

Annual Review of Economics

The Past and Future of Economic Growth: A Semi-Endogenous Perspective

Charles I. Jones

Graduate School of Business, Stanford University, Stanford, California, USA;
email: chad.jones@stanford.edu

Annu. Rev. Econ. 2022. 14:125–52

First published as a Review in Advance on
April 6, 2022

The *Annual Review of Economics* is online at
economics.annualreviews.org

<https://doi.org/10.1146/annurev-economics-080521-012458>

Copyright © 2022 by Annual Reviews.
All rights reserved

JEL codes: J11, O40

Keywords

economic growth, ideas, population, innovation

Abstract

The nonrivalry of ideas gives rise to increasing returns, a fact celebrated in Paul Romer's recent Nobel Prize. An implication is that the long-run rate of economic growth is the product of the degree of increasing returns and the growth rate of research effort; this is the essence of semi-endogenous growth theory. This review interprets past and future growth from a semi-endogenous perspective. For 50+ years, US growth has substantially exceeded its long-run rate because of rising educational attainment, declining misallocation, and rising (global) research intensity, implying that frontier growth could slow markedly in the future. Other forces push in the opposite direction. First is the prospect of "finding new Einsteins": How many talented researchers have we missed historically because of the underdevelopment of China and India and because of barriers that discouraged women inventors? Second is the longer-term prospect that artificial intelligence could augment or even replace people as researchers. Throughout, the review highlights many opportunities for further research.

ANNUAL
REVIEWS **CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

1. INTRODUCTION

For the past 150 years, GDP per person in the United States has increased exponentially at a relatively stable rate of around 2 percent per year. How do we understand the sustained growth of average incomes in a rich country like the United States? This article describes one potential answer to this question, from the perspective of semi-endogenous growth (SEG) theory.

The essence of SEG theory is simple. Romer (1990), in a contribution that was recognized with a Nobel Prize in 2018, emphasized that ideas are nonrival, or, more colloquially, infinitely usable. That is, once an idea is invented, it is feasible for the idea to be used by one person or by one thousand people or by one billion people simultaneously. Most goods are rival: Suppose we raise the productivity of one worker by giving them a computer. If we want to similarly raise the productivity of a million such workers, we need a million computers. But ideas are different. The computer code for a spreadsheet program only needs to be invented once. It can then be used by one million workers simultaneously.

This infinite usability of ideas means that income per person depends on the aggregate stock of ideas, not the stock of ideas per person. Contrast that with capital, where income per person depends on capital per person. Each new idea can be used by any number of people at the same time, so every improvement has the potential to benefit everyone. Faster computers, a new pharmaceutical innovation, a new technique that makes rice more resistant to drought, better solar panel efficiency: Each of these new ideas can be used simultaneously by any number of workers and therefore raises the productivity and income of any number of workers.

Where do ideas come from? The history of innovation is very clear on this point: new ideas are discovered through the hard work and serendipity of people. Just as more autoworkers will produce more cars, more researchers and innovators will produce more new ideas.

The surprise is that we are now done; that is all we need for the SEG model. People produce ideas and, because of nonrivalry, those ideas raise everyone's income. This means that income per person depends on the number of people searching for ideas (i.e., researchers). But then the growth rate of income per person depends on the growth rate of researchers, which is in turn ultimately equal to the growth rate of the population.

This article develops the point more clearly in a formal model and then explores the past and future of economic growth using this framework. Section 2 lays out the mathematical model that parallels the description above. Section 3 conducts a growth accounting exercise for the United States to make a key point: Although SEG theory implies that the entirety of long-run growth is ultimately due to population growth, this is far from being true historically, say for the past 75 years. Instead, population growth contributes only around 20 percent of US economic growth since 1950. Rising educational attainment, declining misallocation, and rising (global) research intensity account for more than 80 percent of growth. Importantly, this statement is true even in the semi-endogenous framework itself.

Section 4 then reviews the literature and makes two broader comments. First, the SEG setup is more widely used than is commonly appreciated. For example, works by Krugman (1979) and Melitz (2003) fit in this class. Second, there are many current research projects using fully endogenous growth models that would be interesting to revisit using the semi-endogenous framework.

Sections 5 and 6 then turn to the future of economic growth. Section 5, building on the growth accounting exercise, highlights various reasons economic growth could slow in the future, including the end of historical transition dynamics and slowing, or even negative, population growth. On the flip side, Section 6 points to possible changes that could keep growth from slowing or even increase the growth rate. One of these forces is "finding new Einsteins," for example in China or India or among populations such as women who are historically underrepresented among innovators. Another force is automation and artificial intelligence: If machines can replace

people in the production of ideas, it is possible for growth to speed up. Finally, Section 7 concludes by discussing some of the most important questions related to long-run economic growth that remain unresolved.

2. A SIMPLE MODEL

The economic environment for a simple SEG model is given in **Table 1**. The basic production function for consumption goods is $Y_t = A_t^\sigma L_{yt}$, where A_t is the stock of ideas and L_{yt} is the amount of labor used to make consumption goods. This production function already embodies the key insight of Romer (1990): The nonrivalry of ideas gives rise to increasing returns. The basic argument is easy to see. If we wish to double the production of any given good, one way to do it is to completely replicate the production setup: We build an identical factory on identical land and use identical machines and workers. If everything is replicated precisely, we should double output. Notice that in this replication process, it is only the rival factors of production that need to be replicated. The idea (design, blueprint, etc.) is nonrival, so the same set of instructions is used by the new factory. This means that there are constant returns to scale to the rival inputs. But as long as the marginal product of another idea is positive, there must be increasing returns to the ideas and the rival factors taken together. In Equation 1, this is captured by the constant returns to labor, the only rival input, and increasing returns to ideas and labor together, with $\sigma > 0$ measuring the overall degree of increasing returns.

Equation 2 is the SEG equation. New ideas are produced using researchers R_t and the existing stock of knowledge. The parameter λ allows for the possibility of duplication effects, so that doubling the number of researchers at a point in time may potentially less than double the innovation rate; however, any $\lambda > 0$, including $\lambda = 1$, is allowed. The parameter $\beta > 0$ captures the rate at which ideas—that is, proportional improvements in productivity—are getting harder to find.

In my own work I go back and forth in terms of how I specify the idea production function. An equivalent formulation is $\dot{A}_t = R_t^\lambda A_t^\phi$, where $\phi < 1$. Comparing these two formulations, one sees that they are the same, with $\beta \equiv 1 - \phi$. The formulation with β is nice in that it preserves a useful convention that parameter values should be specified so that they are positive. The formulation with ϕ , however, is valuable in other ways. For example, in the case of $\phi > 0$, past discoveries increase the productivity of researchers in producing new ideas, whereas when $\phi < 0$, past discoveries make it harder to generate new ideas. In much of the quality ladder literature (mentioned shortly below), ideas are thought of as proportional improvements in productivity, so the β formulation is more relevant.

The remainder of the equations in **Table 1** describe the resource constraints. There is an exogenous population L_t that grows exponentially at an exogenous rate $n > 0$, and this population is divided into researchers and people who work to make consumption goods. How to make this division is the only economic allocation in the simple model. We assume a basic rule of thumb allocation with constant research intensity: A fixed fraction \bar{s} of the population works in research,

Table 1 The economic environment of the simple model

	Equation	Equation no.
Final good	$Y_t = A_t^\sigma L_{yt}$	1
Ideas	$\dot{A}_t/A_t = R_t^\lambda A_t^{-\beta}$	2
Resource constraint	$R_t + L_{yt} = L_t$	3
Population growth	$L_t = L_0 e^{nt}$	3
Allocation	$R_t = \bar{s}L_t$	4

so that $1 - \bar{s}$ make consumption goods. Jones (2005) considers letting markets or a social planner allocate resources, but the key points can be made most easily with exogenous research intensity.

2.1. The Quality Ladder Alternative

This setup can easily be recast as a quality ladder model à la Grossman & Helpman (1991) and Aghion & Howitt (1992). In that literature, ideas represent proportional improvements in productivity, viewed as steps up the quality ladder. We briefly outline here how this interpretation can be incorporated. First, let productivity be given by $A_t = q^{N_t}$, where $q > 1$ is the step size of the quality ladder and N_t corresponds to how far up the quality ladder (i.e., how many rungs) the economy has traversed. Each step raises productivity by a constant percentage.

The innovation rate in the economy is denoted $z_t = R_t^\lambda A_t^{-\beta}$; this is the rate at which the economy takes steps up the quality ladder, i.e., $\dot{N}_t = z_t$. The presence of $\beta > 0$ (i.e., it gets progressively harder to take steps up the quality ladder) makes this a SEG model; this is basically the structure proposed by Segerstrom (1998). Notice that combining $A_t = q^{N_t}$ and $\dot{N}_t = z_t = R_t^\lambda A_t^{-\beta}$ means that $\dot{A}_t/A_t = \log q \times R_t^\lambda A_t^{-\beta}$, which is essentially the idea production function in **Table 1**. We will stick with the simpler framework summarized in the table, but it is straightforward to give a quality ladder interpretation to this setup.

2.2. Balanced Growth

We now characterize the balanced growth path, defined as a situation in which all variables grow at constant exponential rates (possibly including zero). To begin, note that output per person, $y \equiv Y/L$, is given by

$$y_t = A_t^\sigma (1 - \bar{s}). \quad 1.$$

Output per person is proportional to the aggregate stock of ideas: Because the ideas are nonrival, they can be used by any number of workers and therefore each idea benefits everyone in this simple economy.

Next, we turn to solving for the stock of ideas. Let g_x denote the growth rate of any variable x along the balanced growth path. Solving the idea production function $\dot{A}_t/A_t = R_t^\lambda A_t^{-\beta}$ for A_t along a balanced growth path (denoted by $*$) then gives

$$\begin{aligned} A_t^* &= \left(\frac{1}{g_A}\right)^{1/\beta} R_t^{\lambda/\beta} \\ &= \left(\frac{1}{g_A}\right)^{1/\beta} (\bar{s}L_t)^{\lambda/\beta}. \end{aligned} \quad 2.$$

If $\lambda = \beta = 1$, then the stock of ideas A_t is directly proportional to the number of researchers along a balanced growth path. This captures a property of all standard production functions: The more autoworkers we have, the more cars we can produce. Similarly, the more researchers we have, the more ideas we can produce. The extent to which $\lambda \neq 1$ and $\beta \neq 1$ allows for the possibility of different returns to scale in the idea production function.

Combining Equations 1 and 2 gives output per person along the balanced growth path,

$$\begin{aligned} y_t^* &= (1 - \bar{s}) \left(\frac{1}{g_A}\right)^{\sigma/\beta} R_t^\gamma \\ &= (1 - \bar{s}) \left(\frac{1}{g_A}\right)^{\sigma/\beta} \bar{s}^\gamma L_t^\gamma, \end{aligned} \quad 3.$$

where $\gamma \equiv \lambda\sigma/\beta$ measures the overall degree of increasing returns to scale in the economy. Output per person is proportional to the number of researchers raised to the power γ . The nonrivalry of ideas gives rise to increasing returns to scale: $\sigma > 0$ implies $\gamma > 0$. More researchers means more ideas, and because the ideas are nonrival this means more consumption for each person in this simple economy.

Finally, the growth rate of the economy along the balanced growth path can be solved for by taking logs and derivatives of Equations 1 and 3 to yield

$$g_y = \sigma g_A = \frac{\lambda\sigma n}{\beta} \equiv \gamma n. \quad 4.$$

That is, the long-run growth rate of the economy is proportional to the rate of population growth, where the factor of proportionality is the degree of increasing returns to scale, γ . The more important ideas are to production ($\uparrow\sigma$), or the less duplication ($\uparrow\lambda$), the higher is the long-run growth rate of the economy. Faster growth in research, ultimately because of faster growth in population ($\uparrow n$), also raises the long-run growth rate. Conversely, the harder it is to find new ideas ($\uparrow\beta$), the slower long-run growth is.

2.3. Discussion

We now discuss some natural questions that arise.

2.3.1. Does the semi-endogenous growth framework apply at the country level? A very important point to appreciate from the beginning is that one must be careful in applying this model to country-level data because of the diffusion of ideas. We do not believe that Luxembourg or Singapore grows only because of ideas invented by the researchers in those small countries. Instead, essentially all countries eventually benefit from ideas created throughout the world. In any empirical work on growth, one must be careful to consider these cross-country idea flows. For example, we do not think that Kenya, Japan, and France have their growth rates determined primarily by their own population growth rates.

2.3.2. Can growth in the quality of researchers replace growth in their quantity? What we see in the SEG model is that long-run growth is proportional to the growth rate of the number of researchers. An interesting question to consider is whether growth in the quality of researchers could substitute for growth in their quantity. In other words, suppose we have a constant number of researchers. Is it possible that their human capital will grow exponentially, thereby generating a growth rate that is proportional to the growth rate of their human capital?

The answer to this question relates to Lucas's (1988) model. In particular, Lucas considers a model in which human capital grows endogenously at a rate determined by the fraction of time that individuals spend accumulating human capital. If one were to adopt Lucas's formulation here, then the answer would be yes, quality (human capital) could replace quantity (population growth).

The problem with this answer is that it does not work in general but only for a knife-edge formulation of human capital accumulation. For example, suppose we take an approach to human capital à la Mincer (1974), where each year of education raises a researcher's productivity by 5 percent. Then if researchers always get 20 years of education, they are more productive than unskilled workers, but by a constant amount. That is, here human capital would, say, double the number of quality-adjusted workers but would not lead to exponential growth in the quality of the workers. Human capital has a level effect on quality; it does not lead to permanent growth in the quality of researchers.

One way to nest these two formulations is to suppose that a person's human capital accumulates according to

$$\dot{b}_t = u_t b_t^\psi - \delta b_t,$$

where $u_t \in (0, 1)$ is the fraction of time a person spends accumulating human capital and δ captures the depreciation of human capital. Notice that if $u_t = \bar{u}$ is constant and $\psi < 1$, then this equation has a steady-state level of b , and changes in \bar{u} would change the long-run level of a person's human capital but would not generate sustained growth in b . Lucas (1988) instead assumed $\psi = 1$ so that a constant value of u could generate sustained exponential growth in quality. This is possible, but there is essentially no evidence for this assumption. Does person A learn more than person B per school hour simply because person A's parents are more educated? Perhaps. But if A's parents have twice as much human capital as B's parents, does A learn twice as much per hour as B does? That is what is required here. Also, the linearity associated with $\psi = 1$ implies that different countries with different investment rates in human capital should grow at permanently different rates, with the United States—with its high educational attainment—counterfactually growing faster than other countries in the twentieth century (see Bils & Klenow 2000). Also, rising educational attainment over the twentieth century would lead to rising growth rates, which we do not see in the United States or in other countries.

The bottom line is that human capital can certainly be important—and in Section 3 below we suggest that it accounts for one-fourth of US growth in recent decades—but it is unlikely to be a substitute for population growth in generating exponential growth in the effective number of researchers.

2.3.3. Linearity and the population equation. One might wonder whether this linearity criticism also applies to the SEG approach. After all, this approach assumes a linear differential equation as well, in this case for population. There are two responses to this concern. First, population growth is of course a historical fact, so one can simply take it as given and ask what it implies about growth in GDP per person.

The second response is that the law of motion for population is inherently linear, in a way that the accumulation of physical capital, human capital, and ideas is not. To see why, consider a simple example, a process at the heart of all population growth: cell division. Cells divide and replicate. One becomes two, then four, then eight, and so on. The population of some types of bacteria can double every 20 minutes as long as sufficient nutrients are available. If this rate were unchecked, a colony of bacteria could grow exponentially to equal the mass of the earth itself in two days! Cells, bacteria, and indeed people reproduce in proportion to their numbers, and it is this physical process that underlies the linear differential equation for population growth.¹

2.4. Evidence

Bloom et al. (2020) provide a recent look at idea production functions at different levels of aggregation. More specifically, they study the null hypothesis that $\beta = 0$ —i.e., the basic hypothesis of fully endogenous growth models that proportional improvements in productivity are not getting harder to find. They do this at the aggregate level for the US economy and then at the micro level in various case studies.

Figure 1 summarizes two of their key findings. **Figure 1a** shows the result for the aggregate US economy. In particular, two series are plotted: a time series of total factor productivity (TFP)

¹For more discussion of the linearity critique, readers are referred to Jones (2005) and Growiec (2007).




Figure 1

Productivity growth and research. These graphs, adapted from Bloom et al. (2020), illustrate that proportional improvements in productivity—i.e., ideas—are getting harder to find. Productivity growth rates are stable or declining, while research effort is rising rapidly. These data and the other case studies in Bloom et al.'s paper suggest that $\beta > 0$. Abbreviation: TFP, total factor productivity.

growth (an estimate of \dot{A}_t/A_t if $\sigma = 1$, since we do not observe ideas directly in comparable units) and a time series of research effort, R_t . If $\beta = 0$, then these two series should move together because then $\dot{A}_t/A_t = R_t^\lambda$. But as is well known at this point, the data look very different: TFP growth rates are stable or even declining, while the number of researchers has grown considerably. In other words, the data prefer $\beta > 0$ so that proportional improvements in productivity are getting harder to find and require ever-rising research effort. The aggregate data are roughly consistent with a value of $\beta \approx 3$ if we assume $\lambda = 1$. Notice that in terms of $\phi = 1 - \beta$, this implies $\phi \approx -2$; in other words, even unit improvements in productivity are getting harder to achieve.

Figure 1b shows the same kind of evidence but for an individual technology—in this case, the famous Moore's law. The law (really a stylized fact) states that the number of transistors that can fit on a computer chip doubled on average every two years between 1971 and 2010. (The

classic picture of this doubling makes a great example of the use of a log scale for teaching.)² By the so-called rule of 70, this is equivalent to saying that the growth rate of chip density was a stable 35% per year during this 50-year period; this is the \dot{A}_t/A_t series plotted in **Figure 1b**. But how was this steady exponential growth accomplished? It turns out that research effort devoted to pushing Moore's law forward grew enormously over this period. Companies like Intel, Samsung, TSMC, and AMD—but also much older companies, like Fairchild Semiconductor and Texas Instruments—together invested ever-growing amounts of research effort into maintaining the steady growth. If $\beta = 0$ had characterized the idea production function for Moore's law, then the growth rate of chip density would have risen by a factor of 18 over the 50 year period, just like the research effort. Instead, growth was stable: By the end of the period, it required 18 times the amount of research effort as in the early 1970s to double chip density. Converting this into an estimate of β reveals that for Moore's law, $\beta \approx 0.2$.

Bloom et al. (2020) conduct this same exercise for other technologies as well, including agricultural productivity for corn, cotton, soybeans, and wheat; for medical technologies such as mortality from cancer and heart disease; and for firm-level data from Compustat. Essentially everywhere they look they find evidence for $\beta > 0$, with values ranging from a low of 0.2 for Moore's law to values of around 6 or 7 for corn and soybeans.³

2.5. Half Lives

Suppose the economy begins with initial conditions that deliver an initial growth rate of g_{A0} , which differs from the steady state. For example, perhaps a permanent increase in \bar{s} just occurred. How many years does it take before the growth rate converges halfway back to the steady-state value g_A^* ? The **Supplemental Appendix A.2** shows that the answer in our simple model is given by

$$t_{1/2}^* = \frac{1}{\beta g_A^*} \ln \left(\frac{g_{A0} + g_A^*}{g_{A0}} \right). \quad 5.$$

Apart from the initial growth rate and the steady-state growth rate, the key parameter that determines the answer is β , the rate at which ideas are becoming harder to find.⁴

We can now plug in various estimates of this parameter to see how it matters and to get some sense of how long the transition dynamics are in SEG models. To do this, let's suppose our steady-state growth rate g_A^* equals 1%, a rough estimate of average TFP growth in the US economy for the last half century.

Table 2 shows the half-life for different values of β and different initial growth rates g_{A0} . Recall that Bloom et al. (2020) estimate $\beta = 0.2$ for semiconductors and $\beta = 3$ for the aggregate and for

²An example can be found at https://web.archive.org/web/20170101183057if_/https://en.wikipedia.org/wiki/Moore's_law.

³A third generation of endogenous growth models associated with Dinopoulos & Thompson (1998), Peretto (1998), Young (1998), and Howitt (1999) maintain the assumption of $\beta = 0$ from the original models by Romer (1990) and Aghion & Howitt (1992) but assume this holds at the level of an individual product line or firm instead of the aggregate economy. A rise in the extensive margin—i.e., the number of product lines—can potentially offset the rise in aggregate research, so that the number of researchers per product line is stable and leads to stable economic growth. One of the motivations for Bloom et al.'s (2020) paper was to provide evidence on this assumption. So the finding of $\beta > 0$ in a wide range of settings raises questions about this alternative class of models.

⁴Trimborn et al. (2008) provide an elegant, fast algorithm that can be used to solve the transition dynamics of general SEG models in which allocations are determined by optimization rather than the assumption made here that allocations are constant over time (see also Lin & Shampine 2018).

Table 2 Half-life estimates (years)

β	$g_{A0} = 2\%$	$g_{A0} = 4\%$
0.2	203	112
1	41	22
3	14	7
5	8	4

The table reports the number of years it takes the basic model to transition halfway from an initial growth rate of g_{A0} to the steady-state value g_A^* . We assume $g_A^* = 1\%$ and a fixed research intensity for these calculations.

the average across the different cases they studied. The message from **Table 2** is that half-lives are very sensitive to the rate at which ideas are getting harder to find, β . For example, with $\beta = 3$, it takes around a decade for the economy to move halfway back to steady state, whereas for the semiconductor estimate of $\beta = 0.2$, the half life can be as high as a century or two (though a higher baseline g_A^* would lower these values). These calculations are consistent with the transition dynamics reported by Atkeson & Burstein (2019) and Atkeson et al. (2019), which allow the allocation of research to change endogenously along the transition path.

3. ACCOUNTING FOR HISTORICAL GROWTH

The preceding section laid out the basic SEG framework and argued that it helps us understand long-run growth. This section uses that framework to make an important point, which at first may seem contradictory but in fact is made by the SEG model itself: Whereas the ultimate source of long-run economic growth is population growth, historical growth accounting suggests that population growth accounts for only around 20% of US growth in recent decades. The remaining 80% is accounted for by rising educational attainment, increases in the fraction of the population devoted to R&D, and declines in misallocation. The way to reconcile these points is this: In the SEG framework, changes in these other factors have level effects in the long run rather than growth effects. However, over any historical period, these level effects could potentially be large and significant; transition dynamics are slow. This is what we observe in the US experience in recent decades. Put differently, the simple fact that a change in an allocation has a level effect (instead of a growth effect) does not mean that the growth consequences for several decades are not large.

To see this more formally, let's add physical capital, human capital, and labor force participation to the SEG framework. Consumption/investment goods Y are produced according to

$$Y_t = K_t^\alpha (Z_t b_t L_t)^{1-\alpha}, \quad 6.$$

where K is physical capital (which accumulates in the standard way), b is human capital per person, and

$$Z_t \equiv A_t M_t \quad 7.$$

is TFP, which captures both the stock of ideas, A_t , and a misallocation term, M_t . Think of M as incorporating the effects from the misallocation of inputs at a more disaggregated level that aggregate up to reduce TFP; we discuss this term in detail below. Relative to the model discussed in the previous section, we are setting $\sigma = 1 - \alpha$ to simplify the exposition.

Letting P_t denote population while L_t is employment, manipulating Equation 6 to solve for income per person, $y_t \equiv Y_t/P_t$, gives

$$y_t = \left(\frac{K_t}{Y_t}\right)^{\frac{\alpha}{1-\alpha}} A_t M_t b_t \ell_t (1 - s_t), \quad 8.$$

where $\ell_t \equiv L_t/P_t$ is the employment-population ratio. Finally, when growth rates are constant, the stock of knowledge is proportional to the number of researchers, raised to a power γ that measures the overall degree of increasing returns; we obtain

$$A_t^* = R_t^\gamma = (s_t L_t)^\gamma. \quad 9.$$

This equation comes from Equation 2 above, where we are assuming the educational attainment of researchers is constant and are setting the factor of proportionality to one for simplicity.

There are two ways to approach the growth accounting: (a) in the long run and (b) historically. In the long run, notice that the capital-output ratio must be constant, so the contribution from K/Y will be zero. For human capital per worker, b , I find it simplest to think about the effects of educational attainment, so that we have $b_t = e^{\psi u_t}$, where u_t is years of educational attainment. If working lives are finite, then the amount of time people spend in school on average is bounded from above (and certainly the fraction of time devoted to human capital is bounded), so educational attainment also leads to level effects but no long-run growth effects. The fraction of people working, ℓ_t , and the fraction of people working as researchers, s_t , are between zero and one, so these variables must be constant in the long run. Finally, for the misallocation term, if the distortions are constant, then M_t will be eventually constant as well. Moreover, in the long run M_t is bounded from above by one: At best there can be zero misallocation, so that resources are optimally allocated and $M = 1$. Reducing misallocation can generate level effects in the long run but not growth effects (given a rate of population growth). In the end, this leaves us only with the growth rate of the labor force, which equals population growth in the long run. Therefore, even in this extended SEG framework, long-run growth is proportional to population growth and everything else has level effects in the long run, as can be seen in Equations 8 and 9.

However—and this is a really important point to appreciate—this does not mean that historical growth is entirely due to population growth. Instead, to the extent that K/Y , educational attainment, labor-force participation, misallocation, and R&D intensity have changed over time, these changes can contribute to growth over any given historical period. The point that we develop now is that, as a historical matter, such level effects account for something like 80 percent of US economic growth.

To see this more formally, take logs and differences of Equations 8 and 9 to get

$$\begin{aligned} \underbrace{d \log y_t}_{\text{GDP per person}} &= \underbrace{\frac{\alpha}{1-\alpha} d \log \frac{K_t}{Y_t}}_{\text{capital-output ratio}} + \underbrace{d \log b_t}_{\text{educational attainment}} + \underbrace{d \log \ell_t}_{\text{emp-pop ratio}} \\ &+ \underbrace{d \log(1-s_t)}_{\text{goods intensity}} + \underbrace{d \log M_t + d \log A_t}_{\text{TFP growth}}, \end{aligned} \quad 10.$$

where

$$\text{TFP growth} \equiv \underbrace{d \log M_t}_{\text{misallocation}} + \underbrace{d \log A_t}_{\text{ideas}} = \underbrace{d \log M_t}_{\text{misallocation}} + \underbrace{\gamma d \log s_t}_{\text{research intensity}} + \underbrace{\gamma d \log L_t}_{\text{LF growth}}. \quad 11.$$

In the equations above, emp-pop ratio stands for employment-population ratio, and LF indicates labor force. Jones (2002) uses equations like these to conduct a growth accounting exercise for the United States for the period 1953–1993, whereas Fernald & Jones (2014) update the calculation through 2007. Because of the global financial crisis and in the interest of saving space, I will not update the accounting exercise to an even later year. However, both of those papers ignored changing misallocation. So I present instead a back-of-the-envelope version of the accounting




Figure 2

Historical growth accounting. The figure shows a growth accounting exercise for the United States since the 1950s using Equations 10 and 11. Details can be found in the main text. Abbreviations: pp, percentage points; TFP, total factor productivity.

here that includes a rough estimate of gains from changing misallocation. For more details on the facts that are discussed in the remainder of this section, readers are referred to Jones (2016).

To begin, consider the pie chart on the left side of **Figure 2**. Growth in GDP per person, y , has averaged something like 2% per year since 1950, and this pie chart uses Equation 10 to decompose this 2% growth into its components. First, the capital-output ratio has been remarkably steady over time, essentially contributing nothing to growth. Second, the $1 - s_t$ term contributes essentially nothing as well: Measures of s_t are so small that we have $1 - s_t \approx 1$ over time. We now turn to the nonzero components of the equation.

3.1. Human Capital and Labor Force Participation

This brings us to educational attainment. A wonderful stylized fact documented by Goldin & Katz (2008) is that educational attainment throughout the twentieth century rose by slightly under one year per decade, that is, from 4 years in 1900 to 13 years by the end of the century. A standard Mincerian return to education is 5% or 7%. This means that rising educational attainment increased GDP per worker by something like 5% each decade or about 0.5% each year. This is a large number: Something like 0.5% of our 2% per year growth in the twentieth century was due to rising educational attainment! If life expectancy has an upper bound, then educational attainment cannot continue to rise forever. And in fact, in the last two decades, we have seen educational attainment for each cohort flatten out: Roughly 85% of kids graduate from high school, roughly 33% graduate from four-year colleges, and these numbers have leveled off since the 1990s. This nicely illustrates both sides of our main point: Historically, rising educational attainment has contributed a large amount to growth, but in the long run, educational attainment per person seems likely to level off and to contribute nothing to growth.

Another similar demographic change is rising labor force participation due to the entry of women. Since 1950, the employment-population ratio has risen from around 55% to around 62%, or by around 0.2% per year.

Subtracting the 0.5% due to rising educational attainment and the 0.2% due to rising labor force participation from the 2% growth leaves us with TFP growth of around 1.3% per year. (Because of slowing growth since around 2003, more recent numbers would be lower; note also that this TFP measure is in labor-augmenting units.)

3.2. Misallocation

The next step in the growth accounting is to further decompose this 1.3% TFP growth. And this is where one of the most important insights of the growth literature in the past 15 years comes into play. The insight—due in large part to Restuccia & Rogerson (2008) and Hsieh & Klenow (2009)—is that misallocation at the micro level aggregates up into TFP differences. This is a great insight for many reasons. For example, it helps us to understand the enormous differences in levels of TFP across countries. It also helps us to understand how TFP levels in countries like Italy and Spain could decline for two decades: Clearly the stock of ideas did not decline, so rising misallocation is the logical explanation. The bottom line is that the growth literature now has two complementary explanations for TFP: misallocation and ideas. And both can be important.

Applied to the topic at hand, the question is, How has changing misallocation contributed to TFP growth in recent decades in the US economy? I do not think we have a firm answer to this question. Bils et al. (2020) find that allocative efficiency in US manufacturing has gotten worse over time, although they suggest that at least two-thirds of this change may be due to measurement error. On the other side, Hsieh et al. (2019a) find that improvements in the allocation of talent associated with declines in discrimination against women and Black Americans have reduced misallocation substantially. They suggest that growth in output per worker between 1960 and 2010 was higher by 0.3 percentage points per year because of this declining misallocation. And of course there may be other changes in misallocation that we have not taken into account. Motivated by this evidence we pencil in a 0.3 percentage point contribution from declining misallocation, while recognizing that there is a lot of uncertainty about this number.

3.3. Ideas

Subtracting the 0.3% contribution from declining misallocation from the 1.3% TFP growth leaves us with 1% growth due to increases in ideas. Now recall Equation 9, establishing that $A_t^* = R_t^\gamma = (s_t L_t)^\gamma$. There are several ways to proceed. If we have a measure of R_t that we believe in, we can use the growth in A_t^* and the growth in R_t to recover an estimate of γ . Jones (2002) does this type of exercise and finds that R_t is growing faster than population (i.e., that s_t is growing substantially) and is consistent with estimates of $\gamma = 1/3$, though other values are certainly possible. There is basically an identification problem here: We have one moment (the contribution of ideas to growth) and two things to estimate (the degree of increasing returns, γ , and the growth rate of research intensity, s_t). In general, one suspects that research effort is undercounted. For example, the effort that goes into starting a new firm should arguably be counted as research, but most of the time it is not. Most measured R&D is in manufacturing, which accounts for less than 10% of US employment; surely research is important to the rest of the economy as well. There is also the question of how much of research growth that comes from Germany, Japan, China, and so on should enter into the calculation. We will not resolve those questions here, which is another reason this is a back-of-the-envelope kind of calculation.

Here, we simply impose $\gamma = 1/3$ and assume that labor force growth is 1% per year. We then recover the growth rate of research intensity, s , as 2% per year. Put differently, if $\gamma = 1/3$ and employment grows at 1% per year, then the growth contribution of $\uparrow L$ is 0.33% per year, leaving a growth contribution from $\uparrow s$ of 0.67% per year, so that the two add to 1%.

3.4. Remarks

We learn a very important lesson from the growth accounting in **Figure 2**. Even in this SEG framework in which population growth is the only potential source of growth in the long run, other

factors explain more than 80% of US growth in recent decades: The contribution of population growth is 0.3% out of the 2% growth we observe. In other words, the level effects associated with rising educational attainment, declining misallocation, and rising research intensity have been overwhelmingly important for the past 50+ years.

4. CONNECTING TO THE BROADER LITERATURE

This section connects the SEG framework to the literature more broadly and makes three main points. First, we note that many existing papers that are not commonly thought of as adopting a SEG framework actually are. Second, we discuss important extensions to heterogeneity in idea production functions in different industries. Third, we highlight research done in the fully endogenous growth setup that could productively be reexamined using SEG models.

4.1. Semi-Endogenous Growth Is More Common than You Might Think

The first point to make is that the SEG setup features broadly in the literature, and perhaps even more broadly than one notices on first inspection. The reason is that essentially any model that possesses increasing returns to scale is likely a SEG model. Romer (1990) provides the most solid microfoundations for justifying increasing returns, but any model that assumes or estimates increasing returns will likely imply that the growth rate is an increasing function of the population growth rate. The classic learning-by-doing models of Arrow (1962) and Frankel (1962) contain this result, for example. Other classic references include works by Phelps (1966), Nordhaus (1969), Judd (1985), Jones (1995), Kortum (1997), and Segerstrom (1998); Jones (2005) provides a detailed discussion of the historical background.⁵

Other models not traditionally thought of as SEG models also share this feature. International trade models are a great example. Krugman (1979), Grossman & Helpman (1989), and Melitz (2003) all feature increasing returns so that income per person is an increasing function of the labor force. Adding population growth to those models leads to the key result that the growth rate of the economy is the product of the degree of increasing returns and the rate of population growth. It is tempting to think that the classic model by Hopenhayn (1992) of industry equilibrium with entry and exit will also fall into this category. However, it does not: Hopenhayn carefully structures his model to feature constant returns to scale so that a competitive equilibrium exists and is efficient. He does this by using Lucas's (1978) span-of-control setup so that within each firm there are decreasing returns to scale; the fixed cost of entry then leads firms to the bottom of their U-shaped average cost curve and upgrades the firm-level decreasing returns to economy-level constant returns. Still, the basic point here is that the expansive literature building on Melitz (2003) typically fits in the class of SEG models.

Another key setup in international trade is Eaton & Kortum's (2002) model. The basic model does not focus on entry and instead takes the T_i scale parameters of the Fréchet distributions as given. However, the natural interpretation of these parameters is that they scale with population, as shown by Kortum (1997) (see especially Eaton & Kortum 2002, section 5.2). Under this interpretation, this structure also is one of SEG. Ramondo et al. (2016) and Arkolakis et al. (2018) explore some of these implications more fully.

Closed-economy versions of these models of firm dynamics have been used recently to study demographics, aging, and dynamism. Models based on Melitz (2003) fit broadly into the SEG

⁵Other papers focused on SEG include those by Eicher & Turnovsky (1999), Li (2000, 2003), and Cozzi (2020). Further works are discussed below.

framework. Models based on Hopenhayn (1992) feature constant returns as described above. But it is easy to go back and forth between these frameworks, and there may be gains to writing the models in both ways. Examples from the Hopenhayn tradition include the works of Karahan et al. (2019) and Hopenhayn et al. (2018). Engbom (2019) is more in the Aghion & Howitt (1992) tradition, whereas Luttmer (2011) and Sterk et al. (2021) follow the Melitz tradition.

More on point, Peters & Walsh (2021) use a rich SEG setup to explore the implications of slowing population growth for firm dynamics and growth. They find that declining population growth generates lower entry, reduced creative destruction, increased concentration, rising markups, and lower productivity growth, all facts that we see in the firm-level data.

Finally, models of economic geography also emphasize increasing returns and therefore could have SEG implications if considered over time (see Redding & Rossi-Hansberg 2017 for a recent survey).

4.2. Heterogeneity and Multisector Models

An underexplored area in the growth literature is sectoral heterogeneity. Different industries have very different growth rates: Why? The literature on structural transformation often studies this from the perspective of agriculture, manufacturing, and services. Examples include works by Kongsamut et al. (2001), Ngai & Pissarides (2007), Acemoglu & Guerrieri (2008), and Herrendorf et al. (2014). This branch of the literature is often concerned with the important question of how heterogeneity at the sectoral level can be reconciled with stable growth at the aggregate level.

A much smaller literature considers an equally interesting question: Why do different industries have different TFP growth rates? Ngai & Samaniego (2011) study this question in a SEG framework in which different industries have different parameters in their idea production functions. In the context of the notation used earlier, different industries may have different β parameters: It may be relatively easy to find new ideas in the semiconductor industry and in agriculture but much harder to find them in education and construction. This question is ripe for reexamination with more recent data, especially given the interest in heterogeneity more broadly. Research by Bloom et al. (2020) and Sampson (2020) can be read as pushing in this direction, but there is clearly scope for much more work here.

4.3. Novel Places Where the Semi-Endogenous Growth Framework Could Be Applied

There are many interesting papers in the recent growth literature that are not written in a SEG framework. It would be interesting to reformulate these models in the SEG setting to see how the results change. I should emphasize that this is not meant in any way to be a criticism of the papers that follow: All research involves trade-offs, and sometimes fully endogenous growth models are more convenient and easier to solve, for example. This is only to highlight that such papers represent a rich and exciting area for further research.

To start, let me give an example from my own research. Jones & Kim (2018) study a Schumpeterian model of top income inequality. A mechanism that helps the paper generate a rich set of results is a Lucas (1988)-style equation for the growth of entrepreneurial incomes: There is a linear differential equation built in. This is useful because models that generate Pareto distributions often do so by specifying a process with exponential growth and a constant death rate. Pareto inequality is then roughly equal to the ratio of the growth rate to the death rate, and having this growth rate be endogenously determined gives a rich theory of inequality. But why should that differential equation be linear?

A large class of recent growth papers consider an idea-driven growth model in a fully endogenous growth framework. For example, Klette & Kortum (2004) build a model of firm dynamics designed to match a bunch of facts about innovation and growth at the firm level. Acemoglu et al. (2012) study climate change and economic growth when research can be focused on clean versus dirty technologies. Acemoglu & Restrepo (2018) consider the effects of automation. Akcigit & Kerr (2018) highlight differences in innovation between large and small firms, both in quantity and in the nature of the innovations.

These models assume that knowledge spillovers are large and homogeneous across firms or sectors. How different would the results be if the idea production functions incorporated an empirically estimated degree of knowledge spillovers? And what if this parameter is allowed to vary across industries? Each of these papers could be extended to the SEG setting to study these questions.

A related literature is carefully designed so that firm-level data on employment can be used to inform us about different components of growth. Works by Aghion et al. (2019a), Garcia-Macia et al. (2019), Hsieh et al. (2019b) are in this category. They have a clever structure whereby different components of growth—incumbent variety growth, incumbent creative destruction, improvements in incumbent process efficiency, and entry, for example—are exogenous parameters that are recovered from data on employment. These papers remain agnostic about where these growth parameters come from, but giving the papers SEG microfoundations would be valuable.

Another class of models in which SEG could be usefully applied is the recent literature on technology diffusion. Models such as those by Alvarez et al. (2013), Lucas & Moll (2014), and Perla & Tonetti (2014) use very special assumptions so that the diffusion of ideas across agents (e.g., people or firms) can itself generate sustained growth. For example, there is a continuum of agents each of whom has an idea drawn from an unbounded distribution. Therefore, there is not a single best idea in the world at any given point in time—there is always someone else with a better idea. Sampson (2016) and Perla et al. (2021) apply such a setting to international trade and find that the dynamic gains from trade may be 3 to 10 times larger than the standard static gains. Benhabib et al. (2021) relax the special assumptions mentioned above and study technology diffusion in an endogenous growth model in which ideas get invented. Buera & Oberfield (2020) instead use a SEG framework to study technology diffusion and international trade. Buera & Lucas (2018) provide a recent overview of this literature.

5. WHY GROWTH MIGHT BE SLOWER IN THE FUTURE

The next two sections are organized around two topics. First, we consider a set of reasons related to the question of why growth might be slower in the future. Then, in the next section, we consider the opposite: why future growth might not be slower and might even be faster.

On the possibility of slowing growth, I highlight three points. The first is a natural implication of the accounting exercise that we explored in Section 3. The second comes from studying the time path of research effort in the advanced countries of the world: In short, the growth rate of research effort appears to be slowing. Finally, we highlight that population growth has been declining around the world and explore the implications of this slowdown for future economic growth.

5.1. Growth Accounting and a Future Slowdown

In Section 3, we used growth accounting to explain the sources of US economic growth over the past several decades. The stylized picture that emerged is that although the only source of long-run growth in living standards is population growth, historically more than 80% of US

economic growth is due to other factors. These factors include rising educational attainment (perhaps 25% of growth), reductions in misallocation (likely more than 15% of growth), and increases in the fraction of the population devoted to research (perhaps 45% of growth).

The point here is that this framework strongly implies that, unless something dramatic changes, future growth rates will be substantially lower, at least at some point. The sources other than population growth are inherently transitory, and once these sources have run their course, all that will remain is the 0.3 percentage point contribution from population growth. In other words, a natural implication of this framework is that long-run growth will be γn . With our estimate of $\gamma = 1/3$ and $n = 1\%$, the implication is that long-run growth in living standards will be 0.3% per year rather than 2% per year—an enormous slowdown!

The reason for this is worth emphasizing again. Educational attainment rose from around 7 years per person for the cohorts born in the 1880s to almost 14 years per person for the cohorts born in the 1970s. This is a much larger increase than the rise in life expectancy, so people have been spending a larger fraction of their lifetimes in school; and given declining retirement ages, this time represents an even larger and growing fraction of their working lives. In the long run, however, this fraction must level out: We cannot spend more than 100% of our time in school. Moreover, this leveling out already seems to be occurring: The cohorts born in more recent decades seem to be getting fewer than 14 years of education.⁶ The same argument applies to the fraction of the labor force devoted to research. Historically, it has risen quite substantially, but in the long run, this fraction must level out. Finally, misallocation can decline only so much: Once we have an allocation of resources that is optimal, there is no more growth to be had from reducing misallocation.

None of this is to say that there are no ways around these conclusions, at least for a while. First, these trends could certainly continue into the future, say for the next 50 or even 100 years (more on this shortly). Second, there is substantial mismeasurement in human capital and research. For example, educational attainment is only one form of human capital. Maybe we have been shifting from on-the-job accumulation of human capital toward education, so that the fraction of time spent accumulating human capital has not risen since 1900. And maybe, despite many different measurements, research intensity measured correctly is not rising. Finally, perhaps misallocation has not declined, or it has a long way to go. The main point is that a natural way of looking at the data on growth in recent decades through the lens of a SEG model suggests that a substantial slowdown in growth may lie in our future.

5.2. Slowing Growth in Research Effort?

Our next point suggests that another dimension of this slowdown is already occurring. Borrowing from Lewis Carroll, consider the Red Queen interpretation of SEG: We have to keep running at a constant speed in order to maintain steady growth in knowledge.⁷ That is, our research effort must continue to grow at its historic rate. What we see in this section is that the growth rate of research effort seems to be slowing down, at least in advanced countries.

Figure 3 plots measures of full-time equivalent research employment from the Main Science and Technology indicators of the Organisation for Economic Co-operation and Development

⁶Readers are referred to Goldin & Katz (2008). Jones (2016) updates the data on this and the other facts discussed in this section. Grossman et al. (2017) study a model in which the fraction of time in human capital investment optimally rises forever, asymptoting at 100%.

⁷The reference is to the Red Queen's race in Lewis Carroll's *Through the Looking Glass*: "Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!" (Carroll 2019, 24–25).




Figure 3

Research employment in the United States, OECD countries, and the world. “World” is the OECD countries plus China and Russia. Average annual growth rates of research employment are reported for each region for the first and second parts of the time frame. Abbreviation: OECD, Organisation for Economic Co-operation and Development. Data from OECD (2021).

(OECD) for three different aggregates: the United States, the OECD countries, and the world (where “world” means the OECD countries plus China and Russia). The reason for plotting the different aggregates is related to the diffusion of ideas. In particular, because of technology diffusion, growth in the United States or in any other country benefits from ideas created throughout the world. A better study would carefully model technology diffusion to trace through how research effort in the rest of the world affects US economic growth. But this is a difficult issue to get right, so we will take the shortcut of looking at research effort at different levels of aggregation.⁸

The key point of **Figure 3** is that the growth rate of research effort appears to be slowing. For the United States, research employment grew at 3.4% per year between 1981 and 2003 but slowed to 2.1% per year afterward. For the OECD, there is a similar slowdown: Research employment grew at 4.1% per year before 2003 but only 2.8% per year afterward. Adding in the large increase in research employment in China is enough to mitigate most of this slowdown, as shown in the “world” number. However, it is unclear how to properly weight the Chinese research effort (i.e., how much of it is catching up historically versus pushing the frontier forward), so for now we will focus on the US and OECD numbers.

The slowdown in US research employment growth is 38%, whereas for the OECD the slowdown is slightly smaller, at 32%. Recall that in a SEG model, if these research employment growth numbers are treated as permanent, then they imply an equal slowdown in long-run growth. In other words, the slowdown in the growth rate of research effort implies an equally large slowdown in long-run growth.

5.3. Slowing Population Growth

A final reason that economic growth rates might be slowing in the future is that population growth—the long-run driving force behind SEG—is itself slowing throughout the world.

⁸Eaton & Kortum (1996, 1999) trace the international diffusion of ideas using patent renewal data. Redoing and enhancing their research would be an excellent topic for future work.




Figure 4

Population growth around the world. The graph shows average annual rates of population growth for countries classified according to their 2018 World Bank income grouping. Each data point corresponds to a 5-year period. Data from United Nations (2019).

Figure 4 shows that this fact has been true in high- and middle-income countries since at least 1965 and characterizes low-income countries since 1990. Slowing rates of population growth are another reason future growth may be even lower than in recent decades.

However, the demographic data are actually even more pessimistic than this suggests. In their recent book entitled *Empty Planet*, Bricker & Ibbitson (2019) use a rich body of demographic research to suggest that global population growth in the future may not only fall to zero but also be negative. For example, the natural rate of population growth—i.e., births minus deaths, ignoring immigration—is already negative in Japan and in many European countries such as Germany, Italy, and Spain (United Nations 2019).

Figure 5 shows the total fertility rate—the average number of live births per woman—for a selection of countries in the 2015–2020 period. Population growth is positive when each woman has more than two (surviving) children during her lifetime. What we see in the figure is that many countries—in fact, the entirety of the high-income countries taken as a group—already feature fertility rates below this level. That is, fertility rates throughout much of the world are consistent with long-run declines in population rather than with a stable population.

What do idea-based growth models predict will happen if population growth is negative instead of positive? Jones (2020) provides a detailed analysis of the possibilities, but the basic point is easy to demonstrate. Consider the key idea production function used above,

$$\frac{\dot{A}_t}{A_t} = R_t^\lambda A_t^{-\beta}, \quad 12.$$

and assume that $R_t = R_0 e^{-\eta t}$, where $\eta > 0$ indicates the rate at which the number of researchers is declining rather than growing.

Substituting $R_t = R_0 e^{-\eta t}$ into the idea production function gives

$$\frac{\dot{A}_t}{A_t} = R_0^\lambda A_t^{-\beta} e^{-\lambda \eta t}, \quad 13.$$




Figure 5

Total fertility rate around the world. The total fertility rate is the average number of live births a hypothetical cohort of women would have over their reproductive lives if they were subject during their whole lives to the fertility rates of a given period and if they were not subject to mortality. Each data point corresponds to the 5-year period 2015–2020. The dashed line at 2.1 shows the fertility rate consistent with zero population growth. Abbreviation: HIC, high-income countries. Data from United Nations (2019).

and it turns out that this equation is easy to solve. Integrating yields

$$A_t = \begin{cases} A_0 \left(1 + \frac{\beta g_{A0}}{\lambda \eta} (1 - e^{-\lambda \eta t})\right)^{1/\beta} & \text{if } \beta > 0 \\ A_0 \exp\left(\frac{g_{A0}}{\lambda \eta} (1 - e^{-\lambda \eta t})\right) & \text{if } \beta = 0 \end{cases}. \quad 14.$$

Taking the limit as $t \rightarrow \infty$, we find the interesting result that when population growth is negative, the stock of knowledge is bounded and converges to the finite level A^* , where we have

$$A^* = \begin{cases} A_0 \left(1 + \frac{\beta g_{A0}}{\lambda \eta}\right)^{1/\beta} & \text{if } \beta > 0 \\ A_0 \exp\left(\frac{g_{A0}}{\lambda \eta}\right) & \text{if } \beta = 0 \end{cases}. \quad 15.$$

That is, when population growth is negative, both the SEG model ($\beta > 0$) and the fully endogenous growth model ($\beta = 0$) imply that the stock of knowledge—and therefore income per person and overall living standards—converges to some finite value.

Jones (2020) calls this situation the empty planet scenario: If population growth is negative, these idea-driven models predict that living standards stagnate for a population that vanishes! This is a stunningly negative result, especially when compared to the standard result we have been examining throughout the article. In the usual case with positive population growth, living standards rise exponentially forever for a population that itself rises exponentially. Whether we live in an expanding cosmos or an empty planet depends, remarkably, on whether the total fertility rate is above or below a number like 2 or 2.1.

Historically, fertility rates have been quite high, but they have declined over time. Women used to have 4 or 5 kids on average, then 3, then 2, and now even fewer in many places. From an individual family's standpoint, there is nothing special about whether we settle at 2.2 kids per woman or 1.9. But taking into account the macroeconomics of the problem means that these options ultimately lead to wildly different outcomes.

5.4. Summing Up

In conclusion, there are several reasons to worry that future growth will be slower in the long run. Many of the sources of growth that have been operating historically—including rising educational attainment, rising research intensity, and declining misallocation—are inherently limited and cannot go on forever. The key source of sustained growth in the semi-endogenous setting is population growth. But that has been slowing historically around the world, and current fertility patterns are more consistent with a declining population than with positive population growth. In the extreme, this could even lead to the stagnation of living standards for a vanishing population.

6. WHY GROWTH MIGHT NOT BE SLOWER AND COULD BE FASTER

Although the preceding section laid out many reasons to be pessimistic about the future of economic growth, there are two broad reasons for optimism. The first is broadly related to the allocation of talent and the so-called missing Einsteins problem. The second is the possibility for automation and artificial intelligence: What if people can be gradually replaced or augmented in producing new ideas?

6.1. Finding Einsteins

The world contains more than 7 billion people. However, according to the OECD's main science and technology indicators, the number of full-time equivalent researchers in the world appears to be lower than 10 million (OECD 2021). In other words, something in the order of one or two out of every thousand people in the world are engaged in research. Even allowing for massive mismeasurement of R&D, the point is that we are a long way from hitting any constraint that we have run out of people to hunt for new ideas. There is ample scope for substantially increasing the number of researchers over the next century, even if population growth slows or is negative. I see three ways this “finding new Einsteins” can occur.

6.1.1. The rise of China, India, and other countries. The United States, Western Europe, and Japan together count about 1 billion people, about one-seventh of the world's population. China and India each have this many people. As economic development proceeds in China, India, and throughout the world, the pool from which we may find new talented inventors will multiply. How many Thomas Edisons and Jennifer Doudnas have we missed out on among these billions of people because they lacked education and opportunity?

One can easily imagine the global population of researchers increasing by a factor of 3 or even 7 as the world develops. In the SEG setting, this would have a long-run level effect of $3^\gamma \approx 1.44$ or $7^\gamma \approx 1.91$, taking a benchmark value of $\gamma = 1/3$. If half that effect occurs over a century, this could easily amount to an extra 0.2–0.4 percentage points of growth each year.

6.1.2. Finding new Doudnas: women in research. Another huge pool of underutilized talent is women. Brouillette (2021) uses patent data to document that in 1976 less than 3% of US inventors were women. In 2016, their share was still less than 12%. He estimates that eliminating the barriers that lead to this misallocation of talent could raise economic growth in the United States by up to 0.3 percentage points per year over the next century.

6.1.3. Other sources of within-country talent. Bell et al. (2019) document that the extent to which people are exposed to inventive careers in childhood has a large influence on who becomes an inventor. They show that exposure in childhood is limited for girls, people of certain races, and

people in low-income neighborhoods, even conditional on math test scores in grade school, and refer to these missed opportunities as “lost Einsteins.” So the opportunities to expand the talent for research are not limited only to developing countries.

6.2. Automation and Artificial Intelligence

Another potential reason for optimism about future growth prospects is the possibility of automation, both in the production of goods and in the production of ideas.

6.2.1. Zeira’s (1998) model of automation and growth. Zeira (1998) provides an elegant model of automation and economic growth. In its simplest form, suppose the production function is

$$Y = AX_1^{\alpha_1} X_2^{\alpha_2} \dots X_n^{\alpha_n}, \quad \text{where } \sum_{i=1}^n \alpha_i = 1. \quad 16.$$

Zeira thought of the X_i ’s as intermediate goods, but we follow Acemoglu & Autor (2011) and refer to them as tasks; both interpretations have merit. Before it is automated, a task can be produced one-for-one by labor. After automation, one unit of capital can be used instead. We have

$$X_i = \begin{cases} L_i & \text{if not automated} \\ K_i & \text{if automated} \end{cases}. \quad 17.$$

If K is sufficiently large and if K and L are assigned to these tasks optimally, the production function can be expressed (up to an unimportant constant) as

$$Y_t = A_t K_t^\alpha L_t^{1-\alpha}, \quad 18.$$

where the exponent α reflects the overall share and importance of tasks that have been automated. For the moment, we treat α as a constant and explore comparative statics when more tasks are automated.

To close the model, embed this setup in a standard neoclassical growth model with a constant investment rate. The share of factor payments going to capital is given by α , and the long-run growth rate of $y \equiv Y/L$ is

$$g_y = \frac{g}{1-\alpha}, \quad 19.$$

where g is the (for now exogenous) growth rate of A . An increase in automation will therefore increase the capital share α and, because of the multiplier effect associated with capital accumulation, will increase the long-run growth rate.

Zeira emphasizes that automation has been going on since at least the Industrial Revolution. And certainly the twentieth century was the century of automation: Assembly lines, cars, trucks, airplanes, forklifts, computers, machine tools, and even robots are key examples of the extensive automation that has occurred.

But the implication of this simple version of the Zeira model is that the capital share of factor payments and the growth rate itself should have risen along with automation. This prediction is strongly rejected by the data. Instead, Kaldor’s (1961) stylized fact that growth rates and capital shares are relatively stable over time is a good characterization of the US economy for the bulk of the twentieth century, certainly through 1980 and perhaps up to 2000 (see, e.g., Jones 2016). Zeira’s framework, then, needs to be improved so that it is consistent with historical evidence.

Acemoglu & Restrepo (2018) provide one approach to solving this problem. Their rich environment allows constant elasticity of substitution (CES) production and endogenizes the number of tasks as well as automation. In particular, they suppose that research can take two different directions: discovering how to automate an existing task or discovering new tasks that can be used in production. In their setting, α reflects the fraction of tasks that have been automated. This leads them to emphasize one possible resolution to the empirical shortcoming of Zeira: Perhaps we are inventing new tasks just as quickly as we are automating old tasks. The fraction of tasks that are automated could be constant, leading to a stable capital share and a stable growth rate.⁹

Aghion et al. (2019b) provide an alternative explanation. Suppose tasks are complementary in production, with an elasticity of substitution less than one. Then automation and capital accumulation push in opposite directions. As above, automation by itself tends to increase the capital share. However, because the elasticity of substitution is less than one, the input that becomes more scarce—labor here, since capital gets accumulated—sees its factor share rise. This is essentially a form of Baumol’s (1967) cost disease. The increase in the fraction of the economy that is automated over time is just offset by a decline in the share of GDP associated with the automated sectors, such as manufacturing or agriculture. Economic growth is determined not by what we are good at but rather by what is essential and yet hard to improve. Labor gets concentrated on fewer and fewer tasks, but those tasks are essential, and therefore the labor share can remain high.

6.2.2. Automating idea production. Even more intriguing possibilities arise from considering the automation of tasks in idea production (see Aghion et al. 2019b, Agrawal et al. 2019). Microscopes, computers, DNA sequencing machines, and the Internet itself are examples of automation that enhance the production of ideas. Artificial intelligence would be an even more extreme example.

Let’s apply the Zeira structure to the idea production function. We have

$$\dot{A}_t = A_t^{1-\beta} X_1^{\alpha_1} X_2^{\alpha_2} \cdots X_n^{\alpha_n} \quad 20.$$

$$= A_t^{1-\beta} K_t^\alpha R_t^{1-\alpha}. \quad 21.$$

To see the possibilities most easily, assume a simple production function for goods with no automation; allowing automation there as well only enhances the results that follow,

$$Y_t = A_t L_{yt}. \quad 22.$$

Then, close the rest of the model with a standard capital accumulation equation and positive population growth, $R_t + L_{yt} = L_0 e^{mt}$.

In the long run, the capital-output ratio is constant, and so the idea production function becomes

$$\dot{A}_t = \kappa A_t^{1-(\beta-\alpha)} L_t, \quad 23.$$

where κ is a constant that depends on the capital-output ratio and the share of labor devoted to research. Along the balanced growth path, the growth rate of A is then equal to

$$g_A = \frac{n}{\beta - \alpha}. \quad 24.$$

⁹Other important contributions to this recent literature include those by Peretto & Seater (2013), Hemous & Olsen (2016), and Acemoglu & Restrepo (2020).

Two important conclusions follow from this setup. First, an increase in the automation of tasks in idea production ($\uparrow\alpha$) causes the growth rate of the economy to increase. Second, if the fraction of tasks that are automated (α) rises to reach the rate at which ideas are getting harder to find (β), we get a singularity! In particular, once $\alpha \geq \beta$, the model exhibits sufficiently strong increasing returns that there is no balanced growth path. Instead, the growth rate rises rapidly over time and the economy reaches infinite knowledge and income in finite time, assuming that is possible. Alternatively, and we return to this below, if the total number of possible ideas that can be discovered is finite, then the economy reaches the maximum possible knowledge stock \bar{A}_{max} in finite time. In particular, notice that if $\beta = 2/3$, for example, then once two-thirds of the research tasks have been automated, growth explodes. The model with ideas does not require automation to be complete in order for explosive growth to occur. Moreover, allowing automation to occur in the goods production function as well only reinforces this possibility.

Yet this model, too, is not without problems. In particular, as we noted at the start of this section, the automation of the idea production function has been occurring for the past hundred years or more—consider the massive improvement in the tools for conducting research, including computers, lasers, laboratories, the internet, etc. Yet growth rates in the United States have not increased. Of course we do not know the counterfactual; maybe absent this automation, growth rates would have slowed already. But at the very least, this suggests that this story of automation is incomplete.¹⁰

6.2.3. Artificial intelligence. An extreme version of this model would involve artificial general intelligence (AGI). Consider what would happen if machines could replace humans in all tasks. On the one hand, this scenario seems quite far-fetched. On the other hand, according to Davidson (2021), many experts in the field of artificial intelligence believe there is a nontrivial chance of this occurring in the next century.

If we replace all labor with capital in the task model—in both goods production and idea production—then there are increasing returns to factors that can be accumulated because of the nonrivalry of ideas. Growth explodes, and simple math reveals a singularity in which knowledge and incomes go to infinity in finite time.

However, this possibility assumes it is possible for A_t to go to infinity: that there are ideas out there with arbitrarily high productivity. If instead one were to think that there exists a best idea with productivity \bar{A}_{max} , then this maximum productivity would be reached in finite time. In that case, the (automated) goods production function would become $Y_t = \bar{A}_{max}K_t$ and the model would behave like the classic AK model of growth theory. Described in words, this result sounds like the plot of a science fiction novel: Machines invent all possible ideas, leading to a maximum productivity, and then they fill the universe rearranging matter and energy to exponentially expand the size of the artificial intelligence.

7. KEY UNRESOLVED QUESTIONS FOR FUTURE RESEARCH

We conclude this essay by summarizing some of the most important outstanding questions related to economic growth.

7.1. How Large Is the Degree of Increasing Returns to Scale?

One of the fundamental contributions of Romer (1990) was the insight that the nonrivalry of ideas means that production is characterized by increasing returns to scale. The magnitude of the

¹⁰Aghion et al. (2019b) follow the Baumol approach discussed above and assume research tasks involve an elasticity of substitution of less than one to help address this problem.

degree of increasing returns is fundamental and plays a key role in shaping the answer to many practical questions. Knowing its value is essential for calibrating models and answering questions such as, What is the value of the sustainable long run growth rate, and what is the optimal top income tax rate? There is remarkably little work aimed at measuring the degree of increasing returns associated with the nonrivalry of ideas, despite the importance of this parameter.¹¹

7.2. What Is the Social Rate of Return to R&D?

To what extent do we underinvest in research activities to create new ideas? What are the best policies for closing this gap? What is the role for basic versus applied research, or publicly funded versus privately funded research? These questions have been the subject of a huge literature in economics, dating back to Griliches (1957). Most existing estimates are large, but the quality of the identification in many of these papers is suspect. One of the best-identified papers is by Bloom et al. (2013), who use variation in R&D tax credits across time and across US states and estimate very high values for the social return to R&D, in the order of 50%. Readers are referred to Hall et al. (2010), Azoulay et al. (2018), and Jones & Summers (2020) for other recent contributions.

7.3. How Can We Best Measure Ideas?

A long tradition in economics uses patents, patent citations, and the stock market reaction to patents to measure the production of new ideas (Griliches 1984, Hall et al. 2005, Budish et al. 2015, Kogan et al. 2017). Yet patents have problems as well. More than 70% of US corporate patents are in manufacturing, a sector employing less than 10% of the labor force (Autor et al. 2020). Most firms do not patent, so patents capture only one part of idea production. Argente et al. (2020) use machine learning algorithms to link patents to a database of consumer goods and argue that at least half of product innovation in that sector comes from firms that never patent. Moreover, because of changing laws and policies over time, a patent granted in 1980 may be very different from a patent granted in 2020, making comparisons over time difficult (Kortum & Lerner 1998).

The growth literature often uses TFP growth to proxy for innovation, but this too has problems. For example, as is now widely appreciated, misallocation is another determinant of TFP, so changes in misallocation will show up as TFP growth, having nothing to do with innovation. The recent changes to the System of National Accounts to treat software, entertainment products, and R&D as intellectual property products is a welcome improvement in measurement. However, the current practice of the US Bur. Labor Stat. (2021) is to treat these investments as aggregating up to an intangible capital stock that is placed inside the constant returns to scale production function, just like machine tools and buildings. That is, it treats these nonrival intellectual property products as rival. But this means that an important part of the contribution of intellectual property products to growth is incorrectly subtracted out when computing TFP growth. More research is needed on this point, but the standard multifactor productivity growth series from the US Bureau of Labor Statistics surely misses an important part of innovation.

7.4. Better Growth Accounting

Back in Section 3, we presented a stylized accounting of US economic growth since the 1950s. A much richer version of this growth accounting would be valuable. How much of US growth is

¹¹Arkolakis et al. (2020) use European immigration to the United States between 1880 and 1920 and estimate a range of values of 0.7 to 1.3. Peters (2019) uses the pseudorandom settlement of East Germans into West Germany after World War II and finds a value of 0.89. These estimates are larger than the 1/3 estimate assumed in the growth accounting exercise above, perhaps because cross-sectional estimates include reallocations from one region to another that may wash out in the aggregate.

due to ideas discovered in other countries, and vice versa? What is the contribution from changing misallocation over time?

8. CONCLUSIONS

I hope this article convinced the reader that there are many exciting opportunities for future research on economic growth. This force has lifted billions of people out of poverty. It holds the promise of equally great advances in the future.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

I am grateful to Jean-Felix Brouillette, Pete Klenow, Michael Peters, Charlotte Siegmann, Chris Tonetti, and seminar participants at the Atkinson Lecture at Oxford University for many helpful discussions and suggestions.

LITERATURE CITED

- Acemoglu D, Aghion P, Bursztyn L, Hemous D. 2012. The environment and directed technical change. *Am. Econ. Rev.* 102(1):131–66
- Acemoglu D, Autor D. 2011. Skills, tasks and technologies: implications for employment and earnings. In *Handbook of Labor Economics*, Vol. 4, ed. O Ashenfelter, D Card, pp. 1043–171. Amsterdam: Elsevier
- Acemoglu D, Guerrieri V. 2008. Capital deepening and nonbalanced economic growth. *J. Political Econ.* 116(3):467–98
- Acemoglu D, Restrepo P. 2018. The race between man and machine: implications of technology for growth, factor shares, and employment. *Am. Econ. Rev.* 108(6):1488–542
- Acemoglu D, Restrepo P. 2020. Robots and jobs: evidence from US labor markets. *J. Political Econ.* 128(6):2188–244
- Aghion P, Bergeaud A, Boppart T, Klenow PJ, Li H. 2019a. Missing growth from creative destruction. *Am. Econ. Rev.* 109(8):2795–822
- Aghion P, Howitt P. 1992. A model of growth through creative destruction. *Econometrica* 60(2):323–51
- Aghion P, Jones BF, Jones CI. 2019b. Artificial intelligence and economic growth. In *The Economics of Artificial Intelligence: An Agenda*, ed. A Agrawal, J Gans, A Goldfarb, pp. 237–82. Chicago: Univ. Chicago Press
- Agrawal A, McHale J, Oettl A. 2019. Finding needles in haystacks: artificial intelligence and recombinant growth. In *The Economics of Artificial Intelligence: An Agenda*, ed. A Agrawal, J Gans, A Goldfarb, pp. 149–74. Chicago: Univ. Chicago Press
- Akcigit U, Kerr WR. 2018. Growth through heterogeneous innovations. *J. Political Econ.* 126(4):1374–443
- Alvarez FE, Buera FJ, Lucas RE. 2013. *Idea flows, economic growth, and trade*. NBER Work. Pap. 19667
- Argente D, Hanley D, Baslandze S, Moreira S. 2020. *Patents to products: product innovation and firm dynamics*. Work. Pap. 2020-4, Fed. Reserve Bank Atlanta, Atlanta, GA
- Arkolakis C, Lee SK, Peters M. 2020. *European immigrants and the United States' rise to the technological frontier*. Work. Pap., Yale Univ., New Haven, CT
- Arkolakis C, Ramondo N, Rodriguez-Clare A, Yeaple S. 2018. Innovation and production in the global economy. *Am. Econ. Rev.* 108(8):2128–73
- Arrow KJ. 1962. The economic implications of learning by doing. *Rev. Econ. Stud.* 29:153–73
- Atkeson A, Burstein A. 2019. Aggregate implications of innovation policy. *J. Political Econ.* 127(6):2625–83
- Atkeson A, Burstein A, Chatzikonstantinou M. 2019. Transitional dynamics in aggregate models of innovative investment. *Annu. Rev. Econ.* 11:273–301

- Autor D, Dorn D, Hanson GH, Pisano G, Shu P. 2020. Foreign competition and domestic innovation: evidence from US patents. *Am. Econ. Rev. Insights* 2(3):357–74
- Azoulay P, Graff Zivin JS, Li D, Sampat BN. 2018. Public R&D investments and private-sector patenting: evidence from NIH funding rules. *Rev. Econ. Stud.* 86(1):117–52
- Baumol WJ. 1967. Macroeconomics of unbalanced growth: the anatomy of urban crisis. *Am. Econ. Rev.* 57:415–26
- Bell AM, Chetty R, Jaravel X, Petkova N, Van Reenen J. 2019. Who becomes an inventor in America? The importance of exposure to innovation. *Q. J. Econ.* 134(2):647–713
- Benhabib J, Perla J, Tonetti C. 2021. Reconciling models of diffusion and innovation: a theory of the productivity distribution and technology frontier. *Econometrica* 89(5):2261–301
- Bils M, Klenow PJ. 2000. Does schooling cause growth? *Am. Econ. Rev.* 90:1160–83
- Bils M, Klenow PJ, Ruane C. 2020. *Misallocation or mismeasurement?* NBER Work. Pap. 26711
- Bloom N, Jones CI, Van Reenen J, Webb M. 2020. Are ideas getting harder to find? *Am. Econ. Rev.* 110(4):1104–44
- Bloom N, Schankerman M, Van Reenen J. 2013. Identifying technology spillovers and product market rivalry. *Econometrica* 81(4):1347–93
- Bricker D, Ibbitson J. 2019. *Empty Planet: The Shock of Global Population Decline*. New York: Crown
- Brouillette JF. 2021. *Women inventors and economic growth*. Work. Pap., Stanford Univ., Stanford, CA
- Budish E, Roin BN, Williams H. 2015. Do firms underinvest in long-term research? Evidence from cancer clinical trials. *Am. Econ. Rev.* 105(7):2044–85
- Buera FJ, Lucas RE. 2018. Idea flows and economic growth. *Annu. Rev. Econ.* 10:315–45
- Buera FJ, Oberfield E. 2020. The global diffusion of ideas. *Econometrica* 88(1):83–114
- Carroll L. 2019. *Through the Looking Glass*. N.p.: SDE Class.
- Cozzi G. 2020. *Semi-endogenous or fully endogenous growth? Both essential, one survives*. Work. Pap., Univ. St. Gallen, St. Gallen, Switz.
- Davidson T. 2021. Could advanced AI drive explosive economic growth? *Open Philanthropy*, June 25. <https://www.openphilanthropy.org/could-advanced-ai-drive-explosive-economic-growth>
- Dinopoulos E, Thompson P. 1998. Schumpeterian growth without scale effects. *J. Econ. Growth* 3(4):313–35
- Eaton J, Kortum SS. 1996. Trade in ideas patenting and productivity in the OECD. *J. Int. Econ.* 40(3–4):251–78
- Eaton J, Kortum SS. 1999. International technology diffusion: theory and measurement. *Int. Econ. Rev.* 40:537–70
- Eaton J, Kortum SS. 2002. Technology, geography, and trade. *Econometrica* 70(5):1741–79
- Eicher TS, Turnovsky SJ. 1999. Non-scale models of economic growth. *Econ. J.* 109(457):394–415
- Engbom N. 2019. *Firm and worker dynamics in an aging labor market*. Work. Pap., New York Univ., New York
- Fernald JG, Jones CI. 2014. The future of US economic growth. *Am. Econ. Rev. Pap. Proc.* 104(5):44–49
- Frankel M. 1962. The production function in allocation and growth: a synthesis. *Am. Econ. Rev.* 52:995–1022
- Garcia-Macia D, Hsieh CT, Klenow PJ. 2019. How destructive is innovation? *Econometrica* 87(5):1507–41
- Goldin C, Katz LF. 2008. *The Race Between Education and Technology*. Cambridge, MA: Belknap Press
- Griliches Z. 1957. Hybrid corn: an exploration in the economics of technological change. *Econometrica* 25(4):501–22
- Griliches Z, ed. 1984. *R&D, Patents and Productivity*. Chicago: Univ. Chicago Press
- Grossman GM, Helpman E. 1989. Product development and international trade. *J. Political Econ.* 97(6):1261–83
- Grossman GM, Helpman E. 1991. *Innovation and Growth in the Global Economy*. Cambridge, MA: MIT Press
- Grossman GM, Helpman E, Oberfield E, Sampson T. 2017. Balanced growth despite Uzawa. *Am. Econ. Rev.* 107(4):1293–312
- Growiec J. 2007. Beyond the linearity critique: the knife-edge assumption of steady-state growth. *Econ. Theory* 31(3):489–99
- Hall BH, Jaffe A, Trajtenberg M. 2005. Market value and patent citations. *RAND J. Econ.* 36(1):16–38
- Hall BH, Mairesse J, Mohnen P. 2010. Measuring the returns to R&D. In *Handbook of the Economics of Innovation*, Vol. 2, ed. BH Hall, N Rosenberg, pp. 1033–82. Amsterdam: North-Holland
- Hemous D, Olsen M. 2016. *The rise of the machines: automation, horizontal innovation and income inequality*. Work. Pap., Univ. Zurich, Zurich, Switz.

- Herrendorf B, Rogerson R, Valentinyi A. 2014. Growth and structural transformation. In *Handbook of Economic Growth*, Vol. 2, ed. P Aghion, SN Durlauf, pp. 855–941. Amsterdam: Elsevier
- Hopenhayn HA. 1992. Entry, exit, and firm dynamics in long-run equilibrium. *Econometrica* 60(5):1127–50
- Hopenhayn HA, Neira J, Singhania R. 2018. *From population growth to firm demographics: implications for concentration, entrepreneurship and the labor share*. NBER Work. Pap. 25382
- Howitt P. 1999. Steady endogenous growth with population and R&D inputs growing. *J. Political Econ.* 107(4):715–30
- Hsieh CT, Hurst E, Jones CI, Klenow PJ. 2019a. The allocation of talent and U.S. economic growth. *Econometrica* 87(5):1439–74
- Hsieh CT, Klenow PJ. 2009. Misallocation and manufacturing TFP in China and India. *Q. J. Econ.* 124(4):1403–48
- Hsieh CT, Klenow PJ, Nath IB. 2019b. *A global view of creative destruction*. NBER Work. Pap. 26461
- Jones BF, Summers LH. 2020. *A calculation of the social returns to innovation*. NBER Work. Pap. 27863
- Jones CI. 1995. R&D-based models of economic growth. *J. Political Econ.* 103(4):759–84
- Jones CI. 2002. Sources of U.S. economic growth in a world of ideas. *Am. Econ. Rev.* 92(1):220–39
- Jones CI. 2005. Growth and ideas. In *Handbook of Economic Growth*, Vol. 1B, ed. P Aghion, SN Durlauf, pp. 1063–111. Amsterdam: North-Holland
- Jones CI. 2016. The facts of economic growth. In *Handbook of Macroeconomics*, Vol. 2, ed. JB Taylor, H Uhlig, pp. 3–69. Amsterdam: Elsevier
- Jones CI. 2020. *The end of economic growth? Unintended consequences of a declining population*. NBER Work. Pap. 26651
- Jones CI, Kim J. 2018. A Schumpeterian model of top income inequality. *J. Political Econ.* 126(5):1785–826
- Judd KL. 1985. On the performance of patents. *Econometrica* 53(3):567–85
- Kaldor N. 1961. Capital accumulation and economic growth. In *The Theory of Capital*, ed. F Lutz, D Hague, pp. 177–222. New York: St. Martin's Press
- Karahan F, Pugsley B, Sahin A. 2019. *Demographic origins of the startup deficit*. NBER Work. Pap. 25874
- Klette TJ, Kortum S. 2004. Innovating firms and aggregate innovation. *J. Political Econ.* 112(5):986–1018
- Kogan L, Papanikolaou D, Seru A, Stoffman N. 2017. Technological innovation, resource allocation, and growth. *Q. J. Econ.* 132(2):665–712
- Kongsamut P, Rebelo S, Xie D. 2001. Beyond balanced growth. *Rev. Econ. Stud.* 68(4):869–82
- Kortum SS. 1997. Research, patenting, and technological change. *Econometrica* 65(6):1389–419
- Kortum SS, Lerner J. 1998. Stronger protection or technological revolution: What is behind the recent surge in patenting? *Carnegie-Rochester Conf. Ser. Public Policy* 48(1):247–304
- Krugman PR. 1979. Increasing returns, monopolistic competition, and international trade. *J. Int. Econ.* 9(4):469–79
- Li CW. 2000. Endogenous versus semi-endogenous growth in a two-R&D-sector model. *Econ. J.* 110(462):C109–22
- Li CW. 2003. Endogenous growth without scale effects: comment. *Am. Econ. Rev.* 93(3):1009–17
- Lin HC, Shampine LF. 2018. R&D-based calibrated growth models with finite-length patents: a novel relaxation algorithm for solving an autonomous FDE system of mixed type. *Comput. Econ.* 51(1):123–58
- Lucas RE Jr. 1978. On the size distribution of business firms. *Bell J. Econ.* 9:508–23
- Lucas RE Jr. 1988. On the mechanics of economic development. *J. Monet. Econ.* 22(1):3–42
- Lucas RE Jr., Moll B. 2014. Knowledge growth and the allocation of time. *J. Political Econ.* 122(1):1–51
- Luttmer EG. 2011. On the mechanics of firm growth. *Rev. Econ. Stud.* 78(3):1042–68
- Melitz MJ. 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6):1695–725
- Mincer J. 1974. *Schooling, Experience, and Earnings*. New York: Columbia Univ. Press
- Ngai LR, Pissarides CA. 2007. Structural change in a multisector model of growth. *Am. Econ. Rev.* 97(1):429–43
- Ngai LR, Samaniego R. 2011. Accounting for research and productivity growth across industries. *Rev. Econ. Dyn.* 14(3):475–95
- Nordhaus WD. 1969. An economic theory of technological change. *Am. Econ. Assoc. Pap. Proc.* 59:18–28
- OECD (Organ. Econ. Coop. Dev.). 2021. *Main science and technology indicators*. Data Set, Organ. Econ. Coop. Dev., Paris, accessed March 15, 2021. https://stats.oecd.org/Index.aspx?DataSetCode=MSTI_PUB

- Peretto PF. 1998. Technological change and population growth. *J. Econ. Growth* 3(4):283–311
- Peretto PF, Seater JJ. 2013. Factor-eliminating technical change. *J. Monet. Econ.* 60(4):459–73
- Perla J, Tonetti C. 2014. Equilibrium imitation and growth. *J. Political Econ.* 122(1):52–76
- Perla J, Tonetti C, Waugh ME. 2021. Equilibrium technology diffusion, trade, and growth. *Am. Econ. Rev.* 111(1):73–128
- Peters M. 2019. *Market size and spatial growth—evidence from Germany’s post-war population expulsions*. Work. Pap., Yale Univ., New Haven, CT
- Peters M, Walsh C. 2021. *Population growth and firm dynamics*. Work. Pap., Yale Univ., New Haven, CT
- Phelps ES. 1966. Models of technical progress and the Golden Rule of Research. *Rev. Econ. Stud.* 33(2):133–45
- Ramondo N, Rodriguez-Clare A, Saborio-Rodriguez M. 2016. Trade, domestic frictions, and scale effects. *Am. Econ. Rev.* 106(10):3159–84
- Redding SJ, Rossi-Hansberg E. 2017. Quantitative spatial economics. *Annu. Rev. Econ.* 9:21–58
- Restuccia D, Rogerson R. 2008. Policy distortions and aggregate productivity with heterogeneous plants. *Rev. Econ. Dyn.* 11:707–20
- Romer PM. 1990. Endogenous technological change. *J. Political Econ.* 98(5):S71–102
- Sampson T. 2016. Dynamic selection: an idea flows theory of entry, trade, and growth. *Q. J. Econ.* 131(1):315–80
- Sampson T. 2020. *Technology gaps, trade and income*. Work. Pap., Lond. Sch. Econ, London
- Segerstrom P. 1998. Endogenous growth without scale effects. *Am. Econ. Rev.* 88(5):1290–310
- Sterk V, Sedlacek P, Pugsley B. 2021. The nature of firm growth. *Am. Econ. Rev.* 111(2):547–79
- Trimborn T, Koch KJ, Steger TM. 2008. Multi-dimensional transitional dynamics: a simple numerical procedure. *Macroecon. Dyn.* 12(3):301–19
- UN. 2019. *World Population Prospects 2019*. Data Files, United Nations, New York. <https://population.un.org/wpp/>
- US Bur. Labor Stat. 2021. Multifactor productivity. *U.S. Bureau of Labor Statistics*. <https://www.bls.gov/mfp/>
- Young A. 1998. Growth without scale effects. *J. Political Econ.* 106(1):41–63
- Zeira J. 1998. Workers, machines, and economic growth. *Q. J. Econ.* 113(4):1091–117