

NBER WORKING PAPER SERIES

THE NEW KALDOR FACTS:
IDEAS, INSTITUTIONS, POPULATION, AND HUMAN CAPITAL

Charles I. Jones
Paul M. Romer

Working Paper 15094
<http://www.nber.org/papers/w15094>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2009

Prepared for a session at the January 2009 annual meeting of the American Economic Association on "The secrets of growth: What have we learned from research in the last 25 years?" We are grateful to Daron Acemoglu, Philippe Aghion, Francesco Caselli, Arthur Chiang, Steve Davis, Rob Feenstra, Penny Goldberg, Claudia Goldin, Avner Greif, Peter Howitt, Pete Klenow, Ron Lee, Bill Nordhaus, Douglass North, Stephen Parente, Lant Pritchett, Andres Rodriguez-Clare, David Romer, and David Weil for helpful suggestions. Jones thanks the National Science Foundation (SES-0720994) for financial support. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2009 by Charles I. Jones and Paul M. Romer. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The New Kaldor Facts: Ideas, Institutions, Population, and Human Capital
Charles I. Jones and Paul M. Romer
NBER Working Paper No. 15094
June 2009
JEL No. O1,O3,O4

ABSTRACT

In 1961, Nicholas Kaldor used his list of six "stylized" facts both to summarize the patterns that economists had discovered in national income accounts and to shape the growth models that they were developing to explain them. Redoing this exercise today, nearly fifty years later, shows how much progress we have made. In contrast to Kaldor's facts, which revolved around a single state variable, physical capital, our six updated facts force consideration of four far more interesting variables: ideas, institutions, population, and human capital. Dynamic models have uncovered subtle interactions between these variables and generated important insights about such big questions as: Why has growth accelerated? Why are there gains from trade?

Charles I. Jones
Graduate School of Business
Stanford University
518 Memorial Way
Stanford, CA 94305-4800
and NBER
chad.jones@stanford.edu

Paul M. Romer
Stanford Institute for Economic Policy Research
Stanford University
Stanford, CA 94305
and NBER
paul.romer@stanford.edu

1. Introduction

If you live from paper to paper, seminar to seminar, and conference to conference, it is easy to lose faith in scientific progress. The risk is perhaps most pronounced at job talks in areas outside our field of interest. “Isn’t this what they were arguing about last year?” Or “last decade?”

In any assessment of progress, as in any analysis of macroeconomic variables, a long-run perspective helps us look past the short-run fluctuations and see the underlying trend. In 1961, Nicolas Kaldor stated six now famous “stylized” facts. He used them to summarize what economists had learned from their analysis of 20th-century growth and also to frame the research agenda going forward (Kaldor, 1961):

1. Labor productivity has grown at a sustained rate.
2. Capital per worker has also grown at a sustained rate.
3. The real interest rate or return on capital has been stable.
4. The ratio of capital to output has also been stable.
5. Capital and labor have captured stable shares of national income.
6. Among the fast growing countries of the world, there is an appreciable variation in the rate of growth “of the order of 2–5 percent.”

Redoing this exercise nearly 50 years later shows just how much progress we have made. Kaldor’s first five facts have moved from research papers to textbooks. There is no longer any interesting debate about the features that a model must contain to explain them. These features are embodied in one of the great successes of growth theory in the 1950s and 1960s, the neoclassical growth model. Today, researchers are now grappling with Kaldor’s sixth fact and have moved on to several others that we list below.

One might have imagined that the first round of growth theory clarified the deep foundational issues and that subsequent rounds filled in the details. This is not what we observe. The striking feature of the new stylized facts driving the research agenda today is how much more ambitious they are. Economists now expect that economic theory should inform our thinking about issues that we once ruled out of bounds as important but too difficult to capture in a formal model.

Here is a summary of our new list of stylized facts, to be discussed in more detail below:

1. **Increases in the extent of the market.** Increased flows of goods, ideas, finance, and people — via globalization as well as urbanization — have increased the extent of the market for all workers and consumers.
2. **Accelerating growth.** For thousands of years, growth in both population and per capita GDP has accelerated, rising from virtually zero to the relatively rapid rates observed in the last century.
3. **Variation in modern growth rates.** The variation in the rate of growth of per capita GDP increases with the distance from the technology frontier.
4. **Large income and TFP differences.** Differences in measured inputs explain less than half of the enormous cross country differences in per capita GDP.
5. **Increases in human capital per worker.** Human capital per worker is rising dramatically throughout the world.
6. **Long-run stability of relative wages.** The rising quantity of human capital relative to unskilled labor has not been matched by a sustained decline in its relative price.

In assessing the change since Kaldor developed his list, it is important to recognize that Kaldor himself was raising expectations relative to the initial neoclassical model of growth as outlined by Solow (1956) and Swan (1956). When the neoclassical model was being developed, a narrow focus on physical capital alone was no doubt a wise choice. The smooth substitution of capital and labor in production expressed by an aggregate production function, the notion that a single capital aggregate might be useful, and the central role of accumulation itself were all relatively novel concepts that needed to be explained and assimilated. Moreover, even these small first steps toward formal models of growth provoked substantial opposition.

The very narrow focus of the neoclassical growth model sets the baseline against which progress in growth theory can be judged. Writing in 1961, Kaldor was already intent on making technological progress an endogenous part of a more complete model of growth. The tip-off about his intention is the inclusion of his final fact, which cited the variation in growth rates across countries and which morphs here into our Fact 3

about the increase in the the variation in growth with distance form the technology frontier.

Growth theorists working today have not only completed this extension but also brought into their models the other endogenous state variables excluded from consideration by the initial neoclassical setup. Ideas, institutions, population, and human capital are now at the center of growth theory. Physical capital has been pushed to the periphery.

Kaldor had a model in mind when he introduced his facts. So do we. Providing a conceptual outline of this model is a subsidiary goal of this paper. In the near term, we believe that this model should capture the endogenous accumulation of and interaction between three of our four state variables: ideas, population, and human capital. For now, we think that progress is likely to be most rapid if we follow the example of the neoclassical model and treat institutions the way the neoclassical model treated technology, as an important force that enters the formalism but which evolves according to a dynamic that is not explicitly modeled. Out on the horizon, we can expect that current research on the dynamics of institutions and politics will ultimately lead to a simple formal representation of endogenous institutional dynamics as well.

2. The Formal Foundations of Modern Growth Theory

One of the great accomplishments of neoclassical growth theory is that it produced a single model that captured the first five of Kaldor's facts. Five different models, one for each fact, would have been a much less significant intellectual achievement. General equilibrium modeling in the sense of Solow and Swan allows the use of simple functional forms, but still insists on a unified framework where different specialized lines of inquiry can be integrated to reveal the connections between seemingly unrelated observations.

One could create a different partial equilibrium model for each of our six facts. For example, increases in the extent of the market might be explained by exogenous changes in transport costs and exogenous changes in laws that limited trade in goods and incoming foreign direct investment. One could also model increases in rates of growth as the result of exogenous changes in institutions. This kind of approach misses

the deep connections that the applied general equilibrium approach uncovers.

The other great accomplishment of the neoclassical model, one which has reverberated throughout macroeconomics ever since, was its explicit microeconomic foundations. Reduced form expressions that are not derived from fundamentals restrict the evidence one can bring to bear on the model, limit its utility in policy analysis, and narrow its focus. External increasing returns are a recurring example of this kind of reduced form approach. Growth models that rely on this assumption typically assume that the underlying spillovers stop at a national border. The resulting model may seem useful for studying trends in a closed economy but is mute or unreliable when used to address questions about international trade or direct foreign investment.

The key to progress in growth theory has always been a tractable description of production possibilities based on a production function and a small list of inputs. Modern growth theory has added a stock of ideas and a stock of human capital to the familiar inputs of physical capital and workers. It has also returned to classical concerns about the endogenous dynamics of population. We turn now to an overview of these variables as well as the institutions that shape their evolution.

Of the three state variables that we endogenize, ideas have been the hardest to bring into the applied general equilibrium structure. The difficulty arises because of the defining characteristic of an idea, that it is a pure nonrival good. A given idea is not scarce in the same way that land or capital or other objects are scarce; instead, an idea can be used by any number of people simultaneously without congestion or depletion.

Because they are nonrival goods, ideas force two distinct changes in our thinking about growth, changes that are sometimes conflated but are logically distinct. Ideas introduce scale effects. They also change the feasible and optimal economic institutions. The institutional implications have attracted more attention but the scale effects are more important for understanding the big sweep of human history.

The distinction between rival and nonrival goods is easy to blur at the aggregate level but inescapable in any microeconomic setting. Picture, for example, a house that is under construction. The land on which it sits, capital in the form of a measuring tape, and the human capital of the carpenter are all rival goods. They can be used to build this house but not simultaneously any other. Contrast this with the Pythagorean Theorem, which the carpenter uses implicitly by constructing a triangle with sides in

the proportions of 3, 4 and 5. This idea is nonrival. Every carpenter in the world can use it at the same time to create a right angle.

Of course, human capital and ideas are tightly linked in production and use. Just as capital produces output and forgone output can be used to produce capital, human capital produces ideas and ideas are used in the educational process to produce human capital. Yet ideas and human capital are fundamentally distinct. At the micro level, human capital in our triangle example literally consists of new connections between neurons in a carpenter's head, a rival good. The 3-4-5 triangle is the nonrival idea. At the macro level, one cannot state the assertion that skill-biased technical change is increasing the demand for education without distinguishing between ideas and human capital.

The scale effect of ideas follows immediately from nonrivalry. The value of the 3-4-5 triangle increases in proportion to the number of houses under construction. If there are many different people in different regions who are building houses, there are efficiency gains to be had from connecting them together so that the idea can be utilized everywhere as soon as it has been discovered somewhere. The 3-4-5 triangle might be shared in written or spoken form, or used to make T-squares that are shipped to all other locations. No matter how it is communicated and reused, nonrivalry by itself creates strong incentives for economic integration among the largest possible group of people. As we will argue in the next section, this is the best candidate explanation for Fact 1, the relentless pressure for expansion in the extent of the market.

Another consequential idea, one that was discovered much more recently, is oral rehydration therapy: if you add some salt and minerals and, critically, a little sugar to water, you can rehydrate a child who would otherwise die from diarrhea. This simple idea now saves millions of lives each year. Like the idea of a 3-4-5 triangle, it is much easier to teach this idea once it has been discovered than to have each person go out and try to rediscover it on his or her own.

In this context, the key point is that nonrivalry did not come to the forefront of growth theory because of attention to Fact 1 and interest in the extent of the market. Rather, this concept emerged from the study of Fact 2, the acceleration of growth over time. It was the commitment to the methods of applied general equilibrium theory that uncovered the surprising link to the extent of the market and the gains from trade.

The discussion of the carpenter already suggests important connections between nonrival ideas, investment in human capital, and scale or population. Of the four state variables needed to make sense of the new list of facts, this leaves only institutions. As we have indicated above, we are far away from a simple formal model that can describe how institutions evolve. However, a model that allows for ideas suggests that institutions are more complicated and matter much more than the neoclassical model once suggested.

As just one example, recall that the increasing returns to scale that is implied by nonrivalry leads to the failure of Adam Smith's famous invisible hand result. The institutions of complete property rights and perfect competition that work so well in a world consisting solely of rival goods no longer deliver the optimal allocation of resources in a world containing ideas. Efficiency in use dictates price equal to marginal cost. But with increasing returns, there is insufficient output to pay each input its marginal product; in general, price must exceed marginal cost somewhere to provide the incentive for profit maximizing private firms to create new ideas.¹ This tension is at the heart of the problem: a single price cannot simultaneously allocate goods to their most efficient uses *and* provide the appropriate incentives for innovation.

An important unresolved policy question is therefore the optimal design of institutions that support the production and distribution of nonrival ideas. In practice, most observers seem to agree that some complicated mix of secrecy, intellectual property rights that convey partial excludability, public subsidies through the institutions of science, and private voluntary provision is more efficient than any corner solution like that prescribed for rival goods. We are, however, very far from results we could derive from first principles to guide decisions about which types of goods are best served by which institutional arrangement.

In what follows, a key point is that the institutions that underlie the production and distribution of new ideas have evolved over time, dramatically so in the last couple of centuries. Growth driven by the discovery of new ideas, both in the recent past and in the foreseeable future, must be understood to occur in the context of institutions such as universities and laws supporting intellectual property rights that have been evolving,

¹Boldrin and Levine (2008) make the point that without any markup, innovation can sometimes arise through a process of voluntary provision of a pure public good, but this observation has the potential to capture only a small part of the innovation that takes place in the private sector of a modern economy.

that are almost surely not yet optimal, and that will continue to change in the future.

3. The New Stylized Facts

What then are the new facts that growth models should explain? We return to our list.

Fact 1: Increases in the extent of the market.

Increased flows of goods, ideas, finance, and people — via globalization as well as urbanization — have increased the extent of the market for all workers and consumers.²

Figure 1 illustrates two manifestations of this fact. For recent decades, it shows the rising share of world trade (imports plus exports) and foreign direct investment in world GDP. World trade as a share of GDP has nearly doubled since 1960. But this surely masks a much larger increase in economic integration. One piece of evidence supporting this argument is that FDI as a share of world GDP has increased by a factor of 30, from less than a tenth of a percent of GDP in 1965 to 2.8 percent of GDP in 2006.

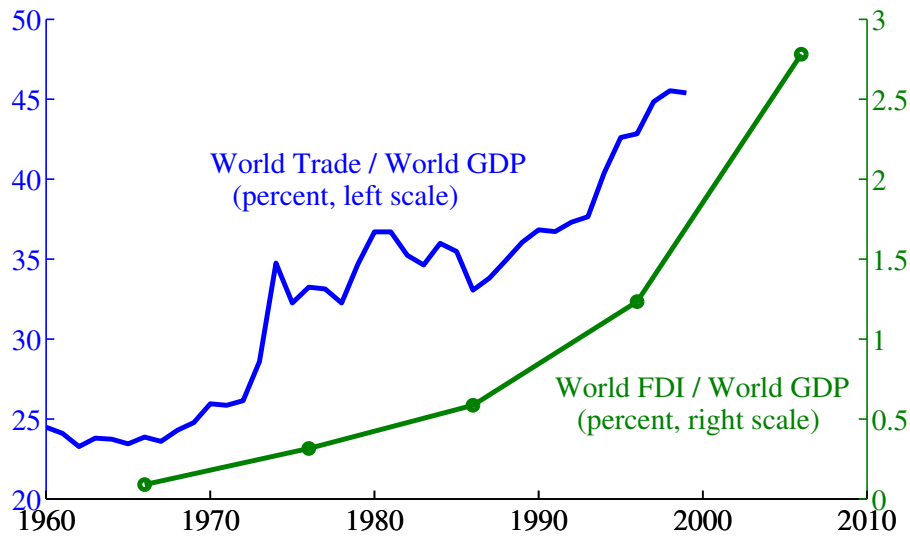
While trade and FDI are key facets of the rising extent of the market, the fact itself is even broader and includes the flow of ideas and people, within as well as across borders. International flows of ideas are indicated by cross-country patent statistics. For example, in 1960, 83% of patents granted by the U.S. Patent and Trademark Office were to domestic entities. In recent years, that fraction has fallen to about 50%.³ Within countries, urbanization rates have risen sharply. The fraction of the world's population living in cities increased from 29.1% in 1950 to 49.4% in 2007 and is projected to rise even further to 69.6% by 2050 (United Nations, 2008). Finally, with the rise of the World Wide Web, information flows both across and within countries have exploded.

All these are facts that we largely take for granted but which cry out for explanation. Why is it that a country with 300 million residents that is as geographically diverse as the United States can still enjoy gains from trade with the rest of the world? Or why should the Western Hemisphere as a whole benefit from trade with the other half of the

²Examples of research that contributed to documenting this fact include Krugman (1995) and Feenstra (1998).

³See the National Science Foundation's *Science and Engineering Indicators*.

Figure 1: The Rise in Globalization



Note: World trade is the sum of world exports and imports as a share of world GDP from the Penn World Tables 6.1. FDI as a share of GDP is from the World Bank's *World Development Indicators*.

globe? When nonrival goods are present, the gains from trade are not exhausted until everyone is connected to everyone else.

Robert Lucas (1988) emphasized a similar point in the context of cities. Why is it that so many people pay high rent to be near so many other people who pay high rent? The same point applies to trade flows and FDI. We take it for granted that trade is good, but rarely stop to ask why. The developers of new trade theory explained why the pattern of trade took the form that it did, with trade flows between countries that were similar at the aggregate level (Krugman, 1979; Helpman and Krugman, 1985). Like Lucas, Krugman (1991) also drew the parallel between spatial patterns of production and models of trade. All these models turn on the observation that each individual is better off if she can interact with others just like her. The explanation for this positive association between individual welfare and the size of the market is often buried in a functional form with increasing returns to scale or some benefit from variety, which is limited by fixed costs. The claim here is that the deeper explanation turns on nonrivalry.

To make this point, start with a simple model with a stock of rival physical goods X , which could be natural resources. Add to these L individuals who can supply labor

and a set of ideas A . If sugar, salt, and water are typical rival goods, an idea is a formula like that for oral rehydration therapy that lists just the right proportions to turn them into a life-saving medicine. If total output takes the form $Y = F(A, X, L)$, the standard replication argument for rival goods implies that there are constant returns to X and L together, so output per capita is $Y/L = F(A, X/L, 1)$.

The most obvious point, which goes back at least to Malthus, is clear from this expression. With any fixed set of ideas A , average welfare has to be decreasing in the size of the population L . As L increases, the per capita endowment of the rival goods falls. Growth can still take place, at least for a while, if the stock of ideas provides instructions for converting some components of X (iron ore perhaps) into other components (machine tools.) Nevertheless, as long as these transformations are governed by convex production possibilities, an increase in L has an unambiguously negative effect on average welfare.

Now suppose that the stock of ideas is proportional to L . (For example, this would be true in the long run if each person has a fixed probability of discovering a new idea and there is constant population growth.) Then output per capita becomes $Y/L = F(L, X/L, 1)$. An increase in L is harmful because of its effect on per capita supplies of the rival goods X but beneficial because of the increase in the stock of nonrival goods A . In principle, it is possible that as L increases, either effect could dominate. In practice, urbanization, increased trade, globalization in all of its forms, and the positive trend in per capita income all point in the same direction. In the long run, the benefits of a larger population which come from an increase in the stock of available ideas decisively dominates the negative effects of resource scarcity. In such a world, any form of interaction that lets someone interact with many others like her and share in the ideas they discover is beneficial, and the benefit need not be exhausted at any finite population size.

We can add some numbers to this argument by considering a practical question of current relevance. China's population is roughly equal to that of the United States, Europe, and Japan combined. Over the next several decades, the continued economic development of China might plausibly double the number of researchers throughout the world pushing forward the technological frontier. What effect will this have on incomes in countries that share ideas with China in the long run?

To address this question, we can use functional forms to make the argument from above concrete. Write output per person y , as an increasing function of the number of people with whom an individual can share ideas, L :

$$y = mL^\gamma \tag{1}$$

For example, if the function F is Cobb-Douglas, γ is equal to the difference between the exponent on the nonrival input A and the sum of the exponents on rival goods in the vector X . This exponent γ measures the degree of increasing returns, and m is a factor of proportionality that captures the effects of institutions, human capital, and other accumulation processes.

Many fully specified idea-based growth models deliver an expression like that in equation (1).⁴ This is most precisely true in the semi-endogenous growth models of Jones (1995), Kortum (1997), and Segerstrom (1998), for example, in which the scale effect from nonrivalry shows up in the level of output per capita but not in the long run growth rate. The first-generation idea-based models of Romer (1990), Aghion and Howitt (1992), and Grossman and Helpman (1991) had even stronger scale effects in the sense that a larger population implies a permanently faster rate of growth, which of course means higher levels. For those models, the calibration exercised offered here can be thought of as a rough lower bound on potential benefits from an increase in scale.

Based on equation (1), one could try to calibrate the value of γ by comparing per capita income growth to growth in the size of the relevant idea-generating population. An obvious problem with this approach, however, is that m is not constant: changes in institutions, education, and research intensity, for example, are important over any reasonable time period from which we can take data. Jones (2002) discusses how to handle this problem and recover the value for γ . There, the level of total factor productivity is proportional to the number of researchers raised to the power γ , which is just another version of equation (1). In the long run, the value of γ must be such that these two series have the same time trend; in other words, γ must equal the ratio of TFP growth to the growth rate of the effective number of researchers. In the United States between

⁴For a discussion of the main exceptions, see Jones (1999).

1960 and 1993, Jones (2002) documents that TFP growth was about 1% per year and the effective number of researchers contributing to ideas used in the United States grew at about 4% per year, suggesting a value of $\gamma \approx 1/4$.

One can, of course, quibble with these numbers. For example, by ignoring the gains associated with rising life expectancy and mismeasured quality, it is possible that growth is actually twice as fast, suggesting that γ should be doubled.⁵ Alternatively, perhaps all people — not just those labeled formally as “researchers” by statisticians — contribute to idea production. In this case, one might argue that the growth rate of research inputs is only about 1% per year rather than 4%, which would raise γ by a factor of 4. Taking these arguments together, one could make a case for a value of γ as high as 1 or even 2 instead of $1/4$.


What does this imply about the gains from economic integration resulting from China’s development in coming decades? If the number of researchers effectively doubles, then per capita income in the United States and in other countries benefiting from China’s ideas could rise by a factor of anywhere from $2^{1/4} = 1.2$ to $2^1 = 2$ or even $2^2 = 4$, depending on the value of γ . Even the smallest of these numbers — the 20% gain associated with $\gamma = 1/4$ — is large by the standards of the conventional trade literature.⁶ It is also large compared to most estimates of the entire cost to the United States of reductions in carbon emissions, so even if we assume (unrealistically) that growth in China is to blame for all the mitigation we do, we are still better off because of that growth. And the gains could be substantially larger.

The key lesson from the first fact is that there are powerful incentives for connecting as many people as possible into trading networks that make all ideas available to everyone. This incentive, we claim, is the deep explanation for all the different processes that are making it easier for ever more people to connect with each other.

⁵For example, Nordhaus (2003) suggests that gains in life expectancy in the 20th century are roughly just as valuable as gains in non-health consumption.

⁶Eaton and Kortum (2002) and Ramondo and Rodriguez-Clare (2009) study these gains in detailed models of trade and multinational production and report relatively small benefits. We thank Andres Rodriguez-Clare and Lant Pritchett for suggestions that motivate this example.

Figure 2: Population and Per Capita GDP over the Very Long Run



Note: Population and GDP per capita for “the West,” defined as the sum of the United States and 12 western European countries. Both series are normalized to take the value 1.0 in the initial year, 1 A.D. Source: Maddison (2008).

Fact 2: Accelerating growth.

For thousands of years, growth in both population and per capita GDP has accelerated, rising from virtually zero to the relatively rapid rates observed in the last century.⁷

This fact is documented in Figure 2, which shows population and per capita GDP for “the West” — an amalgam of the United States and twelve western European countries for which Maddison (2008) reports data going back 2,000 years. Plotted on a linear scale, the by-now-familiar “hockey stick” pattern would be highlighted, where both population and per capita GDP appear essentially flat for nearly two thousand years and then rise very sharply in the last two centuries. We’ve chosen to plot these two series on a logarithmic scale instead to emphasize the point that the rates of growth — the slopes of the two series — have themselves been rising over time.

⁷Examples of research that contributed to documenting and analyzing this fact include Romer (1986), Lee (1988), Kremer (1993), Galor and Weil (2000) and Clark (2007).

Kremer (1993) documents this fact for world population, going back 1 million years in history. The growth rate of population is zero to several decimal places for hundreds of thousands of years. For example, from 25,000 B.C. to 1 A.D., the population grows at an average annual rate of only 0.016%. Nevertheless, the power of exponential growth is such that the level of population increased substantially, from 3.3 million to 170 million. The growth rate in recent decades is approximately 100 times faster.

Nordhaus (1997) provides a tangible example linking accelerating growth to ideas in his famous “price of light” calculation. Between 38,000 B.C. and 1750 B.C., the real price of light fell by a total of about 17%, based on the transition from animal or vegetable fat to sesame oil as a fuel. The use of candles and whale oil reduced the price by a further 87% by the early 1800s, an average annual rate of decline of 0.06% per year. Between 1800 and 1900, the price of light fell at an annual rate that was 38 times faster, 2.3%, with the introduction of the carbon filament lamp. And then in the 20th century, the price of light has fallen at the truly remarkable pace of 6.3% per year with the use of tungsten filaments and fluorescent lighting. New ideas are very clearly at the heart of this accelerating productivity growth.⁸

If Fact 1 is fundamentally about the static gains from scaling given ideas across ever larger markets, Fact 2 is its dynamic counterpart. More people lead to more ideas. For most of human history, more ideas made it possible for the world to support more people. In a dynamic version of the model used to explain Fact 1, this simple feedback loop generates growth rates that increase over time.

To see this, consider a simple model based on Lee (1988) and Kremer (1993). Interpret the rival good X as land, which is fixed and normalized to one ($X = 1$). Suppose consumption Y is produced using ideas A , land X , and labor L :

$$Y_t = A_t X^\beta L_t^{1-\beta}. \quad (2)$$

Assume that each person discovers α ideas as a side effect of other activities:

$$\dot{A}_t = \alpha L_t. \quad (3)$$

⁸These calculations are based on Table 1.4 in Nordhaus (1997).

We can close the model in two ways, depending on what assumption we make regarding the evolution of population. To begin, we use the Malthusian assumption that the number of people jumps immediately to whatever level is consistent with subsistence output, \bar{y} , which we normalize to one:

$$Y_t/L_t = \bar{y} = 1. \quad (4)$$

Combining equations (2) and (4), it is easy to see that the level of the population at any point in time will be proportional to the number of ideas raised to some power. And substituting this result into (3) yields

$$\dot{A}_t = \alpha A_t^{1/\beta}. \quad (5)$$

Because the land share β is less than one, the exponent in this equation is greater than one: the growth rate of ideas accelerates over time. And since population is a power function of the number of ideas, population growth similarly accelerates over time. In fact, if these forces are unchecked, it is easy to show that both ideas and population become infinite in a finite amount of time. Such is the power of faster-than-exponential growth!


Of course, it is biologically impossible for the population growth rate to become infinite in finite time. At some point, human fertility can't keep up. This leads to a second way to close the model. One can relax the assumption that population adjusts instantaneously to drive income down to subsistence and replace it with an economic model of fertility. At low rates of growth and low levels of income, desired fertility may be able to keep up with new ideas, so per capita income remains close to the subsistence level and population growth accelerates. Eventually, though, fertility and population growth level out. Growth in per capita income then accelerates until it reaches modern rates.

The acceleration of population growth and per capita growth are striking bits of time series evidence that support the feedback between population and ideas. There is no comparable cross sectional evidence at the level of individual nations because economic integration lets countries of widely varying sizes draw from a shared pool of ideas. Until India reformed its trade and investment laws, workers in tiny Hong Kong may have had access to more ideas than workers in much of India.

Searching for a case where economic integration does not pollute the cross-section, Kremer (1993) and Diamond (1997) document their famous example of regions isolated by the melting of the polar ice caps at the end of the last ice age. Five regions that began as similar hunter-gatherer societies around 10,000 B.C. look extremely different when the world is re-integrated by the time of Columbus, around 1500 A.D. These outcomes correlate perfectly with the land areas and initial populations, just as the increasing returns associated with nonrivalry would predict. The Eurasian-African continent is the region using large ships to explore the rest of the world, the Americas contains the Maya and Aztec civilizations with their cities and calendars, Australia has advanced somewhat with the boomerang, polished stone tools, and fire-making technologies, Tasmania remains a primitive hunter-gatherer society, and the people of the tiny Flinders Island have died out completely.

Looking forward, virtually all demographic projections call for the number of humans on earth to reach a maximum in this century, which may lead to a slowing of growth in technology. Many forces could offset this change. The effective number of people with whom each individual can share ideas could grow through more intense integration. The total number of people living in cities will continue to grow long after total population has begun its decline. Barring some drastic political setback, the trade and communication links between all these cities will also grow tighter still. Rising levels of human capital per capita could make the average individual better at discovering and sharing ideas. If new institutions change incentives, the fraction of the available human capital that is devoted to producing and sharing ideas could go up fast enough to offset the decline in the total. These forces have been operating in much of the OECD during the last half century and there is much room for each to have big effects as policies and levels of human capital in places like China and India come to resemble those in the OECD. For all these reasons, it is quite possible that growth at the technological frontier could continue for the foreseeable future and who knows, might even increase yet again in this century compared to the last. Nevertheless, this century will mark a fundamental phase shift in the growth process. Growth in the stock of ideas will likely no longer be supported by growth in the total number of humans.

Figure 3: Growth Variation and Distance from the Frontier



Source: Penn World Tables 6.1.

Fact 3: Variation in modern growth rates.

The variation in the rate of growth of per capita GDP increases with the distance from the technology frontier.⁹

Figure 3 illustrates this fact by showing one of the more familiar graphs in the growth literature, the “triangle” plot of the average annual growth rate between 1960 and 2000 against initial per capita GDP. At the frontier, the United States is one of the richest countries in the world and exhibits steady growth at a rate of about 2% per year. The variation of growth rates is much smaller for the richest countries than for the poorest. Both rapid catch-up growth and tremendous lost opportunities can be seen in the growth experiences among the poor.

One of the main reasons the variance far from the frontier can be so high is that the rate at which rapid catch-up growth can occur is now faster than it has ever been. For example, between 1950 and 1980, growth in Japan averaged 6.5% per year. More recently, China’s catch-up growth has been even faster, averaging 8.2% between 1980 and

⁹Examples of papers that contributed to documenting and analyzing this fact include Romer (1987), DeLong (1988), Lucas (1988) and Barro (1991).

2004; per capita GDP in China has risen by a factor of seven in just a single generation. By comparison, the most rapid growth in the world between 1870 and 1913 occurred in Argentina, at a rate that averaged less than 2.5% per year.¹⁰ Comin and Hobijn (2008) show this same increase in data on the take up of individual technologies in different countries. Adoptions lags have been growing shorter over time.

Yet this rapid catch-up is far from the norm. Ethiopia and Nicaragua provide two examples. Growth in Ethiopia has been slow and unsteady. In 1950, Ethiopia was 34 times poorer than the United States. By 2003, however, this ratio had risen to 50. For Nicaragua, the situation is even worse, as per capita GDP has actually declined over the last half century.

Before discussing this fact in the context of growth theory, we state the next closely-related fact.

Fact 4: Large Income and TFP differences.

Differences in measured inputs explain less than half of the enormous cross country differences in per capita GDP.¹¹

As shown in Figure 4, per capita GDP in the poorest countries of the world is about 1/50th of that in the United States. Just as Solow (1957) documented a large total factor productivity (TFP) “residual” in accounting for growth over time, there is a large TFP residual — when measured in a conventional neoclassical setup — in accounting for differences in levels of per capita GDP across countries. In both cases, the residual is quantitatively at least as important as the measured factor inputs. The differences in income and TFP across countries are large and highly correlated: poor countries are poor not only because they have less physical and human capital per worker than rich countries, but also because they use their inputs much less efficiently.¹²

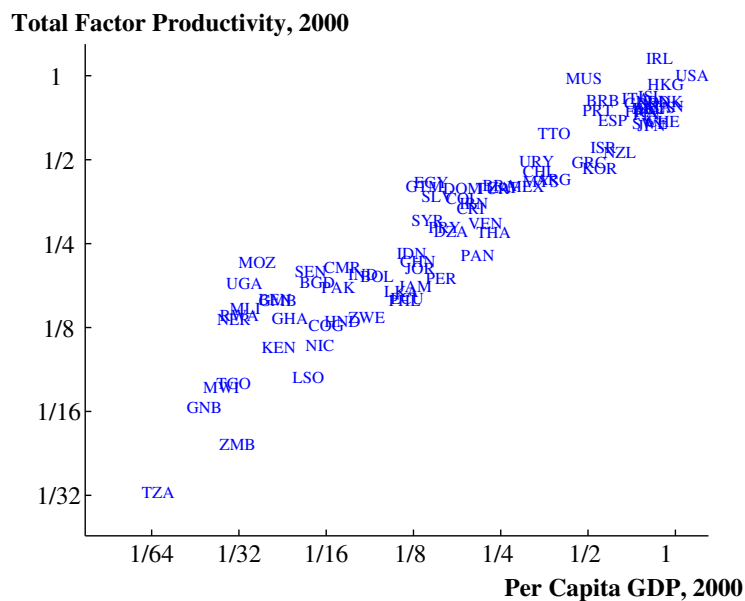
Facts 3 and 4 are closely related: there are enormous income differences across countries, but these gaps can occasionally be closed with remarkable speed. And while

¹⁰China’s recent growth is measured using the Penn World Tables 6.2. The fact about Argentina is taken from the 2008 update of Angus Maddison’s data.

¹¹Examples of papers that contributed to documenting this fact include Mankiw, Romer and Weil (1992), Islam (1995), Caselli, Esquivel and Lefort (1996), Klenow and Rodriguez-Clare (1997), Hall and Jones (1999), and Clemens, Montenegro and Pritchett (2008). See Caselli (2005) for an overview.

¹²The TFP measure is constructed following the methodology in Hall and Jones (1999). It corresponds to the value of A in a production function $Y = K^\alpha (AhL)^{1-\alpha}$, with $\alpha = 1/3$ and h measured as average efficiency units of labor per worker in a Mincerian schooling approach.

Figure 4: Large Income and TFP Differences



Note: Both TFP and per capita GDP are normalized so that the U.S. values are 1.0. TFP is reported in “labor-augmenting” form and is constructed following the methodology of Hall and Jones (1999) using the Penn World Tables 6.1 and the education data of Barro and Lee (2000).

ideas (in our reading of the literature) are a generally accepted explanation for economic growth in the frontier countries, the role of ideas in explaining economic development — the rising rates of catch-up growth and large income differences across countries — is less widely appreciated.

The textbook explanation for the rapid catch-up growth that we see in Japan, South Korea, and China is transition dynamics in a neoclassical growth model. A significant problem with this explanation, of course, is that it is based on a closed-economy setting where capital cannot flow across countries to equate marginal products. But international capital flows seem important in practice; for example, Caselli and Feyrer (2007) show that marginal products of capital are remarkably similar across countries. While the textbook transition dynamics — driven by diminishing returns to capital accumulation — are elegant and easy to explain, they are most likely not especially relevant to catch-up growth in practice. Instead, one of two directions is more promising. It could be that some kind of adjustment costs in an open economy provide the diminishing

returns that govern transition dynamics (Barro, Mankiw and Sala-i-Martin, 1995). Alternatively, catch-up growth could be associated with the dynamics of idea flows and technology adoption.¹³ The fact that rates of catch-up growth have been rising over time is difficult to understand in a pure neoclassical framework but is a natural occurrence in a world of ideas, where the technology frontier is relevant to what can be achieved. For some reason — perhaps because the diminishing returns to capital in a production function are much easier to measure than the process of technology adoption — the idea explanation for catch-up growth is not as well-established as it probably should be.

There is very broad agreement that differences in institutions must be the fundamental source of the wide differences in growth rates observed for countries at low levels of income and for the low income and TFP levels themselves. In any model, bad institutions will distort the usage of rival inputs like labor and capital (Banerjee and Duflo, 2005; Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009). Our point is that one must allow for the possibility that they also distort the adoption and utilization of ideas from leading nations. The potential for ideas to diffuse across countries can significantly amplify the effects of institutions.

The interaction between institutions and idea flows is easy to illustrate in familiar contexts. For example, until 1996, opponents successfully used the local permit process to keep Wal-Mart from building stores or distribution centers in Vermont. This kept powerful logistics ideas like cross-docking that Wal-Mart pioneered from being used to raise productivity in retailing in the state. Such nonrival ideas must have been at least partly excludable. This is why Wal-Mart was willing to spend resources developing them and why competitors were not able to copy them. All this fits comfortably in the default model of endogenous discovery of ideas as partially excludable nonrival goods.

Some economists have suggested that even though ideas should be treated as partially excludable nonrival goods in the country at the technological frontier where new ideas are discovered, they should nevertheless be treated as pure public goods when we consider questions of development. This might sound reasonable in the abstract, but at the micro-level of a specific idea it can seem absurd. It would imply that institutions

¹³See Nelson and Phelps (1966), Parente and Prescott (1994), Romer (1994), Howitt (2000), Klenow and Rodriguez-Clare (2005), and Lucas (2009).


could keep cross-docking out of Vermont, at a time when it was available in Japan and North Korea.

The pure public good approach also makes it very difficult to address one of the most striking episodes in modern history. By something like 1300 A.D., China was the most technologically advanced country in the world, with a large integrated population. According to the Lee model, it should have persisted indefinitely as the world technology leader. The explosive dynamics of the virtuous circle between population and ideas suggests that such technological leads should never be lost. Only a remarkable and persistent failure of institutions can explain how China fell so far behind Europe. A model in which institutions can stifle innovation could explain why China lost the lead, but it takes a model in which institutions can also stop inflows of ideas from the rest of the world to explain why, for more than 500 years, ideas developed in the west were not more systematically adopted in China.

If incentives, and hence institutions, can influence the rate of discovery of partially excludable nonrival goods, they must also be able to influence the spread and utilization of these goods. We see this at the level of entire nations. We also see wide variation in the utilization of specific technologies like cross-docking. For example, Baily, Gersbach, Scherer and Lichtenberg (1995) note the use of different technologies in Germany, Japan, and the United States in manufacturing steel (the use of minimills or modern integrated mills) and beer (mass production equipment versus small breweries). Comin and Hobijn (2004) document that many different technologies are utilized with widely varying intensities throughout the world, including electricity, personal computers, and shipping technologies.

There is much interesting work that tries to explain why inefficient institutions can persist and why efficient institutions can be difficult to establish (Acemoglu and Robinson, 2006; Greif, 2006; North et al., 2009). As that work progresses, it seems reasonable to insist that models of growth and development allow for the possibility that political and regulatory institutions like those in used in Vermont can sometimes be used in different countries to keep out technologies like cross-docking. At the other extreme, it also seems reasonable to allow for the possibility that in some countries (think of Haiti or Somalia as extreme cases) ideas like cross-docking are sometimes not brought into a country because its institutions cannot assure even the most basic elements of

Figure 5: Years of Schooling by Birth Cohort, United States



Source: Goldin and Katz (2007), Figure 7.


personal security and protection of private property.

Fact 5: Rising human capital.

Human capital per worker is rising throughout the world.

One of Kaldor's main stylized facts was the sustained growth in capital per worker over time. Fact 5 is the human capital counterpart. Figure 5 documents the sustained increase in educational attainment over time in the U.S. economy. The cohort born in 1920 obtained just over 10 years of education, while the cohort born in 1980 went to school for 14 years. Another way to look at education is by the average years of educational attainment for the entire labor force in a given year. By this measure (not shown), educational attainment has, until recently, increased by about one year per decade. Assuming a Mincerian return to education of 6% per year, this increase contributes about 0.6 percentage points per year to U.S. growth, a significant fraction of our 2% per capita growth. The slowdown by birth cohort shown in Figure 5 suggests that average years of education in the labor force is destined to slow in the future, erasing this 0.6 percentage points in the decades to come.

Figure 6: The U.S. College and High School Wage Premiums



Source: Goldin and Katz (2008), Table D1.

Another of Kaldor's facts was that the rise in the capital-labor ratio occurred without a decline in the real interest rate. A natural interpretation of these two facts is that technological progress kept the marginal product of capital from declining. A similar phenomenon occurs with human capital, as shown in our last stylized fact.

Fact 6: Long-run stability of relative wages.

The rising quantity of human capital relative to unskilled labor has not been matched by a sustained decline in its relative price.

Figure 6 shows the college and high school wage premiums in the United States. While there is (interesting) variation in these premiums over time, one of the main things that stands out is this: despite the large increases in educational attainment by some people in the United States, the wage premiums associated with college and with high school show no tendency to decline. The standard interpretation of this fact is that skill-biased technical change has shifted out the relative demand for highly-educated workers, more than offsetting the downward pressure on the wage premium that is associated with the increase in their relative supply (Katz and Murphy, 1992). Similar facts apply more generally throughout the world, as reviewed by Goldberg and Pavcnik

(2007).

A fascinating question to consider is why technological change might be skill-biased in this fashion. Acemoglu (1998), emphasizing nonrivalry and the interaction between scale and incentives, argues that a key determinant of the direction of technical change is the number of people for whom the new technology will be useful. The rising supply of highly-educated labor tilts technical change in its own direction. Another related possibility is tied to the rising extent of the market cited in Fact 1. In particular, when research is conducted primarily in the advanced countries, increases in the extent of the market to developing countries will raise the return to ideas and therefore the wages of the people who produce them.

4. Conclusion

In 1961, Nicholas Kaldor used his list of six “stylized” facts to summarize what was then known about economic growth and to shape the direction of research in the future. Re-doing this exercise today, nearly fifty years later, reveals just how much progress has been made. Whereas Kaldor’s original facts were explained almost entirely using the neoclassical growth model, the facts we highlight reveal the broader reach of modern growth theory. To capture these facts, a growth model must consider the interaction between ideas, institutions, population, and human capital. Two of the major facts of growth — the extraordinary rise in the extent of the market associated with globalization and the acceleration over the very long run — are readily understood as reflecting the defining characteristic of ideas, their nonrivalry. The next two major facts — the enormous income and TFP differences across countries as well as the stunning variation in growth rates for countries far behind the technology frontier — testify to the importance of institutions and institutional change. Our final two facts parallel two of Kaldor’s original observations, but whereas his emphasis was on physical capital, the emphasis in modern growth theory is on human capital. Human capital per worker is rising rapidly, and this occurs despite no systematic trend in the wage premium associated with education.

These facts also reveal important complementarities among the key endogenous variables. The virtuous circle between population and ideas accounts for the acceler-

ation of growth. Institutions may have their most important effects on cross-country income differences by hindering the adoption and utilization of ideas from throughout the world. Institutions like public education and the university system are surely important for understanding the growth in human capital. And institutions are themselves ideas — inventions that shape the allocation of resources — and the search for better institutions is unending. Finally, the rising extent of the market, which raises the return to ideas and therefore to the human capital that is a fundamental input to idea production, may help explain why the college wage premium has not fallen systematically despite the huge increases in the ratio of college graduates to high school graduates.

Such complementarities exemplify the value of the applied general equilibrium approach. They are the fundamental reason why we seek a unified framework for understanding growth. Going forward, the research agenda will surely include putting ingredients like those we have outlined in this paper together into a single formal model. Further out on the horizon, one may hope for a successful conclusion to the ongoing hunt for a simple model of institutional evolution. Combining that with the unified approach to growth outlined here would surely constitute the economics equivalent of a grand unified theory — a worthy goal by which we may be judged when future generations look back fifty years from now and quaintly revisit our “ambitious” list of stylized facts.

References

- Acemoglu, Daron, "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *Quarterly Journal of Economics*, 1998, 113, 1055–1089.
- and James A. Robinson, *Economic Origins of Dictatorship and Democracy* 2006.
- Aghion, Philippe and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica*, March 1992, 60 (2), 323–351.
- Baily, Martin Neil, Hans Gersbach, F. M. Scherer, and Frank R. Lichtenberg, "Efficiency in Manufacturing and the Need for Global Competition," *Brookings Papers on Economic Activity. Microeconomics*, 1995, pp. 307–358.
- 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 J., "Economic Growth in a Cross Section of Countries," *Quarterly Journal of Economics*, May 1991, 106, 407–443.
- and Jongwha Lee, "International Data on Educational Attainment: Updates and Implications," April 2000. CID Working Paper No. 42, Harvard University.
- , N. Gregory Mankiw, and Xavier Sala-i-Martin, "Capital Mobility in Neoclassical Models of Growth," *American Economic Review*, March 1995, 85, 103–115.
- Boldrin, Michele and David K. Levine, "Perfectly Competitive Innovation," *Journal of Monetary Economics*, April 2008, 55 (3), 435–453.
- Caselli, Francesco, "Accounting for Cross Country Income Differences," in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Elsevier, 2005.
- and James Feyrer, "The Marginal Product of Capital," *Quarterly Journal of Economics*, 05 2007, 122 (2), 535–568.
- , 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.
- Clark, Gregory, *A Farewell to Alms: A Brief Economic History of the World*, Princeton University Press, 2007.

- Clemens, Michael A., Claudio E. Montenegro, and Lant Pritchett, "The Place Premium: Wage Differences for Identical Workers across the US Border," July 2008. Harvard Kennedy School of Government working paper.
- Comin, Diego A. and Bart Hobijn, "Cross-Country Technology Adoption: Making the Theories Face the Facts," *Journal of Monetary Economics*, January 2004, 51 (1), 39–83.
- and —, "An Exploration of Technology Diffusion," Harvard Business School Working Paper 08-093 March 2008.
- DeLong, J. Bradford, "Productivity Growth, Convergence, and Welfare: Comment," *American Economic Review*, 1988, 78, 1138–1154.
- Diamond, Jared, *Guns, Germs, and Steel*, New York: W.W. Norton and Co., 1997.
- Eaton, Jonathan and Samuel S. Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, 70 (5), 1741–1779.
- Feenstra, Robert C, "Integration of Trade and Disintegration of Production in the Global Economy," *Journal of Economic Perspectives*, Fall 1998, 12 (4), 31–50.
- Galor, Oded and David Weil, "Population, Technology, and Growth: From the Malthusian Regime to the Demographic Transition," *American Economic Review*, September 2000, 90, 806–828.
- Goldberg, Pinelopi K. and Nina Pavcnik, "Distributional Effects of Globalization in Developing Countries," *Journal of Economic Literature*, March 2007, 45, 39–82.
- 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.
- and —, *The Race between Education and Technology*, Belknap Press, 2008.
- Greif, Avner, *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*, Cambridge University Press, 2006.
- Grossman, Gene M. and Elhanan Helpman, *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press, 1991.
- Hall, Robert E. and Charles I. Jones, "Why Do Some Countries Produce So Much More Output per Worker than Others?," *Quarterly Journal of Economics*, February 1999, 114 (1), 83–116.
- Helpman, Elhanan and Paul R. Krugman, *Market Structure and Foreign Trade*, Cambridge, Mass.: MIT Press, 1985.

- Howitt, Peter, "Endogenous Growth and Cross-Country Income Differences," *American Economic Review*, September 2000, 90 (4), 829–846.
- Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, forthcoming 2009.
- Islam, Nazrul, "Growth Empirics: A Panel Data Approach," *Quarterly Journal of Economics*, November 1995, 110, 1127–1170.
- Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, August 1995, 103 (4), 759–784.
- , "Growth: With or Without Scale Effects?," *American Economic Association Papers and Proceedings*, May 1999, 89, 139–144.
- , "Sources of U.S. Economic Growth in a World of Ideas," *American Economic Review*, March 2002, 92 (1), 220–239.
- Kaldor, Nicholas, "Capital Accumulation and Economic Growth," in FA. Lutz and D.C. Hague, eds., *The Theory of Capital*, St. Martins Press, 1961, pp. 177–222.
- 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.
- 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 —, "Externalities and Growth," in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Amsterdam: Elsevier, 2005.
- 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.
- Krugman, Paul R., "Increasing Returns, Monopolistic Competition, and International Trade," *Journal of International Economics*, November 1979, 9 (4), 469–479.
- , "Increasing Returns and Economic Geography," *Journal of Political Economy*, June 1991, 99 (3), 483–499.

- , “Growing World Trade: Causes and Consequences,” *Brookings Papers on Economic Activity*, 1995, 1, 327–377.
- 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.
- , “Trade and the Diffusion of the Industrial Revolution,” *American Economic Journal: Macroeconomics*, January 2009, 1 (1), 1–25.
- Maddison, Angus, “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.
- Nelson, Richard R. and Edmund S. Phelps, “Investment in Humans, Technological Diffusion, and Economic Growth,” *American Economic Association Papers and Proceedings*, May 1966, 56, 69–75.
- Nordhaus, William D., “Do Real Output and Real Wage Measures Capture Reality? The History of Lighting Suggests Not,” in Timothy F. Bresnahan and Robert J. Gordon, eds., *The Economics of New Goods*, University of Chicago Press, 1997, pp. 29–66.
- , “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.
- North, Douglass C., John J. Wallis, and Barry R. Weingast, *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History*, Cambridge University Press, 2009.
- Parente, Stephen L. and Edward C. Prescott, “Barriers to Technology Adoption and Development,” *Journal of Political Economy*, April 1994, 102 (2), 298–321.
- Ramondo, Natalia and Andres Rodriguez-Clare, “Trade, Multinational Production, and the Gains from Openness,” 2009. Penn State University working paper.
- Restuccia, Diego and Richard Rogerson, “Policy Distortions and Aggregate Productivity with Heterogeneous Plants,” *Review of Economic Dynamics*, October 2008, 11, 707–720.

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

—, "Crazy Explanations for the Productivity Slowdown," in Stanley Fischer, ed., *NBER Macroeconomics Annual 1987*, Cambridge, MA: MIT Press, 1987.

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

—, "New Goods, Old Theory, and the Welfare Costs of Trade Restrictions," *Journal of Development Economics*, 1994, 43, 5–38.

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

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.

Swan, Trevor W., "Economic Growth and Capital Accumulation," *The Economic Record*, November 1956, 32, 334–361.

United Nations, "World Urbanization Prospects 2007," Technical Report, Department of Economic and Social Affairs February 2008.