The Outlook for Long-Term Economic Growth

Charles I. Jones*
Stanford GSB and NBER

August 29, 2023
Jackson Hole Symposium

Abstract

What are the prospects for economic growth in the United States and other advanced countries over the next several decades? U.S. growth for the past 150 years has been surprisingly stable at 2% per year. Growth theory reveals that in the long run, growth in living standards is determined by growth in the worldwide number of people searching for ideas. At the same time, a growth accounting exercise for the United States since the 1950s suggests that many other factors have temporarily contributed to growth, including rising educational attainment and a rising investment rate in ideas. But these forces are inherently temporary, implying that growth rates could slow in the future. This prediction is reinforced by declining population growth rates throughout the world. In contrast, other forces could potentially sustain or even increase growth rates. The emergence of countries such as China and India provides large numbers of people who could search for ideas. Improvements in the allocation of talent — for example, the rise of women inventors — and increased automation through artificial intelligence are other potential tailwinds.

*Prepared for the August 2023 Jackson Hole Symposium panel on "Globalization at an Inflection Point." This paper summarizes work reported in Jones (2022).

Figure 1: Real GDP per Person in the United States

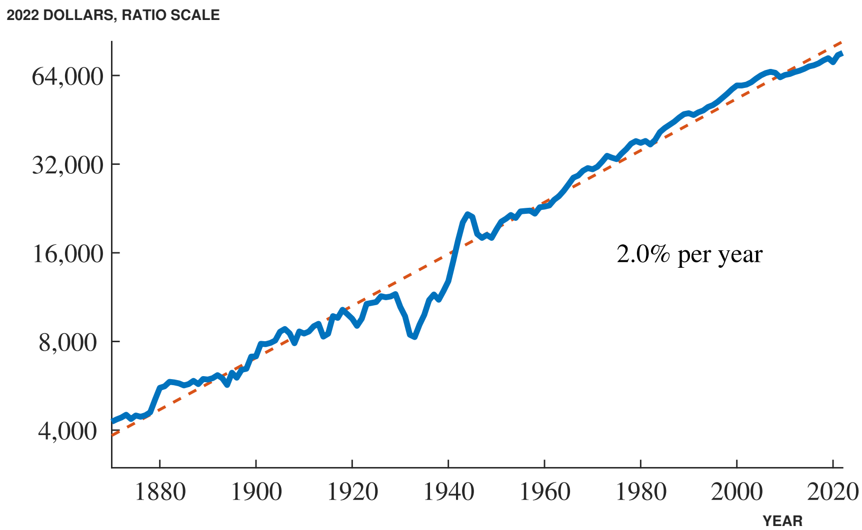

Source: Barro and Ursua (2010) until 1928 and after using the NIPA data from the FRED database (GDPCA and B230RC0A052NBEA).

I. Introduction

What are the prospects for economic growth in the United States and in other advanced countries over the next several decades? Let me begin with one of my favorite charts in economics: the time path of U.S. GDP per person over the past 150 years, shown in Figure 1. The surprising thing about this chart is how well a straight line with a slope of 2% per year fits the logarithm of U.S. living standards. Clearly there are decades with faster than 2% growth as well as decades with slower than 2% growth. But the extent to which the data adhere to a straight line is remarkable.

At some level, then, a simple forecast — which has worked out very well historically — is to project the 2% growth into the future. While the statistical regularity makes this tempting, I'll argue below that the theory of economic growth provides reasons to be cautious about assuming this continuation. There are reasons why growth might be slower as well as reasons why growth might be faster. I'll discuss each of these below in the context of "tailwinds" and "headwinds," after first summarizing the lessons from growth theory.

II. Growth Theory in a Nutshell

Why do living standards tend to rise over time? That is, how do we understand the sustained exponential growth in GDP per person in frontier countries like the United States for more than a century? The key to answering this question was provided by Paul Romer in his 1990 paper (Romer, 1990) and served as the basis for his Nobel Prize in Economics in 2018. The key is that ideas are different from nearly all other goods in economics in that they are *nonrival*, or what I prefer to call *infinitely usable*.

Consider most goods in economics, such as a computer or a barrel of oil or an hour of a surgeon's time. Each of these goods is rival, meaning that it can only be used by one person at a time. If I'm using my computer or an hour of a surgeon's time, those goods are not available for you to use simultaneously. Most goods in economics share this feature, which gives rise to the scarcity that is at the heart of economics.

Ideas, in contrast, are nonrival or infinitely usable. Think about calculus or the latest Covid-19 vaccine or a state-of-the-art machine language algorithm. Each of these are ideas and, once the idea is invented, can potentially be used by any number of people simultaneously: one person's use does not inherently reduce the amount of the idea available for use by others. The Covid-19 vaccines provide an excellent recent example. Once the vaccine has been invented, its design can benefit billions of people without ever having to be reinvented.

The implication is that living standards are tied to the total number of ideas that have ever been invented, not to "ideas per person." The contrast with rival goods like capital is important here. One computer makes one worker more productive. If we want to make a million workers more productive in the same way, we need a million computers. In contrast, one new idea — such as the harnessing of electricity — can make any number of workers more productive.

This gives rise to the following theory of growth:

Income per person* ← *Ideas* ← *People

That is, income per person depends on the total number of ideas ever discovered. Where do these ideas come from? They are invented by researchers, entrepreneurs, and scientists — that is by people. This leads to the somewhat surprising conclusion that living standards depend on the number of people searching for ideas. And therefore the *growth rate* of living standards in the long run depends on the growth rate of the number of people searching for ideas:

Growth in income per person $\leftarrow$ ***Growth in people searching for ideas***

In the long run, the growth in the number of researchers and entrepreneurs is limited by the population growth rate because the share of people doing research must level off. Growth in living standards in countries at the world frontier is ultimately tied to the population growth rate of the countries that produce ideas. A more detailed but still accessible overview of this theory is provided by Jones (2019, 2022).

III. Headwinds

We turn now to five headwinds that might lead to slower growth in frontier countries such as the United States in the coming decades.

1. Growth is already slowing. The first headwind is, of course, that we are already seeing slow growth. Figure 2 shows the level of total factor productivity in the United States since 1990. Two lines are shown, and both suggest slowing growth. For the private business sector as a whole, growth in the years before 2003 averaged 1.1% per year but just over half as much as 0.6% per year since 2003.

One hypothesis often put forward is that many of the so-called "free" goods associated with Google, Facebook, and other tech companies have become increasingly important and are likely not adequately captured in our output measures. A response to this concern is to consider manufacturing, which is traditionally viewed as the sector of the economy that is best measured. The remarkable fact is that the slowdown in TFP growth is much more severe in manufacturing than in the rest of the economy: from 1.6% before 2003 to just 0.4% after. Moreover, there has essentially been zero growth in manufacturing TFP since 2005.

2. Ideas are getting harder to find. The second headwind is a direct implication of the growth theory that we laid out in Section II. In particular, according to this theory, there is an important sense in which ideas are getting harder to find. Now of course we do not have great measures of ideas themselves, and some ideas are huge while others are

Figure 2: U.S. Total Factor Productivity

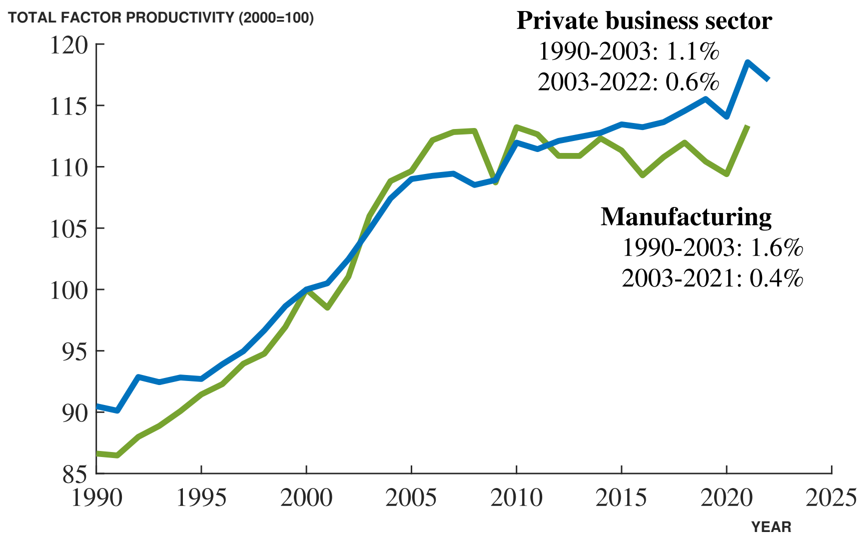

Source: BLS Multifactor Productivity from the FRED database (MFGPROD and MFPPBS).

small. But the sense in which this statement is true is that achieving constant rates of exponential growth requires devoting larger and larger numbers of people to the hunt for ideas.

My favorite example of this phenomenon is Moore's Law, the empirical regularity that the density of computer chips — the number of transistors packed into each CPU — doubled every two years between the 1970s and the 2010s. What makes this such a great example is that Moore's Law is at the heart of the most dynamic sector of modern economies. Through the widespread adoption of computers, smartphones, the internet, and other complementary inventions, Moore's Law is likely responsible for a large fraction of economic growth in recent decades. As is well known, a process that doubles every T years corresponds to constant exponential growth at $70/T$ percent per year. So another way of stating Moore's Law is that the density of computer chips rose at a stable rate of 35% per year for at least half a century.

How was this growth achieved? As documented carefully in Bloom, Jones, Van Reenen and Webb (2020), the global research effort devoted to pushing Moore's Law forward by companies like Intel, AMD, Nvidia, Samsung, TSMC, and ASML — but also historically by IBM, Motorola, AT&T, and Texas Instruments — rose dramatically over this period,

Figure 3: Research Employment in Select Countries

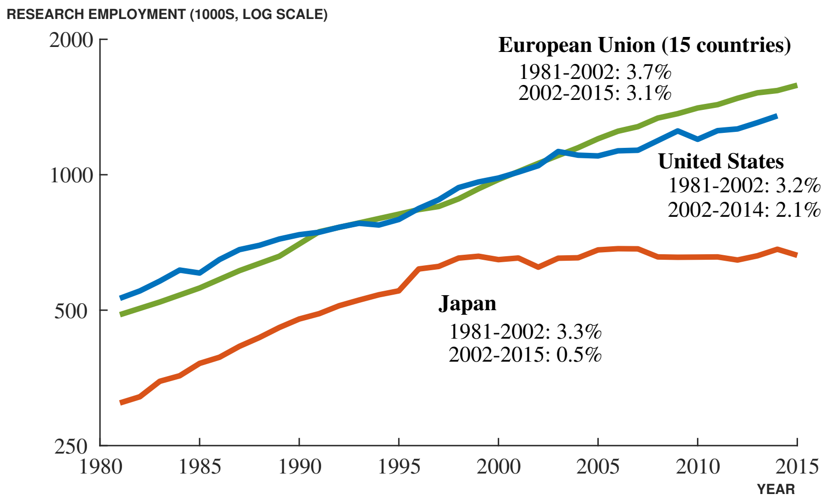

Source: Main Science and Technology Indicators, OECD (2017).

doubling roughly every decade. By the 2010s, it took 18 times more researchers than in the 1970s to generate the doubling of semiconductor chip density. It is the ever-increasing number of people searching for ideas that sustains constant exponential growth.

An immediate implication of this theory is that if the growth rate of those searchers were to decline, exponential growth rates in living standards or in the power of computer chips would also slow. What does the data look like?

Figure 3 shows overall research employment in select countries and regions around the world. Interestingly, in each of these cases, the growth rate of research employment slowed after 2002 relative to the two decades before. In other words, one possible explanation for the slowdown in productivity growth in the U.S. and in other countries is that the growth rate of people searching for ideas also appears to have slowed.

3. Investment rates in infinitely usable ideas have been rising historically. Figure 4 provides another measure of research effort by showing the U.S. investment share of GDP in intellectual property products, which includes privately-funded R&D, publicly-funded R&D, as well as computer software and digital entertainment products like songs and movies. This investment rate in ideas has risen from around 1% of GDP in the

Figure 4: U.S. Investment in Infinitely Usable Ideas

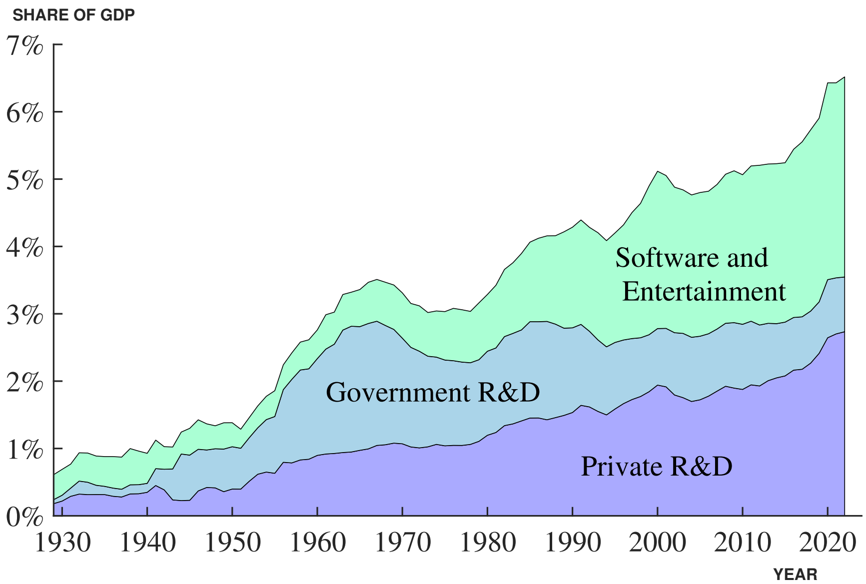

Source: U.S. intellectual property products investment and components, National Income and Product Accounts via FRED.

1930s to more than 6% of GDP in recent years. On the one hand, it is great that we (and other countries) are investing ever larger shares of GDP in these infinitely usable ideas. On the other hand, though, this inherently suggests a future slowdown in growth: our stable growth rate of 2% per year has been achieved while investing an increasing share of GDP in intellectual property products. At some point, this share has to stop rising, and when it does, the implication is that this past source of growth will be unavailable. The growth accounting exercise in the next point quantifies this slowdown.

4. Educational attainment is stagnating. During the 20th century, educational attainment in the United States increased substantially, by just under one year per decade when averaged over the entire labor force (Goldin and Katz, 2008). However, the educational attainment of recent cohorts has risen much more slowly and nearly stagnated, as discussed in Autor, Goldin and Katz (2020).

Figure 5 quantifies this effect using a growth accounting exercise conducted by Jones (2022). With a Mincerian return to education of roughly 6%, each year of educational attainment for the adult labor force should raise labor productivity by around 6%. Be-

Figure 5: Historical Growth Accounting

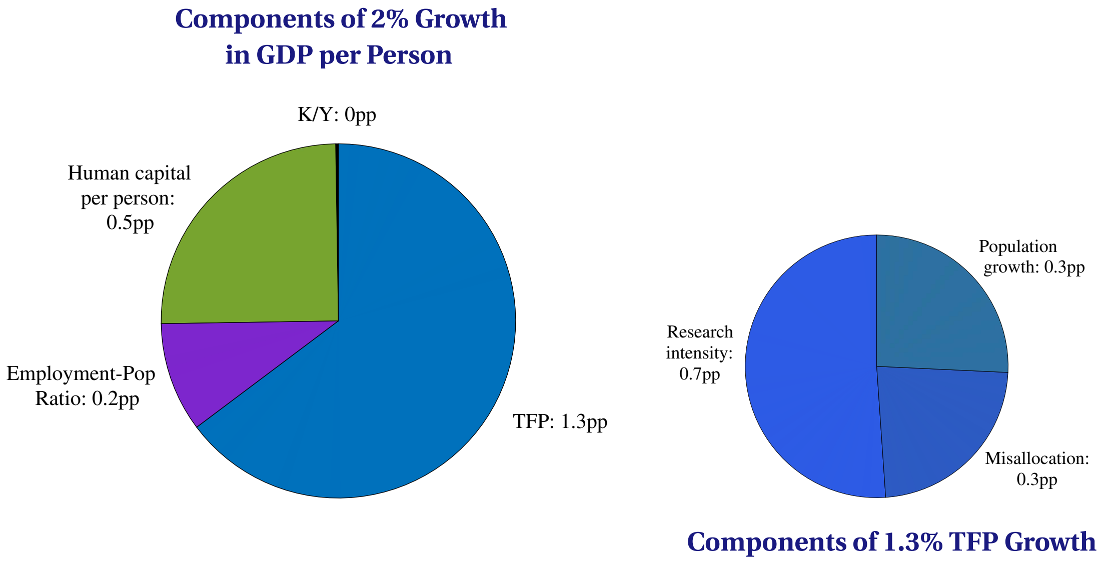

Note: The figure shows a stylized growth accounting exercise for the United States since the 1950s. Source: Jones (2022).

cause this change occurred each decade historically, we divide by 10 to get the annual growth contribution which works out to be about 0.5 percentage points per year, or fully one quarter of the 2% U.S. growth rate. (This is the "human capital per person" slice of the pie chart.) If educational attainment continues to stagnate, the implication is that this half-a-percentage point contribution to economic growth would disappear.

The broader growth accounting in Figure 5 is also helpful for judging the contribution of the rising investment share of intellectual property products. In particular, in this accounting, rising "research intensity" accounts for 0.7 percentage points of the 2% growth. When research intensity levels out at some point in the future, this component of growth would vanish as well. There is more uncertainty in the magnitude of the contribution than is suggested in this chart, but the numbers still give a helpful sense of what could happen.

5. Population growth is slowing and may turn negative. According to the theory discussed in Section II, 100% of growth in living standards must be due to population growth in the long run. Interestingly, that need not be true historically because the "long run" may not yet have been reached. An important finding of the growth accounting exercise shown in Figure 5 is that since the 1950s, only about 15 percent (0.3 percentage points) of growth is due to population growth. Other forces have contributed 85 percent of the total. These include the rise in educational attainment and the rise in research intensity just discussed, but also the rising employment-population ratio and a decline in misallocation (which will be discussed further below). At some level, this is encouraging because it suggests that other economic forces can raise growth rates for more than half a century. However, the implication of growth theory is that these other forces are inherently transitory. The implication is that the long-run component of growth may be something like 0.3%, much slower than the 2% growth we've experienced historically.

In fact, the numbers are even more pessimistic than this because population growth rates themselves are slowing throughout the world. This was driven home to me by the Bricker and Ibbitson (2019) book, *Empty Planet*. These authors observed that fertility rates all around the world have been declining. In fact, for the high-income countries as a whole, the total fertility rate — the number of children women have over their lifetime on average — is now 1.7. Simply to keep populations constant, total fertility rates need to be just over two. In other words, in rich countries as a whole, observed fertility rates are already consistent with a declining population rather than with a growing population. It is distinctly possible that global population will level off and then start to decline over the next century. The implication for growth theory is that living standards could stagnate rather than continue to grow exponentially: if the number of people searching for ideas declines over time rather than rising, economic growth eventually comes to an end (Jones, 2020).

IV. Tailwinds

The preceding discussion of headwinds highlights numerous reasons why growth rates in frontier countries such as the United States might slow in coming decades. However, there are at least three important tailwinds that will push against a slowdown and could possibly even increase growth rates, at least for a while.

1. The rise of China and India. The production of new ideas is a *global* phenomenon: ideas created anywhere have a strong tendency to benefit people all around the world. The "infinite usability" does not stop at national borders. In this context, the economic development of China and India is noteworthy. Each country has a population as large as the United States, Western Europe, and Japan combined. In 1970, these economies were so poor and so far from the world technological frontier that their populations could not meaningfully contribute to producing new ideas that would raise living standards in the U.S. and Europe. However, decades of rapid economic growth mean that this is no longer true, especially in China already but also in India in the future. For example, in 2013–2016, Tsinghua University produced more of the 10 percent most highly cited papers in STEM than any other university in the world (The Economist, 2018). So even if population growth rates are slowing around the world, global research effort could continue to rise in the next several decades as researchers and entrepreneurs in China and India join the search for ideas.

2. Improving the allocation of talent. Sandra Day O'Connor — the first women Supreme Court Justice in the United States — graduated from Stanford Law School in 1952 with the third best academic record in her class. The only job she could get in the private sector was as a legal secretary (Biskupic, 2006). According to Hsieh, Hurst, Jones and Klenow (2019), a similar fact was true more broadly for high-skilled occupations: in 1960, 94 percent of doctors, lawyers, and managers in the United States were white men. By 2010, this share had fallen to 60 percent. Part of what was going on during that fifty-year period was the Sandra Day O'Connor story writ large. The allocation of people was increasingly based on talent and comparative advantage rather than being distorted by various barriers and social norms. Hsieh, Hurst, Jones and Klenow (2019) show that 40% of growth in income per person and 20% of growth in income per worker over these 50 years was due to the improved allocation of talent.

According to Brouillette (2023), 4% of inventors based on patent statistics were women in 1976, rising to 12% by 2020. Bell, Chetty, Jaravel, Petkova and Van Reenen (2019) show that the extent to which people are exposed to inventive careers in childhood has a substantial influence on who becomes an inventor. Exposure in childhood is limited for girls, people of certain races, and people in low-income neighborhoods. So the opportunities to expand the talent for research are not only limited to China and India and other developing countries. How many future Steve Jobs and Jennifer Doudnas are waiting to realize their potential?

3. Artificial intelligence. The final tailwind is perhaps the most uncertain but also has the greatest upside potential. The recent emergence of ChatGPT and other large language models indicates dramatic advances in artificial intelligence. Machines are increasingly able to substitute for humans in various tasks. We've argued that a lack of talented people to search for new ideas is an impediment to future growth. What if machines can replace people in this task as well? Aghion, Jones and Jones (2019) show that in models like those discussed in the first part of this paper, it is at least possible for growth rates to rise if A.I. can partially or fully replace people in generating ideas.

However, that paper also emphasizes various bottlenecks that can limit the extent of these effects. For example, automation has been going on since the Industrial Revolution. The steam engine, electricity, internal combustion engines, tractors, and semiconductors are all examples of amazing new technologies that helped automate various parts of the economy. And yet as shown back in Figure 1, none of these technologies accelerated growth during the past 150 years. The development of a new general purpose technology every few decades may be precisely what kept the 2% trend going for so long. Perhaps A.I. is just the latest amazing technology that will postpone a slowdown for several more decades and permit 2% growth to continue a bit longer.

So while it is conceptually possible for A.I. to raise growth rates, it is far from certain. Theory says it is possible while history gives reasons for caution.

V. Concluding Thoughts

Because ideas are infinitely usable, living standards in any country depend on the total stock of ideas that have ever been invented throughout the world. Ideas are discovered by people, so living standards are tied to the global number of people searching for ideas. In growth rates, this means that the growth rate of living standards in the long run depends on population growth.

Historically, other factors have been important. In the United States since the 1950s, perhaps 85 percent of growth has been due to other factors such as rising investment rates in ideas, rising educational attainment, the increase in the employment-population ratio, and the improvement in the allocation of talent. But each of these forces is inherently temporary: the fraction of GDP devoted to investment in ideas will someday level off and educational attainment is already beginning to stagnate. Moreover, population growth rates themselves are slowing around the world and even potentially turning negative. A long list of headwinds confront future economic growth, suggesting that growth in the next several decades could be slower than in the past half century.

On the other hand, there are tailwinds that could offset these forces. The rise of China and India and the improved allocation of talent throughout the world mean that there are many more people with the potential to become the next Steve Jobs or Jennifer Doudna. Artificial intelligence appears to be a new general purpose technology, perhaps on par or even exceeding electricity and the semiconductor. The widespread application of A.I. could stimulate economic growth in the coming decades, though there are substantial uncertainties around this possibility.

References

Aghion, Philippe, 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.

Autor, David, Claudia Goldin, and Lawrence F. Katz, "Extending the Race between Education and Technology," Working Paper 26705, National Bureau of Economic Research January 2020.

Barro, Robert J. and José F. Ursua, "Barro-Ursua Macroeconomic Data," 2010. Downloaded from https://scholar.harvard.edu/barro/publications/barro-ursua-macroeconomic-data.

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.

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

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, April 2020, *110* (4), 1104–44.

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," 2023. Stanford University manuscript.

Goldin, Claudia and Lawrence F. Katz, *The Race between Education and Technology*, Belknap Press, 2008.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow, "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Jones, Charles I., "Paul Romer: Ideas, Nonrivalry, and Endogenous Growth," *The Scandinavian Journal of Economics*, 2019, *121* (3), 859–883.

___ , "The End of Economic Growth? Unintended Consequences of a Declining Population," 2020. NBER Working Paper No. 26651.

___ , "The Past and Future of Economic Growth: A Semi-Endogenous Perspective," *Annual Review of Economics*, August 2022, *14*, 125–152.

OECD, "Main Science and Technology Indicators," 2017.

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

The Economist, "Tsinghua University may soon top the world league in science research," November 17 2018.