is to rewrite the production function in a way that assigns any induced capital accumulation to TFP. This is accomplished by dividing both sides of the production function by Y_t^{α} and solving for Y_t to get

$$\frac{Y_t}{L_t} = \left(\frac{K_t}{Y_t}\right)^{\frac{\alpha}{1-\alpha}} \frac{H_t}{L_t} \cdot Z_t \tag{6}$$

where $Z_t \equiv (A_t M_t)^{\frac{1}{1-\alpha}}$ is total factor productivity measured in labor-augmenting units. To understand why this equation assigns the induced capital accumulation to TFP, notice that in the steady state of a neoclassical growth model, the capital-output ratio K/Y is proportional to the investment rate and independent of TFP. Hence the contributions from productivity and capital deepening are separated in this version, in a way that they were not in equation (5). This was the equation on which we based our growth accounting, and we will use the same equation for development accounting.

The Penn World Tables, starting with Version 8.0, contains all the information needed to conduct the simplest form of development accounting as in equation (6). That data set contains measures of the economy's stock of physical capital and measures of human capital that are based on educational attainment data from Barro and Lee (2013) and Mincerian returns to education of 13.4 percent for the first 4 years, 10.1 percent for the second 4 years, and 6.8 percent for all additional years, as in Hall and Jones (1999). We conduct our growth accounting exercise using this data and equation (6), assuming

$\alpha = 1/3$.[24]

Table 5 shows the basic development accounting exercise using the Penn World Tables data for a sample of countries.[25] To see how the accounting works, consider the row for Mexico. According to the Penn World Tables, Mexico has a GDP per worker that is about 1/3 that in the United States in 2010. This 1/3 number is the product of the next three terms in the row, following the formula in equation (6).

Remarkably little of the difference is due to physical capital: the capital-output ratio in Mexico is about 87 percent of that in the United States. Because of diminishing returns, though, it is the square root ($\alpha/(1-\alpha) = 1/2$ when $\alpha = 1/3$) of this difference that matters for income, and $\sqrt{0.87} \approx 0.93$, so differences in physical capital only lead to about a 7 percent difference in GDP per worker between the U.S. and Mexico.

In the next column, we see a larger contribution from human capital. In 2010, people aged 15 and over in Mexico had on average around 8.8 years of education according to Barro and Lee (2013). In contrast, educational attainment in the United States was 13.2. The difference is 4.6 years of schooling. With a return to each year of education of around 7 percent, this implies about a 32 percent differnce due to education. The entry from human capital for Mexico is 0.76, and $1/0.76 \approx 1.32$, consistent with this reasoning.

The implied difference in TFP between the U.S. and Mexico is then $0.338/(0.931 \times 0.760) \approx 0.477$. Put another way, GDP per worker was 3 times higher in the U.S. than in Mexico. A factor of $1.07 \times 1.32 \approx 1.4$ of this difference is due to inputs, meaning a factor of 2.1 was due to TFP, since $1.4 \times 2.1 \approx 3$. From these numbers, one can also see easily how the "Share due to TFP" column is calculated. For each 2 parts due to inputs, 3 parts are due to TFP, hence the share due to TFP is calcuated as 60 percent (i.e. as 3/(2+3)).

More generally, several key findings stand out from Table 5. First, the capital-output ratio is remarkably stable across countries. Its average value is very close to one, and even the poorest country in the table, Malawi, is reported by the Penn World Tables to have a capital-output ratio very close to the U.S. value. So differences in physical

[24] In fact, the Penn World Tables 8.0 contains its own measure of TFP as well. However, this measure is based on Cobb-Douglas production functions in which the exponents on capital vary across countries. In this sense, the measure of TFP that is reported there is incomplete: countries with the same inputs and the same level of TFP will have different outputs!

[25] Data on all countries can be obtained in the data files available on the author's web page; see the file "DevelopmentAccounting.log."

Table 5: Basic Development Accounting, 2010

	GDP per worker, y	Capital/GDP $(K/Y)^{\alpha/(1-\alpha)}$	Human capital, h	TFP	Share due to TFP
U.S.	1.000	1.000	1.000	1.000	...
Hong Kong	0.854	1.086	0.833	0.944	48.9%
Singapore	0.845	1.105	0.764	1.001	45.8%
France	0.790	1.184	0.840	0.795	55.6%
Germany	0.740	1.078	0.918	0.748	57.0%
U.K.	0.733	1.015	0.780	0.925	46.1%
Japan	0.683	1.218	0.903	0.620	63.9%
South Korea	0.598	1.146	0.925	0.564	65.3%
Argentina	0.376	1.109	0.779	0.435	66.5%
Mexico	0.338	0.931	0.760	0.477	59.7%
Botswana	0.236	1.034	0.786	0.291	73.7%
South Africa	0.225	0.877	0.731	0.351	64.6%
Brazil	0.183	1.084	0.676	0.250	74.5%
Thailand	0.154	1.125	0.667	0.206	78.5%
China	0.136	1.137	0.713	0.168	82.9%
Indonesia	0.096	1.014	0.575	0.165	77.9%
India	0.096	0.827	0.533	0.217	67.0%
Kenya	0.037	0.819	0.618	0.073	87.3%
Malawi	0.021	1.107	0.507	0.038	93.6%
Average	0.194	0.978	0.694	0.286	64.3%
1/Average	5.146	1.022	1.440	3.496	70.4%

Computed using the Penn World Tables 8.0 for the year 2010 assuming a common value of $\alpha = 1/3$. The product of the three input columns equals GDP per worker. The penultimate row, "Average," shows the geometric average of each column across 134 countries. The "Share due to TFP" column is computed as described in the text. The 70.4% share in the last row is computed looking across the columns, i.e. as approximately 3.5/(3.5+1.5).

capital contribute almost nothing to differences in GDP per worker across countries. Caselli and Feyrer (2007) document a closely-related fact in great detail: the marginal product of capital (which here is proportional to the inverse of the capital-output ratio) is very similar in rich and poor countries.[26]

Second, the contribution from educational attainment is larger, but still modest. For example, countries like India and Malawi only see their incomes reduced by a factor of 2 due to educational attainment. Loosely speaking, the poorest countries of the world have 4 or 5 years of education, while the richest have 13. Eight years of education with a Mincerian return of around 10 percent leads to a 0.8 difference in logs, and $exp(0.8) \approx 2$.

Finally, the implication of these first two points is that differences in TFP are the largest contributor to income differences in an accounting sense. Figure 29 shows the levels of TFP plotted against GDP per worker for 134 countries in 2010. The two series are highly correlated at 0.96. And the differences in TFP are very large: the Democratic Republic of the Congo is about 128 times poorer than the United States and its TFP is about 64 times lower than the U.S. level.

The large contribution from TFP is verified by the last column of Table 5, where the share explained by TFP ranges from just under 50 percent for Singapore and Hong Kong to more than 90 percent for Malawi. To understand the "Share due to TFP" column, consider the last row of Table 5. According to that row, the average country in the 134-country sample is just over 5 times poorer than the United States. Essentially none of this difference (a factor of 1.022) is due to differences in K/Y, while a factor of 1.44 is due to differences in educational attainment. Taken together, this means a factor of $1.022 \times 1.44 \approx 1.5$ is due to inputs, leaving a factor of about 3.5 attributed to TFP. We then compute the "Share due to TFP" as $3.5/(1.5+3.5) \approx 70\%$, as shown in the last entry in Table 5. The rest of the shares are computed in an analogous way. For example, for Malawi, about a factor of 2 is due to inputs and a factor of 26 is due to TFP, meaning that $26/28 \approx 93\%$ is due to TFP.

More generally, the share across all 134 countries is shown in Figure 30. There, a systematic pattern is obvious. In the poorest countries of the world, well over 80 percent of the difference in GDP per worker relative to the United States is due to TFP

[26]The general lack of correlation between the capital-output ratio and GDP per person is discussed by Inklaar and Timmer (2013).

Figure 29: Total Factor Productivity, 2010

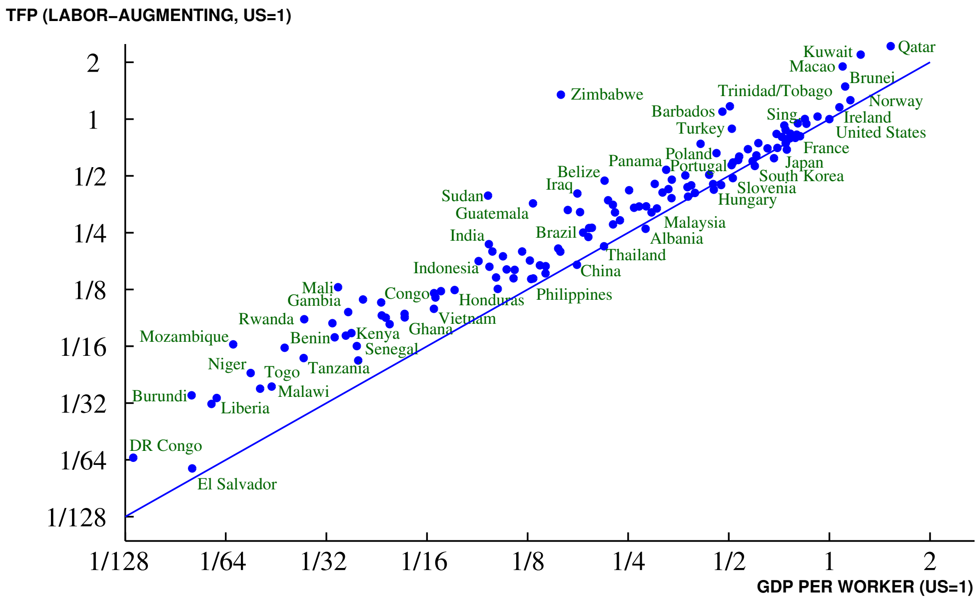

Source: Computed using the Penn World Tables 8.0 assuming a common value of $\alpha = 1/3$.

differences. Moving across the graph to richer countries, one sees that less and less is due to TFP, and for the richest countries as a whole, TFP contributes around 50 percent of the differences.

4.5. Understanding TFP Differences

The basic finding that TFP differences account for the bulk of income differences across countries has led to a large body of research designed to explain what these differences are. This is exemplified by the title of a famous paper by Prescott (1998): "Needed: A Theory of TFP."

In the last 15 years, this challenge has been approach in two ways. First, several papers have improved our measures of inputs in various ways, chipping away at the contribution of the "measure of our ignorance." Second, the literature on misallocation has emerged to provide in general the kind of theory that Prescott was seeking. The remainder of this section will review the efforts to improve input measurement, while the next several sections will consider misallocation and its implications.

Figure 30: The Share of TFP in Development Accounting, 2010

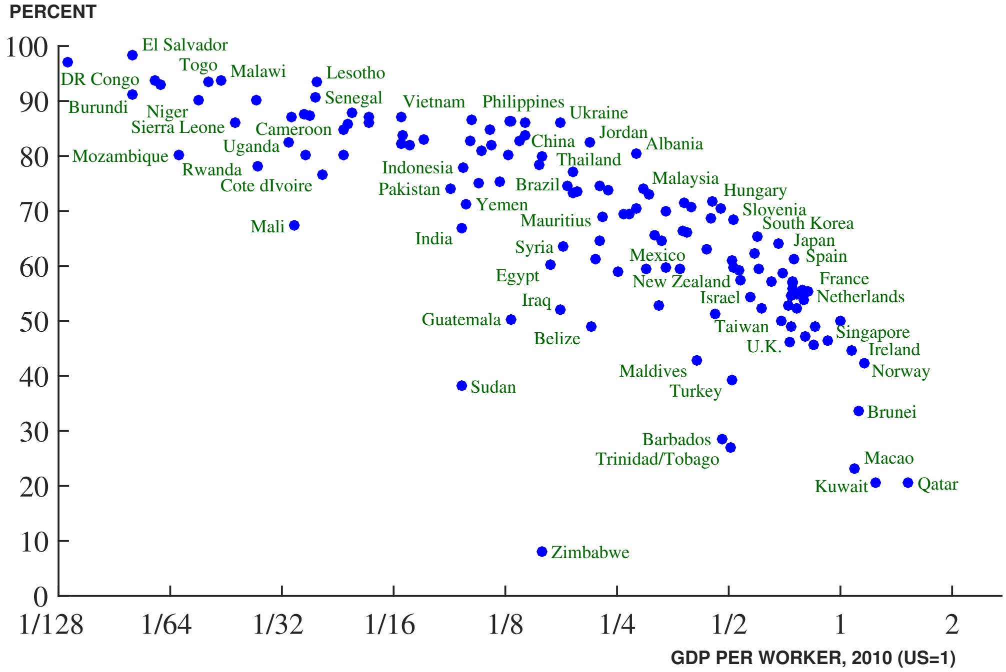

Source: Computed as described in the text and in Table 5 using the Penn World Tables 8.0 assuming a common value of $\alpha = 1/3$.

Caselli (2005) provides a detailed survey and analysis of the state of development accounting as of 2005. The interested reader should certainly look there to get up to speed. Caselli reviews progress on many dimensions: measuring the quality of education using test scores (Hanushek and Kimko, 2000); considering differences in the experience of the labor force across countries (Klenow and Rodriguez-Clare, 1997); sectoral differences in productivity, especially agriculture (Restuccia, Yang and Zhu, 2008); differences in labor productivity due to health (Weil, 2007); differences in the quality of capital (Caselli and Wilson, 2004); and the potential role of non-neutral productivity (Caselli and Coleman, 2006).

Much additional progress has been made in the decade since Caselli's review was published, especially with respect to misallocation, as discussed in the next section. But there has also been much valuable work on measuring the inputs into development accounting. Lagakos, Moll, Porzio, Qian and Schoellman (2012) use household survey data from 35 countries to show that the returns to experience vary substantially across countries, with poorer countries typically having much flatter age-earnings profiles. Incorporating differential returns to experience in development accounting boosts the importance of the human capital term by about 50 percent. Hendricks and Schoellman (2014) use data on immigrants from 50 source countries into 11 different host countries to improve our measurement of labor quality differences, providing another boost to the human capital term — of about 30 percent. Ben Jones (2014) observes that the standard approach outlined above treats workers with different levels of education as perfect substitutes, up to differences in efficiency units. He proposes instead a generalized aggregator and shows that the traditional perfect-substitutes case delivers a lower bound for the role of human capital. If workers with different human capital are less than perfect substitutes, the share of income differences explained by human capital can rise dramatically.[27]

[27] This finding is related to an observation made by Caselli and Coleman (2006). They noted that the ratio of "skilled" to "unskilled" workers varies enormously across countries. For example, if we let high school completion be the dividing line, the ratio of skilled to unskilled workers is just 0.025 in Kenya versus 1.8 in the United States — a difference of a factor of 70. If college completion is the dividing line, the factor proportions differ by even more. When workers with different human capital levels are no longer perfect substitutes, this ratio becomes relevant. The difficulty is that it can then imply implausibly large differences in the return to schooling across countries if one is not careful. Caselli and Coleman introduce additional TFP terms so they can match the returns to education, but then the large differences in factor proportions shows up as enormous differences in the non-neutral TFP terms.

4.6. Misallocation: A Theory of TFP

One of the great insights of the growth literature in the last 15 years is that misallocation at the micro level can show up as a reduction in total factor productivity at a more aggregated level. This insight appears in various places, including Banerjee and Duflo (2005), Chari, Kehoe and McGrattan (2007), Restuccia and Rogerson (2008) and Hsieh and Klenow (2009).

As we discussed briefly in the context of misallocation and frontier growth (in Section 2.6.), the essence of this insight is quite straightforward: when resources are allocated optimally, the economy will operate on its production possibilities frontier. When resources are misallocated, the economy will operate inside this frontier. But that is just another way of saying that TFP will be lower: a given quantity of inputs will produce less output.

A simple example illustrates this point. Suppose output is produced using two tasks according to $Y = X_1^{\alpha} X_2^{1-\alpha}$. This could describe a firm, and the tasks could be management and the production line, or this could be the economy as a whole and the tasks could be manufacturing and services. Suppose that each task is accomplished very simply: one unit of labor can produce one unit of either task, and the economy is endowed with L units of labor. Finally, suppose that the allocation of labor is such that a fraction s works in the first task, and the fraction $1-s$ works in the second task. We leave the source of this allocation unspecified: labor could be optimally allocated, or it could be misallocated because of taxes, poor management, information problems, unions, or many other reasons. But given this allocation, there is a reduced-form production function given by

$$Y = M(s)L \quad \text{where} \quad M(s) \equiv s^{\alpha}(1-s)^{1-\alpha} \tag{7}$$

In other words, total factor productivity in this economy is $M(s)$, which depends on the allocation of labor in the economy.[28] Moreover, it is easy to see that the output-maximizing allocation of labor in this example has $s^* = \alpha$, and any departure of the allocation from s^* will reduce total factor productivity. This is the essence of the literature on misallocation and TFP.

[28]One could easily assume both capital and labor are used to produce each X so that the result in equation (7) would be $Y = M(s)K^{\alpha}L^{1-\alpha}$, which makes the connection between $M(s)$ and TFP even more apparent.

This insight is at the heart of our best candidate explanations for answering the question of why some countries are so much richer than others. Development accounting tells us that poor countries have low levels of inputs, but they are also remarkably inefficient in how they use those inputs. Misallocation provides the theoretical connection between the myriad of distortions in poor economies and the TFP differences that we observe in development accounting.

The remainder of this section explores various facts related to misallocation.

4.7. Institutions and the Role of Government

Countries are a natural unit of analysis for growth economists for the simple reason that national borders are the places where different political and economic institutions begin and end. It has long been conjectured that differences in these institutions are fundamental determinants of long-run economic success. But what is the evidence for such a claim? How do we know that the income differences we see across countries are not primarily driven by differences in natural resources or other aspects of geography?

One of the best sources of evidence on this question was provided by Olson (1996). Olson observed that history itself provides us with "natural experiments" that allow us to see the large impact of institutions on economic success. For example, prior to World War II, North and South Korea were not separate countries. As a rough approximation, the north and south of Korea contained people that shared a cultural heritage, a government, institutions, and even a geography. In fact, if anything, North Korea was economically advantaged relative to the South, containing a disproportionate share of electricity production and heavy industry.[29] After the Korean War ended in 1953, North and South Korea were divided and governed according to very different rules. The resulting economic growth of the next half century was dramatically different, as illustrated in Figure 31. The picture in this figure was taken by an astronaut on the International Space Station in early 2014 and shows North and South Korea at night.

[29] "[Under Japanese rule before World War II], the colonial industries were unevenly distributed between South Korea and North Korea. Heavy and chemical industries were concentrated in the North, while many light industries, such as textile, food, printing and wood, were located in the South. In 1940, North Korea's share of the total production in the metal industry was 96 percent, and 82 percent for the chemical industry. Also, 92 percent of the total electricity production originated from the North in 1945 (Lee, D-G, 2002: 39). Thus, in 1945 when Japan withdrew from Korea and when Korea was divided into two separate political regimes, the South Korean economy in general and industry in particular were severely crippled." Yang (2004), p. 16.

Figure 31: Korea at Night

Note: North Korea is the dark area in the center of the figure, between China to the north and South Korea to the south. Pyongyang is the isolated cluster in the center of the picture. Source: http://commons.wikimedia.org/wiki/File:North_and_South_Korea_at_night.jpg

South Korea is brightly lit, while North Korea is almost completely dark, barely indistinguishable from the ocean. Whatever was different between North and South Korea after 1953 apparently had an enormous influence on their long-term economic success.

Similar "natural" experiments can be observed in East and West Germany after World War II, Hong Kong and southeastern China, and across the Rio Grande between Mexico and Texas. These examples make clear that something malleable matters for economic success even if they do not specifically identify what that something is.

Another fascinating piece of evidence comes from Acemoglu, Johnson and Robinson (2002) and is illustrated in Figure 32, the so-called "reversal of fortune." Restricting our attention to former European colonies, economic success 500 years ago is *negatively* correlated with economic success today. That is, the places that were most successful 500 years ago, as measured by population density or urbanization, are on average comparatively poor today.

A classic example of this phenomenon, highlighted by Engerman and Sokoloff (1997),

Figure 32: The Reversal of Fortune

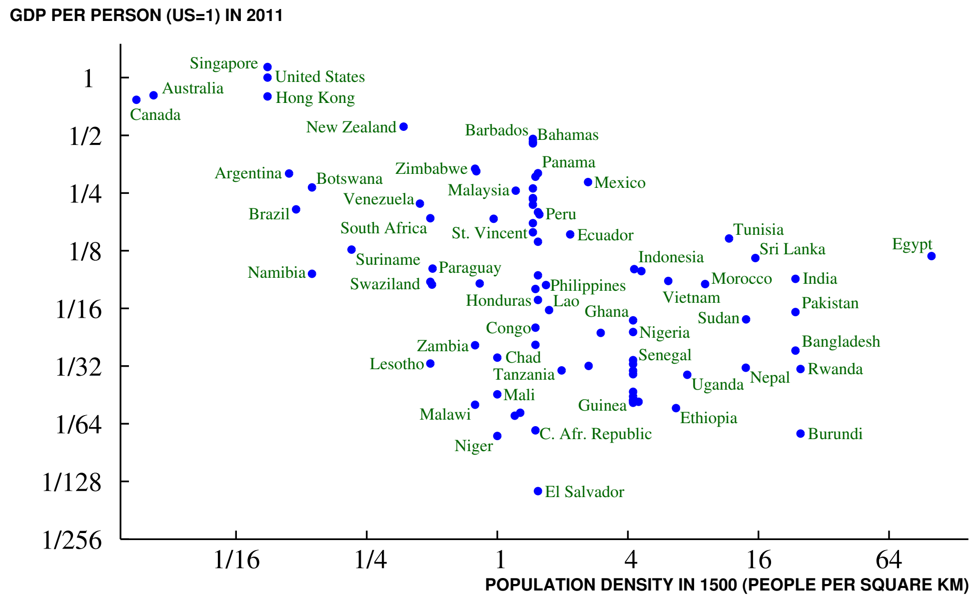

Note: Source: Population density is from Acemoglu, Johnson and Robinson (2002) and GDP per person is from the Penn World Tables 8.0.

is the New World:

> [Latin America] began with — by European standards of the time — vast supplies of land and natural resources per person and were among the most prosperous and coveted of the colonies in the seventeenth and eighteenth centuries. Indeed, so promising were these other regions that Europeans of the time generally regarded the thirteen British colonies on the North American mainland and Canada as of relatively marginal economic interest — an opinion evidently shared by Native Americans who had concentrated disproportionately in the areas the Spanish eventually developed. Yet, despite their similar, if not less favorable, factor endowments, the United States and Canada ultimately proved to be far more successful than the other colonies in realizing sustained economic growth over time. (pp. 260–261)

These examples suggest that economic success is not permanently given, for example by geographic endowments, but rather can be changed by the rules that are put in place. Engerman and Sokoloff (1997), Acemoglu, Johnson and Robinson (2002), and others suggest that the institutions adopted by Europeans in response to these initial conditions influenced subsequent growth. In places that were already economically successful in 1500, Europeans tended to set up "extractive" institutions to transfer the economic gains back to Europe. In contrast, Europeans themselves migrated to places that were sparsely populated in 1500, setting up "European" institutions that were conducive to long-run economic success.[30]

Dell (2010) provides a detailed analysis of the long-reaching nature of one such institution in Peru. The "mita" was a forced labor system conscripting one seventh of the adult male population in a region of Peru to work in local silver and mercury mines between 1573 and 1812. A regression discontinuity analysis reveals that today — Two hundred years later after this system ended — consumption is lower inside the former mita by 25 percent, educational attainment is lower, and the region is less well-connected by roads and infrastructure.

4.8. Taxes and Economic Growth

One of the most obvious and readily quantified measures of government involvement in the economy is taxes. It is easy to write down models in which governments that tax heavily reduce the long-run success of their economies. The facts, however, are not so clear.

Figure 33 shows the growth rate of real GDP per person in the United States since 1980 as well as the total government tax revenues (including state and local) as a share of GDP, updating a graph first highlighted by Stokey and Rebelo (1995). The stunning fact that emerges from this graph is that taxes have increased enormously, from around 10 percent of GDP in 1929 to more than 30 percent of GDP at their peak in 2000. But as we already noted earlier, growth rates over the 20th century were remarkably stable — if anything, they were higher after 1950 than before.

Figure 34 shows a related fact by looking across the countries of the world: tax revenues as a share of GDP are *positively* correlated with economic success, not negatively correlated.

Of course, these are just simple correlations, and they have an obvious interpreta-

[30]For example, see also Hall and Jones (1999), Acemoglu, Johnson and Robinson (2001), and Acemoglu and Robinson (2012).

Figure 33: Taxes and Growth in the United States

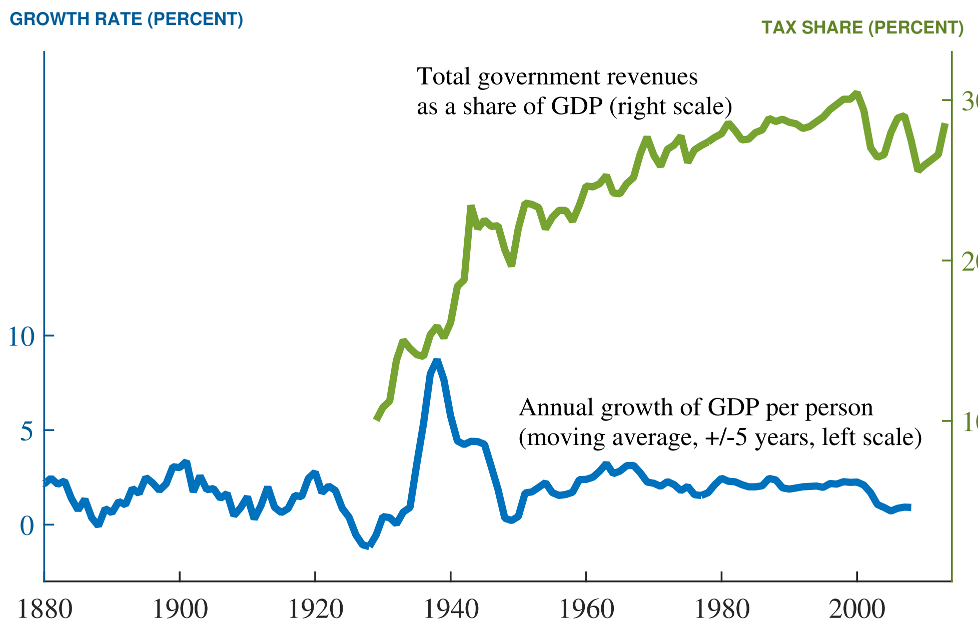

Note: This graph updates a similar figure in Stokey and Rebelo (1995). Total government current receipts are from NIPA Table 3.1 via the FRED database and include federal, state, and local revenues. Real GDP per person is constructed as in Figure 1. The growth rate is smoothed by taking a moving average across the 5 years before and after the relevant date.

Figure 34: Tax Revenues as a Share of GDP

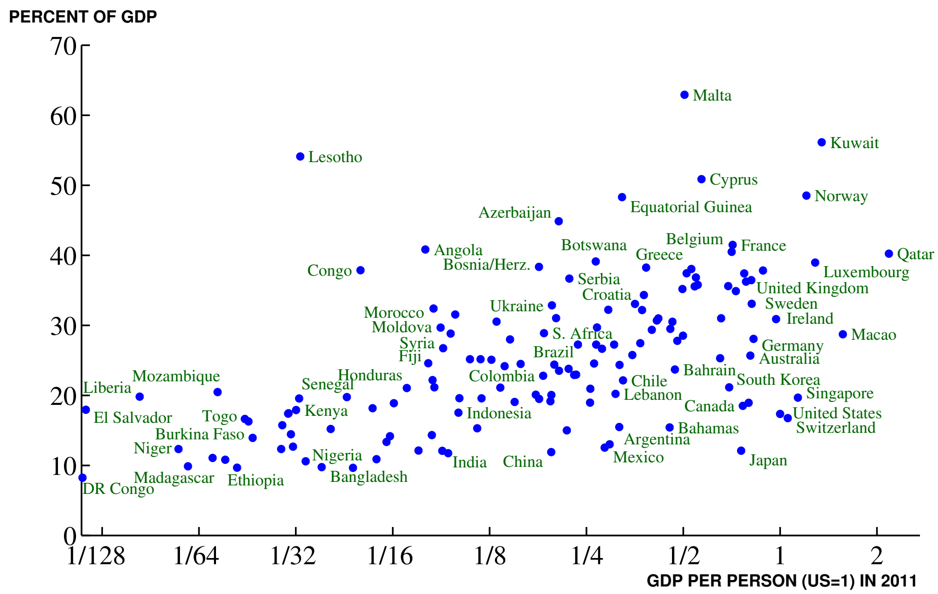

Note: Tax revenue is averaged for the available years between 2000 and 2014, is for the central government only, and includes receipts for social insurance programs. This is an updated graph of a figure from Acemoglu (2005). Source: The World Bank, *World Development Indicators*. GDP per person is from the Penn World Tables 8.0.

tion. Governments do not simply throw the tax revenue that they collect into the ocean. Instead, this revenue — at least to some extent — is used to fund the good purposes that governments serve: providing a stable rule of law, a judicial system, education, public health, highways, basic research, and so on.

4.9. TFPQ versus TFPR

An important realization related to the measurement of either labor productivity or TFP emerged during the last decade. Specifically, to measure true productivity, one needs detailed information on micro level prices. Foster, Haltiwanger and Syverson (2008) introduced the labels "TFPQ" and "TFPR," which will be explained in detail below. This distinction plays a crucial role in Hsieh and Klenow (2009). For more discussion, see the recent survey by Syverson (2011).

To see this point most easily, consider the following setup. The economy consists of a unit measure of heterogeneous varieties that enter the utility function via a CES aggregator:

$$C = \int_0^1 (\alpha_i Y_i)^\rho di \tag{8}$$

where α_i are taste parameters related to each variety and $0 < \rho < 1$ governs the elasticity of substition between varieties.

Each variety is assumed to be produced by different monopolistically-competitive firms using labor:

$$Y_i = A_i L_i \tag{9}$$

where A_i is the (exogenous) productivity with which each variety is producted. Finally, assume labor is homogenous and can be hired by any firm at a wage rate w.

It is well-known that in this kind of setup, monopolistically-competitive firms charge a price p_i for their variety that is a markup over marginal cost:

$$p_i = \frac{1}{\rho} \cdot \frac{w}{A_i}. \tag{10}$$

This implies that sales revenue for each firm is $p_i Y_i = wL_i/\rho$.

Now consider measuring firm-level productivity in this economy. Since we've left capital out of this example, we focus on labor productivity. But exactly the same issues apply to TFP as well.

In general, the data we have available on firms include sales revenues p_iY_i and employment L_i. This leads immediately to an important point: if one does not have data on the firm-level price p_i, then one cannot recover A_i. For example, deflating revenue by an industry-level price deflator is not the same as deflating by p_i because firms are heterogeneous and have different productivity levels. In the absence of firm-level prices, one typically recovers

$$\text{Revenue Productivity, TFPR}_i: \quad \frac{p_iY_i}{L_i} = \frac{w}{\rho}.$$

If firms have identical markups and do not face any distortions, revenue productivity should be equated across heterogeneous firms. Workers have to earn the same wage at each firm, and this equates the marginal revenue product of labor across firms, which is all that is being recovered here.

The same argument applies to total factor productivity, and gives rise to the label TFPR, where the "R" denotes "revenue" TFP. The marginal revenue product of capital should also be equated across firms in a simple model like this one, so weighted averages of the average revenue products of capital and labor — which is what TFPR is — should be equated across firms.

By this point, the reader should realize that in a world of heterogeneous goods, it is not even obvious how to compare "true productivity." How do we compare Ford's productivity in making pickup trucks to Tesla's productivity in making electric cars? Or how do we compare Dell's productivity at making PC's with Apple's productivity at making Macs? Even if we had detailed data on the price of Ford trucks and Tesla cars — even if we recovered the A_i's using these prices — they would not be comparable, because the products are different!

Both of these issues are addressed by having knowledge of the utility function, in this case the C aggregator in equation (8). In particular, knowledge of the utility function allows one to compute the marginal rate of substitution between different products — it tells us how to compare Fords and Teslas or Apples and Dells.

To see how, recall that the demand curve from utility maximization of (8) is

$$p_i = \rho\alpha_i^{\rho}Y_i^{\rho-1}. \tag{11}$$

Therefore, sales revenue for this variety is

$$p_i Y_i = \rho(\alpha_i Y_i)^\rho \tag{12}$$

and therefore we can invert this equation to recover the term that enters utility from sales revenue:

$$\alpha_i Y_i = \left(\frac{p_i Y_i}{\rho}\right)^{1/\rho}. \tag{13}$$

And this is what we require. The α_i terms tell us how to relate Fords and Teslas:[31]

$$\text{True Productivity, TFPQ}_i: \quad \frac{\alpha_i Y_i}{L_i} = \frac{(p_i Y_i)^{1/\rho}}{L_i} = \alpha_i A_i.$$

That is, TFPQ (the "Q" denotes "quantity" TFP) tells us how effective a firm is at taking a unit of labor (in this case, with no capital) and using it to produce Fords or Teslas in comparable units.

Notice that TFPQ is measured in this case using only the data we typically have — sales revenue and employment — together with knowledge of the elasticity of demand.

TFPQ reflects A_i but TFPR is independent of A_i in this simple specification.[32]

4.10. The Hsieh-Klenow Facts

Hsieh and Klenow (2009) use the insight that TFPR should be equated across plants if resources are optimally allocated to quantify the effect of misallocation on aggregate TFP. In particular, they consider misallocation across plants within 4-digit manufacturing industries in the United States, China, and India.

The first part of their approach has already been explained in the previous section: they assume CES demand, and use this elasticity to back out prices and real output from sales revenue. This allows them to measure "true" TFP, called "TFPQ," for each establishment in their data. The average distribution of TFPQ within 4 digit manufacturing that they recover is shown in Figure 35.

As shown in this figure, the distributions within industry of TFPQ in the United States and China are relatively similar, while the distribution is significantly wider in

[31] I'm dropping the $\rho^{1/\rho}$ term in the equation to keep things clear.
[32] In richer settings, TFPR can depend on A_i, for example if there are fixed costs in production.

Figure 35: The Distribution of TFPQ in 4-digit Manufacturing Industries

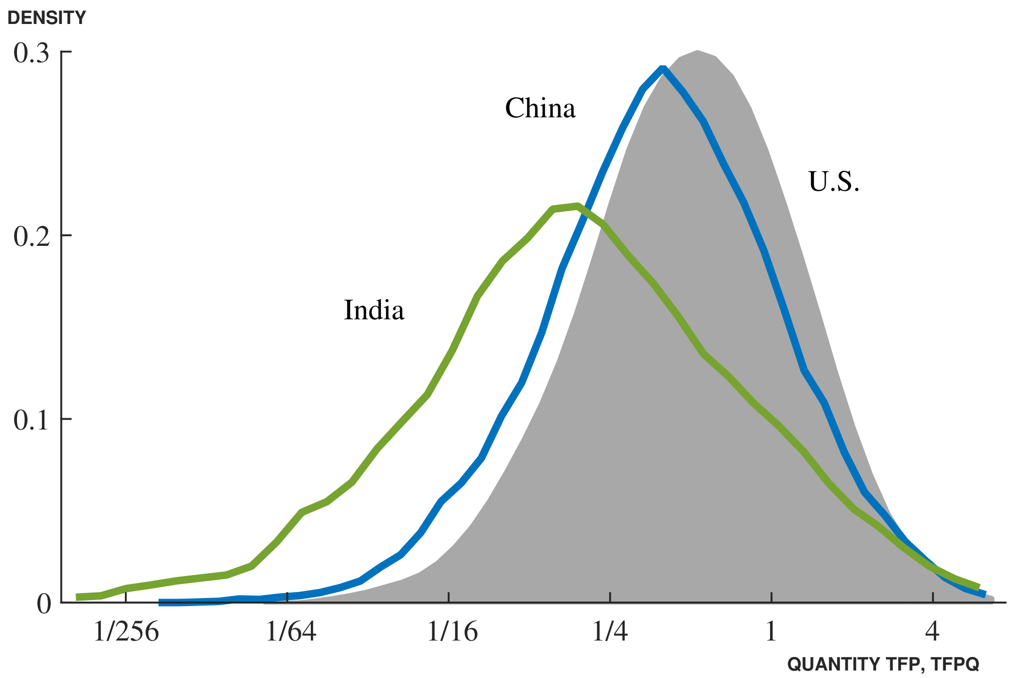

Note: This is the average distribution of TFPQ within 4-digit manufacturing industries for the U.S. in 1997, China in 2005, and India in 1994, computed as described in the text. The means across countries are not meaningful. Source: Hsieh and Klenow (2009); data provided by Chang Hsieh.

India.[33] What is surprising, perhaps, is just how large the differences in TFPQ are within an industry. Hsieh and Klenow (2009) find that the 90-10 ratio of TFPQ across plants is 8.8 in the United States, 22.4 in India, and 11.5 in China. One way of thinking about these large differences is to note that employment differences across plants is very large. Why does a large textile manufacturer employ many more workers than a family shop? One answer is that TFPQ is much higher for the large plant.

Hsieh and Klenow's most valuable contribution, however, is to quantify misallocation. To see how they do this, Consider a plant that produces with a Cobb-Douglas production function, using capital and labor, and that faces distortions τ_K and τ_L in choosing its inputs. These distortions are modeled as if they are taxes, but the literature interprets the distortions much more broadly to include credit market frictions, hiring and firing costs, quantity restrictions, and so on. The profit-maximizing firm will

[33] However, the authors note that small firms are underrepresented in the Chinese data, so this could reflect differences in the sample.

hire capital and labor until the marginal revenue product of these factors equals their net-of-distortion rental price. Written differently, payments to factors will equal the product of the factor exponents and the distortion terms:

$$\alpha_K(1-\tau_K) = \frac{(r+\delta)K_i}{p_i Y_i} \tag{14}$$

and

$$\alpha_L(1-\tau_L) = \frac{wL_i}{p_i Y_i}. \tag{15}$$

Roughly speaking, Hsieh and Klenow observe the right-hand side of these expressions in the data — they observe the share of revenues spent on labor and capital. They then use these observed spending shares to back out the distortions.

The key identification issue in recovering the τ_K and τ_L is this: when a manufacturing plant pays a large fraction of its revenue to labor, is that because it faces a low τ_L, or is that because its technology is labor intensive (a high α_L)? Hsieh and Klenow solve this identification problem by assuming that all firms within a 4-digit industry have common α_K and α_L Cobb-Douglas exponents. Then variation in factor shares across plants reflects distortions rather than technologies. (This is one of the reasons why their approach works well within industries but would run into problems across industries.)

TFPR is a weighted average of the marginal revenue product of capital and the marginal revenue product of labor, relative to the average values in the industry, and with no distortions would take a value of one. It is therefore also equal to a weighted average of the $1-\tau_K$ and $1-\tau_L$ distortions, where the weights are the Cobb-Douglas exponents in the production function.[34]

The average distribution of TFPR within 4-digit manufacturing industries is shown in Figure 36. The first thing to note about this figure is that TFPR is not equal to one for every firm, not even in the United States. One interpretation of this fact is that resources are misallocated even in the United State and to use the deviations from one as a measure of U.S. misallocation. Another interpretation — and these likely both contain elements of truth — is that there is measurement error in the U.S. data, and

[34] For aggregating the τ_K and τ_L into a single index of "TFPR"and to measure the effect of the distortions on TFP — Hsieh and Klenow actually require values for the Cobb-Douglas exponents. They assume the U.S. average shares are undistorted and assume China and India have the same Cobb-Doulgas technology as the U.S.

Figure 36: The Distribution of TFPR in 4-digit Manufacturing Industries

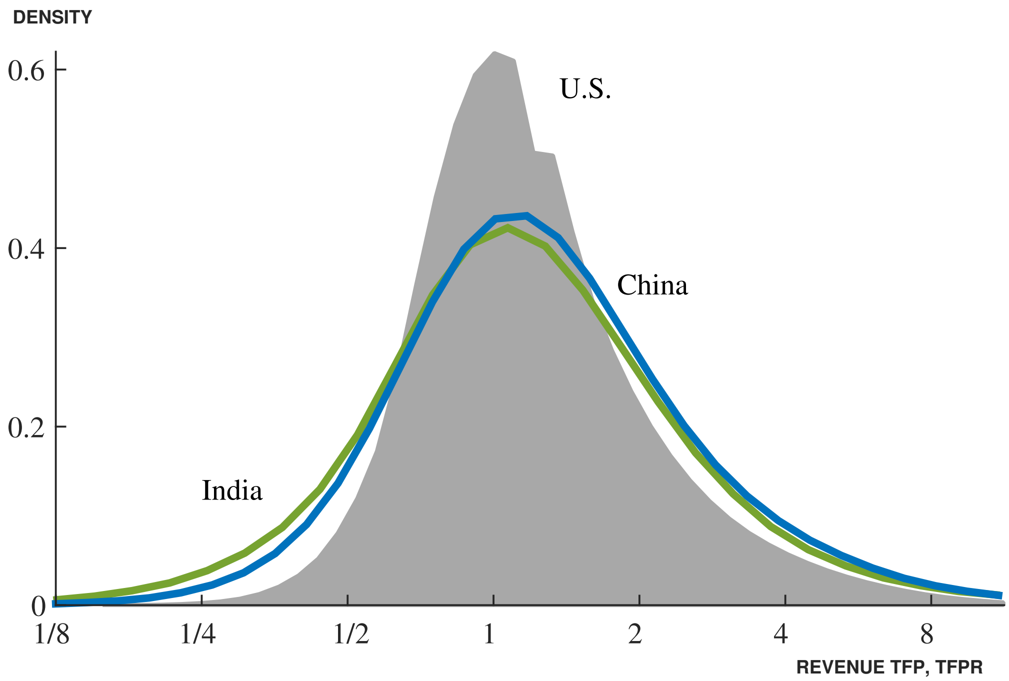

Note: This is the average distribution of TFPR within 4-digit manufacturing industries for the U.S. in 1997, China in 2005, and India in 1994, computed as described in the text. Source: Hsieh and Klenow (2009); data provided by Chang Hsieh.

some of the deviations away from one reflect this measurement error.

The second point to note in Figure 36 is that the dispersion of TFPR in India and China is significantly larger than the dispersion in the United States. To the extent that this does not reflect larger measurement error in India and China, it suggests that the misallocation of capital and labor across establishments within 4-digit industries in China and India is a factor reducing GDP in those economies. Hsieh and Klenow quantify these effects and find that if China and India has the same dispersion of TFPR as the United States, their aggregate TFP would be higher by 30% to 50% in China and 40% to 60% in India. GDP would be higher by approximately twice this amount if capital accumulates in response to the higher TFP.

In a recent follow-up paper, Hsieh and Klenow attempt to understand what could be causing this misallocation. Hsieh and Klenow (2014) looks at how establishments in the United States, India, and Mexico grow as they age. The remarkable finding of this paper is summarized in Figure 37: plants in the United States get much larger as they

Figure 37: Average Employment over the Life Cycle

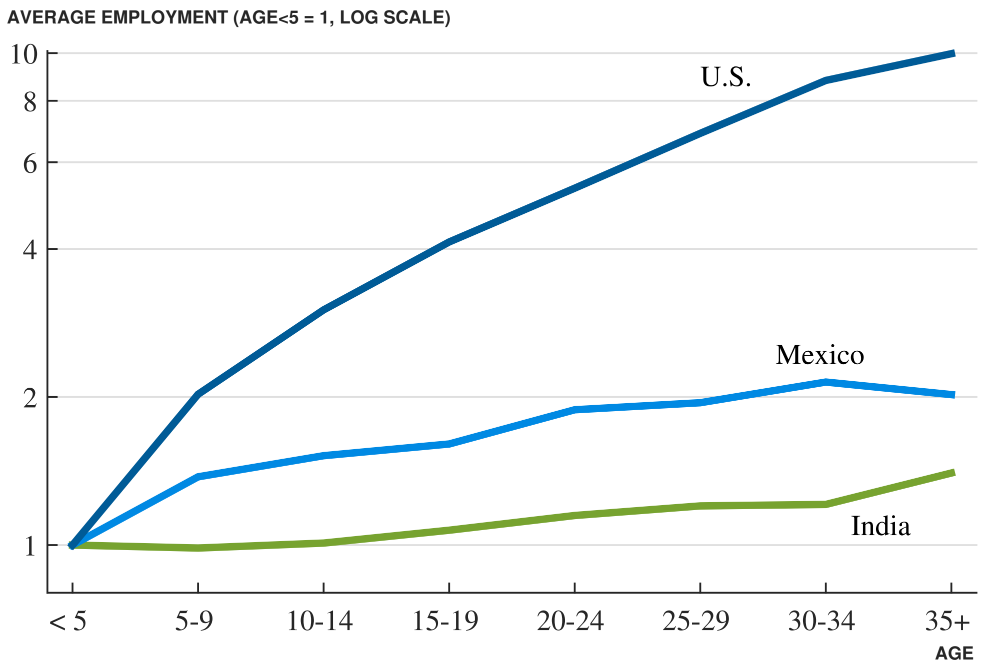

Note: The graph compares average employment per surviving plant in a later year to average employment per operating plant in an earlier year from the same cohort using census data for the manufacturing industry in the United States, Mexico, and India. Source: Hsieh and Klenow (2014); data provided by Chang Hsieh.

age, while this is barely true at all in India.

To be more precise, plants that are more than 35 years old in the U.S. have more than 8 times the employment of plants that are less than 5 years old. In contrast, old plants in Mexico are only twice as large as young plants, while plants in India exhibit even less employment growth. The suggestion, explored in detail in this paper, is that distortions in Mexico and India prevent the most productive plants from growing in size, and this is one cause of the lower aggregate TFP in these economies. They estimate that moving from the U.S. life cycle to the Indian or Mexican life cycle of plant growth could reduce aggregate TFP by about 25 percent.

Motivated by facts like these, a growing number of recent papers explore various kinds of misallocation their effects on TFP. Asker, Collard-Wexler and Loecker (2011) examine the role of volatility and adjustment costs in explaining variation in TFPR and TFPQ. Midrigan and Xu (2014) and Moll (2014) study the extent to which credit market

Figure 38: Technology Adoption is Speeding Up Over Time

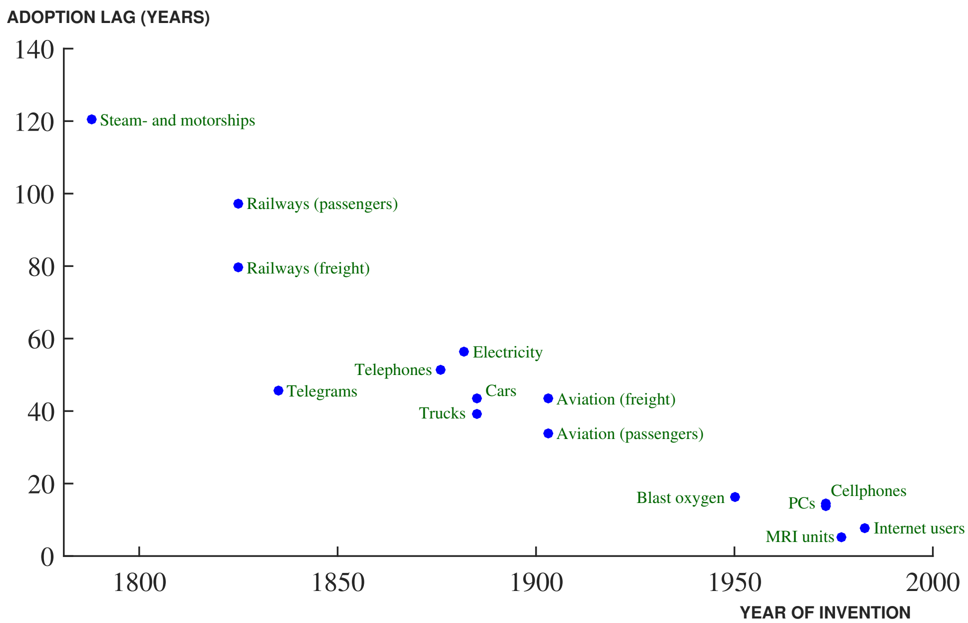

Note: Adoption lags for each country measure the amount of time between when a technology is invented and when it was adopted in the country. The figure reports averages estimated across 166 countries spanning the period 1820 to 2003. Source: Comin and Hobijn (2010)

.

frictions can generate misallocation and TFP losses. Peters (2013) considers the role of heterogeneous markups in accounting for misallocation. Guner, Ventura and Yi (2008) and Garicano, Lelarge and Van Reenen (2014) consider the effect of regulations tied to the size of firms. Akcigit, Alp and Peters (2014a) suggest that incentive problems for managers limit the ability of potentially highly-productive small firms to expand, leading to lower TFP.

4.11. The Diffusion of Ideas

Figure 38 shows our next fact: lags in the adoption of new technologies have declined sharply over the last 200 years. This fact is taken from Comin and Hobijn (2010) and is based on the CHAT ("Cross-country Historical Adoption of Technology") database that these authors previous assembled. The database contains information on the diffusion of more than 100 technologies, in more than 150 countries, since 1800.

For 15 technologies, the graph plots the year of invention of each technology versus

Figure 39: The Number of “Million Cities”

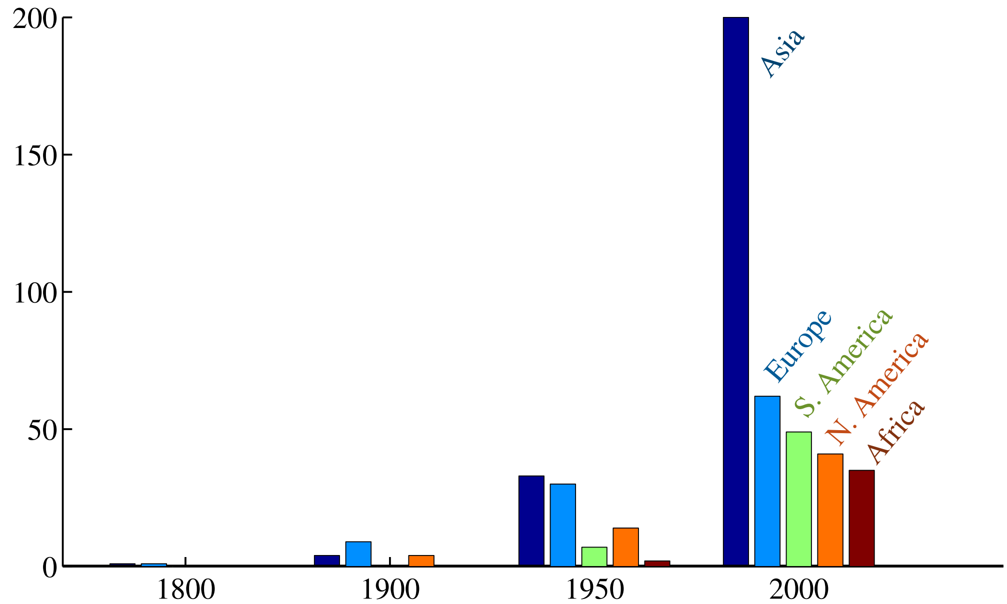

Note: The histogram shows the number of cities on each continent with populations greater than 1 million. Oceania is included with Asia. Source: Satterthwaite (2005), Table 3.

an average adoption lag across 166 countries. More precisely, the adoption lag for each country/technology observation measures the number of years between the date a technology was invented and the date it was adopted in the country.

A strong negative correlation is evident in the graph, suggesting the fact that technology adoption lags have been shrinking over time. Comin and Hobijn (2010) estimate that technologies invented ten years later are on average adopted 4.3 years faster.

Note how misallocation can show up in ideas as well.

4.12. Urbanization

We remarked earlier on the inverse correlation between urbanization in the year 1500 and income per person today. Figure 39 exemplifies the large trend in urbanization over time. The figure shows the number of cities containing more than a million people, by continent, since 1800. The large rise in urbanization over the last two centuries is well known. This figure also emphasizes the extent to which that urbanization is even stronger in Asia than in Europe and North America.

5. Conclusion

Facts we'd like to know?

Things I have neglected

Intermediate goods and IO. Network stuff

Income and democracy? Recent Acemoglu et al paper.

Countries that are "open" are richer and/or grow faster. How best to document?

Globalization??? What are the best globalization facts out there? Spillovers?

Lucas-Moll, Perla Tonetti

How the spread of growth feeds into growth of the Frontier China, India, PhDs and the future of growth. Maybe move this fact to the R&D graph section.

References

Abramovitz, Moses, "Resource and Output Trends in the United States Since 1870," *American Economic Association Papers and Proceedings*, May 1956, *46* (2), 5–23.

Acemoglu, Daron, "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *Quarterly Journal of Economics*, 1998, *113*, 1055–1089.

_ , "Politics and economics in weak and strong states," *Journal of Monetary Economics*, October 2005, *52* (7), 1199–1226.

_ and James A. Robinson, *Why nations fail: the origins of power, prosperity and poverty*, Crown Business, 2012.

_ , Simon Johnson, and James A. Robinson, "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, December 2001, *91* (5), 1369–1401.

_ , _ , and _ , "Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution," *Quarterly Journal of Economics*, 2002, *117* (4), 1231–1294.

_ , Ufuk Akcigit, Nicholas Bloom, and William R. Kerr, "Innovation, Reallocation and Growth," NBER Working Papers 18993, National Bureau of Economic Research, Inc April 2013.

Aghion, Philippe and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica*, March 1992, *60* (2), 323–351.

Akcigit, Ufuk, Harun Alp, and Michael Peters, “Lack of Selection and Limits to Delegation: Firms Dynamics in Developing Countries,” Technical Report 2014. University of Pennsylvania manuscript.

_ , Murat Alp Celik, and Jeremy Greenwood, “Buy, Keep or Sell: Economic Growth and the Market for Ideas,” 2014. University of Pennsylvania manuscript.

Alvaredo, Facundo, Anthony B. Atkinson, Thomas Piketty, and Emmanuel Saez, “The World Top Incomes Database,” 2013. Accessed on October 15, 2013.

Aoki, Shuhei and Makoto Nirei, “Pareto Distributions and the Evolution of Top Incomes in the U.S,” MPRA Paper 47967, University Library of Munich, Germany July 2013.

Asker, John, Allan Collard-Wexler, and Jan De Loecker, “Productivity Volatility and the Misallocation of Resources in Developing Economies,” NBER Working Papers 17175, National Bureau of Economic Research, Inc June 2011.

Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez, “Top Incomes in the Long Run of History,” *Journal of Economic Literature*, 2011, *49* (1), 3–71.

Autor, David H., “Skills, Education, and the Rise of Earnings Inequality among the “Other 99 Percent”,” *Science*, 2014, *344* (6186), 843–851.

_ , Frank Levy, and Richard J. Murnane, “The Skill Content Of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, November 2003, *118* (4), 1279–1333.

Banerjee, Abhijit V. and Esther Duflo, “Growth Theory through the Lens of Development Economics,” in Philippe Aghion and Steven A. Durlauf, eds., *Handbook of Economic Growth*, New York: North Holland, 2005, pp. 473–552.

Barro, Robert and Jong-Wha Lee, “A new data set of educational attainment in the world, 19502010,” *Journal of Development Economics*, 2013, *104* (C), 184–198.

Barro, Robert J., “Economic Growth in a Cross Section of Countries,” *Quarterly Journal of Economics*, May 1991, *106*, 407–443.

_ , “Convergence and Modernization Revisited,” NBER Working Papers 18295, National Bureau of Economic Research, Inc August 2012.

_ and Gary S. Becker, “Fertility Choice in a Model of Economic Growth,” *Econometrica*, March 1989, *57* (2), 481–501.

___ and Xavier Sala-i-Martin, "Convergence," *Journal of Political Economy*, 1992, *100* (2), 223–251.

Baumol, William J., "Productivity Growth, Convergence and Welfare: What the Long-Run Data Show," *American Economic Review*, December 1986, *76*, 1072–1085.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu, "The Distribution of Wealth and Fiscal Policy in Economies With Finitely Lived Agents," *Econometrica*, 01 2011, *79* (1), 123–157.

Biskupic, Joan, *Sandra Day O'Connor: How the First Woman on the Supreme Court Became Its Most Influential Justice*, New York, NY: HarperCollins, 2006.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 07 2013, *81* (4), 1347–1393.

Bolt, Jutta and J. L. van Zanden, "The Maddison Project: collaborative research on historical national accounts," *The Economic History Review*, 2014, *67* (3), 627651.

Boppart, Timo, "Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and NonGorman Preferences," *Econometrica*, November 2014, *82*, 2167–2196.

Bourguignon, Francois and Christian Morrisson, "Inequality among World Citizens: 1820-1992 Inequality among World Citizens: 1820-1992," *American Economic Review*, September 2002, *92* (4), 727–744.

Bridgman, Benjamin, "Is Labor's Loss Capital's Gain? Gross versus Net Labor Shares," June 2014. Bureau of Economic Analysis manuscript.

Brynjolfsson, Erik and Andrew McAfee, *Race Against the Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity and Irreversibly Transforming Employment and the Economy*, Digital Frontier Press, 2012.

Buera, Francisco J., Alexander Monge-Naranjo, and Giorgio E. Primiceri, "Learning the Wealth of Nations," *Econometrica*, 2011, *79* (1), 1–45.

Byrne, David M., Stephen D. Oliner, and Daniel E. Sichel, "Is the Information Technology Revolution Over?," *International Productivity Monitor*, Spring 2013, *25*, 20–36.

Caballero, Ricardo J and Adam B. Jaffe, "How High are the Giants' Shoulders?," in Olivier Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual*, Cambridge, MA: MIT Press, 1993, pp. 15–74.

Cagetti, Marco and Mariacristina De Nardi, "Entrepreneurship, Frictions, and Wealth," *Journal of Political Economy*, October 2006, *114* (5), 835–870.

Caselli, Francesco, "Technological Revolutions," *American Economic Review*, March 1999, *89* (1), 78–102.

__ , "Accounting for Cross Country Income Differences," in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Elsevier, 2005.

__ and Daniel J. Wilson, "Importing technology," *Journal of Monetary Economics*, January 2004, *51* (1), 1–32.

__ and James Feyrer, "The Marginal Product of Capital," *Quarterly Journal of Economics*, 05 2007, *122* (2), 535–568.

__ and Wilbur John Coleman, "The World Technology Frontier," *American Economic Review*, June 2006, *96* (3), 499–522.

__ , Gerardo Esquivel, and Fernando Lefort, "Reopening the Convergence Debate: A New Look at Cross-Country Growth Empirics," *Journal of Economic Growth*, September 1996, *1*, 363–390.

Cass, David, "Optimal Growth in an Aggregative Model of Capital Accumulation," *Review of Economic Studies*, July 1965, *32*, 233–240.

Castaneda, Ana, Javier Diaz-Gimenez, and Jose-Victor Rios-Rull, "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, August 2003, *111* (4), 818–857.

Chari, V.V., Pat Kehoe, and Ellen McGrattan, "Business Cycle Accounting," *Econometrica*, May 2007, *75* (3), 781–836.

Christensen, Laurits R., Dianne Cummings, and Dale W. Jorgenson, "Relative Productivity Levels, 1947-1973," *European Economic Review*, May 1981, *16* (1), 61–94.

Clark, Gregory, "The Secret History of the Industrial Revolution," 2001. U.C. Davis mimeo.

__ , "The Industrial Revolution," in "Handbook of Economic Growth," Vol. 2 of *Handbook of Economic Growth*, Elsevier, 2014, chapter 5, pp. 217–262.

Coe, David T. and Elhanan Helpman, "International R&D Spillovers," *European Economic Review*, 1995, *39* (5), 859–887.

Comin, Diego and Bart Hobijn, "An Exploration of Technology Diffusion," *American Economic Review*, 2010, *100* (5), 2031–59.

__ , Danial Lashkari, and Marti Mestieri, "Structural Transformations with Long-run Income and Price Effects," March 2015. Dartmouth College unpublished manuscript.

Cordoba, Juan Carlos and Marla Ripoll, "The Elasticity of Intergenerational Substitution, Parental Altruism, and Fertility Choice," Staff General Research Papers 37766, Iowa State University, Department of Economics August 2014.

Dalgaard, Carl-Johan and Holger Strulik, "Optimal Aging And Death: Understanding The Preston Curve," *Journal of the European Economic Association*, 06 2014, *12* (3), 672–701.

Dell, Melissa, "The Persistent Effects of Peru's Mining Mita," *Econometrica*, November 2010, *78* (6), 1863–1903.

DeLong, J. Bradford, "Productivity Growth, Convergence, and Welfare: Comment," *American Economic Review*, 1988, *78*, 1138–1154.

Denison, Edward F., *Why Growth Rates Differ*, Washington, DC: The Brookings Institution, 1967.

Diamond, Jared, *Guns, Germs, and Steel*, New York: W.W. Norton and Co., 1997.

Doepke, Matthias, "Child mortality and fertility decline: Does the Barro-Becker model fit the facts?," *Journal of Population Economics*, 06 2005, *18* (2), 337–366.

Elsby, Michael WL, Bart Hobijn, and Ayşegül Şahin, "The Decline of the U.S. Labor Share," *Brookings Papers on Economic Activity*, 2013, *2013* (2), 1–63.

Engerman, Stanley L. and Kenneth L. Sokoloff, "Factor Endowments, Institutions, and Differential Paths of Growth Among New World Economies," in Stephen Haber, ed., *How Latin America Fell Behind*, Stanford, CA: Stanford University Press, 1997.

Fernald, John, "Productivity and Potential Output Before, During, and After the Great Recession," in "NBER Macroeconomics Annual 2014, Volume 29," University of Chicago Press, 2014.

Fernald, John G. and Charles I. Jones, "The Future of US Economic Growth," *American Economic Review Papers and Proceedings*, May 2014, *104* (5), 44–49.

Foster, Lucia, John Haltiwanger, and Chad Syverson, "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, March 2008, *98* (1), 394–425.

Galor, Oded, "From Stagnation to Growth: Unified Growth Theory," *Handbook of Economic Growth*, 2005, *1*, 171–293.

_ and David N. Weil, "The Gender Gap, Fertility, and Growth," *American Economic Review*, June 1996, *86* (3), 374–387.

__ and David Weil, "Population, Technology, and Growth: From the Malthusian Regime to the Demographic Transition," *American Economic Review*, September 2000, *90*, 806–828.

Garicano, Luis, Claire Lelarge, and John Van Reenen, "Firm size distortions and the productivity distribution: Evidence from France," Technical Report, LSE manuscript 2014.

Goldin, Claudia and Lawrence F. Katz, "Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing," *Brookings Papers on Economic Activity*, 2007, *2*, 135–165.

Greenwood, Jeremy, Ananth Seshadri, and Guillaume Vandenbroucke, "The Baby Boom and Baby Bust," *American Economic Review*, March 2005, *95* (1), 183–207.

__ and Guillaume Vandenbroucke, "The Baby Boom and Baby Bust: O.E.C.D. Fertility Data," 2004. http://guillaumevdb.net/BabyBoom-data.pdf.

__ , Zvi Hercowitz, and Per Krusell, "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, June 1997, *87* (3), 342–362.

Griliches, Zvi, "Productivity Puzzles and R&D: Another Nonexplanation," *Journal of Economic Perspectives*, 1988, *2*, 9–21.

__ , "The Search for R&D Spillovers," *Scandinavian Journal of Economics*, 1992, *94*, 29–47.

__ , "Productivity, R&D and the Data Constraint," *American Economic Review*, March 1994, *84* (1), 1–23.

Grossman, Gene M. and Elhanan Helpman, *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press, 1991.

Guner, Nezih, Gustavo Ventura, and Xu Yi, "Macroeconomic Implications of Size-Dependent Policies," *Review of Economic Dynamics*, October 2008, *11* (4), 721–744.

Haines, Michael, "Fertility and Mortality in the United States," in Robert Whaples, ed., *EH.Net Encyclopedia*, 2008.

Hall, Robert E. and Charles I. Jones, "The Productivity of Nations," November 1996. NBER Working Paper No. 5812.

__ and __ , "Why Do Some Countries Produce So Much More Output per Worker than Others?," *Quarterly Journal of Economics*, February 1999, *114* (1), 83–116.

__ and __ , "The Value of Life and the Rise in Health Spending," *Quarterly Journal of Economics*, February 2007, *122* (1), 39–72.

Hansen, Gary D. and Edward C. Prescott, "Malthus to Solow," *American Economic Review*, September 2002, *92* (4), 1205–1217.

Hanushek, Eric A. and Dennis D. Kimko, "Schooling, labor-force quality, and the growth of nations," *American Economic Review*, December 2000, *90* (5), 1184–1208.

Heisenberg, Werner, *Physics and Beyond, Encounters and Conversations*, New York: Harper & Row, 1971.

Hemous, David and Morten Olsen, "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality," 2014 Meeting Papers 162, Society for Economic Dynamics 2014.

Hendricks, Lutz and Todd Schoellman, "Human Capital and Development Accounting: New Evidence from Immigrant Earnings," 2014 Meeting Papers 702, Society for Economic Dynamics 2014.

Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi, "Growth and Structural Transformation," in "Handbook of Economic Growth," Vol. 2 of *Handbook of Economic Growth*, Elsevier, 2014, chapter 6, pp. 855–941.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

_ and _ , "The Life Cycle of Plants in India and Mexico," *Quarterly Journal of Economics*, 2014, *129* (3), 1035–1084.

_ , Erik Hurst, Charles I. Jones, and Peter J. Klenow, "The Allocation of Talent and U.S. Economic Growth," 2013. Stanford University, unpublished paper.

Inklaar, Robert and Marcel P. Timmer, "Capital, Labor and TFP in PWT 8.0," July 2013. Groningen Growth and Development Center unpublished paper.

Islam, Nazrul, "Growth Empirics: A Panel Data Approach," *Quarterly Journal of Economics*, November 1995, *110*, 1127–1170.

Jones, Benjamin F., "The Human Capital Stock: A Generalized Approach," *American Economic Review*, November 2014, *104* (11), 3752–77.

Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, August 1995, *103* (4), 759–784.

_ , "On the Evolution of the World Income Distribution," *Journal of Economic Perspectives*, Summer 1997, *11*, 19–36.

— , "Was an Industrial Revolution Inevitable? Economic Growth Over the Very Long Run," *Advances in Macroeconomics*, 2001, *1* (2), Article 1. http://www.bepress.com/bejm-/advances/vol1/iss2/art1.

— , "Sources of U.S. Economic Growth in a World of Ideas," *American Economic Review*, March 2002, *92* (1), 220–239.

— , "Growth and Ideas," in Philippe Aghion and Steven A. Durlauf, eds., *Handbook of Economic Growth*, New York: North Holland, 2005, pp. 1063–1111.

— and Dean Scrimgeour, "A New Proof of Uzawa's Steady-State Growth Theorem," *The Review of Economics and Statistics*, February 2008, *90* (1), 180–182.

— and Jihee Kim, "A Schumpeterian Model of Top Income Inequality," October 2014. Stanford University manuscript.

— and John C. Williams, "Measuring the Social Return to R&D," *Quarterly Journal of Economics*, November 1998, *113* (4), 1119–1135.

— and Paul M. Romer, "The New Kaldor Facts: Ideas, Institutions, Population, and Human Capital," *American Economic Journal: Macroeconomics*, January 2010, *2* (1), 224–245.

— and Peter J. Klenow, "Beyond GDP: Welfare across Countries and Time," April 2015. Stanford University unpublished manuscript.

Jones, Larry and Michele Tertilt, "An Economic History of Fertility in the U.S.: 1826-1960," in Peter Rupert, ed., *Handbook of Family Economics*, Elsevier, forthcoming.

Jones, Larry E., Alice Schoonbroodt, and Michele Tertilt, "Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?," in "Demography and the Economy" NBER Chapters, National Bureau of Economic Research, Inc, May 2010, pp. 43–100.

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.

Karabarbounis, Loukas and Brent Neiman, "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

Katz, Lawrence and Kevin Murphy, "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics*, February 1992, *107* (1), 35–78.

King, Robert G. and Ross Levine, "Capital Fundamentalism, Economic Development, and Economic Growth," *Carnegie-Rochester Conference Series on Public Policy*, 1994, *40*, 259–292.

Klenow, Peter J. 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.

— and — , "Extenalities and Growth," in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Amsterdam: Elsevier, 2005.

Koopmans, T. C., "On the Concept of Optimal Economic Growth," in "The Econometric Approach to Development Planning," Amsterdam: North Holland, 1965.

Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, *65* (6), 1389–1419.

Kremer, Michael, "Population Growth and Technological Change: One Million B.C. to 1990," *Quarterly Journal of Economics*, August 1993, *108* (4), 681–716.

Krusell, Per, Lee Ohanian, Jose-Victor Rios-Rull, and Giovanni Violante, "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, September 2000, *68* (5), 1029–1053.

Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman, "Experience Matters: Human Capital and Development Accounting," NBER Working Papers 18602, National Bureau of Economic Research, Inc December 2012.

Lee, Ronald D., "Induced Population Growth and Induced Technological Progress: Their Interaction in the Accelerating Stage," *Mathematical Population Studies*, 1988, *1* (3), 265–288.

Lucas, Robert E., "On the Mechanics of Economic Development," *Journal of Monetary Economics*, 1988, *22* (1), 3–42.

Lucas, Robert E., "Some Macroeconomics for the 21st Century," *Journal of Economic Perspectives*, Winter 2000, *14* (1), 159–168.

Lucas, Robert E., "The Industrial Revolution: Past and Future," *Lectures on Economic Growth*, 2002, pp. 109–188.

— , "Ideas and Growth," *Economica*, February 2009, *76* (301), 1–19.

— and Benjamin Moll, "Knowledge Growth and the Allocation of Time," *Journal of Political Economy*, February 2014, *122* (1), 1–51.

Maddison, Angus, "Per Capita Output in the Long Run," *Kyklos*, 1979, *32*, 412–419.

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

— , "Statistics on World Population, GDP and Per Capita GDP, 1–2006 AD," 2008. Downloaded on December 4, 2008 from http://www.ggdc.net/maddison/.

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.

Manuelli, Rodolfo E. and Ananth Seshadri, "Human Capital and the Wealth of Nations," *American Economic Review*, 2014, *104* (9), 2736–62.

Midrigan, Virgiliu and Daniel Yi Xu, "Finance and Misallocation: Evidence from Plant-Level Data," *American Economic Review*, February 2014, *104* (2), 422–58.

Mokyr, Joel, *The Lever of Riches*, New York: Oxford University Press, 1990.

Moll, Benjamin, "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?," *American Economic Review*, 2014, *104* (10), 3186–3221.

Murphy, Kevin M. and Robert H. Topel, "The Value of Health and Longevity," *Journal of Political Economy*, October 2006, *114* (5), 871–904.

Nordhaus, William D., "The Health of Nations: The Contribution of Improved Health to Living Standards," in Kevin M. Murphy and Robert Topel, eds., *Measuring the Gains from Medical Research: An Economic Approach*, Chicago: University of Chicago Press, 2003, pp. 9–40.

OECD, *Frascati Manual 2002*, OECD Publishing, 2002.

Oeppen, Jim and James W. Vaupel, "Broken Limits to Life Expectancy," *Science*, May 2002, *296* (5570), 1029–1031.

Olson, Mancur, "Big Bills Left on the Sidewalk: Why Some Nations are Rich, and Others Poor," *Journal of Economic Perspectives*, Spring 1996, *10* (2), 3–24.

Peters, Michael, "Heterogeneous Mark-ups, Growth, and Endogenous Misallocation," 2013. Yale University manuscript.

Piketty, Thomas, *Capital in the Twenty-first Century*, Harvard University Press, 2014.

— and Emmanuel Saez, "Income Inequality In The United States, 1913–1998," *Quarterly Journal of Economics*, February 2003, *118* (1), 1–39.

— , — , and Stefanie Stantcheva, "Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities," *American Economic Journal: Economic Policy*, February 2014, *6* (1), 230–71.

Pomeranz, Kenneth, *The Great Divergence: China, Europe, and the Making of the Modern World Economy*, Princeton University Press, 2009.

Prescott, Edward C., "Needed: A Theory of Total Factor Productivity," *International Economic Review*, August 1998, *39* (3), 525–51.

— , "Why Do Americans Work So Much More than Europeans?," *Quarterly Review*, July 2004, pp. 2–13.

Pritchett, Lant, "Divergence: Big Time," *Journal of Economic Perspectives*, Summer 1997, *11* (3), 3–17.

Quah, Danny, "Empirical Cross-Section Dynamics in Economic Growth," *European Economic Review*, 1993, *37*, 426–434.

Ramey, Valerie A. and Neville Francis, "A Century of Work and Leisure," *American Economic Journal: Macroeconomics*, July 2009, *1* (2), 189–224.

Ramsey, Frank, "A Mathematical Theory of Saving," *Economic Journal*, 1928, *38*, 543–559.

Restuccia, Diego and Guillaume Vandenbroucke, "The Evolution of Education: A Macroeconomic Analysis," *International Economic Review*, 2013, *54* (3), 915–936.

— and Richard Rogerson, "Policy Distortions and Aggregate Productivity with Heterogeneous Plants," *Review of Economic Dynamics*, October 2008, *11*, 707–720.

— , Dennis Tao Yang, and Xiaodong Zhu, "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis," *Journal of Monetary Economics*, 2008, *55* (2), 234–250.

Romer, Paul M., "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, October 1986, *94*, 1002–1037.

— , "Endogenous Technological Change," *Journal of Political Economy*, October 1990, *98* (5), S71–S102.

Rosenberg, Nathan, *Exploring the Black Box: Technology, Economics, and History*, New York: Cambridge University Press, 1994.

Saez, Emmanuel and Gabriel Zucman, "Wealth Inequality in the United States since 1913," July 2014. U.C. Berkeley slides.

Sala-i-Martin, Xavier, "The World Distribution of Income: Falling Poverty and... Convergence, Period," *Quarterly Journal of Economics*, May 2006, *121* (2), 351–397.

Satterthwaite, David, "The Scale of Urban Change Worldwide 1950-2000 and its Underpinnings," May 2005. International Institute for Environment and Development unpublished paper.

Schlicht, Ekkehart, "A Variant of Uzawa's Theorem," *Economics Bulletin*, February 2006, *5* (6), 1–5.

Segerstrom, Paul, "Endogenous Growth Without Scale Effects," *American Economic Review*, December 1998, *88* (5), 1290–1310.

Simon, Julian L., *The Ultimate Resource*, Princeton, NJ: Princeton University Press, 1981.

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.

Stokey, Nancy L. and Sergio Rebelo, "Growth Effects of Flat-Rate Taxes," *Journal of Political Economy*, June 1995, *103*, 519–550.

Syverson, Chad, "What Determines Productivity?," *Journal of Economic Literature*, June 2011, *49* (2), 326–65.

Weil, David N., "Accounting for The Effect of Health on Economic Growth," *Quarterly Journal of Economics*, 2007, *122* (3), 1265–1306.

Weitzman, Martin L., "Recombinant Growth," *Quarterly Journal of Economics*, May 1998, *113*, 331–360.

Whelan, Karl, "A Two-Sector Approach to Modeling U.S. NIPA Data," *Journal of Money, Credit, and Banking*, August 2003, *35* (4), 627–656.

Yang, Jonghoe, "Colonial Legacy and Modern Economic Growth in Korea: A Critical Examination of their Relationships," *Development and Society*, June 2004, *33* (1), 1–24.

Zeira, Joseph, "Workers, Machines, And Economic Growth," *Quarterly Journal of Economics*, November 1998, *113* (4), 1091–1117.