Comment

Samuel Kortum, *Yale and NBER*
Unni Pillai, *SUNY*

John Fernald's paper establishes a disheartening yet intriguing fact. The rapid productivity growth that began in the mid-1990s appears to have run its course. We have now returned to the slow growth of the 1970s. The Great Recession is not to blame. Rather, it appears that information technology (IT) is the central player in this recent slowdown, as well as in the speedup of the 1990s. In this comment, we use data for microprocessors to reflect on how a single small sector might play such an outsized role in driving aggregate productivity.

A Recap

Labor productivity in the US economy grew much faster during the period 1995–2003 than subsequently from 2004 to 2013. Fernald brings a wealth of national-, state-, and industry-level data to examine the reasons behind the slowdown in labor productivity growth after 2003. Using a growth accounting approach, he shows that the slowdown in total factor productivity (TFP) and in capital deepening in the aggregate economy both contributed to slower labor productivity growth.

While one would expect that the Great Recession played a major role in this recent productivity slowdown, Fernald shows that it started well before the recession and both labor productivity and TFP rebounded within a few quarters after the recession hit in 2007. After the rebound, both trended along the slower growth path that started around 2004. Fernald also considers the argument that the boom and bust in the housing market contributed to the slowdown in productivity growth. Using state-level data, he shows that changes in housing prices had only a negligible impact on labor productivity.

978-0-226-26873-6/2015/2014-0102$10.00

After excluding these other explanations, Fernald argues that the most likely culprit is slower technological change in IT. The slowdown in productivity during 2004–2013 was actually a waning of the IT-induced surge in productivity that occurred during the mid-1990s to early in the twenty-first century.

Fernald makes a convincing case. First, he uses industry data to show that the growth rate of TFP in IT-producing industries (especially computers and electronics manufacturing) jumped during 1995–2000 compared to pre-1995 values. A number of other studies have reported the jump in TFP growth in IT-producing industries, which Jorgenson (2001) traces to the speedup in technological progress in the semiconductor industry. TFP in IT-producing industries dropped after 2000, but TFP growth in industries that use IT intensively (chemicals and petroleum manufacturing, airline transportation, wholesale trade, motion pictures, broadcasting and telecommunication) jumped during 2000–2003. These two increases, related to production and use of IT, drove the increase in aggregate TFP growth during the entire period 1995–2003. TFP growth in IT-using industries dropped after 2003, setting the stage for the productivity slowdown well before the Great Recession.

Evidence from Microprocessors

We examine Fernald's IT-based account of changes in aggregate productivity growth using data on microprocessors, a driving force behind IT productivity growth. While Fernald and others use measures of TFP growth (calculated as a residual) to infer changes in the rate of technological progress, we ask whether more direct measures of technology in the industry also point to the primacy of IT.

Technological progress in the semiconductor industry is synonymous with Moore's law, that the number of transistors on a chip doubles every two years (Moore 1965). Figure 1 illustrates the remarkable stability of Moore's law over more than four decades, during which the number of transistors on a chip increased by six orders of magnitude.[1] If technological advance in the industry were solely determined by transistors (the basic electronic component on a chip), then the stability apparent in figure 1 would seem to rule out Fernald's claim that changes in technological progress in IT drove changes in aggregate productivity growth.

<AQ: Captions needed for figures? I placed "Figure 1" etc indicators, OK?>

But, it is well known that performance of a semiconductor chip depends not only on the number of transistors but also on their physical

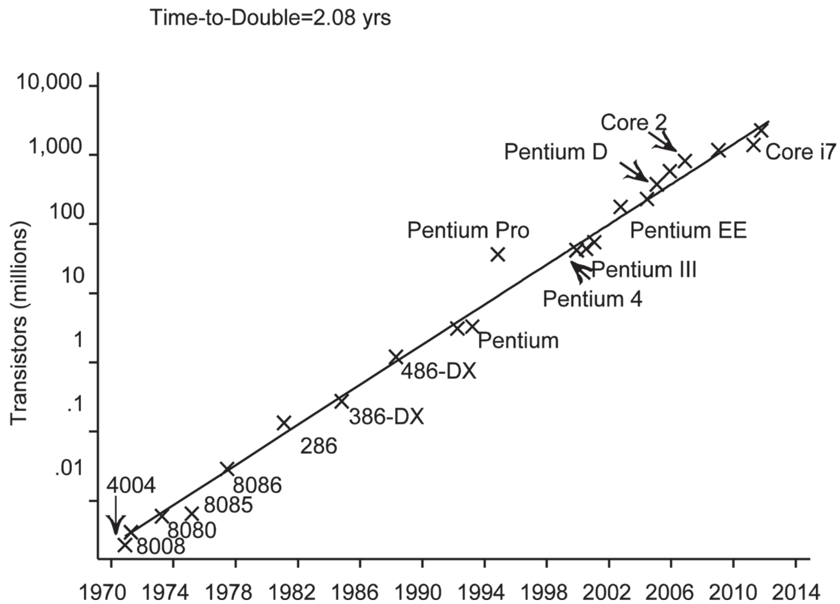

Fig. 1.

size, with smaller transistors operating more efficiently. The industry has reduced the size of transistors by adopting newer vintages of capital equipment. Adoption of each new capital vintage has led transistors to shrink by a factor of about 0.7, decreasing the dimensions from around 6,000 nanometers in 1974 to just 22 by 2012. Figure 2 shows that the time interval between the adoption of these vintages has shortened, from over three years in the pre-1993 period to around two years since then. Jorgenson (2001) noted this fact in his AEA Presidential Address, calling it a speedup in the product cycle for semiconductors.

The speedup in the product cycle would likely show up as a jump in calculated TFP in industries that produce semiconductors or that use them directly (the IT-producing industries), providing direct evidence in support of Fernald's account, which he based on the surge in TFP in IT-producing industries. Figure 2 provides no evidence in support of the following slowdown, however. The product cycle during 2004–2012 remained very close to two years, the same as during the 1993–2003 period.

Of course other factors, besides the number of transistors and the size of transistors on a microprocessor, affect productivity in IT-producing or using industries. To explore these other factors, we look at data on

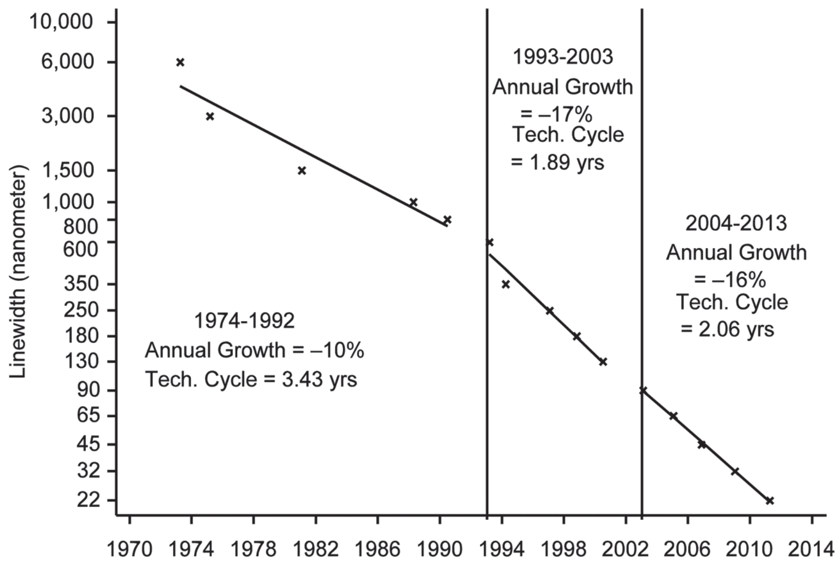

Fig. 2.

the performance of microprocessors, a good proxy for the value that users derive from them. Figure 3 plots the performance numbers over time, showing not only the speedup after 1993, but also the slowdown after 2003. The growth rate is quite similar in the pre-1993 and post-2003 periods, and about twice as fast in between.[2]

Pillai (2013) offers an explanation for the post-2003 slowdown based on a technological shock in the microprocessor industry. The key is to note that during the latter two time periods (1993–2003 and 2004–2013), the transistor dimensions have decreased at about 17% per year (figure 2), and the number of transistors have increased at twice that rate at 34% (figure 1, depicting Moore's law of doubling every two years).[3] But the same increase in the number of transistors translated to much less improvement in performance during 2004–2013 than in 1993–2003. Pillai (2013) argues that engineering problems around the year 2004 forced Intel to abandon the trajectory of previous improvements in microprocessor design that it had followed in the past, and to adopt a new paradigm of design improvements (based on multicore processing) that turned out to be less effective. The lower slope in 2004–2013, shown in figure 3, reflects this inability of new designs to make use of the additional transistors that Moore's law, in essence, made available.

How does the speedup and slowdown in figure 3 compare with the

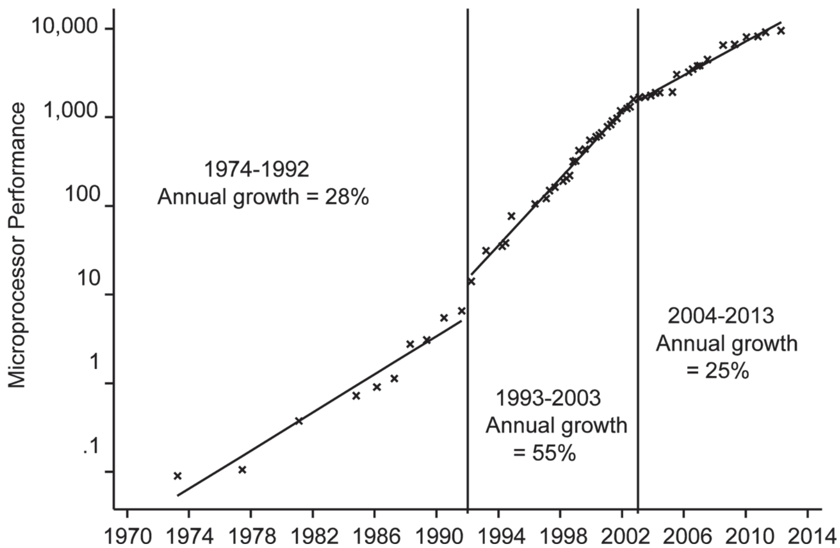

Fig. 3.

speedup and slowdown in measured TFP growth rates? Table 1 compares the growth rate in performance to TFP growth in the semiconductor industry calculated by Byrne, Oliner, and Sichel (2013). The slowdown in TFP growth (18 percentage points) is less than the slowdown in measured performance (30 percentage points), perhaps reflecting the fact that improvements in other semiconductor products (like memory chips) did not slow down as much as microprocessors. Overall, however, the microprocessor data does provide support that technological capabilities measured directly through performance followed the same pattern as measures of TFP in IT-related industries.[4]

Back to Growth Accounting

Taking the performance gains in figure 3 as a proxy for TFP growth in the semiconductor industry, what is the contribution of this single industry to aggregate TFP growth? Charles Hulten (1978) derived a formula that we can use to answer this question. He showed that an industry's contribution to aggregate TFP growth is the product of the industry's TFP growth and the Domar weight (named for Evsey Domar). The Domar weight for an industry is its gross production rela-

Table 1
Microprocessor Performance and Calculated TFP Growth in the Semiconductor Industry

Period	Microprocessor Performance Growth (%)	Semiconductor Industry TFP Growth (%)
1974–1992	28	26 (1974–1995)
1993–2003	55	44 (1996–2004)
2004–2013	25	26 (2005–2013)

Notes: The first column is calculated by the authors. The second column is taken from Byrne et al. (2013), for the time intervals shown in parentheses.

Table 2
Semiconductor Contribution to Aggregate TFP Growth

Period	Aggregate TFP Growth (%)	Microprocessor Performance Growth (%)	Semiconductor Domar Weight (%)	Semiconductor Contribution to TFP Growth (%)
1974–1992	0.64	28	0.39	0.11
1993–2003	1.29	55	0.80	0.44
2004–2013	0.70	25	0.52	0.13

Notes: Note that the time periods in Fernald's paper (1974–1995, 1996–2004, 2004–2013) are slightly different from those in the table. The aggregate TFP growth (second column) was recalculated using data provided by Fernald. Microprocessor performance picks up around 1993, which was why we chose 1993 as the beginning year of the speedup, instead of 1996 as in Fernald.

tive to aggregate value added. It does not matter for the Domar weight that semiconductors are primarily used as intermediates and it would not matter if the semiconductor industry had little value added. What does matter for the Domar weight is overall spending on domestically produced semiconductors. Table 2 shows that the Domar weight for semiconductors (taken from Byrne et al. [2013]) hit a peak of 0.8 in the period 1993–2003. When multiplied by the huge growth rates in microprocessor performance, this one industry can clearly have a sizable impact on aggregates.

Using Hulten's formula, which was also employed in Byrne et al. (2013), table 2 reports the contribution of the semiconductor industry to aggregate TFP growth. This contribution rises from 0.1% prior to 1992 to nearly half a percent per year during the period of rapid TFP growth from 1993 to 2003. After 2003, the contribution declines to 0.13. The jump in the contribution of semiconductors to aggregate TFP growth accounts for nearly half of the jump in aggregate TFP growth during the period 1993–2003, as measured by Fernald.

Conclusion

John Fernald has provided new estimates of US productivity growth as well as a comprehensive analysis of why productivity growth has slowed in the last decade. The evidence points to IT, both as the source of more rapid productivity growth in the mid-1990s and the cause of slower growth since 2003. Our examination of data on microprocessors corroborates Fernald's conclusion and highlights an intriguing fact, that technological change in a single industry can have a substantial impact on aggregate productivity.

A number of questions are raised by this analysis. First, as shown in table 2, a large part of the decline in the contribution of semiconductors to aggregate TFP growth came about because of a fall in the Domar weight (rather than simply a decline in semiconductor TFP growth itself). The declining Domar weight is almost surely due to a shift of semiconductor production overseas. Yet, this shift could occur even as technological change in the industry continues to be driven by research done in the United States. We will need to think more carefully about the interpretation of productivity growth as the US economy becomes increasingly integrated with the rest of the world.

Second, how should we think about productivity growth going forward if for the past 20 years it has been so dependent on technological change in a single industry? The transistor scaling that has driven technology in the industry is continuing at the same rate as it has been since the early 1990s. While performance growth rates for existing devices could fail to keep up with this scaling, new innovations could certainly sprout up that make use of the tremendous opportunities provided by the ability to miniaturize all sorts of devices.

Endnote

Authors can be contacted at samuel.kortum@yale.edu and usadasivanpillai@sunycnse.com. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see http://nber.org/chapters/c13408.ack.

1. Figure 1 uses only data for microprocessors. Moore (1965) based his law on memory chips, prior to the invention of the microprocessor.

2. All three figures shown here are updated versions from Pillai (2013), which also describes the data sources.

3. Halving the size of a transitor allows four times as many to be crammed into a given area on a chip. If chips remain the same size over time, then we expect the slope of figure 1 to be about twice (in absolute value) the slope of figure 2. This prediction is borne out in the post-1993 period.

4. Of course, we should expect similar findings from measures of TFP and direct measures of technological improvement. Byrne et al. (2013) use data on price declines to compute the growth rate of TFP. Aizcorbe and Kortum (2005) show, in a model with technological change embodied in vintages of capital (as depicted in figure 2), that price declines of individual microprocessors reflect the rate of technological progress.

References

Aizcorbe, Ana, and Samuel Kortum. 2005. "Moore's Law and the Semiconductor Industry: A Vintage Model." *Scandinavian Journal of Economics* 107:603–30.

Byrne, David, Stephen Oliner, and Daniel Sichel. 2013. "Is the Information Technology Revolution Over?" *International Productivity Monitor* 25:20–36.

Hulten, Charles. 1978. "Growth Accounting with Intermediate Inputs." *Review of Economic Studies* 45:511–18.

Jorgenson, Dale. 2001. "Information Technology and the US Economy." *American Economic Review* 91:1–32.

Moore, Gordon. 1965. "Cramming More Components onto Integrated Circuits." *Electronics* 38:114–17.

Pillai, Unni. 2013. "A Model of Technological Progress in the Microprocessor Industry." *Journal of Industrial Economics* 61:877–912.