The Productivity Growth Slowdown in Advanced Economies

By Charles I. Jones[135]

Abstract

Perhaps the most remarkable fact about economic growth in recent decades is the slowdown in productivity growth that occurred around the year 2000. This slowdown is global in nature, featuring in many countries throughout the world. In this discussion, I summarize some important characteristics of the slowdown and consider recent insights from the growth literature on its possible causes.

1 Introduction

Many economic problems can be solved by economic growth. Most obviously, economic growth is almost definitionally related to improvements in living standards. Less obviously, problems with government budget constraints – and the issues related to sovereign defaults and even high inflation – may be relieved by economic growth. All of this makes the global slowdown in productivity growth since around the year 2000 one of the most significant economic issues confronting the advanced economies of the world today.

I begin my discussion by summarizing some of the key facts related to this slowdown. Next, I review the insights offered by the growth literature concerning how we might understand slowing productivity growth. Finally, I conclude by offering some thoughts on the prospects for growth in advanced economies in the coming decades.

1.1 Basic Facts

Chart 1 shows the level of total factor productivity (TFP) for the U.S. economy since 1990, first for the private business sector and then for manufacturing. Splitting the sample in half reveals the slowdown. Between 1990 and 2003, TFP growth in the private business sector averaged 1.2 percent per year. Since 2003, however, the growth rate has slowed to 0.7 percent.

---

[135] Graduate School of Business, Stanford University. I am grateful to John Fernald, Loukas Karabarbounis, Pete Klenow, and Ricardo Reis for helpful comments and discussions in preparing this paper.

Chart 1

U.S. Total Factor Productivity

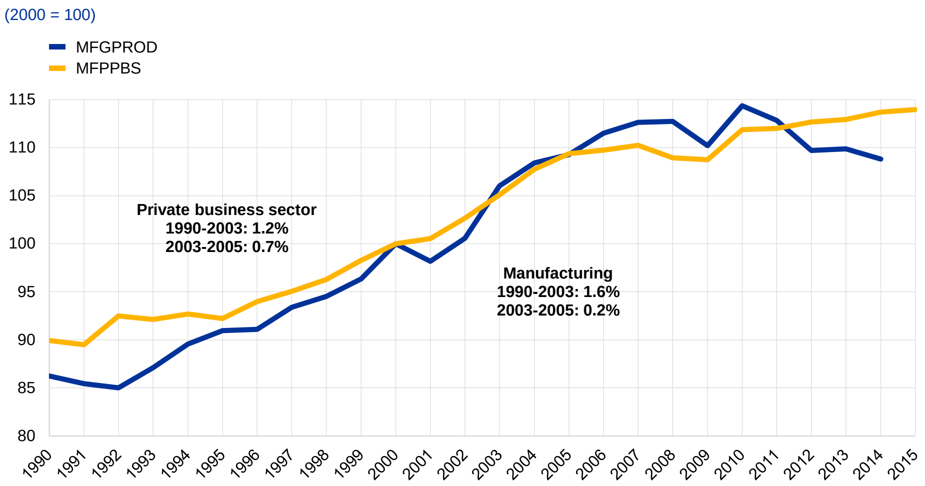

Sources: Bureau of Labor Statistics, “Multifactor Productivity Trends,” downloaded from the Federal Reserve Economic Data database of the FRB St. Louis on June 12, 2017.
Note: Average annual growth rates over various periods are reported in the graph.

A common hypothesis for the slowdown is increasing mismeasurement, particularly associated with the “free” goods provided by IT firms like Google and Facebook. Byrne, Fernald, and Reinsdorf (2016) and Syverson (2017) conclude that the slowdown is so large relative to the importance of the “free” sector that mismeasurement is likely a small part of the explanation. Aghion et al. (2017) use firm-level employment dynamics to assess the mismeasurement of growth due to creative destruction. They find that growth is understated by around 0.5 percentage points per year, but they do not find a substantial change in mismeasurement over time.

The second line in Chart 1 – TFP in the manufacturing sector – reinforces this point. Manufacturing is one of the better-measured sectors in the U.S. economy and a traditional stronghold for productivity growth. Yet this sector exhibits an even sharper slowdown than the aggregate economy: from an annual growth rate of 1.6 percent between 1990 and 2003 to just 0.2 percent since 2003. The growth slowdown, then, appears to be real and economically important.

Chart 2 reports levels of total factor productivity for the U.S., France, Italy, Germany, and Spain using data from the Penn World Tables. These data correspond to the aggregate economy in each country, including the government sector.

Chart 2

Total Factor Productivity in Select Countries

(2000 = 100)

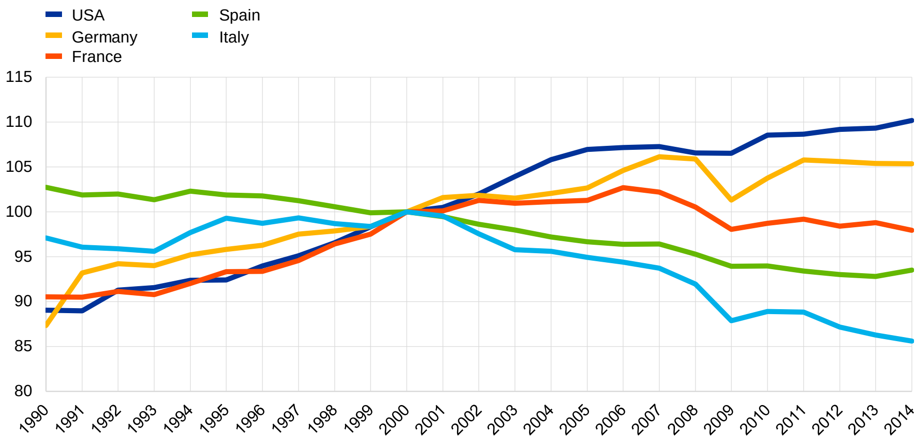

Sources: Author's calculation using the Penn World Tables 9.0.
Notes: Based on the series "rtfpna".

The U.S. productivity slowdown is once again apparent in the chart. More remarkable, however, is that productivity growth since 2000 is even slower in the European economies. TFP is about 5 percent higher in Germany in 2014 than it was in 2000. But TFP levels in the remaining three countries are actually lower than they were in 2000: in France by about 2 percent, in Spain by about 5 percent, and in Italy by more than 10 percent. Therefore, it is not merely a slowdown in productivity growth that needs to be explained: we also need to understand how it is possible for some advanced economies to be substantially less productive today than they were in the year 2000.

1.2 How Growth Theory Can Explain These Facts

Modern growth theory views productivity as being determined by two main economic forces: ideas and misallocation.[136] In the seminal models put forward by Romer (1990) and Aghion and Howitt (1992), productivity growth results from the discovery of new ideas. A slowdown in productivity growth could then occur in two possible ways. We could be moving from a regime in which ideas were relatively easy to discover to one in which ideas are harder to discover, or we could – in some sense to be made more precise below – be investing fewer resources in the search for new ideas.

The second key force impacting productivity that has been highlighted in the growth literature is misallocation. Indeed, the insight that the misallocation of resources at

[136] A third possibility is suggested by Lucas (1988), namely that productivity can be determined by human capital investments. This explanation requires care in that many determinants of human capital are already incorporated into the inputs in the total factor productivity calculations, including formal schooling and demographics associated with age, gender, and occupation. However, human capital accumulation that is "within cell" so to speak – such as on-the-job training or learning-by-doing – could show up as productivity growth.

the micro level leads to declines in TFP at higher levels of aggregation is one of the key insights from the growth literature in the past decade. Restuccia and Rogerson (2008) emphasized this point conceptually, while Hsieh and Klenow (2009) studied its quantitative importance in accounting for TFP differences between the U.S., China, and India. Applied to the question at hand, a slowdown in TFP growth – or even a decline in TFP – could result from a systematic increase in the misallocation of resources at the microeconomic level.

The next two sections explore these possibilities in more detail.

2 Ideas and the Productivity Slowdown

At the heart of virtually all models of economic growth based on the discovery of new ideas is an equation like the following:

Economic growth	=	Idea TFP	×	Number of researchers

That is, the growth rate of some economic entity (a firm, a product, or even the economy itself) arises through the discovery of new ideas. These ideas in turn are produced by research effort ("Number of Researchers") multiplied by their research productivity ("Idea TFP").

My latest research project – Bloom, Jones, Van Reenen, and Webb (2017) – studies this equation empirically at various levels of aggregation. For example, we've looked at the aggregate economy, at Moore's Law for the density of computer chips, at the productivity of various types of agriculture, at a range of medical innovations, and at firm-level data in Compustat. What we find everywhere is that growth rates are relatively stable, sometimes rising slightly, sometimes falling slightly, sometimes constant. But research effort increases dramatically. The implication is that research productivity – the TFP in the idea production function – is falling dramatically.

Two examples illustrate this point. First, for the aggregate U.S. economy, it is well known that growth rates of GDP per person are relatively stable over time; if anything, growth has slowed a bit recently. In contrast, research effort has risen by more than a factor of 20 since the 1930s. The implication is that Idea TFP has fallen by more than a factor of 20: it is more than 20 times harder today to generate exponential growth in the U.S. economy than it was in the 1930s.

A similar finding occurs with respect to computer chip density and Moore's Law. Recall that Moore's Law is the stylized fact that the density of transistors on computer chips doubles every two years, and this doubling time has been remarkably stable back to the 1970s. A constant doubling time of course corresponds to constant exponential growth, in this case at a rate of about 35 percent per year. How has this growth been achieved? Bloom et al. (2017) measure the research input in a variety of ways, but no matter how we do it, the finding is the same: the research required today to double computer chip density is much higher – perhaps by a factor of 15 or more – than it was in the 1970s. Once again, exponential growth is getting harder to achieve over time, and the way it is achieved is by committing ever increasing quantities of research effort to the endeavor.

The conclusion of that paper is that ideas are systematically getting harder to find, a conclusion reminiscent of Gordon (2016). Interestingly, however, the analysis suggests that this has always been true! In other words, there is nothing in our work arguing that anything has changed over time. It has always been getting harder and harder to find new ideas, and the exponential growth we see is the outcome of throwing more resources into research. So this particular story does not seem to be able to account for slowing productivity growth in advanced economies.

It does, though, suggest a possible alternative explanation. One model consistent with these facts might be called a “Red Queen” model of economic growth, after a character in Lewis Carroll’s Through the Looking Glass: “Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!” Constant exponential growth in research effort offsets its declining productivity. Continued growth in research inputs is required to generate exponential growth in the economy, and if the growth rate of research effort were to slow, the outcome would be a slower rate of exponential productivity growth.[137]

The next two graphs examine how resources devoted to innovation have changed over time. Chart 3 takes advantage of the recent revisions to the U.S. National Income and Product Accounts to show investment in “intellectual property products” (IPP) as a share of GDP. IPP includes public and private investment in research and development but also incorporates spending on other nonrival goods, including computer software, books, music, and movies.

[137] See Jones (1995) and Kortum (1997) for models along these lines. This literature is reviewed in Jones (2005).

Chart 3

U.S. Intellectual Property Products Investment

(share of GDP)

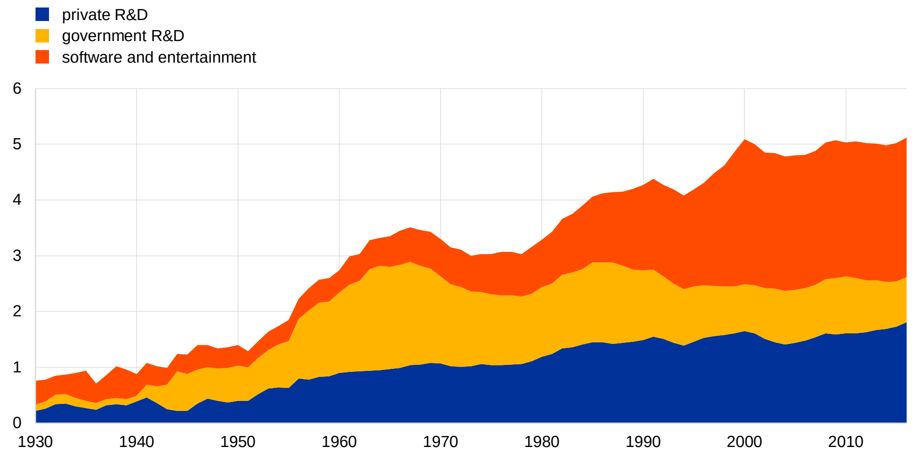

Sources: U.S. Bureau of Economic Analysis via FRED.

Several important facts are evident in Chart 3. First, there was a general rise in R&D spending as a share of GDP up until the mid-1960s. The rise in private R&D has continued though the present time, but public R&D as a share of GDP has been declining in general since the first astronaut set foot on the moon. Since 1980, investment in software and entertainment products has risen sharply. Taken as a whole, investment in intellectual property products has generally been rising over time as a share of GDP. However, since about 2000, the investment rate has been more stable.

The evidence from this chart about the Red Queen model is somewhat mixed. Private R&D appears to be growing in intensity. Private plus public R&D shows no trend or change in trend since around 1980. Finally, the total IPP investment rate does appear to have changed behavior in the last two decades, stabilizing rather than continuing to rise.

It is important to appreciate, of course, that the search for new ideas is a global phenomenon. The United States benefits from ideas created all over the world, as do all other countries. Chart 4 shows research employment in the European Union, Japan, and in the United States since 1981 and reveals an interesting finding.

Chart 4

Research Employment in the E.U., the U.S., and Japan

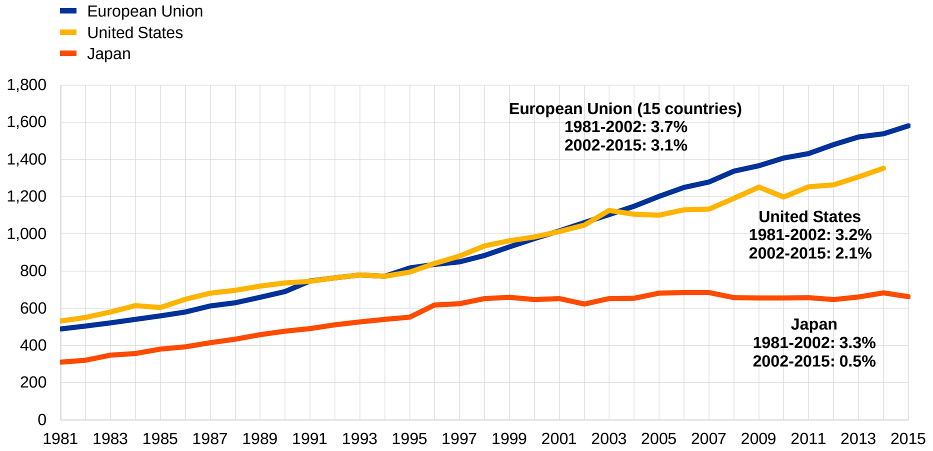

Sources: OECD Main Science and Technology Indicators.
Note: Average annual growth rates over various periods are reported in the graph.

All three measures of research employment show a slowdown in research effort. After growing at 3.7 percent per year between 1981 and 2002, for example, research employment in the E.U. grew at only 3.1 percent between 2002 and 2015. The slowdown in research in the United States is even larger: from an average growth rate of 3.2 percent before 2002 to just 2.1 percent after. Finally, the slowdown in Japan is especially noteworthy. Research employment grew at 3.3 percent per year before 2002 but just 0.5 percent after. Remarkably, research employment in Japan has been relatively stable since the late 1990s. According to the logic of the Red Queen model, this slowdown in the growth rate of research effort should translate into a slowdown in productivity growth. Slowing research effort, then, is a possible explanation for the productivity slowdown. This hypothesis should be explored further in future work.

3 Misallocation

A key contribution of the growth literature of the past decade is a renewed appreciation of the importance of the misallocation of resources in the determination of aggregate TFP. The level of TFP is simply a measure of how successful a given basket of aggregate inputs, including capital, labor, and human capital, are in production. When resources are misallocated at the microeconomic level, it is easy to show that this will reduce TFP at a more aggregated level; for example, see Restuccia and Rogerson (2008) and Hsieh and Klenow (2009) for precise statements along these lines. This misallocation could occur across sectors, across firms within a sector, or even within a firm itself. For example, a housing bubble will provide agents in the economy with mistaken price signals that will lead to too many housing being built at the expense of other economic activity. This misallocation across sectors would reduce aggregate TFP. Alternatively, privileged access to credit by some firms at the expense of others could result in the misallocation of resources within a narrowly defined sector. This would reduce the TFP of that sector and therefore also contribute to lower aggregate TFP. Finally, poor management practices could result in the misallocation of resources within a firm. In some ways, this is a natural explanation for TFP differences across firms: as has been long appreciated, firms are in some respects like centrally planned economies, and the absence of markets inside firms can make it particularly difficult to allocate resources optimally. Good management practices contribute to good allocations, while bad management practices lead to misallocation and reduce firm-level TFP.

Misallocation is a particularly appealing hypothesis for explaining the decline in aggregate TFP that we've seen in economies like Italy and Spain. It is not plausible that these economies are forgetting knowledge, so the idea-based models have a hard time explaining a decline in TFP. But an increased misallocation of resources – perhaps because of the housing bubble and frictions associated with the financial crisis – offers a possible explanation. And of course an increase in misallocation in the United States could contribute to slowing productivity growth there as well.

A number of papers have attempted to measure the macroeconomic consequences of misallocation. These papers build on a simple insight. In particular, the efficient allocation of resources requires that the value of the marginal product of labor, for example, be equated across firms. Any deviation from this equality implies that reallocation can improve productivity by moving labor from a place where it has a low marginal product to a place where it has a high marginal product. Similar statements apply to capital or intermediate inputs or any other input. Therefore, variation in the marginal revenue product of an input across firms is one summary statistic commonly reported in the literature on misallocation. Moreover – and relying on important structural assumptions such as the precise form of production functions – one can calculate the implied gain in output that could be achieved by eliminating the variation in marginal revenue products across firms or plants. With this approach, Hsieh and Klenow (2009) argue that total factor productivity in the manufacturing sector in China and India would be about 50 percent higher if the marginal revenue products of capital and labor were equated across firms within relatively narrow 4-digit industries.

More recently, research has turned to measuring the extent to which misallocation has changed over time. Reis (2013) noted that productivity growth slowed sharply in Portugal and several other economies of southern Europe after they joined the euro. He proposed that the misallocation of the large capital inflows that followed the adoption of the common currency could be part of the explanation. This hypothesis has been explored in more detail using micro data in recent research. Gopinath, Kalemli-Ozcan, Karabarbounis, and Villegas-Sanchez (2017) use the Orbis database to study misallocation in the manufacturing sector in Spain and several other countries of southern Europe. They show that the dispersion in the marginal revenue product of capital increased significantly in Spain between 1999 and 2012, potentially accounting for a 7-12 percent decline in manufacturing TFP relative to trend. With less complete data, they document similar results for Italy and Portugal. Interestingly, they do not find a substantial increase in misallocation in France, Germany, and Norway.

Other authors with access to administrative data have now gone further and studied misallocation across firms throughout the economy, not just in manufacturing. These papers provide fairly consistent support for the misallocation hypothesis as an important factor in explaining slow and even declining TFP in southern Europe.

Bils, Klenow, and Ruane (2017) examine changing misallocation within the manufacturing industries in the United States and document two important findings. The first is perhaps surprising: applying the basic techniques of Hsieh and Klenow (2009) to U.S. manufacturing over time, they find a large increase in misallocation between 1978 and 2007. In particular, their basic calculation suggests that U.S. manufacturing efficiency was only 2/3 of its maximum level in 1978 and that this fell to just 1/3 by the mid-2000s. Taken at face value, this implies that rising misallocation could be an important contributor to the slowdown in manufacturing productivity growth.

Their second finding, however, casts doubt on their first. In particular, it has long been appreciated that measurement error is a significant problem for judging the importance of misallocation. The key contribution of Bils, Klenow, and Ruane (2017) is to provide a statistical technique that controls for measurement error. In particular, they argue that additive measurement error that is constant can be “differenced out” by looking at changes in inputs and outputs over time in order to compute marginal products. They use this basic insight to show how even time-varying measurement error can be controlled for, at least when it is of an additive form. Applying this correction to U.S. manufacturing, they suggest that allocative efficiency is roughly 80 percent of its optimal level and, importantly, that the degree of misallocation is relatively stable over time. The implication is that measurement error in U.S. manufacturing data has been rising over time.

Several important directions for future work are implied. First, what accounts for the increase in measurement error over time in the U.S. data? Second, given that measurement error is obviously important, are there techniques that can control for multiplicative measurement error (which is arguably more natural) and not just additive measurement error? Finally, do the results in the various European countries look different if one adjusts for measurement error?

4 Future Prospects

What about the prospects for productivity growth in advanced countries in the coming decades? I offer three observations.

First, if one of the explanations for slowing productivity growth at the frontier is a slowdown in the growth rate of research inputs, then a natural question to ask is “Why?”. What economic forces are behind the slowdown in research, and can these be reversed in the future? I do not know the answer to this question, but it is clearly important. One observation, based on Chart 3, is that public funding for basic research has declined significantly when expressed as a share of GDP.

Understanding the role that this may have played in the latest productivity growth slowdown strikes me as quite important.

Second, I already observed that the search for new ideas is a global phenomenon. Because of the diffusion of technologies, ideas invented anywhere in the world ultimately benefit people in countries that are open to the international flow of ideas. Along these lines, the rise of research in China and India as these economies develop may be central in driving the future growth rate of global research effort. Each of these economies by themselves has as many people as the United States, Western Europe, and Japan combined. How many Edisons and Einsteins have we missed in the previous century because of the lack of economic development in China and India? And how many more will they contribute in coming decades as their development proceeds?

Finally, to the extent that increases in misallocation have reduced total factor productivity in some of the countries in Europe, it should be appreciated that this is a “level effect” and not a “growth effect.” In other words, productivity growth will surely resume in these countries eventually. Nevertheless, once that resumption occurs, we should not stand by and be satisfied. If misallocation permanently reduces the level of productivity relative to trend by 15 percent, the damage is done. Understanding the economic forces that have caused such a loss is critical so that the misallocation can be reversed.

References

Aghion, Philippe, Bergeaud, Antonin, Boppart, Timo, Klenow, Peter, and Li, Huiyu (2017), “Missing Growth from Creative Destruction”, Stanford University manuscript.

Aghion, Philippe and Howitt, Peter (1992), “A model of growth through creative destruction”, *Econometrica*, Vol. 60, Issue 2, pp. 323-351.

Bartelsman, Eric, Lopez-Garcia, Paloma and Presidente, Giorgio (2017), “Factor Reallocation in Europe”, Tinbergen Institute manuscript.

Bils, Mark, Klenow, Peter J., and Ruane, Cian (2017), “Misallocation or Mismeasurement”, Stanford University manuscript.

Bloom, Nicholas, Jones, Charles I., Van Reenen, John, and Webb, Michael (2017), “Are Ideas Getting Harder to Find?”, Stanford University manuscript.

Borio, Claudio, Kharroubi, Enisse, Upper, Christine and Zampolli, Fabrizio (2016), “Labour reallocation and productivity dynamics: financial causes, real consequences”, *BIS Working Papers* 534, Bank for International Settlements.

Byrne, David M., Fernald, John G., and Reinsdorf, Marshall B. (2016), “Does the United States Have a Productivity Slowdown or a Measurement Problem?”, *Brookings Papers on Economic Activity*, Vol. 47(1 (Spring), pp. 109-182.

Calligaris, Sara. (2015), “Misallocation and Total Factor Productivity in Italy: Evidence from Firm-Level Data”, *Labour*, Vol. 29, pp. 367-393.

Cette, Gilbert, Fernald, John and Mojon, Benoît, (2016), “The pre-Great Recession slowdown in productivity”, *European Economic Review*, Vol. 88(C), pp. 3-20.

Dias, Daniel A., Robalo Marques, Carlos, and Richmond, Christine (2016), “Misallocation and Productivity in the Lead Up to the Eurozone Crisis”, *Journal of Macroeconomics*, Volume 49, pp. 46-70.

García-Santana, Manuel, Moral-Benito, Enrique, Pijoan-Mas, Josep and Ramos, Roberto (2016), “Growing Like Spain: 1995-2007”, *Banco de Espana Working Paper* No 1609.

Gopinath, Gita, Kalemli-Ozcan, Sebnem, Karabarbounis, Loukas, and Villegas-Sanchez, Carolina (2017), “Capital Allocation and Productivity in South Europe”, *Quarterly Journal of Economics*, forthcoming.

Gordon, Robert J. (2016), *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*, Princeton University Press.

Jones, Charles I. (1995), “R&D-based models of economic growth”, *Journal of Political Economy*, Vol. 103, Issue 4, pp. 759-784.

Jones, Charles I. (2005), “Growth and Ideas”, *Handbook of Economic Growth, Volume* 1B, pp. 1063-1111.

Kortum, Samuel S. (1997), “Research, Patenting, and Technological Change”, *Econometrica* Vol. 65, pp. 1389-1419.

Lucas, Robert E. Jr. (1988), “On the Mechanics of Economic Development”, *Journal of Monetary Economics* Vol. 22 (July), pp. 3-42.

Reis, Ricardo (2013), “The Portugese Slump and Crash and the Euro Crisis”, *Brookings Papers on Economic Activity*, Vol. 44, Issue 1 (Spring), pp. 143-210.

de Ridder, Marteen (2017), “Investment in Productivity and the Long-Run Effect of Financial Crises on Output”, University of Cambridge manuscript.

Romer, Paul M. (1990), “Endogenous Technological Change”, *Journal of Political Economy* Vol. 98, pp. S71-S102.

Syverson, Chad (2017) “Challenges to Mismeasurement Explanations for the US Productivity Slowdown”, *Journal of Economic Perspectives*, 31(2): 165-86.
DOI: 10.1257/jep.31.2.165.