

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

Charles I. Jones*

Stanford GSB and NBER

June 1, 2021 — Version 0.5

Preliminary, Comments Appreciated

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 paper interprets past and future growth from a semi-endogenous perspective. For 50+ years, U.S. 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 paper highlights many opportunities for further research.

*Draft in preparation for the *Annual Review of Economics*. I am grateful to Pete Klenow, Michael Peters, and Chris Tonetti for many helpful discussions and suggestions.

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 paper describes one potential answer to this question: the perspective from semi-endogenous growth theory.

The essence of semi-endogenous growth 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 her 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 semi-endogenous model of economic growth. People produce ideas and, because of nonrivalry, those ideas raise everyone's income. This means that income per person depends on the number of 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.

The rest of this paper develops this 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 words you've just read. Section 3 conducts a growth accounting exercise for the United States to make a key point: despite the fact that semi-endogenous growth theory implies that the entirety of long-run growth is ultimately due to population growth, this is far from true historically, say for the past 75 years. Instead, population growth contributes only around 20 percent of U.S. 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 growth framework itself.

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

Section 5 and 6 then turn to the future of economic growth. Section 5, building on the growth accounting exercise, highlights various reasons why 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.

Table 1: The Economic Environment of the Simple Model

Final good	$Y_t = A_t^\sigma L_{yt}$	(1)
------------	---------------------------	-----

Ideas	$\frac{\dot{A}_t}{A_t} = R_t^\lambda A_t^{-\beta}$	(2)
-------	--	-----

Resource constraint	$R_t + L_{yt} = L_t$	(3)
---------------------	----------------------	-----

Population growth	$L_t = L_0 e^{nt}$	(4)
-------------------	--------------------	-----

Allocation	$R_t = \bar{s} L_t$	(5)
------------	---------------------	-----

2. A Simple Model

The economic environment for a simple semi-endogenous growth 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 there are constant returns to scale to the rival inputs. But as long as the marginal product of another idea is positive, there must then be increasing returns to ideas and the other 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 semi-endogenous growth equation. New ideas are produced using researchers 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 — proportional improvements in productivity — are getting harder to find.

I go back and forth in my own papers 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 nice 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), ideas are thought of as proportional improvements in productivity, so the β formulation is more relevant.

The remainder of the equations in Table 1 lay out the resource constraints. There is an exogenous population L_t that grows exponentially at 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 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.

The Quality Ladder Alternative. This setup can easily be re-cast as a quality ladder model, a la Grossman and Helpman (1991) and Aghion and 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 (“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 $\beta > 0$ — it gets progressively harder to take steps up the quality ladder — makes this a semi-

endogenous growth model; this is basically the structure of Segerstrom (1998). Notice that combining $A_t = q^{N_t}$ and $\dot{N}_t = z_t = R_t^\lambda A_t^{-\beta}$ means that $\frac{\dot{A}_t}{A_t} = \log q \cdot R_t^\lambda A_t^{-\beta}$, 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.

2.1 Balanced Growth

We can 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

$$y_t = A_t^\sigma (1 - \bar{s}) \quad (6)$$

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.

Now, 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 $\frac{\dot{A}_t}{A_t} = R_t^\lambda A_t^{-\beta}$ for A_t along a BGP (denoted by a “*”) then gives

$$\begin{aligned} A_t^* &= \left(\frac{1}{g_A} \right)^{1/\beta} R_t^{\lambda/\beta} \\ &= \left(\frac{1}{g_A} \right)^{1/\beta} (\bar{s} L_t)^{\lambda/\beta}. \end{aligned} \quad (7)$$

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 is a property of all standard production functions: the more autoworkers we have, the more cars a factory can produce. Similarly, the more researchers we have, the more ideas we 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 (6) and (7) gives output per person along the balanced growth

path:

$$\begin{aligned} y_t^* &= (1 - \bar{s}) \left(\frac{1}{g_A} \right)^{\sigma/\beta} R_t^\gamma \\ &= (1 - \bar{s}) \left(\frac{1}{g_A} \right)^{\sigma/\beta} \bar{s}^\gamma L_t^\gamma \end{aligned} \quad (8)$$

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 (6) and (8) to yield

$$g_y = \sigma g_A = \frac{\lambda\sigma n}{\beta} \equiv \gamma n. \quad (9)$$

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 is long-run growth.

2.2 Discussion

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 grow 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 popula-

tion growth rates.

Can growth in the quality of researchers replace growth in their quantity? What we see in the semi-endogenous growth 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 had a constant number of researchers. Is it possible that their human capital could 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 the Lucas (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 the Lucas formulation here, then the answer is yes, quality (human capital) could replace quantity (population growth).

The problem with this answer is that it does not work in general, only for a knife-edge formulation of human capital accumulation. For example, suppose we take a Mincer (1974) approach to human capital, 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 an unskilled worker, but by a constant amount. That is, human capital here 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

$$\dot{h}_t = u_t h_t^\psi - \delta h_t$$

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 h , and changes in \bar{u} would change the long-run level of a person's human capital but would not generate sustained growth in h . 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. Do I learn more per hour in school simply because my parents were more educated? Perhaps. But if your parents have twice as much human capital as my parents, do you learn twice as much per hour as I do? That's 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 U.S. — with its high educational attainment — growing faster than other countries in the 20th century; see Bils and Klenow (2000). Also, rising educational attainment over the 20th century would lead to rising growth rates, which we do not see in the U.S. or in other countries.

The bottom line is that human capital can certainly be important — and in Section 3 below we will suggest that it accounts for 1/4 of U.S. 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.

Linearity and the Population Equation. One might wonder whether this linearity criticism also applies to the semi-endogenous growth approach. After all, it assumes a linear differential equation as well, in 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 are 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 2, then 4, then 8 and so on. The population of some types of bacteria can double every twenty 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 number, and it is this physical process that underlies the linear differential equation for population growth.¹

2.3 Evidence

Bloom, Jones, Van Reenen and Webb (2020) provide a recent look at idea production

¹For more discussion of the linearity critique, see Jones (2005) and Growiec (2007).

Figure 1: Productivity Growth and Research



Note: These graphs based on Bloom, Jones, Van Reenen and Webb (2020) illustrate that proportional improvements in productivity — “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 that paper suggest $\beta > 0$.

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 U.S. economy and then at the micro level in various case studies.

Figure 1 summarizes two of their key findings. The left panel shows the result for the aggregate U.S. economy. In particular, two series are plotted: a time series of TFP growth (an estimate of $\frac{\dot{A}_t}{A_t}$) and a time series of research effort, R_t . If $\beta = 0$, then these two series should move together because then $\frac{\dot{A}_t}{A_t} = R_t^\lambda$. 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.

The right panel of Figure 1 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 “rule of 70,” this is equivalent to saying that the growth rate of chip density was a stable 35% per year for this fifty-year period; this is the $\frac{\dot{A}_t}{A_t}$ series plotted in the right panel. 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 $\beta = 0$ 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. But instead, growth was stable: by the end of the period, it required 18 times the amount of research effort as in the early 1970s in order to double chip density. Converting this into an estimate of β reveals that for Moore’s Law, $\beta \approx 0.2$.

Bloom, Jones, Van Reenen and Webb (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 the low of 0.2 for Moore’s Law to values around 6 or 7 for corn and soybeans.³

2.4 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 half-

²For example, see https://web.archive.org/web/20170101183057if_/https://en.wikipedia.org/wiki/Moore's_law

³A third generation of endogenous growth models associated with Dinopoulos and Thompson (1998), Peretto (1998), Young (1998), and Howitt (1999) maintain the assumption of $\beta = 0$ from the original Romer (1990) and Aghion and Howitt (1992) models but assume this holds at the level of an individual product line or firm instead of for the aggregate economy. A rise in the extensive margin — the number of product lines — can potentially offset the rise in aggregate research, so that researchers per product line is stable and leads to stable economic growth. One of the motivations for the Bloom, Jones, Van Reenen and Webb (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.

way back to the steady state value g_A^* ? Appendix A shows that the answer in our simple model is given by

$$t_{1/2}^* = \frac{1}{\beta g_A^*} \ln \left(\frac{g_{A0} + g_A^*}{g_{A0}^2} \right) \quad (10)$$

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 semi-endogenous growth 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 U.S. 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, Jones, Van Reenen and Webb (2020) estimate $\beta = 0.2$ for semiconductors and $\beta = 3$ for the aggregate and for the average across the different cases they studied. The message from Table 2 is that half lives can be very long. For example, with $\beta = 3$, it takes more than a century for the economy to move halfway back to steady state. And for the semiconductor estimate of $\beta = 0.2$, the half life is more than 1700 years! These numbers are consistent with the long transition dynamics reported in Atkeson and Burstein (2019) and Atkeson, Burstein and Chatzikonstantinou (2019), which allow the allocation of research to change endogenously along the transition path.

A robust finding, then, is that transition dynamics in semi-endogenous growth models are very long. This has important implications. Shocks to the economy and changes in policy may have very long-lasting effects. Is it possible that the transition dynamics associated with the space program of the 1960s or even the extensive R&D conducted during World War II are still playing out in the U.S. and global economies today?

3. Accounting for Historical Growth

The preceding section laid out the basic semi-endogenous growth framework and argued that it helps us understand long-run growth. This section uses that framework

⁴Trimborn, Koch and Steger (2008) provide an elegant, fast algorithm that can be used to solve the transition dynamics of general semi-endogenous growth models in which allocations are determined by optimization rather than the assumption made here that allocations are constant over time. See also Lin and Shampine (2018).

Table 2: Half-Life Estimates (Years)

β	$g_{A0} = 2\%$	$g_{A0} = 4\%$
0.2	2159	1721
1	432	344
3	144	115
5	86	69

Note: The table reports the number of years it takes the basic model to transition half way 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.

to make an important point, which at first may seem contradictory but in fact is made from the semi-endogenous growth model itself: *while the ultimate source of long-run economic growth is population growth, historical growth accounting suggests that population growth accounts for only around 20% of U.S. 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 semi-endogenous growth 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 U.S. experience in recent decades. Put differently, just because 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 semi-endogenous growth framework. Consumption/investment goods Y are produced according to

$$Y_t = K_t^\alpha (Z_t h_t L_{Yt})^{1-\alpha} \quad (11)$$

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

$$Z_t \equiv A_t M_t \quad (12)$$

is total factor productivity, 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 which aggregate up to reduce TFP; we discuss this term in detail below. Relative to the model of 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 (11) to solve for income per person, $y_t \equiv Y_t/P_t$, gives

$$y_t = \left(\frac{K_t}{Y_t} \right)^{\frac{\alpha}{1-\alpha}} A_t M_t h_t \ell_t (1 - s_t) \quad (13)$$

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:

$$A_t^* = R_t^\gamma = (s_t L_t)^\gamma \quad (14)$$

This equation comes from (7) above, where we are assuming the educational attainment of researchers is constant and setting the factor of proportionality to one for simplicity.

There are two ways to approach the growth accounting: (i) in the long run, and (ii) 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, h , I find it simplest to think about the effects of educational attainment, so that $h_t = e^{\psi u_t}$ where u_t is years of educational attainment. At least if lives are finite, then the fraction of time people spend in school on average is bounded from above, so educational attainment also leads to level effects but no long-run growth effects. The fraction of people working (ℓ_t) and the fraction 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 constant eventually 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 only leaves us with the growth rate of the labor force, which equals

population growth in the long run. Therefore, even in this extended semi-endogenous growth 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 (13) and (14).

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 U.S. economic growth.

To see this more formally, take logs and differences of these two equations to get:

$$\underbrace{d \log y_t}_{\text{GDP per person}} = \underbrace{\frac{\alpha}{1-\alpha} d \log \frac{K_t}{Y_t}}_{\text{Capital-Output ratio}} + \underbrace{d \log h_t}_{\text{Educational att.}} + \underbrace{d \log \ell_t}_{\text{Emp-Pop ratio}} + \underbrace{d \log(1-s_t)}_{\text{Goods intensity}} + \underbrace{d \log M_t + d \log A_t}_{\text{TFP growth}} \quad (15)$$


where

$$\text{TFP growth} \equiv \underbrace{d \log M_t}_{\text{Misallocation}} + \underbrace{d \log A_t}_{\text{Ideas}} = \underbrace{d \log M_t}_{\text{Misallocation}} + \underbrace{\gamma d \log s_t}_{\text{Research intensity}} + \underbrace{\gamma d \log L_t}_{\text{LF growth}} \quad (16)$$

Jones (2002) uses equations just like these to conduct a growth accounting exercise for the U.S. for the period 1953 to 1993, while Fernald and 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 instead, I present a “back-of-the-envelope” version of the accounting here that includes a rough estimate of gains from changing misallocation. Also, for more details on the facts that are discussed in the remainder of this section, see 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 (15) to decompose this 2% growth into its components. First, the capital-

Figure 2: Historical Growth Accounting



Note: The figure shows a growth accounting exercise for the United States since the 1950s using equations (15) and (16). See the main text for details.

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 $1 - s_t \approx 1$ over time. We now turn to the non-zero components of the equation.

Human capital and labor force participation. This brings us to educational attainment. A wonderful stylized fact documented by Goldin and Katz (2008) is that educational attainment throughout the 20th century rose by slightly under one year per decade, e.g. 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 20th 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've seen educational attainment for each cohort flatten out: roughly 85% of kids graduate from high school, roughly 1/3 graduate from 4-year colleges, and these numbers have levelled 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 contribution 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 2% leaves us with TFP growth of around 1.3% per year. (Because of slowing growth since around 2003, more recent numbers would be lower; also note that this TFP measure is in labor-augmenting units.)

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 and Rogerson (2008) and Hsieh and 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. And it helps us to understand how TFP levels in countries like Italy and Spain can decline for two decades; clearly the stock of ideas did not decline, so the misallocation explanation is likely very important here. 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 U.S. economy?” I don’t think we have a firm answer to this question. Bils, Klenow and Ruane (2020) find that allocative efficiency in U.S. manufacturing has gotten worse over time, although they suggest that at least 2/3rds of this change may be due to measurement error. On the other side, Hsieh, Hurst, Jones and Klenow (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.

Ideas. Subtracting the 0.3% contribution from declining misallocation from the 1.3% TFP growth leaves us with 1.0% growth due to increases in ideas. Now, recall equation (14): $A_t^* = R_t^\gamma = (s_t L_t)^\gamma$. There are several ways to proceed. If we have a measure of R_t that we believe, we can use the growth in A^* and 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 idea contribution 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. There is also the question of how much of research growth that comes from Germany, Japan, China, etc., should enter into the calculation. We will not resolve those questions here, which is another reason why this is a “back of the envelope” kind of calculation.

Here, we simply impose $\gamma = 1/3$ and assume 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 the two add to 1.0%.

Remarks. We learn a very important lesson from the growth accounting in Figure 2. Even in this semi-endogenous growth framework in which population growth is the only potential source of growth in the long run, other factors explain more than 80% of U.S. growth in recent decades: the contribution of population growth is 0.3% or less 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 semi-endogenous growth (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 SEG 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 re-examined using SEG models.

SEG is more common than you might think. The first point to make is that the semi-endogenous growth 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 features increasing returns to scale is likely a semi-endogenous growth model. Romer (1990) provides the most solid microfoundations for justifying increasing returns, but in fact any model that assumes or estimates increasing returns will likely feature the result 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 Phelps (1966), Nordhaus (1969), Judd (1985), Jones (1995), Kortum (1997), and Segerstrom (1998); see Jones (2005) for a detailed discussion of the historical background.⁵

Other models not traditionally thought of as semi-endogenous growth models also share this feature. International trade models are a great example of this. Krugman (1979), Grossman and Helpman (1989), and Melitz (2003) each feature increasing returns so that income per person is an increasing function of the labor force. Adding population growth to those models would lead 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 Hopenhayn (1992) model of industry equilibrium with entry and exit will also fall into this category. However, it does not: Hopenhayn carefully structures his model so it features constant returns to scale so that a competitive equilibrium exists and is efficient. He does this by using the Lucas (1978) span-of-control setup so that within each firm there is decreasing returns to scale; the fixed cost of entry then leads firms to the bottom of their U-shaped average cost curve

⁵Other papers focused on semi-endogenous growth include Eicher and Turnovsky (1999), Li (2000), Li (2003), and Cozzi (2020). More are discussed below.

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 semi-endogenous growth models.

Another key setup in international trade is the Eaton and Kortum (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 in Kortum (1997); see especially Section 5.2 of Eaton and Kortum (2002). Under this interpretation, this structure also is one of semi-endogenous growth. Ramondo, Rodriguez-Clare and Saborio-Rodriguez (2016) and Arkolakis, Ramondo, Rodriguez-Clare and Yeaple (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 semi-endogenous growth framework. Models based on Hopenhayn (1992) feature constant returns to scale 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 Karahan, Pugsley and Sahin (2019) and Hopenhayn, Neira and Singhania (2018). Engbom (2019) is more in the Aghion and Howitt (1992) tradition, while Sterk, Sedlacek and Pugsley (2021) follows the Melitz tradition.

More to the point, Peters and Walsh (2021) use a rich semi-endogenous 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 to scale and therefore typically would have semi-endogenous growth implications if considered over time. See Redding and Rossi-Hansberg (2017) for a recent survey.

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 includes Kongsamut, Rebelo and

Xie (2001), Ngai and Pissarides (2007), Acemoglu and Guerrieri (2008), and Herrendorf, Rogerson and Valentinyi (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 the interesting question: why do different industries have different TFP growth rates? Ngai and Samaniego (2011) study this question in a semi-endogenous growth framework in which different industries have different parameters of the idea production function. 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 in the education and construction industries. This question is ripe for re-examination with more recent data, especially given the interest in heterogeneity more broadly. Bloom, Jones, Van Reenen and Webb (2020) and Sampson (2020) can be read as pushing in this direction, but there is clearly scope for much more work here.

Novel places where the SEG Framework could be applied. There are many interesting papers in the recent growth literature that are not written in a semi-endogenous growth framework. It would be interesting to reformulate these models in a semi-endogenous growth 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 tradeoffs, 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 and Kim (2018) study a Schumpeterian model of top income inequality. One of the mechanisms 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 there. This is useful because models that generate Pareto distributions often do so by specifying a process with exponential growth and a constant “death rate” or exit 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 and Kortum (2004) build a model of firm dynamics designed to match a bunch of facts about innovation and growth at the firm level. Acemoglu, Aghion, Bursztyn and Hémous (2012) study climate change and economic growth when research can be focused on clean versus dirty technologies. Acemoglu and Restrepo (2018) consider the effects of automation. Akcigit and 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” in some sense and homogeneous across firms or sectors, and this may drive the quantitative results that are derived. How different would the results be if the idea production functions incorporate 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 semi-endogenous growth 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. Aghion, Bergeaud, Boppart, Klenow and Li (2019a), Garcia-Macia, Hsieh and Klenow (2019), Hsieh, Klenow and Nath (2019b) are in this category. They have a clever structure so that 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 semi-endogenous growth microfoundations would be valuable.

Another class of models in which semi-endogenous growth could be usefully applied is the recent literature on technology diffusion. Models such as Alvarez, Buera and Lucas (2013), Lucas and Moll (2014), and Perla and 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 no “best” idea in the world at a given point in time — there is always someone else with a better idea. Sampson (2016) and Perla, Tonetti and Waugh (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, Perla and Tonetti (2021) relax the spe-

cial assumptions mentioned above and study technology diffusion in an endogenous growth model in which ideas get invented. Buera and Oberfield (2020) instead use a semi-endogenous growth framework to study technology diffusion and international trade. Buera and 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 “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 U.S. economic growth over the past several decades. The stylized picture that emerged is that while the only source of long-run economic growth in living standards is population growth, historically more than 80 percent of U.S. economic growth is due to other factors. These factors include rising educational attainment (perhaps 25 percent of growth), reductions in misallocation (likely more than 15 percent of growth), and increases in the fraction of the population devoted to research (perhaps 45 percent of growth).

The point to emphasize here is that this framework strongly implies that, unless something dramatic changes, future growth rates will be substantially lower. In particular, all 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. In the long-run, however, this fraction must level out — we cannot spend more than 100 percent of our time in school. Moreover, this leveling-out already seems to be occurring. In particular, the cohorts born in more recent decades also 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 appears to have 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 and improving the allocation of resources.

None of this is to say that there aren't ways around these conclusions. Certainly there is substantial mismeasurement in human capital and research. For example, educational attainment is only one form of human capital. Maybe we've 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. And perhaps misallocation has not declined, or has a long way to go. But the main point is that a natural way of looking at the data on growth in recent decades through the lens of a semi-endogenous growth model suggests that a substantial slowdown in growth may lie ahead.

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 semi-

⁶See Goldin and Katz (2008). Jones (2016) updates the data on this and the other facts discussed in this section.

endogenous growth: 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 OECD's Main Science and Technology indicators for three different aggregates: the United States, the OECD countries, and the "world" (where "world" here 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 U.S. economic growth. But this is a difficult issue to get right, so we will take the short-cut 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 afterwards. For the OECD, there is a similar slowdown: research employment grew at 4.1% per year before 2003 but only 2.8% per year after. Adding in the large increase in research employment in China is enough to mitigate most of this slowdown, as shown in the "world" number. But it is unclear how to properly weight Chinese research efforts (how much is historically catch-up versus pushing the frontier forward), so for now we will focus on the U.S. and OECD numbers.

The slowdown in U.S. research employment growth is 38 percent while for the OECD the slowdown is slightly smaller at 32 percent. Recall that in a semi-endogenous growth 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.

⁷From 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!"

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

Figure 3: Research Employment in the U.S., OECD, and World



Note: “World” is the OECD 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. The data for Russia start in 1994, so we assume the values for 1991 to 1993 are equal to the 1994 value (research employment in Russia was declining in the 1990s). Source: OECD Main Science and Technology Indicators (2021).


5.3 Slowing Population Growth

A final reason why economic growth rates might be slowing in the future is that population growth — the long-run driving force behind semi-endogenous growth — is itself slowing throughout the world. 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 is another reason why 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 and 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 may actually 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

Figure 4: Population Growth around the World




Note: Average annual rates of population growth for countries classified according to their 2018 World Bank income grouping. Each data point corresponds to a five-year period. Source: United Nations (2019).

— for a selection of countries in the 2015–2020 period. Adjusting for mortality, population growth is positive when each woman has two (or slightly more) children during her lifetime. What we see in the figure is that many countries — and 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:

$$\frac{\dot{A}_t}{A_t} = R_t^\lambda A_t^{-\beta} \quad (17)$$

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.

Figure 5: The Total Fertility Rate around the World

Note: The total fertility rate is the average number of live births a hypothetical cohort of women would have over their reproductive life 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 five-year period 2015–2020. Source: United Nations (2019).

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

$$\frac{\dot{A}_t}{A_t} = R_0^\lambda A_t^{-\beta} e^{-\lambda \eta t} \quad (18)$$

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

$$A_t = \begin{cases} A_0 \left(1 + \frac{\beta g A_0}{\lambda \eta} (1 - e^{-\lambda \eta t})\right)^{1/\beta} & \text{if } \beta > 0 \\ A_0 \exp\left(\frac{g A_0}{\lambda \eta} (1 - e^{-\lambda \eta t})\right) & \text{if } \beta = 0 \end{cases} \quad (19)$$

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*

$$A^* = \begin{cases} A_0 \left(1 + \frac{\beta g A_0}{\lambda \eta}\right)^{1/\beta} & \text{if } \beta > 0 \\ A_0 \exp\left(\frac{g A_0}{\lambda \eta}\right) & \text{if } \beta = 0 \end{cases} \quad (20)$$

That is, when population growth is negative, both the semi-endogenous growth 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 — converge 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 paper. 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 were 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 ultimately lead to stunningly different outcomes.

5.4 Summing Up

To conclude this section, there are several reasons to worry that future growth will be slower in the long run. Many of the sources of growth historically — including rising educational attainment, rising research intensity, and declining misallocation — are inherently limited and cannot go on forever. The key source of sustainable 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

While 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 possibilities 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 less than 10 million. In other words something on the order of one or two out of every thousand people in the world is 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.

The rise of China, India, and other countries. The United States, Western Europe, and Japan together have about 1 billion people, or only about 1/7th 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 semi-endogenous growth 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 to 0.4 percentage points of growth each year.

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 percent of U.S. inventors were women. Even as of 2016 the share was less than 12 percent. 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.

Other sources of within-country talent. Bell, Chetty, Jaravel, Petkova and Van Reenen (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 only limited 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.

The Zeira (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

$$Y = AX_1^{\alpha_1} X_2^{\alpha_2} \cdot \dots \cdot X_n^{\alpha_n} \quad \text{where} \quad \sum_{i=1}^n \alpha_i = 1. \quad (21)$$

Zeira thought of the X_i 's as intermediate goods, but we follow Acemoglu and 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:

$$X_i = \begin{cases} L_i & \text{if not automated} \\ K_i & \text{if automated} \end{cases} \quad (22)$$

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

$$Y_t = A_t K_t^\alpha L_t^{1-\alpha} \quad (23)$$

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

$$g_y = \frac{g}{1 - \alpha}, \quad (24)$$

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, increase the long-run growth rate.

Zeira emphasizes that automation has been going on at least since the industrial revolution. And certainly the 20th 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 this automation. This prediction is strongly rejected by the data. Instead the Kaldor (1961) stylized fact that growth rates and capital shares are relatively stable over time is a good

characterization of the U.S. economy for the bulk of the 20th century, certainly through 1980 and perhaps up to 2000; for example, see Jones (2016). The Zeira framework, then, needs to be improved so that it is consistent with historical evidence.

Acemoglu and Restrepo (2018) provide one approach to solving this problem. Their rich environment allows 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, Jones and Jones (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 (1967)'s 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.

Automating Idea Production. Even more intriguing possibilities arise from considering the automation of tasks in idea production; see Aghion, Jones and Jones (2019b) and Agrawal, McHale and Oettl (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.

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

Let's apply the Zeira structure to the idea production function:

$$\dot{A}_t = A_t^{1-\beta} X_1^{\alpha_1} X_2^{\alpha_2} \cdot \dots \cdot X_n^{\alpha_n} \quad (25)$$

$$= A_t^{1-\beta} K_t^\alpha R_t^{1-\alpha} \quad (26)$$

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:

$$Y_t = A_t L_{yt} \quad (27)$$

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

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

$$\dot{A}_t = \kappa A_t^{1-(\beta-\alpha)} L_t \quad (28)$$

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

$$g_A = \frac{n}{\beta - \alpha} \quad (29)$$

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 will 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

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 in 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.

¹⁰Aghion, Jones and Jones (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.

7. Key Unresolved Questions for Future Research

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

How large is the degree of IRS? 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 *size* of the degree of increasing returns is a fundamental parameter that plays a key role in shaping the answer to many practical questions. Knowing its value is essential for calibrating the 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

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 Bloom, Schankerman and Van Reenen (2013), which uses variation in R&D tax credits across time and across U.S. states and estimates very high values for the social return to R&D, on the order of 50%. See Hall, Mairesse and Mohnen (2010) for a recent review.

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, ed 1984; Hall et al. 2005; Budish et al. 2015; Kogan et al. 2017). Yet patents have problems as well. Most firms do not patent, so patents capture only one part of idea production. Argente, Hanley, Baslandze and Moreira (2020) use machine learning algorithms to link patents to a database of consumer goods and argue that

¹¹Arkolakis, Lee and Peters (2020) use European immigration to the U.S. between 1880 and 1920 and estimate a range of values of 0.7 to 1.3. Peters (2019) using the pseudo-random 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.

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 than a patent granted in 2020, making comparisons over time difficult (Kortum and 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 U.S. Bureau of Labor Statistics (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 total factor productivity growth. More research is needed on this point, but the standard multifactor productivity growth series from the BLS surely misses an important part of innovation.

Better growth accounting. Back in Section 3, we presented a stylized accounting of U.S. economic growth since the 1950s. A much richer version of this growth accounting would be valuable. How much of U.S. growth is due to ideas discovered in other countries (and vice versa)? What is the contribution from changing misallocation over time and from multiple sources? Given the extremely long half-lives in the basic semi-endogenous growth model, how much of growth in recent decades is due to changes that occurred 50 or 100 years ago — what was the contribution from DARPA, the NSF and the NIH, the migration of top scientists out of Europe surrounding World War II, the migration of talented people from around the world, etc.?

Conclusions. I hope this essay convinces you 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.

A. Appendix: Solving the Key Differential Equation

A key equation in semi-endogenous growth models is the idea production function:

$$\frac{\dot{A}_t}{A_t} = R_t^\lambda A_t^{-\beta}. \quad (30)$$

This differential is easy to integrate and doing so yields useful insights:

$$\int_0^\infty A^{\beta-1} dA = \int_0^\infty R_t^\lambda dt \equiv K_t. \quad (31)$$

The right-hand side of this equation is already interesting, even without imposing any conditions on the time path of R_t . The stock K_t is the cumulative amount of “effective” research that has been conducted as of date t , where “effective” denotes the fact that we are cumulating R_t^λ rather than just R_t to adjust for any duplication effects.

Solving the integral on the left side of equation (31) gives the solution for the stock of knowledge at any date:

$$A_t = \left(A_0^\beta + \beta K_t \right)^{1/\beta} \quad (32)$$

That is, the stock of knowledge at any date t depends on the initial stock A_0 and the cumulative amount of effective research undertaken, K_t . The parameter β is partly the weight on the new research and partly a curvature parameter governing how the sum is taken. Notice that K_t looks somewhat like the “stock of R&D” often used in the productivity literature (Bloom et al., 2013). It is also reminiscent of the stock of research in Kortum (1997). This equation holds at any point in time and for any time series history of research.

A.1 Transition Dynamics

This derivation can be taken further and put to additional use to study the transition dynamics of the basic growth model. For this, we do need to specify a time path of R_t . The most natural one is to assume that R_t is growing at a constant exponential rate. We could proceed for any growth rate, but the natural one is the rate of population growth, n , so we will assume $R_t = R_0 e^{nt}$.

Under this assumption, the cumulative amount of research is given by

$$\begin{aligned} K_t &= \int_0^t R_v^\lambda dv = \frac{1}{\lambda n} (R_t^\lambda - R_0^\lambda) \\ &= \frac{R_0^\lambda}{\lambda n} (e^{\lambda n t} - 1) \end{aligned}$$

Substituting this expression into the solution for A_t in (32) gives

$$\begin{aligned} A_t^\beta &= A_0^\beta + \beta(K_t - K_0) \\ &= A_0^\beta + \frac{\beta R_0^\lambda}{\lambda n} (e^{\lambda n t} - 1) \\ &= A_0^\beta \left[1 + \frac{\beta}{\lambda n} \frac{R_0^\lambda}{A_0^\beta} (e^{\lambda n t} - 1) \right] \\ &= A_0^\beta \left[1 + \frac{g_{A0}}{g_A^*} (e^{\lambda n t} - 1) \right] \end{aligned} \tag{33}$$

where the last equation uses the fact that $g_A^* = \frac{\lambda n}{\beta}$ and $g_{A0} = R_0^\lambda / A_0^\beta$. This gives a simple expression for the stock of ideas at each date when research grows at a constant rate.

A.2 Half Lives

We can now use this last expression to solve for the half life of the transition dynamics. In particular, suppose the economy begins with initial conditions that deliver an initial growth rate of $g_{A0} = R_0^\lambda / A_0^\beta$. How many years does it take before the growth rate has converged half the way to the steady state value g_A^* ? The answer is the value of t such that $g_{At} = \frac{1}{2}(g_{A0} + g_A^*)$.

To solve for this time, notice that we can plug the solution for A_t^β in (33) into the idea production function in (30) to get

$$\begin{aligned} \frac{\dot{A}_t}{A_t} &= \frac{R_t^\lambda}{A_0^\beta \left[1 + \frac{g_{A0}}{g_A^*} (e^{\lambda n t} - 1) \right]} \\ &= \frac{R_0^\lambda e^{\lambda n t}}{A_0^\beta \left[1 + \frac{g_{A0}}{g_A^*} (e^{\lambda n t} - 1) \right]} \\ &= g_{A0} \cdot \frac{e^{\lambda n t}}{1 + \frac{g_{A0}}{g_A^*} (e^{\lambda n t} - 1)} \end{aligned}$$

Now we just set this expression equal to $\frac{1}{2}(g_{A0} + g_A^*)$ to find the half life:

$$\frac{1}{2}(g_{A0} + g_A^*) = g_{A0} \cdot \frac{e^{\lambda n t}}{1 + \frac{g_{A0}}{g_A^*} (e^{\lambda n t} - 1)}$$

Doing a lot of algebra to solve this equation for t gives

$$\begin{aligned} t_{1/2}^* &= \frac{1}{\lambda n} \ln \left(\frac{g_{A0} + g_A^*}{g_{A0}^2} \right) \\ &= \frac{1}{\beta g_A^*} \ln \left(\frac{g_{A0} + g_A^*}{g_{A0}^2} \right). \end{aligned} \quad (34)$$

References

- Acemoglu, Daron and David Autor, “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 4, Elsevier, June 2011, chapter 12, pp. 1043–1171.
- and Pascual Restrepo, “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, June 2018, 108 (6), 1488–1542.
- and —, “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 2020, 128 (6), 2188–2244.
- and Veronica Guerrieri, “Capital Deepening and Nonbalanced Economic Growth,” *Journal of Political Economy*, 2008, 116 (3), 467–498.
- , Philippe Aghion, Leonardo Bursztyn, and David Hemous, “The Environment and Directed Technical Change,” *American Economic Review*, 2012, 102 (1), 131–166.
- Aghion, Philippe and Peter Howitt, “A Model of Growth through Creative Destruction,” *Econometrica*, March 1992, 60 (2), 323–351.
- , Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li, “Missing Growth from Creative Destruction,” *American Economic Review*, August 2019, 109 (8), 2795–2822.
- , Benjamin F. Jones, and Charles I. Jones, “Artificial Intelligence and Economic Growth,” in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, 2019, pp. 237–282.

- Agrawal, Ajay, John McHale, and Alexander Oettl, "Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, 2019, pp. 149–174.
- Akcigit, Ufuk and William R. Kerr, "Growth through Heterogeneous Innovations," *Journal of Political Economy*, 2018, 126 (4), 1374 – 1443.
- Alvarez, Fernando E., Francisco J. Buera, and Robert E. Lucas, "Idea Flows, Economic Growth, and Trade," Technical Report, National Bureau of Economic Research 2013. Working Paper No. 19667.
- Argente, David, Douglas Hanley, Salome Baslandze, and Sara Moreira, "Patents to Products: Product Innovation and Firm Dynamics," FRB Atlanta Working Paper 2020-4, Federal Reserve Bank of Atlanta April 2020.
- Arkolakis, Costas, Natalia Ramondo, Andres Rodriguez-Clare, and Stephen Yeaple, "Innovation and Production in the Global Economy," *American Economic Review*, August 2018, 108 (8), 2128–2173.
- , Sun Kyoung Lee, and Michael Peters, "European Immigrants and the United States' Rise to the Technological Frontier," Technical Report 2020. Yale University manuscript.
- Arrow, Kenneth J., "The Economic Implications of Learning by Doing," *Review of Economic Studies*, June 1962, 29, 153–173.
- Atkeson, Andrew and Ariel Burstein, "Aggregate Implications of Innovation Policy," *Journal of Political Economy*, December 2019, 127 (6), 2625–2683.
- , —, and Manolis Chatzikonstantinou, "Transitional Dynamics in Aggregate Models of Innovative Investment," *Annual Review of Economics*, 2019, 11 (1), 273–301.
- Baumol, William J., "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *American Economic Review*, June 1967, 57, 415–426.
- Bell, Alexander M., Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen, "Who Becomes an Inventor in America? The Importance of Exposure to Innovation," *The Quarterly Journal of Economics*, 2019, 134 (2), 647–713.
- Benhabib, Jess, Jesse Perla, and Christopher Tonetti, "Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier," *Econometrica*, 2021, *forthcoming*.

- Bils, Mark and Peter Klenow, "Does Schooling Cause Growth?," *American Economic Review*, December 2000, 90, 1160–1183.
- , Peter J. Klenow, and Cian Ruane, "Misallocation or Mismeasurement?," Technical Report, National Bureau of Economic Research 2020. Working Paper No. 26711.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, April 2020, 110 (4), 1104–44.
- , Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 07 2013, 81 (4), 1347–1393.
- Bricker, Darrell and John Ibbitson, *Empty Planet: The Shock of Global Population Decline*, Crown Publishing Group, 2019.
- Brouillette, Jean-Felix, "Women Inventors and Economic Growth," 2021. Stanford University manuscript.
- Budish, Eric, Benjamin N. Roin, and Heidi Williams, "Do firms underinvest in long-term research? Evidence from cancer clinical trials," *The American Economic Review*, 2015, 105 (7), 2044–2085.
- Buera, Francisco J and Ezra Oberfield, "The Global Diffusion of Ideas," *Econometrica*, 2020, 88 (1), 83–114.
- Buera, Francisco J. and Robert E. Lucas, "Idea Flows and Economic Growth," *Annual Review of Economics*, 2018, 10, 315–345.
- Cozzi, Guido, "Semi-Endogenous or Fully Endogenous Growth? Both Essential, One Survives," 2020. University of St. Gallen manuscript.
- Davidson, Tom, "Longterm GWP Forecasts," 2021. Open Philanthropy manuscript.
- Dinopoulos, Elias and Peter Thompson, "Schumpeterian Growth without Scale Effects," *Journal of Economic Growth*, December 1998, 3 (4), 313–335.
- Eaton, Jonathan and Samuel Kortum, "Trade in ideas Patenting and productivity in the OECD," *Journal of International Economics*, May 1996, 40 (3-4), 251–278.
- and Samuel S. Kortum, "International Technology Diffusion: Theory and Measurement," *International Economic Review*, 1999, 40, 537–570.
- and —, "Technology, Geography, and Trade," *Econometrica*, 2002, 70 (5), 1741–1779.

- Eicher, Theo S. and Stephen J. Turnovsky, “Non-Scale Models of Economic Growth,” *The Economic Journal*, 1999, 109 (457), 394–415.
- Engbom, Niklas, “Firm and Worker Dynamics in an Aging Labor Market,” 2019. NYU unpublished manuscript.
- Fernald, John G. and Charles I. Jones, “The Future of US Economic Growth,” *American Economic Review Papers and Proceedings*, May 2014, 104 (5), 44–49.
- Frankel, Marvin, “The Production Function in Allocation and Growth: A Synthesis,” *American Economic Review*, December 1962, 52, 995–1022.
- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow, “How Destructive Is Innovation?,” *Econometrica*, 2019, 87 (5), 1507–1541.
- Goldin, Claudia and Lawrence F. Katz, *The Race between Education and Technology*, Belknap Press, 2008.
- Griliches, Zvi, “Hybrid Corn: An Exploration in the Economics of Technological Change,” *Econometrica*, October 1957, 25 (4), 501–522.
- , ed., *R&D, Patents and Productivity*, Chicago: University of Chicago Press, 1984.
- Grossman, Gene M. and Elhanan Helpman, “Product Development and International Trade,” *Journal of Political Economy*, 1989, 97 (6), 1261–1283.
- and —, *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press, 1991.
- Growiec, Jakub, “Beyond the Linearity Critique: The Knife-Edge Assumption of Steady-State Growth,” *Economic Theory*, 2007, 31 (3), 489–499.
- Hall, Bronwyn H, Adam Jaffe, and Manuel Trajtenberg, “Market value and patent citations,” *RAND Journal of economics*, 2005, pp. 16–38.
- Hall, Bronwyn H., Jacques Mairesse, and Pierre Mohnen, “Measuring the Returns to R&D,” in Bronwyn H. Hall and Nathan Rosenberg, eds., *Handbook of the Economics of Innovation*, Vol. 2, North-Holland, 2010, pp. 1033 – 1082.
- Hemous, David and Morten Olsen, “The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality,” 2016. University of Zurich manuscript.
- Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi, “Growth and Structural Transformation,” in “Handbook of Economic Growth,” Vol. 2 of *Handbook of Economic Growth*, Elsevier, 2014, chapter 6, pp. 855–941.

- Hopenhayn, Hugo A., "Entry, Exit, and Firm Dynamics in Long-Run Equilibrium," *Econometrica: Journal of the Econometric Society*, 1992, pp. 1127–1150.
- Hopenhayn, Hugo, Julian Neira, and Rish Singhania, "From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship and the Labor Share," Working Paper 25382, National Bureau of Economic Research December 2018.
- Howitt, Peter, "Steady Endogenous Growth with Population and R&D Inputs Growing," *Journal of Political Economy*, August 1999, 107 (4), 715–730.
- Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, 124 (4), 1403–1448.
- , Erik Hurst, Charles I. Jones, and Peter J. Klenow, "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 2019, 87 (5), 1439–1474.
- , Peter J. Klenow, and Ishan B. Nath, "A Global View of Creative Destruction," Working Paper 26461, National Bureau of Economic Research November 2019.
- Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, August 1995, 103 (4), 759–784.
- , "Sources of U.S. Economic Growth in a World of Ideas," *American Economic Review*, March 2002, 92 (1), 220–239.
- , "Growth and Ideas," in Philippe Aghion and Steven A. Durlauf, eds., *Handbook of Economic Growth*, New York: North Holland, 2005, pp. 1063–1111.
- , "The Facts of Economic Growth," *Handbook of Macroeconomics*, 2016, 2a, 3–69.
- , "The End of Economic Growth? Unintended Consequences of a Declining Population," 2020. NBER Working Paper No. 26651.
- and Jihee Kim, "A Schumpeterian Model of Top Income Inequality," *Journal of Political Economy*, October 2018, 126 (5), 1785–1826.
- Judd, Kenneth L., "On the Performance of Patents," *Econometrica*, May 1985, 53 (3), 567–585.
- Kaldor, Nicholas, "Capital Accumulation and Economic Growth," in F.A. Lutz and D.C. Hague, eds., *The Theory of Capital*, St. Martins Press, 1961, pp. 177–222.
- Karahan, Fatih, Benjamin Pugsley, and Aysegul Sahin, "Demographic Origins of the Startup Deficit," Working Paper 25874, National Bureau of Economic Research May 2019.

- Klette, Tor Jakob and Samuel Kortum, "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, October 2004, 112 (5), 986–1018.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," *The Quarterly Journal of Economics*, 2017, 132 (2), 665–712.
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie, "Beyond Balanced Growth," *The Review of Economic Studies*, 2001, 68 (4), 869–882.
- Kortum, Samuel and Josh Lerner, "Stronger protection or technological revolution: what is behind the recent surge in patenting?," *Carnegie-Rochester Conference Series on Public Policy*, June 1998, 48 (1), 247–304.
- Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, 65 (6), 1389–1419.
- Krugman, Paul R., "Increasing Returns, Monopolistic Competition, and International Trade," *Journal of International Economics*, November 1979, 9 (4), 469–479.
- Li, Chol-Won, "Endogenous vs. Semi-endogenous Growth in a Two-R&D-Sector Model," *Economic Journal*, March 2000, 110 (462), C109–C122.
- , "Endogenous Growth Without Scale Effects: Comment," *American Economic Review*, June 2003, 93 (3), 1009–1017.
- Lin, Hwan C. and L. F. Shampine, "R&D-based Calibrated Growth Models with Finite-Length Patents: A Novel Relaxation Algorithm for Solving an Autonomous FDE System of Mixed Type," *Computational Economics*, January 2018, 51 (1), 123–158.
- Lucas, Robert E., "On the Mechanics of Economic Development," *Journal of Monetary Economics*, 1988, 22 (1), 3–42.
- and Benjamin Moll, "Knowledge Growth and the Allocation of Time," *Journal of Political Economy*, February 2014, 122 (1), 1–51.
- Lucas, Robert E., Jr., "On the Size Distribution of Business Firms," *Bell Journal of Economics*, 1978, 9, 508–523.
- Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, 71 (6), 1695–1725.
- Mincer, Jacob, *Schooling, Experience, and Earnings*, New York: Columbia University Press, 1974.

- Ngai, L. Rachel and Christopher A. Pissarides, “Structural Change in a Multisector Model of Growth,” *American Economic Review*, March 2007, 97 (1), 429–443.
- Ngai, Rachel and Roberto Samaniego, “Accounting for Research and Productivity Growth Across Industries,” *Review of Economic Dynamics*, July 2011, 14 (3), 475–495.
- Nordhaus, William D., “An Economic Theory of Technological Change,” *American Economic Association Papers and Proceedings*, May 1969, 59, 18–28.
- Peretto, Pietro, “Technological Change and Population Growth,” *Journal of Economic Growth*, December 1998, 3 (4), 283–311.
- Peretto, Pietro F. and John J. Seater, “Factor-eliminating technical change,” *Journal of Monetary Economics*, 2013, 60 (4), 459–473.
- Perla, Jesse and Christopher Tonetti, “Equilibrium Imitation and Growth,” *Journal of Political Economy*, February 2014, 122 (1), 52–76.
- , —, and Michael E. Waugh, “Equilibrium Technology Diffusion, Trade, and Growth,” *American Economic Review*, January 2021, 111 (1), 73–128.
- Peters, Michael, “Market Size and Spatial Growth—Evidence from Germany’s Post-War Population Expulsions,” 2019. Yale University manuscript.
- and Connor Walsh, “Population Growth and Firm Dynamics,” 2021. Yale University manuscript.
- Phelps, Edmund S., “Models of Technical Progress and the Golden Rule of Research,” *Review of Economic Studies*, April 1966, 33 (2), 133–45.
- Ramondo, Natalia, Andres Rodriguez-Clare, and Milagro Saborio-Rodriguez, “Trade, Domestic Frictions, and Scale Effects,” *American Economic Review*, October 2016, 106 (10), 3159–84.
- Redding, Stephen J. and Esteban Rossi-Hansberg, “Quantitative Spatial Economics,” *Annual Review of Economics*, 2017, 9 (1), 21–58.
- Restuccia, Diego and Richard Rogerson, “Policy Distortions and Aggregate Productivity with Heterogeneous Plants,” *Review of Economic Dynamics*, October 2008, 11, 707–720.
- Romer, Paul M., “Endogenous Technological Change,” *Journal of Political Economy*, October 1990, 98 (5), S71–S102.
- Sampson, Thomas, “Dynamic Selection: An Idea Flows Theory of Entry, Trade, and Growth,” *The Quarterly Journal of Economics*, 2016, 131 (1), 315–380.

— , “Technology Gaps, Trade and Income,” July 2020. LSE manuscript.

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

Sterk, Vincent, Petr Sedlacek, and Benjamin Pugsley, “The Nature of Firm Growth,” *American Economic Review*, February 2021, *111* (2), 547–79.

Trimborn, Timo, Karl-Josef Koch, and Thomas M. Steger, “Multi-Dimensional Transitional Dynamics: A Simple Numerical Procedure,” *Macroeconomic Dynamics*, 2008, *12* (3), 301–319.

United Nations, “World Population Prospects 2019,” Technical Report 2019.

U.S. Bureau of Labor Statistics, “Multifactor Productivity,” Technical Report 2021. <https://www.bls.gov/mfp/>.

Young, Alwyn, “Growth without Scale Effects,” *Journal of Political Economy*, February 1998, *106* (1), 41–63.

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