THREE

The Economy Is about Firms

Productivity Slowdown and Divergence

DANY BAHAR

It would be impossible to fully understand the global macroeconomic trends on productivity—in particular the marked slowdown across all sectors—without looking in detail at what has been happening at the firm level. After all, economic growth in a country is a reflection of the growth and, in turn, the productivity dynamics among its firms. In what follows, the analysis examines firm-level empirical evidence on productivity and links it to the overall productivity slowdown.

The productivity of a firm reflects how efficiently it can convert input into output. For example, imagine an experiment in which two football manufacturing plants are exactly the same in terms of their inputs: same workers, same equipment, and same amount of leather to be used for making the balls. They should be able to make the same number of balls, shouldn't they? In practice, their output might be quite different due to differences in productivity. If one of these plants is able to produce more footballs than the other one, then the former is more productive than the latter. This is what economists call Total Factor Productivity: what firms produce after taking into account the factors of production used as inputs

This analysis is based on research by Bahar (2018).

in the process, such as labor, machinery, and raw materials. Economists calculate the productivity of a firm by employing the "residual" method: TFP is the difference between the actual output of a firm and its expected output given its inputs. In other words, it is the portion of a firm's output that cannot be explained by its inputs. Moses Abramovitz wrote in 1956 that productivity is a "measure of our own ignorance."[1] In a more intuitive way, a firm is more productive when it adopts new technologies, often embedded in new machines or in experienced workers and managers, that allow the firm to do more with the same resources. As firms become more productive they can produce at lower cost, sell at lower prices, export to foreign markets, gain more market share, and grow in size and in sales.

The slowdown in productivity growth has been recognized by many researchers. In the United States the manufacturing sector experienced a considerable slowdown between 2005 and 2015 as compared to the previous decade. Average annual firm-level TFP growth fell from 2.2 percent in 1995–2004 to 0.4 percent in the following decade (Syverson 2016). Within the manufacturing sector, the hardest hit was taken by firms manufacturing computers and electronics, whose TFP annual growth rate dropped from 10.7 percent during 1995 to 2004 to 3.7 percent from 2005 to 2014. It is argued that the fast-paced productivity growth of firms in the computer and electronics sector during 1995 to 2004—and more broadly, all ICT-using and ICT-producing sectors—is behind the rise in aggregate productivity for the United States during that same period. Yet, Japan and non-English-speaking European countries seem not to have benefited to the same degree from the vast innovation coming out of this sector, as their overall productivity growth did not accelerate during that period and also suffered from a slowdown in the following decade. In fact, slowdowns in productivity are observed in the United States and in twenty-four of the twenty-nine countries in the OECD (Syverson 2016), tracing to before the global recession in 2008, which is also consistent with findings in other studies (see Cette, Fernald, and Mojon 2016).

The post-2005 slowdown, however, goes beyond the manufacturing sector. In the United States, it also happened in other sectors, such as retail and wholesale, as well as the service economy as a whole. A striking pattern seen in the data is that industries that experienced fast-paced productivity growth in the 1995–2004 decade typically slowed down considerably in the following decade, as indicated in figure 3-1 (Baily and Montalbano

FIGURE 3-1 Changes in TFP Growth for Acceleration and Slowdown, Major Sectors

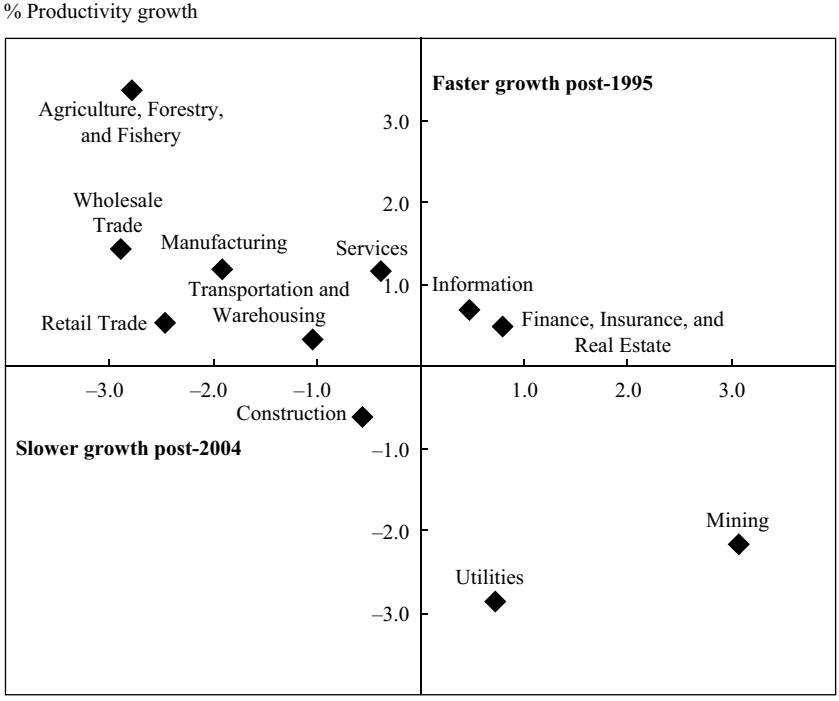

Source: Baily and Montalbano (2016).

2016). Besides firms in the construction, information, financial, insurance, and real estate (FIRE) sectors, which maintained a similar growth rate in both decades, all other sectors "reversed" their growth pace.

The pattern is consistent across many other countries, too. Bahar (2018) explores in detail the evolution of productivity of firms in a global data set and finds a number of interesting facts regarding the distribution of firms' TFP for years 2006 and 2014 by sector. Based on this analysis, two interesting facts arise, visualized in figure 3-2. First, median TFP in year 2014 is slightly lower than in year 2006, across all industries. In addition, dispersion in TFP has increased for most sectors, if not all.

Note that these facts are consistent across all industries, including services, which is frequently neglected because of lack of data. The importance of looking beyond manufacturing, however, is crucial when thinking about

FIGURE 3-2 TFP Distribution by Sector, 2006 and 2014

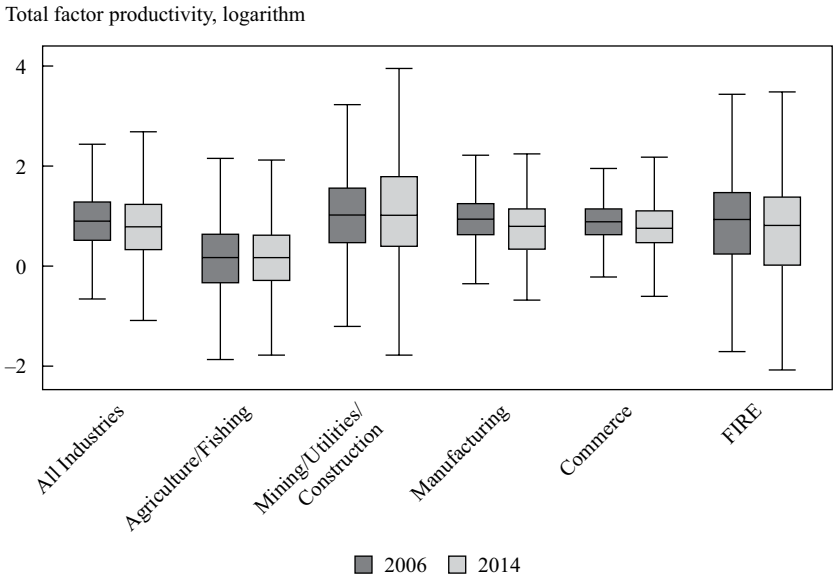

Source: Bahar (2018).

overall productivity growth in large economies such as the United States, Japan, and European nations. The service-oriented nature of these economies could present an important challenge for future growth. This is because the established unconditional convergence in the manufacturing sector does not necessarily hold in the service sector, for which international competition is less of an issue (Rodrik 2011). In fact, some economists have suggested that innovations in services are less relevant in producing dramatic changes in productivity and efficiency, claiming that the productivity of an artist, for example, is not very different today than centuries ago (Baumol and Bowen 1966).

But this view can be challenged. Take, for example, sports. There are shocks in productivity that might not represent the long-run trend. The Jamaican runner Usain Bolt won the gold medal in three consecutive Olympic games (2008, 2012, and 2016) for the hundred-meter run, each time with a slightly different performance (his 2012 time of 9.63 seconds remains his best, as well as the Olympic record). When looking at average performance of runners throughout the 1900s, the improvement has been

dramatic. Before Jim Haines (United States), no Olympic athlete had run one hundred meters in less than ten seconds. In fact, if Bolt and Jesse Owens had raced together, Owens would still have had fourteen feet to go by the time Bolt arrived at the finish line.[2] These improvements probably have little to do with the fact that all humans are faster today than they were decades ago. It is, in fact, a consequence of athletes having improved their training techniques, as well as improvements in technologies that provide them inputs that fuel productivity, such as special clothing, nutrition, or improvement measurement precision.

These same ideas apply to all firms in the service sector that could innovate and use technology-embedded inputs that would make them more productive. The fast-food industry, for example, underwent a number of improvements in the past decades that allowed it to significantly reduce the time between ordering and serving food, incidentally reducing costs and human mistakes. In the retail sector, for example, stores have become more productive by innovating in providing tailored customer service by using data and expanding their platforms online. Innovation and the adoption of technology can improve the way service firms, as well as manufacturing ones, deliver to their customers.

A Framework to Understand Productivity Slowdown

In a nutshell, there are two components in the dynamics linking firm productivity to overall economic growth. First is the improvement in productivity for each firm in the economy over time, known as the "within" component. The second is the growth in size of the most productive firms relative to the least productive ones, known as the "reallocation" component. The reallocation component, in fact, reflects the process through which least productive firms shed labor and other resources—either because they exit the market or simply become smaller—toward the most productive ones. On the aggregate, the speed at which these two processes occur is the key factor that differentiates fast-growing countries from slow-growing ones (McMillan, Rodrik, and Verduzco-Gallo 2014).

Historically, the contribution of each component—within and reallocation—has been different depending on the period and the industry under consideration. A Brookings paper (Baily, Hulten, and Campbell 1992)

studying productivity growth in the United States manufacturing sector shows that, overall, the reallocation effect has been positive and significant during every five-year period from 1972 to 1987. That is, a big chunk of productivity growth in the United States during that period can be attributed to the fact that the most productive firms took over a larger portion of the market share in the overall economy. The within component, however, did not always contribute to overall productivity growth. During the period from 1972 to 1977, for example, manufacturing firms became more productive, but this growth was driven exclusively by firms in the computer and the automobile industry. During the period from 1977 to 1982, firms experienced a decrease in their productivity that was compensated by the reallocation effect. Finally, during the period from 1982 to 1987, the within effect dominated overall productivity growth in manufacturing, and overall productivity was fueled by the reallocation component.

All in all, these two components are essential for overall productivity growth, and they all could play a role in explaining the productivity slowdown of the past decade. This framework is important to understand what could be behind the decreasing productivity growth. Even if some or even most firms experience improvements, overall productivity growth might suffer if workers and other resources flow from the most toward the least productive firms. This is, for example, what happened in Latin American during the period from 1990–2005, when in spite of productivity increases among active industries, overall growth was below potential given flows of workers to least productive industries, often in the informal sector (Pagés 2010). McMillan, Rodrik, and Verduzco-Gallo (2014) expand this decomposition during the same period for other regions and shows a similar case for African nations, whereas countries in Asia experienced unusually high productivity growth of almost 4 percent a year, due both to their industries becoming more productive and the reallocation of resources toward these industries.

In the context of the recent productivity slowdown, all components might play a role. If dynamism in the economy is hurt, then the reallocation and entry/exit components could hinder overall growth. This can be a result of firms not responding effectively to changes in their idiosyncratic productivity. If firms that are least productive are less likely to exit, or if more productive firms fail to attract resources from less productive ones, for example, then overall growth slows down.

Slowdown in Reallocation?

The historic importance of reallocation of labor and the role of entry and exit has been established by many economists. Yet, evidence suggests that dynamism in the United States has been declining in recent decades and, therefore, could play a part in the slowdown that started since the early 2000s. Firms have been less responsive to changes in their individual productivity levels, and this has significant implications for overall growth.

Slowdown in reallocation can play both a positive and a negative role on overall growth. In general, in sectors where it is more likely that any given firm can grow to dominate a big part of the market, dynamism plays an important role. A successful technology start-up with high potential to grow would require resources to flow toward it. A decrease of dynamism in this case would imply that resources aren't flowing toward firms with the highest potential to grow. On the other hand, once resources have moved to the fastest-growing firms, declining dynamism might contribute to overall growth, as happened with the retail sector in the United States during the 1990s and early 2000s. After having already employed a large share of the industry-wide workforce—most likely flowing from small and unproductive mom-and-pop stores—big-box stores kept a fast pace in their productivity growth, contributing to overall industry growth. Thus, reallocation is particularly important in industries where small firms can grow very fast by pioneering innovations, and compete with even the largest firms, such as ICT-producing or ICT-using sectors.

In the United States, declining dynamism in the high-tech industries since the 2000s can explain a significant loss in annual growth up to 2010 (Decker and others 2018). This decline in dynamism is a result of the inability of firms to respond to changes in their productivity, and this could be explained by frictions or high adjustment costs. For example, unnecessary subsidies or high closing costs for a failing enterprise would keep such firms in the market longer, occupying resources that, ideally, could be reallocated to more productive firms in the same sector. Similarly, inflexible labor markets could impede fast-growing firms from hiring more workers when needed and, thus, keep them from responding positively to productivity improvements.

Declining trends of job reallocation in the United States have been common across all industries since the early 2000s, as can be seen in figure 3-3.

FIGURE 3-3 Sectoral Trends in Job Reallocation

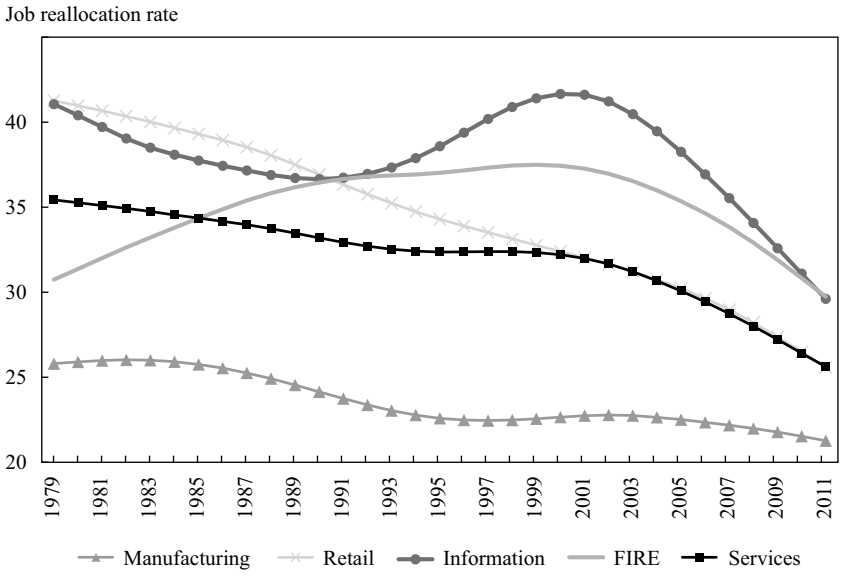

Source: Decker and others (2018).

Both the information and the FIRE industries have experienced a stable, even rising, rate in job reallocation since the 1980s until the early 2000s, after which the rate sharply declined. This post-2000 decline is consistent with the productivity slowdown seen across the economy.

In addition, data shows that young firms in the United States (less than five years old), across all sectors, employ a smaller share of the economy than they did in the early 1980s. If over the past three decades most of the productivity growth had concentrated more among mature firms—rather than the younger firms—then the shifting of resources from young to mature firms would be, in fact, optimal for overall productivity growth. This was, as discussed, the case for retail trade, but not for the other industries. The shifting of resources away from small firms goes hand-in-hand with the productivity slowdown. In fact, employment growth among firms that experienced improvements in productivity has weakened today as compared to the 1980s, both for young and mature firms. This pattern is consistent both for high technology and other plants in the manufacturing sector in the United States.

How important is reallocation in explaining the overall slowdown in productivity as compared to other components? The short answer: not much. A recent paper by Chang-Tai Hsieh and Peter Klenow shows that the decline in dynamism can explain up to 10 percent of the decline in productivity growth in the United States (Hsieh and Klenow 2018). The biggest chunk of what explains a loss in aggregate productivity is, then, the within component.

In fact, as shown by Bahar (2018) and visualized in figure 3-2, firms—on average—reduced their productivity between 2006 and 2014. Yet, firms becoming less productive, on the aggregate, is a difficult concept to digest. Are firms now doing less than they could do with the same resources than before? Some industries are. In the United States alone, for instance, during the 2004–14 period industries such as apparel and leather products, paper products, chemicals, plastics, as well as furniture, among others, experienced negative productivity growth (Baily and Montalbano (2016). Figure 3-2 reflects changes in productivity between 2006 and 2014, just around the Great Recession. Thus, part of this negative growth can be explained as a drop in demand for the industry as a whole, which resulted in reduced sales without immediate changes in the resources the firm employed. This will result in a productivity drop, following the standard measurement techniques. Yet, as noted, the slowdown in productivity preceded the recession in 2007 and, therefore, even if only suggestive, these results support the idea that, even during a period of crisis, reallocation played a positive role. The analysis that follows focuses on understanding firm-level productivity dynamics.

Innovation and Adoption

Simply put, improvements to the productivity of a single firm can be explained in one of two ways: innovation or adoption. Innovation implies the creation of a new and unique method, idea, or product that allows the firm to create more output using the same amount of resources. Adoption, on the other hand, implies that a firm gains access to methods, ideas, or products that were invented by other firms (normally within the same sector) to be able to do more with the same inputs. In fact, innovation is typically done by firms at the frontier who invest large amounts of their budgets in

research and development (R&D) activities, while other, smaller firms grow due to access to previous innovation by other firms in the industry. Both components are crucial; in the absence of innovation there is no productivity growth, and in the absence of adoption not only would most firms not grow (hindering overall productivity growth), there would be less incentive to innovate for frontier firms, given the lack of competition.

The importance of adoption in the process of productivity growth is critical. Even if technologies exist and are available in the country, the process of aggregate growth requires the eventual diffusion of these technologies widely across firms and not only the ones that invented them or adopted them first. An example of an adoption of an already existing technology is the implementation of a customer relationship system that allows the firm to be more efficient in the management of customers, suppliers, and inventory and brings with it the ability to produce more output with the same resources.

The historic trend of both the availability of technologies and their penetration is quite striking. Diego Comin and Martí Mestieri Ferrer, two economists studying the historical diffusion of technologies, find that while diffusion of technologies across countries (what they call the extensive margin) is much faster than in the past, the penetration (the intensive margin) of those technologies within the country has slowed down (Comin and Ferrer, 2013). For instance, it took, on average, forty-five years for the telegraph to reach all countries in the world after it was invented in the 1830s. On the other hand, a newer technology, such as cell phones, took on average only five years to reach all countries after it was invented in the early 1970s. The speed of penetration within other countries relative to Western nations was, however, significantly faster for the telegraph than for the cell phone. In short, as compared to the past century, newer technologies diffuse faster across countries but much slower within countries. This implies that frictions for the adoption of technologies by firms have increased, which would have a direct result in the dispersion of productivity within industries.

Dispersion is precisely what is also documented in figure 3-2. Between 2006 and 2014, not only the median productivity declined but the dispersion in TFP increased as well, across all sectors (see Bahar 2018 for more details on this). This could be a result of a number of developments. First, the weakening of business dynamism; in a highly competitive environment,

firms would be forced by the market to follow "up or out" dynamics. That is, small firms that don't grow because they have been proven to be unproductive would not be able to remain in the market. But in an environment with weak business dynamism, small unproductive firms will remain in the market longer, drawing down the average and increasing the dispersion. Second, the mix of continuous innovation by frontier firms, together with high adoption frictions for the rest of the firms within that industry, are consistent with the results of Comin and Mestieri. If the ability of firms at any part of the distribution to adopt technologies is different from firms at the top, for example, dispersion would increase too.

High productivity dispersion is consistent with the findings of the work by Andrews, Criscuolo, and Gal (2016). Their findings show how productivity grew much faster for firms at the technological frontier as compared to laggard firms within the same industry and that the productivity gap between the two groups of firms widened greatly (see figure 3-4). That implies that large firms, such as Google and Facebook, would become much more productive by the day relative to smaller firms who are not at the frontier, who for a variety of reasons are unable to adopt technologies that would make them grow fast too.

This might seem like an obvious result, but it is not quite, as it contradicts an important belief held by economists about growth: convergence. Since small firms start at a much lower level of productivity than large firms, the latter will tend to grow much faster than the former. But the results by Andrews, Criscuolo, and Gal (2016) show a picture that is consistent with divergence, not convergence. The most productive firms will keep becoming more and more productive relative to the least productive ones, increasing dispersion.

Before digging deeper into the concept of convergence and divergence, there is more to say about dispersion.

The first question to examine is whether dispersion has been increasing systematically or if it has changed in response to booms and recessions. Naturally, business cycles could explain some level of dispersion given heterogeneous changes in firms' response to booms and recessions in terms of investing or, on the contrary, cutting back on investment to adopt new technologies (Kehrig 2015). In fact, the average slowdown in TFP during and after the global recession can be partly attributed to the decline in the speed of adoption of new technologies in response to credit disruptions that

FIGURE 3-4 **Growth at the Frontier versus Laggards**

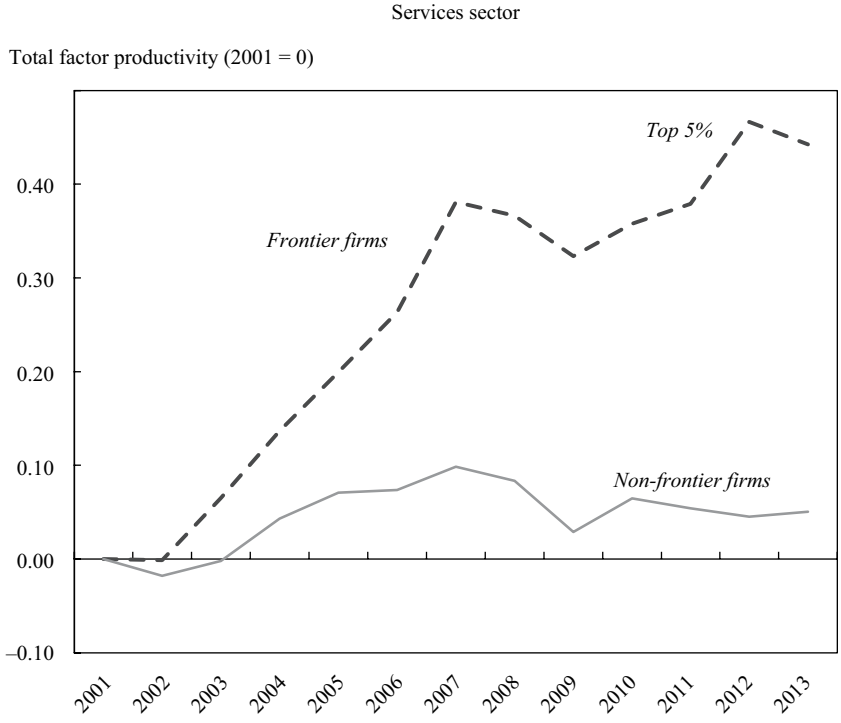

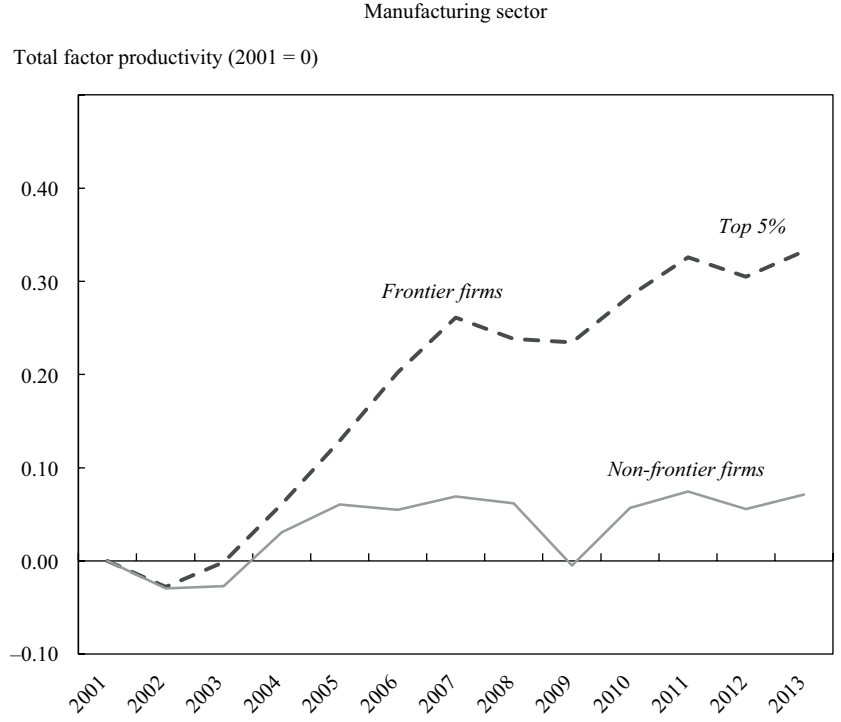

Source: Andrews, Criscuolo, and Gal (2016).

FIGURE 3-5 Productivity Dispersion within Industries Has Been Increasing

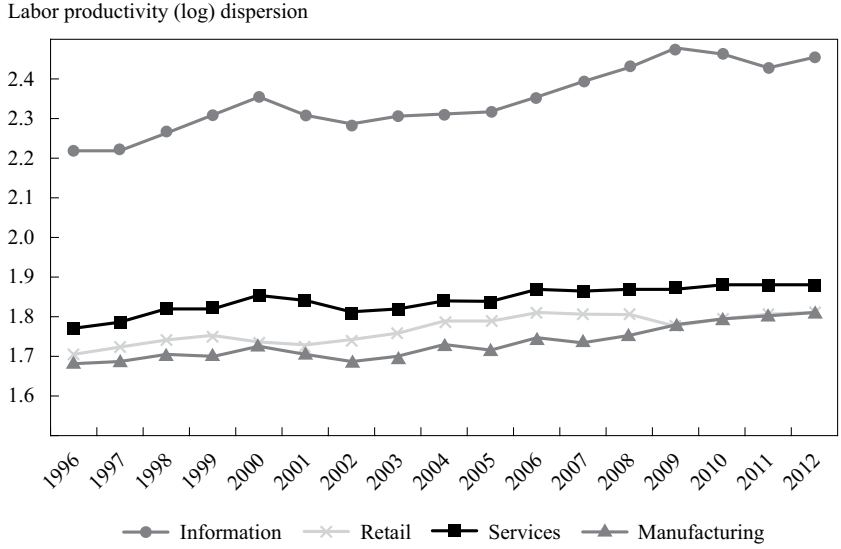

Source: Decker and others (2016).

Note: Y axis does not begin at zero. Data reflect interdecile range of log labor productivity deviated from industry by year means. Sectors are defined on a consistent NAICS basis. Author calculations from the RE-LBD.

have shocked the U.S. economy since the beginning of the worldwide recession in 2007. Yet the slowdown prior to 2007 could be explained by a decline in the ability of R&D investment to bear fruit (Anzoategui and others 2016). Thus, both structural and cyclical factors have played a role in the technology adoption patterns of firms in recent decades and, with it, the trends in productivity dispersion.

Indeed, slowing productivity is a trend that precedes the Great Recession. The fact that productivity dispersion is persistent and large even within narrowly defined firms isn't new, either. For example, research that uses the 1977 U.S. Census of Manufactures show important differences in productivity across plants within a four-digit industry; plants at the 75th percentile were, on average, twice as productive as plants in the 25th percentile in terms of labor productivity.[3] This differential, however, has increased further since then, not only for manufacturing but for all other sectors, as shown

in figure 3-5 from Decker and others (2016). Even though the service sectors have higher levels of productivity dispersion than manufacturing, increasing dispersion across time is common to all. In fact, within manufacturing, increasing dispersion since the 1980s is present for both high-technology firms as well as low-technology ones, and within younger and mature firms. Across the board, productivity dispersion is increasing.

Convergence and Divergence

So what explains increasing dispersion and, perhaps most important, why should we care about it?

As mentioned, in the presence of frictions for adoption, this could affect the typical convergence patterns that economists would expect across firms. Firms at the bottom of the productivity distribution would be expected to grow faster, in relative terms, than those at the top of the distribution. Why? Because the process of adoption is much easier than that of invention. This is, in fact, the essence of convergence. Technology adoption costs, both in terms of resources and time, are smaller than technology discovery costs. As such, less productive firms can enjoy faster productivity growth just by adopting the technologies discovered by those firms at the frontier. The frontier—those firms at the top of the distribution—face a tougher challenge. They have already adopted all the innovations that took them where they are; therefore, to keep growing, they need to lead the innovation process. Even if they are successful, they likely won't be able to grow as fast as those that are only adopting technologies.

One way to examine divergence is to look at each firm's productivity growth trajectory, say, three years down the road, conditional on its initial level. Bahar (2018) found that firms with low initial productivity levels typically experienced faster TFP annual growth over the following three years than firms with a higher level of productivity, consistent with convergence. This process on its own would reduce dispersion, not increase it. Yet, the story doesn't end there. For the most productive firms, this pattern is reversed. It turns out that firms with very high levels of productivity tend to grow faster than their less productive peers, generating a U-shaped relationship between TFP growth and initial productivity levels, as shown in figure 3-6.

FIGURE 3-6 TFP Three-Year Growth Estimate Based on Initial TFP Levels

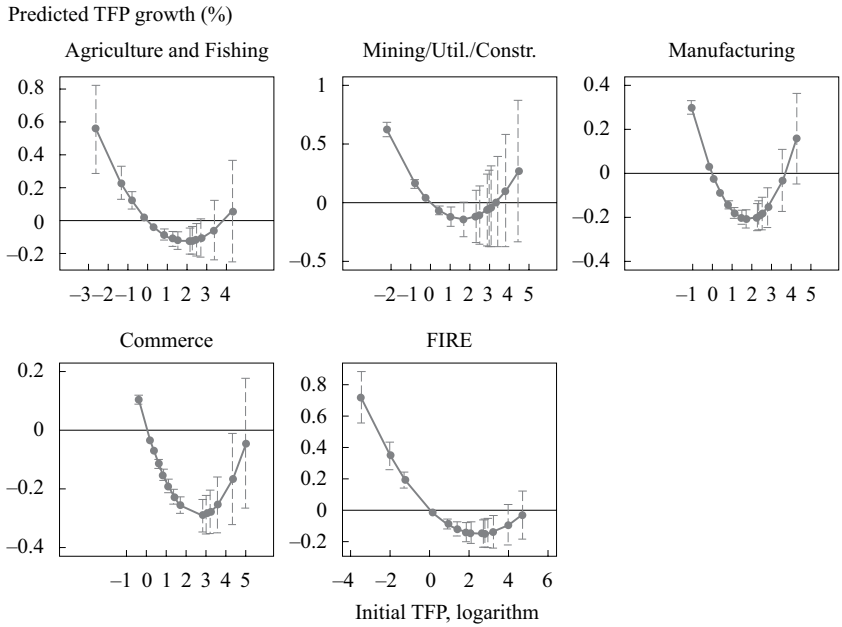

Source: Bahar (2018).

Thus, what we see in the data is a "middle productivity trap" problem. Firms starting low in the productivity scale experience fast-paced growth in TFP, and as they get closer to the productivity frontier, relative growth stagnates. But the top 1 percent keeps growing, and much faster than those in the middle. These dynamics create dispersion. The fast growth of the low productivity firms cannot offset the growth, of the ones at the top when looking at nominal increases. These convergence-divergence dynamics are present across most sectors in the economy, particularly in manufacturing as well as in FIRE sectors, where adoption is key to remain competitive. They are also particularly strong for developing countries.

The fact that some highly productive firms are able to maintain a fast pace when it comes to productivity growth suggests that innovation is, indeed, taking place among these few firms. Yet, these new innovations seem not to be trickling down to other, less productive firms. In other words, there seem to be some friction in the process of technology adoption. What these frictions could be is a key research question, but before we discuss the

FIGURE 3-7 TFP Annual Growth 2008–13 by Initial Dispersion Quartile

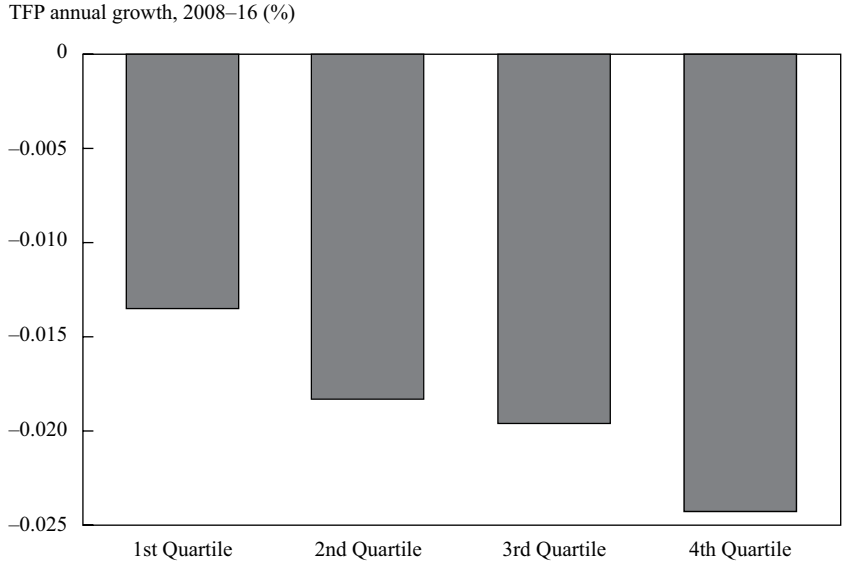

Source: Bahar (2018) and author's calculations.

why, there is an important question that is still up in the air: Does this matter?

It turns out it does. Frictions in the technology adoption process would result in larger productivity dispersion, which in turn has been associated with slower productivity growth. In fact, based on the sample used by Bahar (2018), growth is significantly slower among industries with the highest dispersion (fourth quartile) than for those industries with the smallest dispersion (first quartile) (see figure 3-7).

We can now trace back this productivity slowdown, partly explained by frictions in the adoption process, to economic growth. In the 1950s, Nobel laureate Robert Solow brought to economics a key insight that is highly relevant today: productivity growth is key to sustained economic growth. Countries grow by investing in acquiring more capital or in improving education attainment of their workforce, but the returns to these investments in the long run are limited; thus, without changes in productivity, according to Solow, economic growth would decline. In the long run, it is productivity that matters more. In fact, it has been shown in several studies

that over half of cross-country income differences can be explained by productivity differences (see Hall and Jones 1999).

Concluding Remarks

In the presence of frictions, firms at the frontier will be able to gain larger market shares, which in the long run could turn into weaker competition, affecting output and prices in several markets. The growth of these few firms is not enough to fuel the rest of the economy. These dynamics could affect other trends, such as income inequality, which has also been growing within countries in recent decades.

Overall these patterns present a plausible explanation of the productivity slowdown experienced by most advanced economies since the beginning of the current century, and that has been documented by many. What could stand behind these patterns is out of the scope of this particular chapter, but a plausible factor could be the increasing presence of frictions in technology adoption by lower productivity firms. Under such a possibility, public policy could play an important role in helping to overcome the market failures causing such divergence.

Notes

1. Another widely used definition is "labor productivity," which measures how much output per worker a firm makes. Note that this measure does not take into account that two firms with the same number of workers could differ in the amount of machinery they use or in the way they use their materials. Yet, there usually is a large and positive correlation between TFP and labor productivity.

2. A story in the *New York Times* published on 08/15/2016 details this claim. See: www.nytimes.com/interactive/2016/08/15/sports/olympics/usain-bolt-and-120-years-of-sprinting-history.html.

3. Syverson (2004) and Hsieh and Klenow (2009) find wide dispersion within narrowly defined sectors in the United States, China, and India.

References

Abramovitz, Moses. 1956. "Resource and Output Trends in the U.S. since 1870," *American Economic Review*, May 1956 (Papers and Proceedings), 46 (2), pp. 5–23.

Andrews, Dan, Chiara Criscuolo, and Peter Gal. 2016. "The Best versus the Rest: The Global Productivity Slowdown, Divergence across Firms and the Role of Public Policy," Productivity Working Paper 5 (Paris: OECD).

Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez. 2016. "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence," Working Paper 22005 (Cambridge, Mass.: National Bureau of Economic Research).

Bahar, Dany. 2018. "The Middle Productivity Trap: Dynamics of Productivity Dispersion," *Economics Letters*, 167, pp. 60–66.

Baily, Martin N., Charles Hulten, David Campbell, Timothy Bresnahan, and Richard E. Caves. 1992. "Productivity Dynamics in Manufacturing Plants." *Brookings Papers on Economic Activity: Microeconomics,* pp. 187–267.

Baily, Martin N., and Nicholas Montalbano. 2016. "Why Is U.S. Productivity Growth So Slow? Possible Explanations and Policy Responses," Hutchins Center Working Paper 22 (Brookings Institution).

Baumol, William J., and William G. Bowen. 1966. *Performing Arts: The Economic Dilemma* (New York: Twentieth Century Fund).

Cette, Gilbert, John G. Fernald, and Benoit Mojon. 2016. "The Pre-Great Recession Slowdown in Productivity," Working Paper 2016–08 (Federal Reserve Bank of San Francisco).

Comin, Diego, and Martí Mestieri Ferrer. 2013. "If Technology Has Arrived Everywhere, Why Has Income Diverged?" Working Paper 19010 (Cambridge, Mass.: National Bureau of Economic Research).

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2016 "Declining Business Dynamism: Implications for Productivity?" Hutchins Center Working Paper 23 (Brookings Institution), September.

———. 2018. "Changing Business Dynamism and Productivity: Shocks vs. Responsiveness," Working Paper 24236 (Cambridge, Mass.: National Bureau of Economic Research).

Foster, Lucia, John C. Haltiwanger, and C.J. Krizan. 2001. *New Developments in Productivity Analysis*, chapter 8, 303–72 (University of Chicago Press).

Hall, Robert, and Charles Jones. 1999. "Why Do Some Countries Produce So Much More Output Per Worker Than Others?" *Quarterly Journal of Economics*, 114 (1), pp. 83–116.

Hsieh, Chang-Tai, and Peter Klenow. 2009. "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124 (4), pp. 1403–48.

———. 2018. "The Reallocation Myth," Working Papers 18–19, Center for Economic Studies, U.S. Census Bureau.

Kehrig, Matthias. 2015 (January 23). The Cyclical Nature of the Productivity Distribution. Earlier version: U.S. Census Bureau Center for Economic Studies Paper CES-WP-11-15 (SSRN: https://ssrn.com/abstract=1854401 or http://dx.doi.org/10.2139/ssrn.1854401).

McMillan, Margaret, Dani Rodrik, and Iñigo Verduzco-Gallo. 2014. "Globalization, Structural Change, and Productivity Growth, with an Update on Africa World Development," *World Development*, vol. 63, pp. 11–32.

Pagés, Carmen. 2010. *The Age of Productivity: Transforming Economies from the Bottom Up* (New York: Palgrave Macmillan).

Rodrik, Dani. 2011. "Unconditional Convergence in Manufacturing," *Quarterly Journal of Economics*, 128 (1), pp. 165–204

Solow, Robert. 1956. "A Contribution to the Theory of Economic Growth," *Quarterly Journal of Economics*, 70, (1), pp. 65–94

Syverson, Chad. 2004. "Market Structure and Productivity: A Concrete Example," *Journal of Political Economy*, 112 (6), pp. 1181–22.

———. 2016. "The Slowdown in Manufacturing Productivity Growth." Paper prepared for Chumir-Brookings Project Workshop on the Technology-Productivity-Distribution Nexus, August 1.