DAVID M. BYRNE
Federal Reserve Board

JOHN G. FERNALD
Federal Reserve Bank of San Francisco

MARSHALL B. REINSDORF
International Monetary Fund

Does the United States Have a Productivity Slowdown or a Measurement Problem?

ABSTRACT After 2004, measured growth in labor productivity and total factor productivity slowed. We find little evidence that this slowdown arises from growing mismeasurement of the gains from innovation in information technology–related goods and services. First, the mismeasurement of information technology hardware is significant preceding the slowdown. Because the domestic production of these products has fallen, the quantitative effect on productivity was larger in the 1995–2004 period than since then, despite mismeasurement worsening for some types of information technology. Hence, our adjustments make the slowdown in labor productivity worse. The effect on total factor productivity is more muted. Second, many of the tremendous consumer benefits from the "new" economy such as smartphones, Google searches, and Facebook are, conceptually, nonmarket: Consumers are more productive in using their nonmarket time to produce services they value. These benefits raise consumer well-being but do not imply that market sector production functions are shifting out more rapidly than measured. Moreover, estimated gains in nonmarket production are too small to compensate for the loss in overall well-being from slower market sector productivity growth. In addition to information technology, other measurement issues that we can quantify (such as increasing globalization and fracking) are also quantitatively small relative to the slowdown.

> The things at which Google and its peers excel, from Internet search to mobile software, are changing how we work, play and communicate, yet have had little discernible macroeconomic impact. . . . Transformative innovation really is happening on the Internet. It's just not happening elsewhere.
>
> —Greg Ip (2015)

U.S. productivity data highlight the paradox at the heart of the quotation above. The fast pace of innovation related to information technology (IT) seems intuitive and obvious. Yet productivity growth has been modest, at best, since the early 2000s. In this paper, we examine the hypothesis that the U.S. economy has a growing measurement problem rather than a productivity slowdown (Aeppel 2015; Feldstein 2015; Hatzius and Dawsey 2015). Some components of real output, including the services provided by IT, are indeed poorly measured. Yet for mismeasurement to explain the productivity slowdown, growth must be mismeasured by more than in the past. Although we find considerable evidence of mismeasurement, we find no evidence that the biases have gotten worse since the early 2000s.

We focus especially on IT-related hardware and software, where mismeasurement is sizable, as well as on e-commerce and "free" digital services such as Facebook and Google. More broadly, we identify potential biases to productivity from intangible investment, globalization, and technical innovations in the production of oil and natural gas (for example, fracking). These are all areas where it is plausible that measurement has worsened since the early 2000s. But taken together, our adjustments turn out to make the post-2004 slowdown in labor productivity even larger than measured. The slowdown of business sector total factor productivity (TFP) growth is only modestly affected.

Figure 1 summarizes our quantitative analysis. The solid portions of the bars show the published data on average growth in U.S. business sector labor productivity, or output per hour. Growth was exceptional from 1995 through 2004, but the pace then slowed by more than about 1¾ percent a year.[1] Suppose productivity growth had continued at its 1995–2004 pace of 3¼ percent a year. Then, holding hours growth unchanged, business sector GDP would be $3 trillion (24 percent) larger by 2015 in inflation-adjusted 2009 dollars.[2]

1. Section I and the online appendix discuss data, the timing of the bars in the chart, and the similar pattern in measures of TFP. The online appendixes for this and all other papers in this volume may be found at the *Brookings Papers* web page, www.brookings.edu/bpea, under "Past Editions."

2. In independent work, Syverson (2016) suggests a similar calculation of the missing growth.

Figure 1. Published and Adjusted Data on U.S. Labor Productivity[a]

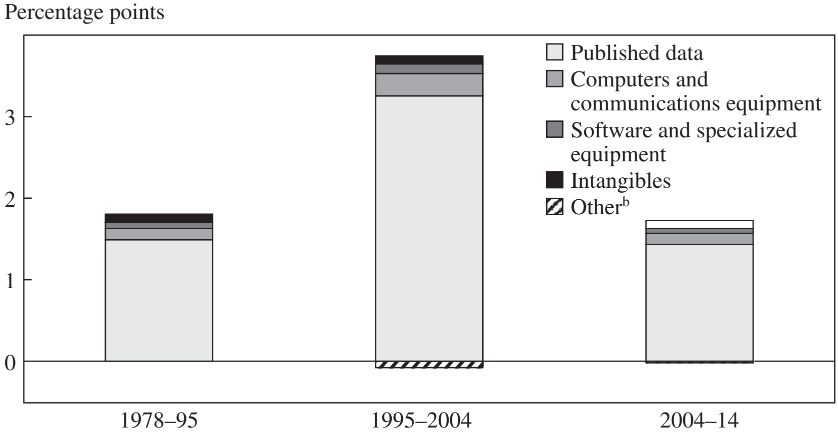

Sources: Fernald (2014); authors' calculations.
a. Shows adjustments to growth in output per hour in the business sector.
b. Comprises Internet access, e-commerce, globalization, and fracking.

We find no evidence that growing mismeasurement related to IT or other factors can fill this gap. In section I, we explore the hypothesis that the slowdown reflects the growing importance of poorly measured industries with low productivity growth, such as health care and other services. These industries are indeed growing as a share of the economy, but holding weights fixed at their 1987 values would make little difference to the slowdown. That most industries show slowing growth matters more than changing weights.

We then turn to biases within specific sectors. Figure 1 shows our adjustments for various biases. We incorporate consistent measurement of quality-adjusted prices for computers and communications equipment; judgmental corrections to prices of specialized information-processing equipment and software; a broader measure of intangible investment than is used in the national accounts; and ballpark adjustments for other issues—Internet access, e-commerce, globalization, and fracking. These adjustments make labor productivity growth since 2004 look better. But the adjustments to account for mismeasurement matter even more in the 1995–2004 period. On balance, therefore, the labor productivity slowdown becomes modestly larger.[3]

3. There are also some sources of *upward* measurement error in growth related to globalization that have become less important. Still, we will usually take "mismeasured" to mean "causing GDP growth to be understated."

In particular, although we find somewhat more mismeasurement of computer and communications equipment prices in the recent period than previously, domestic production of these products has plunged, making this mismeasurement less important for GDP. Although David Byrne, Stephen Oliner, and Daniel Sichel (2015) show that microprocessor (MPU) price declines are substantially understated, this has little immediate implication for productivity; because MPUs are not final products, they only affect GDP through net trade, which is roughly in balance for semiconductors.

The "other" adjustments in figure 1 include improved Internet quality (section III) and e-commerce (section IV), which together add about 5 basis points (bp) more in the post-2004 period than from 1995 to 2004. This adjustment is small, reflecting the conceptual challenges involved in bringing more of the services of Google, Facebook, and the like into market sector GDP. The major cost to consumers of these services is not broadband access, cell phone service, or the phone or computer; rather, it is the opportunity cost of time. This time cost is not consumption of market sector output. It is akin to the consumer surplus obtained from television (an old economy invention) or from playing soccer with one's children. Following Gary Becker (1965), activities that combine market products with the consumer's own time are properly thought of as nonmarket production that uses market goods and services as inputs. As we discuss, a small amount of market output could conceivably be included in final consumption, corresponding to online ad spending; this spending is relatively modest and has little effect on growth in output or productivity. Thus, though the digital services are valuable to households, the possible mismeasurement in these areas makes essentially no difference to market sector labor productivity and TFP growth.[4] That said, to the extent that the effect of innovation on the quality of leisure is outpacing the effect on market activities, market productivity growth might have become a less reliable measure of overall welfare.

These other adjustments also include effects from globalization and fracking (section V). Globalization was most intense in the late 1990s and early 2000s. That caused real import growth to be understated and, correspondingly, artificially boosted measured GDP growth by about 10 basis points (bp) per year during the period from 1995 to 2004. Hence, in figure 1 the "other" bar contributes negatively in the period. Fracking, on the

4. Nordhaus (2006) sketches principles of national accounting for nonmarket as well as market goods and services.

other side, boosts productivity growth by about 5 bp after 2004. Together, these adjustments shave about 10 bp from growth in the 1995–2004 period, and add about 10 bp to growth thereafter.

For TFP, the adjustments are even smaller than for labor productivity. Adjusting equipment, software, and intangibles implies faster GDP growth, but also faster input growth (because effective capital services rise more quickly). After adjusting hardware and software, the aggregate TFP slowdown after 2004 is modestly worse. Adding a broader measure of intangibles—as is done by Carol Corrado, Charles Hulten, and Sichel (2009)—works modestly in the other direction, so our broadest adjustment for investment goods leaves the 1¼ percentage point slowdown in TFP a few basis points worse. The other (non-investment-good) adjustments we make pass directly into TFP; but, on balance, they still leave the slowdown in TFP only modestly attenuated.

In making these points, we draw on a large body of existing research. Before presuming that the measurement problems have gotten worse, it is worth remembering that in the 1990s and early 2000s, much research looked at missing quality improvement, the problem of new goods, and the fact that consumers had an explosion of new varieties. The biases were frequently estimated to be large. For example, VCRs, cell phones, and other similar products were added to the consumer price index (CPI) a decade or so after they appeared, and when their prices had already fallen by 80 percent or so (Gordon 2015; Hausman 1999). The explosion in consumer choice, and the possibilities for so-called mass customization, were documented in the 1990s. At about the same time, the Boskin Commission estimated that omitted quality change in new goods was worth at least 0.5 percent a year (Boskin and others 1998).[5] So again, the issue is not whether there is bias. The question is whether it is larger than it used to be.

The structure of the paper is as follows. Section I lays out motivating facts about the productivity slowdown, including a discussion of the changing industry composition of the U.S. economy. Section II discusses improved deflators for information technology and intangibles, and reworks the growth accounting with alternative capital deflators. We then turn to other issues in sections III, IV, and V that plausibly changed after 2004. Section VI concludes.

5. Some academic research found even larger effects—for example, Bils and Klenow (2001)—while Schultze and Mackie (2002) argued for a smaller number.

I. The Recent Rise and Fall of U.S. Productivity Growth

Three productivity facts frame our subsequent discussion. First, as measured, the growth in business sector labor productivity and TFP increases sharply in the mid-1990s but then slows down after about 2004. Second, the slowdown is broad-based across industries, including in relatively well-measured ones, such as wholesale and retail trade, manufacturing, and utilities. Third, the TFP slowdown is not caused by the rising share of slow-productivity-growth industries.

John Fernald (2015) interprets the slowdown as a "return to normal" following a period of exceptional, broad-based gains from the production and use of information technology. The remaining sections of this paper explore rising mismeasurement as an alternative explanation.[6]

We focus now on TFP, which is defined as a residual: output growth that is not explained (in a proximate sense) by growth in inputs of capital and labor. In the longer run, TFP growth mainly reflects innovation in a broad sense. The online appendix shows that changes in TFP growth have been the proximate driver of changes in labor productivity growth, as theory would suggest. TFP as well as labor productivity slow sharply in the 2004–07 period (before the Great Recession) relative to the late 1990s and early 2000s; the slowdown in growth is statistically significant in formal tests for a change in mean growth.[7]

Figure 2 shows the industry sources of the slowdown in business sector TFP growth from a Bureau of Labor Statistics (BLS) data set. Because of data availability, the subperiods shown are all between 1987 and 2013. We divide the private business economy into four mutually exclusive categories: IT-producing; wholesale and retail trade; other well-measured; and

6. A separate debate is whether the productivity slowdown of the 1970s was itself due to mismeasurement. Griliches (1994) points out that the post-1973 slowdown was concentrated in poorly measured industries. Gordon (2016) argues instead that the post-1973 slowdown reflects the unusual strength of the 1920–70 period rather than anything specific that happened in the 1970s. Relatedly, Fernald (1999) estimates that building the Interstate Highway System substantially boosted productivity growth in the 1950s and 1960s, but then its effects ran their course. Triplett (1999) reviews arguments that the post-1973 slowdown was illusory.

7. A possibly more optimistic perspective on recent developments comes from noting that TFP growth has continued since the Great Recession at its pre-1995 pace. This pace of TFP growth may be normal—it was, perhaps, the 1995–2004 period that was exceptional. Furthermore, in recent years TFP may be more relevant than labor productivity, whose weakness since 2010 partly reflects transitory factors associated with weak capital deepening.

Figure 2. Contributions to U.S. Total Factor Productivity Growth, by Industry Subgroup

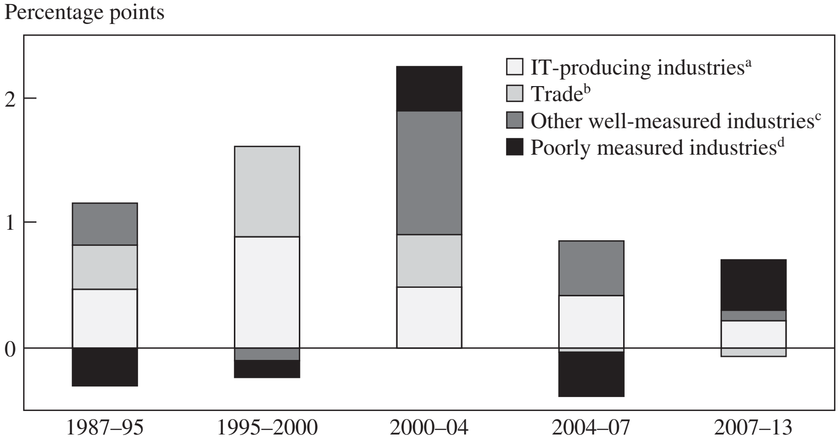

Sources: U.S. Bureau of Labor Statistics; Fernald (2015).

a. Includes computer and electronic product manufacturing, publishing (including software), and computer systems design.

b. Includes wholesale and retail trade.

c. Following Nordhaus (2002), includes manufacturing (excluding IT-producing), agriculture, mining, transportation, utilities, broadcasting, and accommodations.

d. Includes the remaining industries not categorized as IT-producing, trade, or other well-measured industries.

poorly measured.[8] All sectors show somewhat slower growth after 2004, but the slowdown is particularly pronounced for wholesale and retail trade and the other relatively well-measured sectors. After 2000, IT production adds less and less to TFP growth, a situation that we discuss in the next section. After 2004, wholesale and retail trade contribute negatively; this is noteworthy because IT provided a substantial boost to wholesale and retail trade in the preceding periods, in part through industry reorganization. Other (nontrade) well-measured industries contribute less after 2004. Thus, the slowdown is apparent even in areas such as trade and non-IT manufacturing, where measurement has traditionally been considered relatively good. (Of course, even in these industries, unmeasured gains from quality improvements and new goods may be occurring.) Finally, the poorly measured subgroup contributes negatively from 2004 to 2007, but then turns substantially positive from 2007 to 2013; quantitatively, the post-2007 shift

8. “Other well-measured” includes most of manufacturing (except computers and electronics equipment), agriculture, mining, utilities, transportation, broadcasting, and accommodations. Nordhaus (2002) also considers wholesale and retail trade as well measured, but we have broken that out separately.

reflects an increasingly positive contribution from finance and the elimination of a large negative contribution from construction.

The slowdown is also not simply a matter of weights that have been shifting toward poorly measured industries with low TFP growth, such as services. Services have been growing as a share of the economy and are inherently challenging to measure in real terms (Griliches 1994; Triplett and Bosworth 2004). The top panel of figure 3 compares actual TFP growth with a counterfactual where nominal industry value added weights are held constant at their 1987 values.[9] During the periods shown, the growth rates of the two measures are within a few basis points. In other words, shifts in the industry composition of the economy play essentially no role in the productivity speedup in the mid-1990s or slowdown in the 2000s.

Why are the two series so similar? The value added share of services and other relatively poorly measured industries rises about 10 percentage points from 1987 to 2013. For the full sample, TFP growth in these poorly measured industries was about zero, compared with 2 percent annual growth for relatively well-measured industries (including IT hardware and trade). Hence, a back-of-the-envelope guess would be that, by the end of the sample, the fixed-weight index should grow about 20 bp faster, reflecting the annual difference of 2 percentage points in growth times the 10 percentage point shift in weights. Roughly half the shift in weights had occurred by 1998, so the expected effect on the post-2000s slowdown might be 10 bp.

In the top panel of figure 3, the differences are even smaller than this back-of-the-envelope calculation. First, within the groups of well-measured and poorly measured industries, weights shifted toward those with faster TFP growth. These shifts partially offset the broader shift toward services. Second, since 2007, "Baumol's cost disease" (Baumol and Bowen 1966) has reversed—TFP growth in poorly measured services has been *faster* than that in well-measured sectors.

The bottom panel of figure 3 makes this point about weights a different way by showing that the slowdown after the early 2000s is broad-based across industries. The figure shows the change in average annual industry value added TFP growth for 2004–13 relative to 1995–2004. About two-thirds of industries show a slowdown in measured TFP growth after 2004. We get a similar picture if we look at the change from 1995–2004 to 2004–07, so it is not simply a matter of the Great Recession affecting many

9. Value added weighting of value added TFP growth is essentially equivalent to doing so-called Domar weighting of gross output residuals (Domar 1961). The fixed weights are based on nominal expenditures, not quantities. In the data, the rise in the nominal share of services reflects both faster growth in quantities and faster growth in prices.

Figure 3. Aggregate Total Factor Productivity Growth across Selected Business Sector Industries

Actual data and the counterfactual[a]

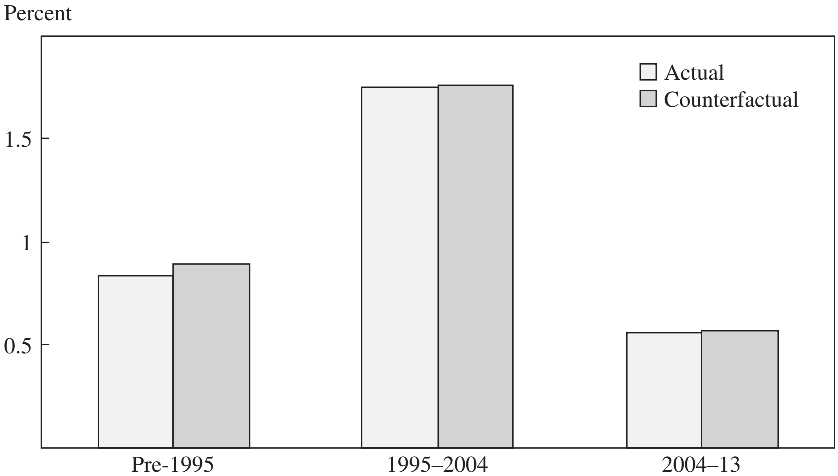

Industries ranked by change in average TFP growth, 1995–2004 to 2004–13[b]

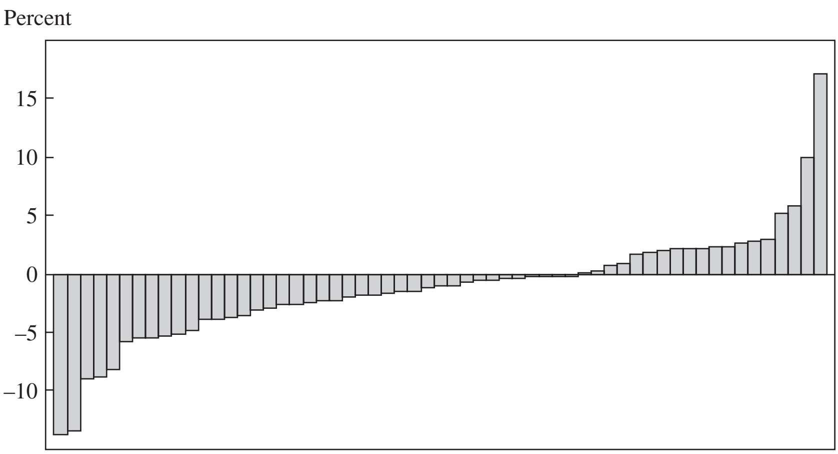

Source: U.S. Bureau of Labor Statistics.

a. The counterfactual assumes that nominal industry value added weights remain fixed at 1987 values.

b. The horizontal axis ranks business sector industries by the change in average annual value added TFP growth from 1995–2004 to 2004–13. Water transportation is omitted from the right side of the figure because of its scale. Growth rates are calculated as 100 times the log change.

industries. We also get a similar picture using labor productivity, so it is not something about capital measurement.

Our results are consistent with previous studies that have found that the shrinking size of well-measured sectors was not a first-order explanation for previous swings in productivity (Baily and Gordon 1988; Sichel 1997).

Why did so many industries show a common slowdown after 2004? The economy plausibly received an exceptional boost from IT in the 1990s and early 2000s that hit many industries. However, by the mid-2000s, the low-hanging fruit of a wave of IT-based innovation (including associated reorganizations) had been plucked. For example, industries along the supply chain from factory to retailing had already been substantially reorganized to reduce inventory, waste, and headcount; and IT-supported efficiencies in middle management and administrative support had been exploited. It is possible that the latest waves of innovation will take time to bear fruit and that we are overlooking nascent IT-based productivity gains in service sectors such as health care and education. But here we sidestep this more challenging question and turn to an alternative hypothesis: that rising mismeasurement might explain the patterns in the data.

II. Growing Mismeasurement of Information Technology?

In this section, we document long-standing challenges in measuring information-processing equipment and software.[10] Correcting for the mismeasurement of these investment goods turns out to make the slowdown in labor productivity and TFP growth even worse after 2004. We also note a rise in uncertainty about these effects: Investment has shifted toward special-purpose information-processing equipment and intangibles, especially software—categories that have proven especially difficult to measure.

After moving roughly sideways in the postwar period through the late 1970s, the official IT investment price index turned downward as the personal computer (PC) era began, and then the rate of decline accelerated sharply, to 6 percent a year on average, during the IT boom of the 1990s and the early 2000s (table 1). Since 2004, the price declines have retreated to a modest rate of 1 percent, coinciding with the decrease in the contribution of

10. Our focus in this section is on the contribution of IT capital services to productivity and its implications for TFP growth. Parallel measurement problems exist for IT consumer durables, which we do not discuss explicitly. However, we account for understatement of GDP from the mismeasurement of IT through our adjustments to domestic production, whether for the consumer or business market.

Table 1. Prices and Weights for Information Technology Investment[a]

Measure	*1947–78*	*1978–95*	*1995–2004*	*2004–14*
IT investment share of business fixed investment	12.2	23.6	30.7	29.3
IT investment price indexes				
National Income and Product Accounts	0.2	–2.2	–6.1	–1.4
Conservative alternative[b]	–1.8	–4.4	–9.2	–4.4
Liberal alternative[c]	–3.9	–6.5	–11.2	–5.9
Share of IT investment				
Computers and peripherals	13.1	22.8	20.8	14.5
Communications equipment	36.9	26.6	22.6	17.0
Other information systems equipment	38.3	26.7	17.3	20.4
Software	11.7	23.9	39.3	48.2
Price deflators				
Computers and peripherals[d]				
National Income and Product Accounts	–18.1	–14.6	–19.3	–6.6
Alternative	–18.1	–19.0	–27.3	–18.6
Communications equipment				
National Income and Product Accounts	1.9	1.4	–5.4	–2.7
Alternative	–3.0	–2.7	–11.2	–10.3
Other information systems equipment				
National Income and Product Accounts	2.3	2.9	–0.6	0.5
Alternative	–1.7	–2.2	–8.9	–4.9
Software[d]				
National Income and Product Accounts	–0.7	–1.2	–1.1	0.1
Alternative	–4.8	–4.4	–2.5	–0.8

Sources: U.S. Bureau of Economic Analysis; Byrne and Corrado (2016).
a. All values are expressed as percents.
b. Incorporates alternative computer and communications equipment prices.
c. Incorporates alternative software and special-purpose equipment prices.
d. Price indexes begin in 1958.

IT production to TFP growth shown in figure 2. This flattening out has led to a revival of interest in measuring IT prices, and some recent studies find that official price statistics have substantially understated price declines in recent years.[11]

11. See research for communications equipment (Byrne and Corrado 2015), computers (Byrne and Pinto 2015; Byrne and Corrado 2016), and microprocessors (Byrne, Oliner, and Sichel 2015).

Figure 4. Information Technology Investment Price Indexes, 1950–2014[a]

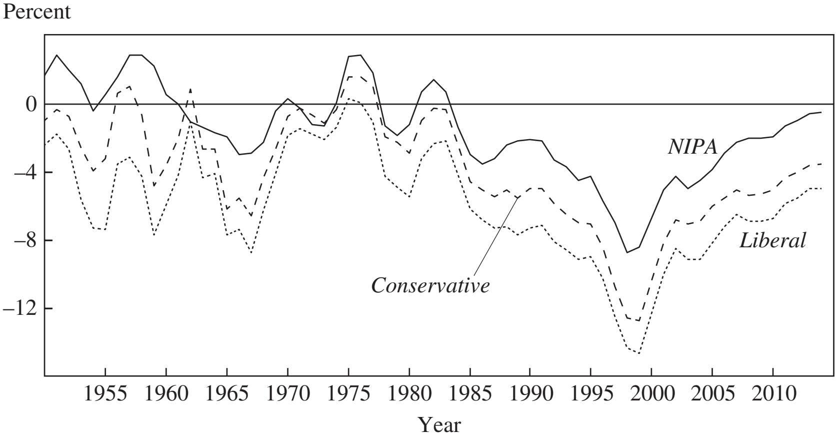

Sources: U.S. Bureau of Economic Analysis; Byrne and Corrado (2016).
a. Data are expressed as 3-year average changes in annual data. The percent change is calculated as 100 times the log change. See the text and the notes to table 1 for definitions of the conservative and liberal indexes.

Has *worsening* price mismeasurement caused a spurious slowdown in official estimates of output and real investment, distorting productivity estimates? Answering this question requires the construction of a fully consistent time series. We employ price indexes developed by Byrne and Corrado (2016), who review the full postwar history of IT price research and construct alternative price indexes for IT investment and production using research not only for recent years but also for earlier periods that may not have been incorporated into the National Income and Product Accounts (NIPA) that are issued by the Bureau of Economic Analysis (BEA).

We provide two alternative price indexes in figure 4. The first, a conservative index, is based solely on research studies that use detailed data sets for specific product classes. We extrapolate these results, as described in Byrne and Corrado (2016), for communications equipment and for computers and peripherals. For the second, liberal, index, we add plausible assumptions about the prices of IT products for which no direct studies are available, namely, other types of information-processing equipment and software. Overall, our alternative indexes suggest substantially faster price declines than those shown in the NIPA *throughout the postwar period.* For some categories (computers and communications equipment), price measurement appears to have worsened, but the importance of these categories in GDP has declined. On balance, the declining importance in GDP dominates, so the bias in GDP growth was larger in the past.

We discuss the component prices briefly here and compare them with the investment prices used in the NIPA.

II.A. Components of IT Investment

COMPUTERS AND PERIPHERALS The official investment price index for computers and peripherals reflects the results of internal BEA research (Cole and others 1986; Dulberger 1989), which led to the adoption of hedonic regression techniques to account for the rapid technological advances embodied in new models of computers and peripherals.[12] For the postwar period, through the early 1980s, BEA prices are consistent with outside studies (Gordon 1990; Triplett 1989). Beginning in the 1990s, the BLS adopted hedonics for computers (but not peripherals) as well, and the BEA now relies on BLS prices as inputs for the NIPA investment deflator (Grimm, Moulton, and Wasshausen 2005). Despite the commitment to quality adjustment in the official statistics, outside research indexes indicate somewhat different price trends beginning in the 1980s.

PERSONAL COMPUTERS Our alternative price index for computers and peripherals diverges from official prices beginning in 1984. For PCs, we adopt an aggregate of the indexes developed in a comprehensive study by Ernst Berndt and Neal Rappaport (2001, 2003), which exhibits declines that are 8 percentage points faster through the early 2000s. The documentation for the BLS hedonic models is not comprehensive enough to allow us to identify the source of the difference in results with confidence.

More recently (since 2004), the BEA index for PCs has slowed dramatically, and some aspects of the sources and methods used raise concerns about the accuracy of this development. The top panel of figure 5 shows the average unit price of PCs sold in the U.S. business market reported by IDC Corporation, which makes no adjustment for quality. The figure also shows the rate of change for the BEA investment price index for PCs. In the late 1990s and early 2000s, the gap between the two series indicates that quality improvements were contributing 15 to 20 percentage points to the fall in constant-quality PC prices. The gap has narrowed since that time, and since

12. With appropriate data on characteristics, hedonic regressions are a useful tool for quality-adjusting prices, but the absence of hedonic adjustment does not necessarily indicate that a price index is biased. Other techniques may also account for quality improvements (Wasshausen and Moulton 2006).

Figure 5. Price Indexes for Personal Computers, 1996–2014[a]

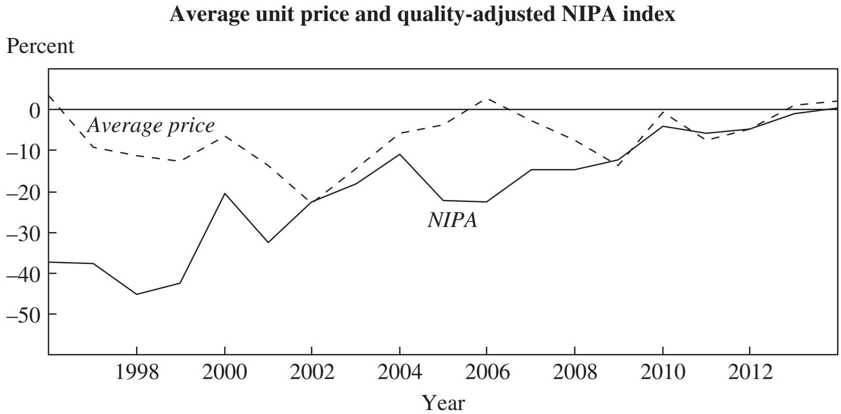

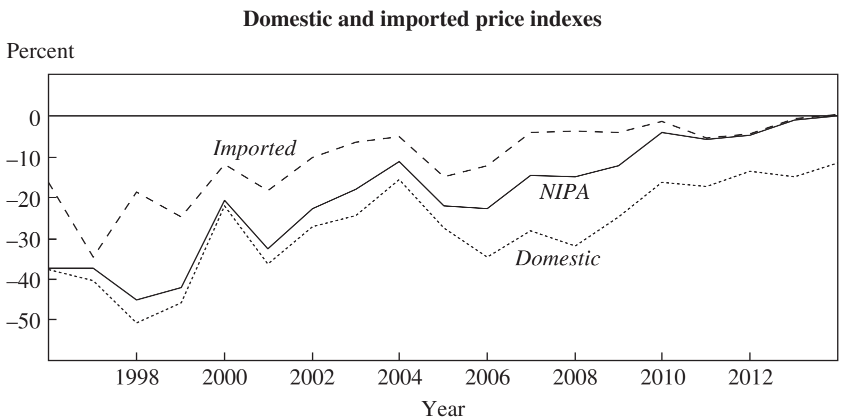

Sources: IDC Corporation; U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics.
a. The percent change is calculated as 100 times the log change.

2010 the two series have been almost identical, implying no improvement in PC quality, holding unit price constant, for the past five years.

Three measurement problems appear to contribute to this implausible result. First, the BEA investment series is the aggregate of a domestic production price index and an import price index that are calculated independently from one another, using different source data (figure 5, bottom panel). As a result, any discount accruing to a business switching from domestically sourced to imported equipment is not reflected in the investment price index—a form of outlet substitution bias akin to omitting from a consumption price index the price savings associated with switching to shopping at Walmart (Reinsdorf 1993; Houseman and others 2011).

Second, the price index for imports falls markedly more slowly than the index for domestic production over a prolonged period—an average annual difference of 14 percentage points since its introduction in 1995. The implied continual rise in the relative price of imported computers is inconsistent with the increase in import penetration from 50 to 90 percent during the same period (Byrne and Pinto 2015). This contradiction suggests that the price mismeasurement is more severe for import prices than for domestic producer prices. Among the possible contributing factors to the relatively flat import price series is the heavy presence of intrafirm (transfer) prices in the index (more than 60 percent of the value of the basket in 2013). These prices may behave differently from arm's-length prices. This may be related to the finding by Emi Nakamura and Jón Steinsson (2012) that a surprisingly high proportion of the items in the import price index sample never experience a price change before exiting the index basket. Also, new models are generally linked into the import price index in a way that would not capture any decline in the quality-adjusted price of the item (Kim and Reinsdorf 2015).

This suggests the producer price index (PPI) would be a more appropriate deflator for investment, though the PPI itself has drawbacks. When quality-adjusting the computer PPIs, the BLS controls primarily for technical features, such as processor clock speed and features associated with changes in production costs (Holdway 2001). Design improvements not clearly tied to costs or not easily identified in technical specifications, such as circuits designed to work more effectively in parallel, may raise the value of the equipment to its user through superior performance without affecting the quality index. Thus, the approach used for quality adjustment in the PPI may lead to an understatement of quality improvements and an overstatement of inflation.

Although we are aware of no research studying computer prices directly in recent periods, Byrne, Oliner, and Sichel (2015) analyze prices for MPUs, the central analytical component of computers. When controls for direct measures of performance were used in their hedonic analysis of MPUs (benchmark scores on a battery of user tasks), their hedonic price index fell more than 20 percentage points faster than a hedonic index controlling for technical features during the 2000–13 period. We infer that the BLS hedonic index may be understating the annual rate of quality improvement for PCs by 4 percentage points—the (rounded) product of the bias in the MPU price index and the share of MPU inputs in the final value of

PCs (15 percent). In our alternative index, we extend the Berndt–Rappaport index with the bias-adjusted PPI.

MULTIUSER COMPUTERS The BLS price index for multiuser computers (such as servers), which is used by the BEA, is quality-adjusted using a hedonic regression as well. Following the same logic used for PCs, we augment the BEA price index beginning in 1993 with an indicator of the average price per computer unit adjusted for MPU performance, which falls markedly faster than the PPI. The performance measure is an average of scores on a suite of benchmark tests developed by Systems Performance Evaluation Corporation (SPEC)—a consortium of industry representatives—to provide reliable comparisons across systems. We blend this price-performance indicator with the PPI, which controls for computer features not accounted for by the SPEC benchmark. We employ a weighted average of the PPI and the price-performance trend to deflate multiuser computers. This alternative index falls 10 percentage points faster than the official BEA price index.

STORAGE EQUIPMENT For storage equipment as well, the PPI that is the basis for the BEA investment price index appears out of alignment with price-performance trends in the industry. From its introduction in 1993 until 2014, the PPI fell 12 percent a year on average, in stark contrast to the price per gigabyte for hard disk drives, currently the dominant technology in the industry, which fell 35 percent per year on average (McCallum 2015). Recent research by Byrne (2015b), employing detailed model-level prices for storage equipment, developed prices that fell at nearly the rate of raw price-per-gigabyte series. We use the Byrne (2015b) index extended backward by the price-per-gigabyte series, with a 4 percentage point bias adjustment.[13]

All told, our alternative index for computers and peripherals falls faster than the NIPA index beginning in the early 1980s, and the gap between the two increases markedly, to 8 percentage points, between 1995 and 2004. The difference between the indexes has been even larger in recent years—an average of 12 percentage points (figure 6, top panel). This substantial gap suggests that additional research is needed to account well for computer investment in the NIPA, and the rising gap makes the issue increasingly important. However, the percentage point slowdown in the alternative

13. Research for the remaining category, peripherals, is sparse. The BEA investment price index fell 12 percent a year on average in the 1990s, but 4 percent on average since then. Aizcorbe and Pho (2005) examine scanner data for eight categories of peripherals for the years 2001–03. Although we note that the geometric mean of price indexes for these categories falls 15 percent per year, we chose not to adjust the peripherals index based on this short time series.

Figure 6. Official and Alternative Computer and Communications Equipment Price Indexes, 1960–2014[a]

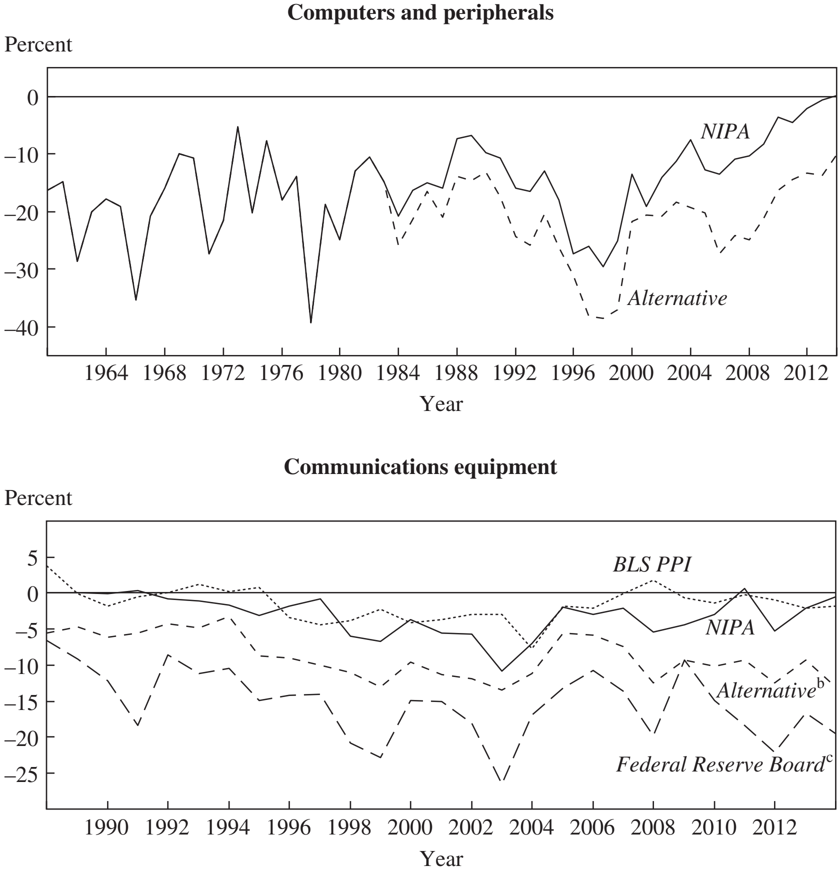

Sources: Byrne (2015b); Byrne and Corrado (2016); Federal Reserve Board; U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics.

a. The percent change is calculated as 100 times the log change.

b. Aggregate of the Federal Reserve Board and the Bureau of Labor Statistics price indexes.

c. Includes selected types of communications equipment.

index is still quite large and returns the rate of price decline to the pace seen before the IT boom of the 1990s.

COMMUNICATIONS EQUIPMENT Official investment prices for communications equipment reflect both BLS producer and import price indexes, and internal BEA research (Grimm 1996). Outside research, including price indexes published by the Federal Reserve Board, is incorporated to some extent as well, and the investment index does fall faster than the PPI for the industry (figure 6, bottom panel). However, a substantial amount of research is not reflected in the NIPA (Byrne and Corrado 2015, 2016). This includes work on transmission and switching equipment in the early postwar era by Kenneth Flamm (1989), as consolidated and augmented by Gordon (1990), and satellite prices constructed by Byrne and Corrado (2015). For more recent years, the BEA investment price index appears inconsistent with new prices for cellular systems, data networking, and transmission developed in Byrne and Corrado (2015) and Mark Doms (2000). Because subindexes are not published for communications equipment investment, it is impossible to analyze the sources of this difference. In any event, technological developments in the field suggest that careful attention needs to be given to account for quality changes, such as fourth-generation cellular systems now capable of delivering video.

Like the computer investment index, the Byrne and Corrado (2016) communications equipment investment index is carefully constructed to match the scope and weighting of the BEA index. All told, the difference between the BEA investment index and the alternative is noteworthy, and the gap is slightly larger in the 2004–14 period than in the 1995–2004 period. Unlike the index for computers and peripherals, the communications equipment index maintains roughly the same pace of decline as during the IT boom.

SPECIAL-PURPOSE ELECTRONICS The remaining components of the BEA's "other information-processing" equipment category form a diverse group of special-purpose types of equipment designed for use in medical, military, aerospace, laboratory, and industrial applications.[14] Examples include magnetic resonance imaging machines, electronic warfare countermeasure devices, and a wide variety of equipment used for monitoring and controlling industrial processes. Technological advances in recent years have been impressive. One well-known example is genomic sequencing, where specialized equipment has contributed to dramatic efficiency gains: The cost of sequencing a human genome has dropped from roughly \$1 million in 2008 to \$1,000 in 2015 (Wetterstrand 2016).[15]

14. Navigational equipment and audiovisual equipment are classified as communications equipment in the BEA investment taxonomy.

15. Although the sequencing of a human genome is not final output, improvements in the tools used to conduct science are the likely foundation of falling prices for health services in the future. Heather and Chain (2016, p. 6) present the history of DNA sequencing equipment, and they note that "over the years, innovations in sequencing protocols, molecular biology and automation increased the technological capabilities of sequencing while decreasing the cost, allowing the reading of DNA molecules that are hundreds of base pairs in length, massively parallelized to produce gigabases of data in one run." On the role of high-performance computing in genetics, also see Stein (2010).

Surprisingly, with the exception of electromedical equipment, which edges down modestly, the PPIs for these products have risen on average since the late 1990s. Differences in market structure (such as the smaller scale of production and the market power of military and medical customers) and the price trends of specialized inputs could cause prices for special-purpose electronics to behave differently from prices for general-purpose electronics like computers (Byrne 2015a). Yet these goods have electronic content comparable to computers, and one might expect the equipment prices to reflect the rapidly falling price of the electronic components used in their production. In our liberal alternative scenario, we remove roughly one-third of the difference between the trend price growth of special-purpose and of general-purpose (computer and communications) electronics.

SOFTWARE Investment in software is deflated in the NIPA by an aggregate of three subindexes: prepackaged, custom, and own-account software. BLS producer prices are available for prepackaged software, and research has been conducted at BEA and by outside researchers into quality-adjusted price trends (Parker and Grimm 2000; Copeland 2013). To deflate investment in prepackaged software, the BEA employs a BLS PPI, with an adjustment reflecting the average difference between the PPI and the BEA's research results. Because direct observation of prices for custom and own-account software has not been possible, investment in these categories of software is deflated by a blend of an input cost index for the industry and the prepackaged software index. In our liberal alternative scenario, we assume that price declines for the other components are understated and deflate own-account and custom software with an index created with one-third weight on prepackaged software and two-thirds weight on existing BEA deflators for the respective categories.[16]

IT INVESTMENT AS A WHOLE All told, declines for the official price index for information technology slow dramatically, from 6 percent a year for the period 1995–2004 to 1 percent a year for 2004–14. Although the alternative index consistently falls faster than the official price, it slows to a similar degree—from 9 percent a year for 1995–2004 to 4 percent a year for 2004–14. The liberal index accelerates as well, and provides essentially the

16. Byrne and Corrado (2016) have added estimates of an alternative price index for software since this paper was written. Their price index accelerates by roughly the same amount (1.9 percent) as the price index we employ (1.7 percent). Consequently, the contribution of IT price mismeasurement to the productivity slowdown would not change if we employed their index. Their price index falls 3 percentage points faster in both periods, implying a somewhat greater contribution to labor productivity of capital deepening and smaller contribution of TFP both before and after 2004, but roughly the same acceleration of TFP.

same picture. Thus, on first examination, increasing mismeasurement does not appear to explain the slowdown in IT price declines when the available research from all periods is considered.

However, it bears emphasis that the composition of IT investment has shifted appreciably toward components for which measurement is more uncertain. Most notably, software investment has gone from 39 percent of IT investment for the period 1995–2004 to 48 percent for 2005–14. Also, special-purpose equipment's share has increased, bringing the share for which measurement is more uncertain to 68 percent. Thus, our confidence in the IT price indexes, even as amended in the alternative indexes, has deteriorated markedly because of compositional shifts.

II.B. Intangibles beyond the NIPA

Conceptually, capital investment represents the use of resources that "reduces current consumption in order to increase it in the future" (Corrado, Hulten, and Sichel 2009, p. 666). Tangible investments in equipment and structures clearly meet this definition. But much intangible spending by businesses and governments also meets this definition. The U.S. national accounts include some intangibles—R&D and artistic originals (history beginning in 1925; introduced in 2013) and software (history beginning in 1960; introduced in 1999)—as final fixed capital formation. However, businesses also undertake considerable other types of spending that have the same flavor—such as training, reorganizations, and advertising.

Corrado, Hulten, and Sichel (2009) and Ellen McGrattan and Edward Prescott (2012) argue that investment spending has increasingly shifted toward intangibles, including those that are not currently counted. Susanto Basu and others (2004) argue that reorganizations associated with IT can explain some of the dynamics of measured U.S. and U.K. aggregate TFP growth.

In the next subsection, we consider the effects of incorporating additional intangibles from Corrado and Kirsten Jäger (2015). Their U.S. intangibles data run from 1997 to 2014. Ordered from largest to smallest estimated values in 2014, their data include investments in organizational capital; branding; training; design; and new finance and insurance products.

II.C. Capital Mismeasurement and TFP

To help interpret the counterfactuals in the next subsection, here we highlight the conceptual reason why capital mismeasurement is unlikely to explain the past slowdown in TFP growth: It affects inputs as well as output, in largely offsetting ways.

Consider a stylized example for a closed economy. Suppose that after some date in the past, we miss q percentage points of true investment growth. This miss could reflect an increase in unmeasured quality improvement (relative to whatever we were missing preceding that date) or an increase in the importance of unobserved intangible investment.

The growing mismeasurement implies that true output and true labor productivity grow at a rate $s_I q$ faster than measured, where s_I is the investment share of output and, by assumption, the good is completely produced domestically. It also implies that true capital input grows more quickly than measured. In a steady state, the perpetual inventory formula implies that capital grows at the same rate as investment, so capital input also grows q percent a year faster.

Thus, the change in TFP growth is the extra output growth less the contribution of the additional capital growth. In a steady state, the change is $(s_I - s_K)q$, where s_K is capital's share in production. In the data (and consistent with dynamic efficiency), $s_I < s_K$. Hence, in a steady state, capital mismeasurement makes true TFP growth *slower,* not faster, than measured.[17]

Of course, this is a steady-state comparison. The initial effect is that output responds more quickly than capital input, so TFP temporarily increases. Also, some domestically produced capital goods are exported, and some goods used for investment are imported. Which effect dominates over particular time frames is thus an empirical question.[18]

II.D. Mismeasurement of Durables Worsens the Slowdown: Evidence from Simulations

We now assess the quantitative importance of the mismeasurement of durable goods. As discussed above, this mismeasurement was large in the past, as well—and domestic production was more important. As a result of both factors, the mismeasurement of productivity appears less important now than in the past. As a result, with consistent measurement, the labor productivity slowdown after 2004 becomes even larger than in the official

17. Though not original to them, Basu and others (2004) make this point in the context of intangible investment. Dale Jorgenson had made this observation to Fernald when software investment was added to the U.S. GDP in 1999.

18. Note, as well, that the slower pace of aggregate TFP growth would be distributed unevenly. Suppose the mismeasurement reflects faster true TFP growth in domestic equipment and software goods. Then TFP growth in the other industries must be slower than measured. Intuitively, this happens because growth in their capital input is more rapid than measured, but growth in their output is the same as measured.

data. For TFP, the adjustments are more modest, but the slowdown is also a touch larger than in the official data.

We begin narrowly, with areas that are most grounded in a consistent methodology over time. This first conservative simulation considers alternative deflators for two categories of equipment for which considerable recent research has been done: computers and peripherals; and communications equipment (see the discussion in section II.A). We also consider alternative deflators for semiconductors. Those are primarily an intermediate input into other types of electronic goods but, because of exports and imports, revised deflators modestly affect final output growth. We then add more speculative adjustments for specialized equipment (NAICS category 3345) and software. Finally, we add estimates of intangibles from Corrado and Jäger (2015).

Given alternative deflators and measures of intangibles, we adjust both output and input (capital services). The online appendix describes the details. Output grows more quickly because of faster growth in domestically produced computers and other types of information-processing equipment. Of course, some of these products are sold to consumers. Hence, the output adjustment also captures the effect on real GDP of consumers' purchases of computers and communications equipment (such as mobile devices). Capital input grows more quickly because of the faster implied growth in investment in computers and other types of information-processing equipment (whether domestically produced or imported).

For semiconductors, the adjustment to output only matters for GDP through its effect on net exports. In a closed economy, an adjustment that raises the true output of semiconductors is exactly offset by higher true intermediate input usage of semiconductors—leaving GDP unchanged. However, in an open economy, semiconductors are exported and imported. We do not have separate adjusted prices for imported versus domestically produced semiconductors, so we assume that any adjustments are proportional.

Column 0 of table 2 shows our baseline from the published data. Measured labor productivity growth (top panel), capital deepening (middle panel), and TFP growth (bottom panel) sped up in the 1995–2004 period, but slowed thereafter. The slowdown in average annual labor productivity growth was about 1¾ percentage points. Some of this slowdown is explained by a reduced pace of capital deepening, leaving a slowdown in TFP growth of about 1¼ percentage points. Labor productivity growth is especially weak after 2010, though the growth accounting attributes this to the lack of capital growth relative to labor. Hence, TFP growth was about equally weak from 2004 to 2010 and from 2010 to 2014.

Table 2. Adjustments to Business Sector Growth Accounting[a]

			Annual percentage point change relative to baseline[b]		
Measure of growth	*Period*	*(0) Published baseline*[c]	*(1) Conservative*[d]	*(2) Liberal*[e]	*(3) Liberal + intangibles*[f]
Labor productivity	1978–95	1.50	0.12	0.21	0.30
	1995–2004	3.26	0.27	0.38	0.49
	2004–14	1.44	0.13	0.19	0.18
	2004–10	1.92	0.17	0.25	0.24
	2010–14	0.71	0.06	0.11	0.10
Capital-to-hours ratio	1978–95	2.20	0.27	0.52	0.66
	1995–2004	3.68	0.54	0.89	1.02
	2004–14	1.80	0.44	0.70	0.55
	2004–10	3.14	0.46	0.74	0.54
	2010–14	–0.22	0.41	0.63	0.58
Total factor productivity	1978–95	0.53	0.04	0.05	–0.01
	1995–2004	1.82	0.09	0.09	–0.08
	2004–14	0.49	–0.04	–0.07	–0.12
	2004–10	0.44	0.00	–0.02	–0.12
	2010–14	0.58	–0.10	–0.14	–0.12

Sources: Fernald (2014); Corrado and Jäger (2015).

a. Averages start in 1978 because of the availability of intangibles data.

b. Each column involves a separate, experimental adjustment to selected components of capital investment. The entries show the percentage point adjustment to business sector growth accounting components, relative to the unadjusted estimates in column 0.

c. Baseline (the business sector) measured as the percent change at an annual rate.

d. Alternative deflators for computers and communications.

e. Column 1, plus alternative deflators for specialized equipment and software.

f. Column 2, plus intangibles from Corrado and Jäger (2015).

Column 1 of table 2 then shows how results change relative to this baseline from adjusting computers, communications equipment, and semiconductors. As the top panel shows, these adjustments do affect labor productivity in a noticeable way. But the increase in the labor productivity growth rate is most pronounced for the 1995–2004 period, at just under 0.3 percentage point. After 2004, the alternative deflators add only a little more than 0.1 percentage point to growth. This reduced effect is due to the declining importance of domestic IT production relative to imports. Domestic production of computer and communications equipment amounted to 2.9 percent of nominal business sector value added in the late 1990s, but only 0.5 percent by 2014. A given amount of mismeasurement of computer and communications equipment therefore would have had a larger effect in the 1990s than today.

The middle panel of the table shows that the adjustments also have a substantial effect on capital services growth. Again, the major adjustment is in the 1995–2004 period, when prices, by any measure, were falling rapidly. The bottom panel shows that the effect on TFP growth is small, but it goes in the direction of exacerbating the post-2004 TFP slowdown. The adjusted TFP is a little stronger than measured in the 1995–2004 period, but a little weaker after 2004.

Column 2 of the table adds more speculative adjustments for specialized equipment and software, as described above. The upward boost to labor productivity is a bit larger in each period than in column 1. But again, the upward boost is larger in the 1995–2004 period than in the post-2004 period—this time by almost 0.2 percentage point. Adjusting capital goods, once again, turns out to exacerbate the slowdown in labor productivity growth. The bottom panel shows that the adjustments also modestly exacerbate the TFP slowdown.

Column 3 of the table adds intangibles from Corrado and Jäger (2015). With intangibles, the adjustments to labor productivity are even larger—but, again, the effects are largest in the 1995–2004 period. Together, the adjustments in column 3 add about 0.5 percentage point to labor productivity relative to the published data for 1995–2004. From 2004 to 2014, the adjustments add only 0.2 percentage point. Thus, the slowdown in labor productivity growth after the adjustments in column 3 is about 0.3 percentage point larger. For labor productivity, then, the adjustments taken together make the productivity slowdown markedly worse.

Other approaches to measuring intangibles—such as the more model-based approach of McGrattan and Prescott (2012)—might yield different results. Still, the results in column 3 suggest that the intangibles route is unlikely to alter the productivity slowdown.

Of course, the slowdown in capital growth, in the middle panel, also becomes much larger. As a result, in the bottom panel, the slowdown of TFP growth is affected by only a few basis points relative to the measured baseline. In particular, the adjustment subtracts 8 bp from TFP growth in the 1995–2004 period but then 12 bp during the 2004–14 period.[19] The

19. The careful reader will note that labor productivity growth for 1995–2004 is about 0.1 percentage point higher in column 3 than column 2, as is capital growth. So why does TFP growth fall, even though the labor productivity effect looks larger than the adjusted contribution of capital (capital's share times capital growth)? The reason is that, with intangibles, capital's share is also adjusted upward, and so the effect on TFP involves not just the adjustment to capital growth but also the adjustment to capital's share multiplied by (the new) capital growth rate. This effect can be a few tenths.

important takeaway is that correcting for capital goods mismeasurement does not resolve the post-2004 slowdown—if anything, it makes it worse.

We also experimented with an aggressive adjustment to software deflators after 2004, whereby true software prices are assumed to fall 5 percent a year faster than measured. This counterfactual captures the hypothesis that measurement has recently gotten worse, because only the post-2004 period is affected. Yet even this aggressive adjustment turns out to have relatively modest effects. The adjustment would add about 0.1 percentage point to labor productivity growth after 2004. Yet capital growth is also higher in this simulation, and TFP is little changed.

The alternative deflators in this section imply faster TFP growth for IT-producing industries, but slower TFP growth for IT-using industries (given that capital input grows more quickly without any adjustment in output growth). Nevertheless, as discussed in the appendix, the alternative deflators do not alter the broad-based nature of the TFP slowdown. With the alternative deflators, TFP growth for industries that produce IT and other types of investment goods slows sharply after 2004, as does TFP growth for other, non-investment-producing industries.

To summarize the takeaways from this section, prices for key capital goods are mismeasured, and this mismeasurement varies over time. However, the effects of mismeasurement on productivity have been less, rather than more, important since 2004. Including intangibles, our adjustments add about 30 bp to the slowdown in labor productivity but make the TFP slowdown only modestly larger.

Thus, if the productivity slowdown after the early 2000s indeed reflects mismeasurement, the source of this mismeasurement is not found in commonly studied IT durable goods. In the remainder of the paper, we find that the growing mismeasurement of Internet services, e-commerce, fracking, and globalization (shown as "other" in figure 1) can fill only a small part of the gap.

III. "Free" Digital Services

The benefits to consumer well-being from online information, entertainment, social connections, and the like are large (Goolsbee and Klenow 2006; Varian 2011; Brynjolfsson and Oh 2014). Nevertheless, these benefits do not change the fact that market sector TFP growth slowed broadly. Under long-standing national accounting conventions, the benefits are largely outside the scope of the market economy; as we discuss, given the small monetary size of the sector, it is very hard to bring many of the benefits inside the market boundary. The largest estimates of the gains are based on models of the time cost of using the Internet as an input into the home-based production of nonmarket services for one's own consumption. The gains from nonmarket production using the consumer's time are conceptually distinct from the gains in market sector output. And regardless of how they are treated, the nonmarket gains are not big enough to offset a significant fraction of the missing $3 trillion a year in business output from the productivity slowdown.

In the standard national accounts approach, none of the output of online service providers whose revenue comes from selling ads is included in the final consumption of households. Rather, their entire output is used for the intermediate consumption of the advertisers.

Drawing on an earlier body of literature on free broadcast television, Rachel Soloveichik (2015b) and Leonard Nakamura and Soloveichik (2015) propose an alternative approach that includes entertainment and information services supported by advertising in household final consumption. This approach prevents artificial changes in GDP when consumers switch between free and subscription-based media. The effect on the GDP growth rate turns out to be minuscule, however, because advertising tends to be a small and relatively stable share of GDP. Further, this alternative approach has no effect on the nominal value added of the business sector by construction, leaving little scope for an effect on business sector productivity. Our "other" category of adjustments in figure 1 therefore adds nothing to productivity growth in any of the periods for ad-supported digital services. Where we can get a small adjustment (about 1 bp from 1995 to 2004, and 4 bp from 2004 to 2014) is for the improved quality of Internet service providers (ISPs) that is not included in the official deflators.

III.A. The Time Cost Approach to Gains from Free Digital Services

The standard approach to measuring gains from new goods considers the difference between the amount of money that consumers would have been willing to pay and the amount that they actually had to pay. Yet the main cost to a user of, say, Facebook, YouTube, or TripAdvisor is the opportunity cost of the user's time. Hence, starting with Austan Goolsbee and Peter Klenow (2006), studies of the gains from free digital services have considered the time costs of using these services, and not only the money costs associated with accessing them.

Time costs are part of Becker's (1965) model of the allocation of time. Suppose the representative consumer has the following utility function:

$$U(Z_I, Z_{TV}, Z_1, Z_2, \ldots).$$

Households benefit from the consumption of (possibly unpriced) services from the Internet, Z_I, from television, Z_{TV}, and from other activities, Z_i, $i \subset \{1, 2, \ldots\}$. The elements of Z_i include meals at home, meals at restaurants, having a clean house, playing soccer, skiing, and so forth.

In this Becker-style model, the Z_i are not the direct purchases of market goods and services. Rather, households combine purchased market goods and services with their own time to generate the actual services they value. They buy a soccer ball (which is part of GDP), and they combine that market purchase with their (leisure) time, and their children's time, to obtain "soccer services." They combine a market purchase of a restaurant meal with several hours of their time. They combine gasoline and a car (both purchased in the market) with their time in order to go on a vacation that they enjoy. They combine a hotel room with their time to get a refreshing night of sleep during this vacation. Broadly, the services take the form

$$Z_i = Z_i(C_i, T_i, Q_i), \qquad i \subset \{I, TV, 1, 2, \ldots\}.$$

Thus, in the household's production function for combining the market purchase with time, playing soccer generates services from the market consumption of a soccer ball, C_i; the time spent playing soccer, T_i; and, possibly, technical change, Q_i.

Now consider a stylized problem that captures the key issues in valuing the Internet. Households seek to maximize their well-being subject to cash and time budget constraints:

$$\text{(1)} \qquad \max \quad U\begin{pmatrix} Z_I(C_I, (1-\tau_I)T_I, Q_I), Z_{TV}(C_{TV}, (1-\tau_{TV})T_I), \\ Z_1(C_1, T_1), Z_2(C_2, T_2), \ldots \end{pmatrix}$$

$$\text{(2)} \qquad \text{s.t.} \quad \textstyle\sum_i P_i C_i + F_I + F_{TV} = WT_{work},$$

$$\text{(3)} \qquad T_{work} + T_I + T_{TV} + \textstyle\sum_i T_i = 1.$$

In the cash budget constraint (equation 2), income is the wage, W, multiplied by time spent working, T_{work}. Households purchase broadband access, C_I, via cable, mobile phone, or another means by paying a fixed or flat cost, F_I, each period. In the time budget constraint (equation 3), total time is normalized to 1; in other words, time spent working is time *not* spent engaged in other activities. The Internet services that they actually value then depend on the time they spend online, T_I, net of a flow "time tax," τ_I, which is proportional to their use of the Internet. For example, they get "free" access to YouTube videos in exchange for spending a proportion of their time watching ads.

As Erik Brynjolfsson and Joo Hee Oh (2014) find, Internet content may get better over time, as captured in quality, Q_I. The quality of Internet content may reflect the growing number of websites available, the number of videos available on YouTube, or whether one's friends are on Facebook. These are conceptually distinct from download speed or other characteristics of one's ISP. And these characteristics conceptually represent a larger *quantity* of C_I. (As we discuss below, not all these characteristics are currently in the implicit deflator for Internet access.)

Television is similar to the Internet. One might pay a fixed cost for watching TV, F_{TV}, as well as paying a time tax, τ_{TV}, again in the form of watching ads. Historically, in the United States, before the inception of cable TV, $F_{TV} = 0$, the entire provision of broadcast TV service was paid for through watching ads. For other types of goods, C_i, the price is P_i.

This formulation illustrates the key issues, but it does make simplifications. For example, it ignores nonwage income, and also durable goods, such as computers, cell phones, TVs, and beds; it assumes that households are unconstrained in their time allocation, so that the marginal opportunity cost of time is the (fixed) wage; and it ignores any extra disutility associated with working or with other activities. Paul Schreyer and W. Erwin Diewert (2014) discuss extensions to Becker's (1965) framework.

It is useful to combine the money and time budget constraints as

$$(4) \qquad \left(\sum_i P_i C_i + F_I + F_{TV}\right) + W\left(T_I + T_{TV} + \sum_i T_i\right) = W.$$

"Full expenditure" in this formula is the sum of market expenditures (the first term in parentheses) and the monetary value of nonmarket expenditures of time (the second term). Some nonmarket expenditures could be on the home-based production of goods and services that are a close substitute for market goods and services, such as cooking and cleaning. Others are for leisure (surfing the Internet for personal reasons, watching TV, playing soccer, and so forth). Some are in the middle, such as Wikipedia, where unpaid content writers create and edit entries for their personal enjoyment, but it substitutes for market encyclopedia services.[20]

20. In "The GNU Manifesto," Richard Stallman (1985) describes his vision that "in the long run, . . . nobody will have to work very hard just to make a living. People will be free to devote themselves to activities that are fun, such as programming." (We thank Hank Farber for pointing us to this quotation.)

The core national accounts measure the prices and quantities that correspond to market activities, which show up in the first term in equation 4. Nevertheless, the importance of nonmarket activities, the second term, has long been recognized. After all, Americans ages 15 and older spend only 15 percent of their total time working, or 24 percent of the time not spent sleeping.[21] Katharine Abraham and Christopher Mackie (2005) and William Nordhaus (2006) discuss the need for nonmarket satellite accounts.

Based on increasing amounts of time spent online, Brynjolfsson and Oh (2014) estimate that the incremental consumer surplus from free digital services is sizable, averaging $25.2 billion for 2002–11, with larger effects in the years after 2005.[22] These incremental gains are the equivalent of adding about 0.3 percentage point a year to business sector output and productivity growth. Adding these gains is not appropriate, however, if the question is the productivity of the economy in producing market goods and services. The gains implied by changes in the allocation of consumers' time are linked to the home-based production of nonmarket services, not market output.

III.B. The Market Production of New Goods

In contrast to the time-based estimates of the value of free digital services, the standard approach used to define the theoretical measure of real GDP implies that only a small amount of extra digital service output is missed, mainly reflecting download speed and other characteristics that are not currently included in the deflators for Internet access and cell phone service.

Real household consumption and real GDP measure changes at the margin, not total amounts of consumer surplus. Hence, even if free digital services belonged in market sector GDP and provided a large amount of consumer surplus, the growth-rate effects would not necessarily be large. What would matter is the incremental consumer surplus from a change in the consumption of the digital services.

For existing goods, the BEA's chained Fisher index of real personal consumption expenditures correctly captures the change in the consumer

21. This is according to the American Time Use Survey (http://www.bls.gov/tus/tables/a1_all_years.xlsx).

22. As a nonprofit institution serving households, Wikipedia's output, about $0.2 billion in 2011, is counted as personal consumption. The $25.2 billion thus overstates the adjustment that could be made to GDP by $0.2 billion.

surplus.[23] For an existing free good, the correct weight on any change in quantity is zero because consumers adjust the quantity consumed of each good (excluding those at a corner solution of zero) so that the value of the marginal unit consumed is proportional to the price.

Conversely, new goods bias can arise even if the good enters at a price of zero. The measurement theory for new goods imagines that the new good previously existed but was offered at the "virtual price" that just drove demand to zero. The area under the demand curve from the virtual price down to the actual price of the good after it entered gives the consumer surplus from the appearance of the new good. Some major free digital services—including Facebook, YouTube, and Google Maps—appeared after the start of the productivity slowdown.

However, because they require Internet access, free digital services are not costless to consume. The price of the required Internet access can be viewed as the price of a bundled commodity, where the free digital services are part of the bundle. With an assumption about the slope and curvature of the demand curve for the bundled commodity, increased spending on Internet access to enjoy the new free services could be used to estimate the gains from this newly available, bundled commodity.

We do not make such an estimate in the present paper, but an indication of its magnitude comes from estimates of welfare gains from Internet access. Shane Greenstein and Ryan McDevitt (2009), for example, use data on the replacement of dial-up Internet access with broadband, and estimate that the uptake of broadband generated an average of $0.3 billion a year in unmeasured consumer surplus for 1999–2003, and an annual average of just over $1 billion for 2004–06. Brynjolfsson and Oh (2014) extend Greenstein and McDevitt's (2009) analysis, adding an adjustment for increased consumption of services per hour, as measured by rising data usage patterns. They find that this "money measure" of the gains from improved ISP services (the part that would be appropriate to add to market sector output) are a little larger, but still small—averaging only $2.7 billion per year (2–3 bp of business output).

This analysis of the monetary mismeasurement applies only to Internet access at home, not mobile access. Using Brynjolfsson and Oh's (2014) data on the improved quality of Internet access, and assuming that the increase in the mobile share since 2004 reflects mobile data that are subject

23. The online appendix shows that the Laspeyres and Paasche quantity indexes that are averaged to obtain the Fisher index are upper and lower bound measures of the relative change in consumer surplus.

to the same unmeasured quality improvement, "true" output and productivity thus rise by 1 bp in the 1995–2004 period and by 4 bp after 2004. We include this adjustment in the "other" category in figure 1.[24]

III.C. An Alternative Treatment of Advertiser-Supported Digital Services

Internet businesses make money in part by creating content that users value. Is it reasonable to exclude this entirely from GDP, just because it does not involve a monetary cost to households? We now consider an alternative that brings some of these otherwise-omitted, advertising-supported digital services into household consumption.

Some free digital services are, in fact, already included in GDP—namely, those provided by nonprofit institutions such as Wikipedia. But most free digital services are supported by advertising.[25] The national accounts treat advertisers as intermediate consumers of the services of a business whose revenue comes entirely from advertising. For example, broadcast television services have long been counted in the national accounts as an intermediate input: Companies buy advertising, so major broadcasting networks such as ABC or NBC are like advertising agencies. Many Internet services have that same treatment: Facebook and Google provide advertising services to businesses, not services consumed by households.

Nakamura and Soloveichik (2015) propose a framework for including ad-supported entertainment and information services in households' consumption that draws on an earlier body of literature on how to treat broadcast television in national accounts. They value the services given to households at their cost of production. This framework is based on the observation that consumers implicitly pay for TV entertainment and information by watching ads (or, in some cases, providing valuable personal information). The time taxes τ_I and τ_{TV} were not included in the cash budget constraint (equation 2) because they do not have an explicit price. But we can express WT_I (the time value associated with the Internet in equation 4) as $W\tau_I T_I + W(1 - \tau_I)T_I$, where the first term is part of a market-oriented barter transaction that can be imputed between households and firms. In this barter transaction, the time that consumers spend viewing

24. We thank Joo Hee Oh for sending these data. Our calculation corresponds to Brynjolfsson and Oh's (2014) equation 14 on the money benefits.

25. Another revenue source for providers of free digital services is the valuable information that the users of these services reveal about themselves, but this is a small revenue source compared with advertising.

ads is a service purchased from households by providing entertainment or information services.

When these "free" entertainment or information services are added to households' consumption, GDP goes up by the value of the extra household consumption. But the national accounts need to balance—someone needs to produce the extra value added. The TV networks or the providers of the digital services have the same inputs of capital and labor, and their measured value added does not change. Instead, on the production side, the rise in GDP can be traced to households' production of "ad-watching services." With no change in the output consumed by advertisers, recording output sold to households requires us to impute an equivalent amount of purchases of services from consumers who view the ads.

This approach is reasonable: It monetizes an implicit barter transaction that consumers undertake with Google and Facebook and other advertising-supported service providers, and it recognizes that consumers value the services they receive. Nonetheless, treating consumers as suppliers of ad-watching services and as consumers of free digital services does not change the business sector's nominal value added; the ad-watching services are outside the boundary of the business sector.[26] On one hand, if the deflators are the same, business sector TFP will also be unaffected because the intermediate inputs of the ad-watching services that are added on the input side of the productivity calculation will exactly offset the "free" entertainment and information services that are added on the output side. On the other hand, it is possible for the deflators to vary in a way that raises business TFP if ad viewing, and the delays caused by the time it takes to download the ads, take up a falling proportion of time spent consuming digital services.

III.D. The Significance of Free Digital Services for Productivity Measures

The effect on the level of GDP from allocating part of output of the providers of free entertainment and information services to household final consumption is limited because advertising is only a small share of GDP.

26. In another project we are working on (Byrne, Fernald, and Reinsdorf ongoing work), we discuss a way to bring the extra value added into the business sector, as opposed to being in the household sector. Their approach requires special treatment of the advertising revenue, so that some of the output that is currently viewed as intermediate consumption by the ad buyers can instead be viewed as consumed by households. This would make business sector nominal and real value added larger, but the effect on TFP growth is still close to zero. A separate issue is that a more explicit agreement for consumers to watch the ads in order to receive the services would be required for the ad watching to qualify as a barter transaction under the international guidelines of the *System of National Accounts* (United Nations and others 2009, para. 3.51 and 3.53, pp. 43–44).

When services to households from traditional print and broadcast media are included along with digital services, the level of U.S. GDP shifts up by about 0.5 percent (Soloveichik 2015a). The effect on the growth rate of real GDP is smaller still. In Nakamura and Soloveichik (2015, table 3), real advertising services have an average growth rate of 2 percent from 2004 to 2013, while the real output of the business sector used in productivity measurement grows at just over 1.5 percent a year. Assuming that the real growth rate of the advertising-supported services was the same as the real growth rate of the advertising and using a share weight of 1.3 percent of business sector output implies an upward revision of less than 1 bp to productivity growth in the slowdown period. But the pre-slowdown adjustment is similar or larger, so this adjustment does not reduce the size of the productivity slowdown. In our benchmark set of "other" adjustments, we round the effect to zero.

How sensitive is this benchmark to the advertising deflator? This deflator may have a new goods bias caused by the emergence of online advertising, if that is a more efficient technology for delivering ads. Soloveichik's (2015a) estimate of the 2012 cost per viewer-hour of an online advertisement is 11 cents, compared with 54 cents for broadcast TV. The lower cost of attracting ad viewers by providing free digital services suggests that the substitution of online advertising for traditional media advertising may involve a productivity gain in ad delivery. Facebook, for example, does not have to pay to acquire content because consumers themselves create the content, making the cost of attracting users to the website quite low.

Suppose the quality-adjusted price for online advertising is half that of traditional media. A unit-value price index would capture the outlet substitution effect as ad buyers switch to online advertising. The market share of online advertising rose from 7 percent in 2004 to 27 percent in 2013 (Nakamura and Soloveichik 2015), implying an average annual growth rate adjustment of −1.9 percent. With a 1.3 percent weight of advertising in the output of business, the implied annual adjustment to productivity growth would be an increase of about 2.5 bp.

Finally, we note that some of the welfare benefits of free digital services involve better choices of where and what to buy. Information from TripAdvisor or Yelp may improve restaurant selection (and even have dynamic spillover effects as bad restaurants improve or exit). In addition, online information and online shopping have expanded the set of available varieties. Moreover, the Internet has also led to new markets for used goods through websites such as eBay and Craigslist. A cost-of-living index that measured the gains from the improved matching of products and product varieties to consumers' preferences and circumstances might show substantial gains (even beyond the e-commerce benefits discussed below). Making more efficient use of what we have raises welfare, but does not represent an outward shift in market output or even the production possibility frontier that is achievable with a given factor endowment.[27]

Divergences between welfare change and real GDP from IT-enabled shifts between market production and nonmarket production also arise in other contexts. For example, tax software has reduced reliance on paid tax preparers, and smartphone apps such as Skype and WhatsApp have reduced spending on phone calls and text messages. Yet it is worth remembering that welfare changes from substitution between nonmarket and market activity are not new. In the early 20th century, for example, paid domestic workers did many tasks that by mid-century had been taken over by the households themselves. Conversely, home appliances such as washing machines served as "engines of liberation" (Greenwood, Seshadri, and Yorukoglu 2005) that dramatically increased women's labor force participation.

Furthermore, though nonmarket and market production are somewhat substitutable in generating consumer welfare, many questions about economic growth require a concept of productivity that covers only the market sector's output and inputs. Imputations for nonmarket output would make the productivity measure more subjective and model-driven, as opposed to data-driven. Gains in nonmarket output and their contribution to welfare, though important, are best treated as a separate concept from productivity change.

IV. E-Commerce and Gains in Variety and Match Quality

E-commerce has grown rapidly in importance, both for business-to-business and business-to-consumer transactions. In this section, we estimate that the growing unmeasured benefits to consumers contribute about 2 bp to the productivity slowdown. Business-to-business e-commerce has made intermediate transactions more efficient, but in principle it does not directly cause the mismeasurement of aggregate productivity. Indirectly, however, it can complicate productivity measurement through its effects on

27. A well-known example—the need to distinguish between increases in welfare and productivity gains—occurs with changes in terms of trade. Favorable shifts in exports and imports increase the opportunity to gain from trade, allowing real consumption to rise as the economy moves to a different point on the production possibility frontier.

outsourcing and the reorganization of production into global supply chains. The next section considers these effects in the context of globalization.

According to the Census Bureau's *Monthly Retail Trade Report,* the share of e-commerce in retail sales has risen about 0.5 percentage point a year since 2000—from 0.9 percent in 2000 to 2.1 percent in 2004, 5.3 percent in 2012, and 7.3 percent in 2015.[28] This steady shift in purchasing patterns reflects the gains to consumers in savings of time and transportation costs, as well as their ability to search over a much broader range of varieties.

For online books, Brynjolfsson, Yu Hu, and Michael Smith (2003) estimate the gains from increases in variety available on Amazon and other websites. They consider obscure book titles as new goods, because these would have been hard to find at brick-and-mortar stores. The compensating variation from a new good with a constant price elasticity of $\alpha < -1$ can be approximated by dividing its post-entry sales by $1 + \alpha$. In 2000, out of \$24.59 billion in total book sales, the authors estimate that \$578 million were from online purchases of obscure titles. Depending on the assumed elasticity, the compensating variation was in the range of \$731 million to \$1.03 billion, or about 3 to 4 percent of total book sales that year.

This approach probably overestimates the gains by assuming a constant demand elasticity and by ignoring losses in consumer surplus from the disappearance of brick-and-mortar bookstores. Robert Feenstra (1994) derives a more conservative formula for the unmeasured gains from net variety growth based on a model with a constant elasticity of substitution $\sigma > 1$. Let λ_t equal 1 minus the share of expenditures in period t going to new varieties, and let λ_0 be 1 minus the share of expenditures in period 0 going to varieties that disappear in period 1. Then the welfare shift from changing the availability of varieties can be calculated by multiplying the constant elasticity of the substitution price index for the continuing varieties by a factor of $(\lambda_t/\lambda_0)^{1/(\sigma-1)}$. The elasticity of substitution between different varieties of the same good is usually high. With $\sigma = 4$ and an assumption of no variety disappearances, the 2.35 percent market share garnered by obscure book titles newly made accessible by the Internet implies a

28. The Census Bureau defines e-commerce as purchases made over the Internet or other electronic network or via email. The e-commerce shares for products that are easy to order online (such as books) are even larger, because some products (such as gasoline and building supplies) presumably involve little e-commerce (https://www.census.gov/retail/mrts/www/data/excel/tsnotadjustedsales.xls). Evans, Schmalensee, and Murray (2016) conjecture that the Census figures underestimate e-commerce.

correction to the price index of –0.8 percent—though, with a relatively low assumption of $\sigma = 3$, the bias becomes 1.2 percent. These gains accumulated during a period of several years, so the annual bias is smaller.

Books, of course, are just one type of good with an increased availability of varieties online. Suppose we view e-commerce itself as a sort of new variety. Using the Census shares and assuming $\sigma = 4$, the correction factor to the price index for retail goods falls 15 bp a year from 2004 to 2014, compared with 8 bp a year from 1995 to 2004 (assuming the online share was zero in 1995). Personal consumption expenditures on goods amount to about 25 percent of the gross value added of business, excluding housing. Using this as a weight on the bias in the retail sales price index implies an upward correction of just under 4 bp a year in business sector productivity after 2004 and about 2 bp a year from 1995 to 2004. Thus, correcting for gains from e-commerce shaves perhaps 2 bp from the productivity slowdown.

V. Fracking and Globalization

Fracking and globalization are two areas where mismeasurement has plausibly contributed in a meaningful way to the slowdown in measured productivity growth. Fracking is a technological innovation that makes it profitable for drillers to access natural resources of an inferior "quality." A back-of-the-envelope calculation suggests that the unmeasured aspects of this innovation have raised true aggregate labor and TFP growth by about 5 bp a year since 2004. For globalization, import-price declines from offshoring and related changes in import sourcing are largely missed, so true import growth is understated in the late 1990s and early 2000s (the time of China's accession to the World Trade Organization); correspondingly, growth in GDP, labor productivity, and TFP are overstated. This globalization adjustment shows up as a negative contribution of about 10 bp from 1995 to 2004 and –2 bp from 2004 to 2014 for the "other" category in figure 1.

V.A. Technological Innovation in Oil and Natural Gas: The Fracking Revolution

In the industry TFP data discussed in section I, the extraction of oil and natural gas performed strongly in TFP during the 2007–13 period. Nevertheless, the standard measure of TFP for mining does not control for variation in the quality of the natural resources being extracted, so it is not a pure measure of technology. Technological innovations that made it possible to

Figure 7. U.S. Production of Oil and Natural Gas, 1995–2014

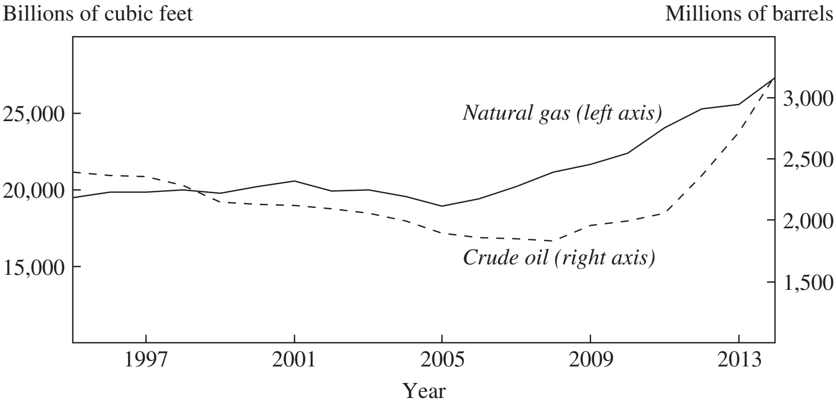

Source: U.S. Energy Information Administration.

extract oil and natural gas from previously uneconomic geologic formations diffused rapidly in the 2000s. This type of technological change is unlikely to be fully reflected in the statistics. Hence, true growth in mining investment in infrastructural capital is almost surely faster than measured. At the same time, a key input (the subsoil reserves component of land) that is not included in the traditional approach to measuring mining productivity fell in quality.

Fracking—originally a cost-effective way to extract natural gas from shale, using horizontal drilling and hydraulic fracturing—was discovered in the late 1990s. During the next decade, this technique was improved and extended to the extraction of oil from shale and other types of low-permeability formations. As a result, the last half of the 2000s saw a remarkable resurgence in the production of oil and natural gas in the United States (figure 7). Import facilities for liquefied natural gas have been hastily repurposed as export facilities, and OPEC has changed its pricing strategy.

The fracked wells are like a new good whose benefits are not counted by conventional measures of TFP. Nordhaus and Edward Kokkelenberg (1999, pp. 63–64) observe that deposits of an exhaustible natural resource vary in their extraction costs. Above some cutoff level of rent (the difference between the extraction cost and the market price of output), extraction does not occur. Suppose that technological progress reduces the unit cost of extraction for all deposits. Now, $\pi > 1$ units can be extracted from any given deposit in period 1, with the same inputs of labor and capital that produced 1 unit in period 0. The output price is set on world markets and does not change, and neither does the cutoff level of rent for extraction to be undertaken. Deposits that were previously uneconomic now begin to be extracted. The level of productivity at the least productive establishment remains constant, though that of the most productive establishment rises from λ_0^{max} to $\lambda_1^{max} = \pi\lambda_0^{max}$. Assuming productivity levels are uniformly distributed across establishments from 1 to λ_0^{max}, and that all establishments are identical in size as measured by inputs, measured productivity growth for the industry, denoted by $\hat{\pi} - 1$, is

$$\hat{\pi} - 1 = \frac{1 + \pi\lambda_0^{max}}{1 + \lambda_0^{max}} - 1 = \frac{\lambda_0^{max}}{1 + \lambda_0^{max}}(\pi - 1).$$

For example, if $\lambda_0^{max} = 2$, only two-thirds of true productivity gains would be measured.

A proper accounting for the quality of land as a factor of production would capture the gains. Deteriorating land quality would imply slower growth of inputs than in the official data, and TFP would grow faster.[29] Careful measurement of land services in mining—and elsewhere—is challenging. In the BLS productivity data, the extraction of oil and natural gas appears to use almost no land, because the value of rights to extract subsoil mineral deposits is included in the services of fixed capital assets (which consist largely of structures). Alternative productivity measures for Australian mining, published by the Australian Bureau of Statistics, imply that roughly half the conventionally measured services of fixed capital assets actually represent services of subsoil natural resources.[30] We assume that this relationship holds for the United States.

Accounting properly for technological progress in the oil and natural gas industries requires not only an assessment of land-quality changes but also quality-adjusting the fixed assets that embody the technological

29. Schreyer, Brandt, and Zipperer (2015) discuss an alternative approach to measuring multifactor productivity (MFP) for mining that includes services of natural resource assets. The Australian Bureau of Statistics publishes an experimental measure of MFP for mining that includes services of subsoil natural resource assets in inputs. In the tables released in December 2015, this raises the estimated growth rate of mining MFP between 2000–01 and 2014–15 from –4 percent a year to –1 percent a year. Similarly, Zheng and Bloch (2014) find that adjusting for inputs of natural resources, declining returns to scale, and capacity utilization raises MFP growth for the mining industry of Australia between 1974–75 and 2007–08 from –0.2 percent a year to 2 percent a year.

30. These data were downloaded at http://www.abs.gov.au/AUSSTATS/abs@.nsf/DetailsPage/5260.0.55.0022014-15.

improvements. These consist primarily of oil and gas wells drilled for exploration or development purposes. The quality adjustment would reflect the cost reduction made possible by better technology while holding constant the mix of deposits being exploited.

In the post-2004 period, the average share of investment in oil and natural gas structures in the value added of business is about 0.9 percent. However, plausibly about half of that—or about 0.5 percent of business output—is the structure itself (which is improving more quickly than measured); the remainder is actually the subsoil asset (where the quality is getting worse). In terms of output (i.e., final investment), suppose there is a fairly large true-quality adjustment to the price index for oil and gas extraction structures of 10 percent a year after 2004. Multiplying this by the roughly 0.5 percent share of business value added implies that the true investment is about 5 bp faster. This goes directly into the "other" portion shown in figure 1 above, boosting true labor productivity in the post-2004 period. For TFP, the question is how much capital is improving and land is deteriorating. As a rough first pass, we assume that the two effects offset each other—leaving measured capital growth about right. In this case, the increment to labor productivity of 5 bp also passes through to aggregate TFP.

V.B. Globalization

Standard techniques for constructing price indexes do not capture the change in the average price paid by the buyers of a product when they alter their purchasing patterns to buy from a different seller. Similarly, import price deflators in the NIPA do not capture changes in the price paid by buyers when they switch from a domestic source to an offshore producer. As a result, when sourcing moves offshore or to a different import-supplying country, real imports are understated and real output is overstated.[31]

This bias was particularly significant in the late 1990s and early 2000s, when the location of many kinds of manufacturing was shifting rapidly from the United States and other countries with high labor costs to emerging market economies. One impetus for this was China's 2001 accession to the World Trade Organization, which coincided with the start of a large shift in the sourcing for many manufactured goods used in the United States to China. Another was a multilateral free trade agreement that reduced tariffs for IT products to zero for an interval of four years ending in 2000, which

31. This problem is examined by Houseman and others (2011) and Mandel (2009).

accelerated international sourcing changes for IT products (Feenstra and others 2013).

Reinsdorf and Robert Yuskavage (forthcoming) use two approaches to estimate the sourcing bias for imported consumer goods in the 1997–2007 period, and they find a bias in the range of 0.8 to 1 percent a year for durable goods, including computers, and about 0.6 percent a year for imported apparel and footwear; after 2007, the effect is small. However, even if we assume that the bias estimate of 1 percent a year for durable goods can be generalized to similar kinds of imported capital goods and that the bias in the apparel index can be generalized to all textile products, the upward bias in business TFP is only 0.1 percent a year because the affected imports have only a small weight in GDP. This globalization adjustment shows up as a negative contribution of 10 bp from 1995 to 2004 for the "other" category in figure 1.

Another aspect of globalization made possible by reduced communications costs is international trade in services over a wire. The number of American jobs that could potentially be offshored to a country with lower wages is large (Blinder 2009), and the offshoring of services could lead to the same sort of upward bias in measures of productivity that is caused by the offshoring of goods. Thus far, however, the effects have been modest; the BEA's input/output accounts show that the imports of business-process services—such as professional, scientific, and technical services, and computer systems services—rose from about 2 percent of total intermediate uses of these services in 1997–98 to about 5 percent in 2010–14.

V.C. The "Sharing" Economy

Nominal GDP includes transactions from the sharing economy, such as car rides via Uber and Lyft.[32] Nevertheless, it is unlikely that the deflator used to compare the new services to previously existing ones correctly measures the decline in the quality-adjusted price experienced by many consumers. Thus, there is probably some (at this point very, very small, but likely growing) downward bias in the growth rate of real GDP.

It would be useful to have official statistics on the nominal output of the various types of services included in the sharing economy. Research indexes of price changes could then be developed to try to calibrate the size of the bias.

32. Where they are in the source data is not clear. However, the Quarterly Services Survey indicates slowing nominal growth of the local transportation measure that includes taxis, which is where one would expect to find the new kinds of local transportation services.

VI. Conclusions

The "productivity paradox 2.0" remains alive: Despite ongoing IT-related innovation, aggregate U.S. productivity growth slowed markedly after about 2004. To investigate this paradox, we propose several adjustments to IT-related hardware, software, and services. The good news is that these adjustments would make recent growth in GDP and investment look modestly better than recorded. The bad news is that they would make the paradox even worse—the slowdown in labor productivity is even larger after our durable goods adjustments, while the slowdown in TFP is not much affected. The reason is that mismeasurement was substantial in the 1995–2004 period, as well as more recently, and rising import penetration for computers and communications equipment means that domestic production (which matters for GDP) has fallen over time.

Moreover, the slowdown was broad-based, which suggests that ongoing innovation in IT is not substantially spilling over into other areas. Other measurement challenges—such as digital services, globalization, and fracking—go in the right direction but are small.

Other evidence also suggests that true underlying growth is relatively modest. First, the U.S. productivity slowdown has been mirrored in many parts of the world (Eichengreen, Park, and Shin 2015; Cette, Fernald, and Mojon 2016; Askenazy and others 2016). This suggests that underlying macroeconomic factors may be driving the slow pace of growth, given the varied sources and methods used across national statistical systems. Chad Syverson (2016) finds that the slowdown across countries is not correlated with IT production or use, again suggesting that the problem is not related to the mismeasurement of IT goods or services. Second, the decline in economic dynamism—both in the form of fewer start-ups and in the slower reallocation of labor resources in response to productivity shocks—supports the idea that productivity-enhancing innovations are diffusing throughout the economy more slowly (Decker and others 2016); also, the subdued pace of investment has slowed the adoption of new technology embodied in capital. And third, Michael Mandel (2016) also finds little evidence of widespread rapid innovation in an analysis of labor market metrics such as occupational employment and help-wanted ads, although there are tremendous occupational changes in some narrow segments of the economy, such as IT and the extraction of oil and natural gas.

If not mismeasurement, why did productivity growth slow? The slowdown predated the Great Recession, which suggests that the event was not the story—or, at least, not the whole story. Given that the 1970s and

1980s had slow growth similar to the period after 2004, the fast-growth 1995–2004 period looks like the anomaly. With the emergence of the Internet, the reorganization of distribution sectors, and IT investments beginning to pay off, many things came together in a short period. With hindsight, this looks like a one-time upward shift in the level of productivity rather than a permanent increase in its growth rate. Looking forward, we could get another wave of the IT revolution. Indeed, it is difficult to say with certainty what may yet come from cloud computing, the Internet of things, and the radical increase in mobility from smartphones. However, we have not yet seen those gains.

Changes in overall welfare are somewhat harder to assess. Transformative gains related to mobile technologies and the Internet clearly raise welfare. Most of these gains properly belong outside the purview of market sector GDP—and proposals to incorporate them into GDP raise concerns. Still, these innovations are valued by households. That said, the available estimates of the welfare gains (based on the value of leisure time) suggest that "free" digital services add the equivalent of perhaps 0.3 percent of GDP a year to well-being; that is small relative to the roughly 1.75 percent slowdown in labor productivity growth in the business sector from 2004 to 2014.

Nevertheless, much is unknown. Matthew Shapiro and David Wilcox (1996) described the area of the quality adjustment of price indexes as "house-to-house" combat in national accounting. The analysis must be done product by product, and statistical agencies are usually playing catch-up.[33] Digital services and the ensuing new modes of delivery of other types of services are particularly challenging. For example, research on the quality-adjusted price indexes associated with changes in the organization of production caused by the digital economy, such as the substitution of Uber and Lyft for traditional taxi services, would be helpful. Satellite accounts could also help to shed light on gains from digital services and shifting of production outside the market boundary by presenting measures of economic activity that extend beyond the market.

Finally, we conclude by touching on the implications for policymakers. Slow productivity growth, if it persists, implies slow future potential growth. The benefits in the nonmarket sector can offset this somewhat vis-à-vis well-being, but they do not help with taxes or the budget.

33. Wasshausen and Moulton (2006) discuss how statistical agencies incorporate quality adjustments.

ACKNOWLEDGMENTS We thank Martin Baily, Erik Brynjolfsson, Carol Corrado, Erwin Diewert, Jean Flemming, Robert Gordon, Peter Klenow, James Stock, and Hal Varian for helpful comments and conversations. We thank Travis Adams, Genevieve Denoeux, and Arthi Rabbane for excellent research assistance. The title of this paper refers to a phrase in a *Wall Street Journal* article by Timothy Aeppel (2015). The views expressed in this paper are those of the authors and should not be attributed to the International Monetary Fund or its managers or executive directors, or to the Federal Reserve Bank of San Francisco, the Federal Reserve System's Board of Governors, or the Federal Reserve System.

References

Abraham, Katharine G., and Christopher Mackie, editors. 2005. *Beyond the Market: Designing Nonmarket Accounts for the United States.* Washington: National Academies Press.

Aeppel, Timothy. 2015. "Silicon Valley Doesn't Believe U.S. Productivity Is Down." *Wall Street Journal,* July 16.

Aizcorbe, Ana, and Yvon Pho. 2005. "Differences in Hedonic and Matched-Model Price Indexes: Do the Weights Matter?" Working Paper no. WP2006-06. Washington: U.S. Department of Commerce, Bureau of Economic Analysis.

Askenazy, Philippe, Lutz Bellmann, Alex Bryson, and Eva Moreno Galbis, editors. 2016. *Productivity Puzzles across Europe.* Oxford University Press.

Baily, Martin Neil, and Robert J. Gordon. 1988. "The Productivity Slowdown, Measurement Issues, and the Explosion of Computer Power." *Brookings Papers on Economic Activity,* no. 2: 347–420.

Basu, Susanto, John G. Fernald, Nicholas Oulton, and Sylaja Srinivasan. 2004. "The Case of the Missing Productivity Growth, or Does Information Technology Explain Why Productivity Accelerated in the United States but Not in the United Kingdom?" *NBER Macroeconomics Annual* 18: 9–82.

Baumol, William J., and William G. Bowen. 1966. *Performing Arts: The Economic Dilemma—A Study of Problems Common to Theater, Opera, Music, and Dance.* New York: Twentieth Century Fund.

Becker, Gary S. 1965. "A Theory of the Allocation of Time." *Economic Journal* 75, no. 299: 493–517.

Berndt, Ernst R., and Neal J. Rappaport. 2001. "Price and Quality of Desktop and Mobile Personal Computers: A Quarter-Century Historical Overview." *American Economic Review* 91, no. 2: 268–73.

———. 2003. "Hedonics for Personal Computers: A Reexamination of Selected Econometric Issues." Paper presented at R&D, Education, and Productivity: An International Conference in Memory of Zvi Griliches (1930–1999), a meeting of the Conference on Research in Income and Wealth, which is administered by the National Bureau of Economic Research, Paris, August 25–27. http://www.nber.org/CRIW/papers/berndt.pdf

Bils, Mark, and Peter J. Klenow. 2001. "Quantifying Quality Growth." *American Economic Review* 91, no. 4: 1006–30.

Blinder, Alan S. 2009. "How Many US Jobs Might Be Offshorable?" *World Economics* 10, no. 2: 41–78.

Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale Jorgenson. 1998. "Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index." In *Getting Prices Right: The Debate over the Consumer Price Index,* edited by Dean Baker. Armonk, N.Y.: M. E. Sharpe.

Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Michael D. Smith. 2003. "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers." *Management Science* 49, no. 11: 1580–96.

Brynjolfsson, Erik, and Joo Hee Oh. 2014. "The Attention Economy: Measuring the Value of Free Digital Services on the Internet." Working paper. https://www.scribd.com/doc/314594157/1-The-Attention-Economy-Measuring-the-Value-of-Free-Goods-on-the-Internet

Byrne, David. 2015a. "Domestic Electronics Manufacturing: Medical, Military, and Aerospace Equipment and What We Don't Know about High-Tech Productivity." FEDS Notes. Washington: Board of Governors of the Federal Reserve System.

———. 2015b. "Prices for Data Storage Equipment and the State of IT Innovation." FEDS Notes. Washington: Board of Governors of the Federal Reserve System.

Byrne, David, and Carol Corrado. 2015. "Recent Trends in Communications Equipment Prices." FEDS Notes. Washington: Board of Governors of the Federal Reserve System.

———. 2016. "ICT Prices and ICT Services: What Do They Tell Us about Productivity and Technology?" Working paper.

Byrne, David M., John G. Fernald, and Marshall Reinsdorf. Ongoing work. "Can We Bring Advertising-Supported Activity into Business-Sector Value Added?" Unpublished manuscript.

Byrne, David M., Stephen D. Oliner, and Daniel E. Sichel. 2015. "How Fast Are Semiconductor Prices Falling?" Working Paper no. 21074. Cambridge, Mass.: National Bureau of Economic Research.

Byrne, David, and Eugenio Pinto. 2015. "The Recent Slowdown in High-Tech Equipment Price Declines and Some Implications for Business Investment and Labor Productivity." FEDS Notes. Washington: Board of Governors of the Federal Reserve System.

Cette, Gilbert, John Fernald, and Benoît Mojon. 2016. "The Pre–Great Recession Slowdown in Productivity." Working Paper no. 2016-08. Federal Reserve Bank of San Francisco. (Forthcoming in *European Economic Review.*)

Cole, Rosanne, Y.C. Chen, Joan A. Barquin-Stolleman, Ellen Dulberger, Nurhan Halvacian, and James H. Hodge. 1986. "Quality-Adjusted Price Indexes for Computer Processors and Selected Peripheral Equipment." *Survey of Current Business* 66, no. 1: 41–50.

Copeland, Adam. 2013. "Seasonality, Consumer Heterogeneity and Price Indexes: The Case of Prepackaged Software." *Journal of Productivity Analysis 39,* no. 1: 47–59.

Corrado, Carol, Charles Hulten, and Daniel Sichel. 2009. "Intangible Capital and U.S. Economic Growth." *Review of Income and Wealth* 55 no. 3: 661–85.

Corrado, Carol, and Kirsten Jäger. 2015. "Wealth and Investment in Mature Societies." Presentation given at Milestone 3: Midterm Conference, Smart Public INTANgibles (SPINTAN), London, April 23–24. http://www.spintan.net/wp-content/uploads/public/Carol-Corrado-Wealth-and-Investment_carol-corrado.pdf

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2016. "Declining Business Dynamism: What We Know and the Way Forward." *American Economic Review* 106, no. 5: 203–07.

Domar, Evsey D. 1961. "On the Measurement of Technological Change." *Economic Journal* 71, no. 284: 709–29.

Doms, Mark. 2005. "Communication Equipment: What Has Happened to Prices?" In *Measuring Capital in the New Economy,* edited by Carol Corrado, John Haltiwanger, and Daniel Sichel. University of Chicago Press.

Dulberger, Ellen R. 1989. "The Application of a Hedonic Model to a Quality-Adjusted Price Index for Computer Processors." In *Technology and Capital Formation,* edited by Dale W. Jorgenson and Ralph Landau. MIT Press.

Eichengreen, Barry, Donghyun Park, and Kwanho Shin. 2015. "The Global Productivity Slump: Common and Country-Specific Factors." Working Paper no. 21556. Cambridge, Mass.: National Bureau of Economic Research.

Evans, David S., Richard Schmalensee, and Scott R. Murray. 2016. "The Census Bureau Needs to Significantly Revise Reporting and Calculation of Its Online and Physical Retail Sales Figures and Commission an Independent Review." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2728918

Feenstra, Robert C. 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review* 84, no. 1: 157–77.

Feenstra, Robert C., Benjamin R. Mandel, Marshall B. Reinsdorf, and Matthew J. Slaughter. 2013. "Effects of Terms of Trade Gains and Tariff Changes on the Measurement of US Productivity Growth." *American Economic Journal: Economic Policy* 5, no. 1: 59–93.

Feldstein, Martin. 2015. "The U.S. Underestimates Growth." *Wall Street Journal,* May 18.

Fernald, John G. 1999. "Roads to Prosperity? Assessing the Link between Public Capital and Productivity." *American Economic Review* 89, no. 3: 619–38.

———. 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Working Paper no. 2012-19. Federal Reserve Bank of San Francisco.

———. 2015. "Productivity and Potential Output before, during, and after the Great Recession." *NBER Macroeconomics Annual* 29: 1–51.

Flamm, Kenneth. 1989. "Technological Advance and Costs: Computers versus Communications." In *Changing the Rules: Technological Change, International Competition, and Regulation in Communications,* edited by Robert W. Crandall and Kenneth Flamm. Brookings Institution Press.

Goolsbee, Austan, and Peter J. Klenow. 2006. "Valuing Consumer Products by the Time Spent Using Them: An Application to the Internet." *American Economic Review* 96, no. 2: 108–13.

Gordon, Robert J. 1990. *The Measurement of Durable Goods Prices.* University of Chicago Press.

———. 2015. "Productivity, Prices, and Measurement." Presentation at Boskin-Fest conference, Stanford University, October 2. https://www.scribd.com/doc/314597445/Productivity-Prices-and-Measurement

———. 2016. *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War.* Princeton University Press.

Greenstein, Shane, and Ryan C. McDevitt. 2009. "The Broadband Bonus: Accounting for Broadband Internet's Impact on U.S. GDP." Working Paper no. 14758. Cambridge, Mass.: National Bureau of Economic Research.

Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. 2005. "Engines of Liberation." *Review of Economic Studies* 72, no. 1: 109–33.

Griliches, Zvi. 1994. "Productivity, R&D, and the Data Constraint." *American Economic Review* 84, no. 1: 1–23.

Grimm, Bruce T. 1996. "A Quality-Adjusted Price Index for Digital Telephone Switches." Working paper. Washington: U.S. Department of Commerce, Bureau of Economic Analysis.

Grimm, Bruce T., Brent R. Moulton, and David B. Wasshausen. 2005. "Information-Processing Equipment and Software in the National Accounts." In *Measuring Capital in the New Economy,* edited by Carol Corrado, John Haltiwanger, and Daniel Sichel. University of Chicago Press.

Hatzius, Jan, and Kris Dawsey. 2015. "Doing the Sums on Productivity Paradox v2.0." *U.S. Economics Analyst* (Goldman Sachs), no. 15/30.

Hausman, Jerry. 1999. "Cellular Telephone, New Products, and the CPI." *Journal of Business and Economic Statistics* 17, no. 2: 188–94.

Heather, James M., and Benjamin Chain. 2016. "The Sequence of Sequencers: The History of Sequencing DNA." *Genomics* 107, no. 1: 1–8.

Holdway, Michael. 2001. "Hedonic Models in the Producer Price Index (PPI)." Washington: U.S. Bureau of Labor Statistics. http://data.bls.gov/cgi-bin/print.pl/ppi/ppicomqa.htm

Houseman, Susan, Christopher Kurz, Paul Lengermann, and Benjamin Mandel. 2011. "Offshoring Bias in U.S. Manufacturing." *Journal of Economic Perspectives* 25, no. 2: 111–32.

Ip, Greg. 2015. "Beyond the Internet, Innovation Struggles." *Wall Street Journal,* August 12.

Kim, Mina, and Marshall B. Reinsdorf. 2015. "The Impact of Globalization on Prices: A Test of Hedonic Price Indexes for Imports." In *Measuring Globalization: Better Trade Statistics for Better Policy* 1, edited by Susan N. Houseman and Michael Mandel. Kalamazoo: Upjohn Institute Press.

Mandel, Michael. 2009. "Growth: Why the Stats Are Misleading." *Bloomberg Businessweek,* June 3.

———. 2016. "Where Is Innovation Falling Short? Using Labor Market Indicators to Map the Successful Innovation Frontier." In *New Entrepreneurial Growth Agenda.* Kansas City: Ewing Marion Kauffman Foundation. http://www.kauffman.org/neg/section-3

McCallum, John C. 2015. "Disk Drive Prices (1955–2015)." http://www.jcmit.com/diskprice.htm

McGrattan, Ellen R., and Edward C. Prescott. 2012. "The Great Recession and Delayed Economic Recovery: A Labor Productivity Puzzle?" In *Government Policies and the Delayed Economic Recovery,* edited by Lee E. Ohanian, John B. Taylor, and Ian J. Wright. Stanford, Calif.: Hoover Institution Press.

Nakamura, Emi, and Jón Steinsson. 2012. "Lost in Transit: Product Replacement Bias and Pricing to Market." *American Economic Review* 102, no. 7: 3277–316.

Nakamura, Leonard, and Rachel H. Soloveichik. 2015. "Valuing 'Free' Media across Countries in GDP." Working Paper no. 15-25. Federal Reserve Bank of Philadelphia.

Nordhaus, William D. 2002. "Productivity Growth and the New Economy." *Brookings Papers on Economic Activity,* no. 2: 211–44.

———. 2006. "Principles of National Accounting for Nonmarket Accounts." In *A New Architecture for the U.S. National Accounts,* edited by Dale W. Jorgenson, J. Steven Landefeld, and William D. Nordhaus. University of Chicago Press.

Nordhaus, William D., and Edward C. Kokkelenberg. 1999. *Nature's Numbers: Expanding the National Economic Accounts to Include the Environment.* Washington: National Academies Press.

Parker, Robert P., and Bruce T. Grimm. 2000. "Recognition of Business and Government Expenditures for Software as Investment: Methodology and Quantitative Impacts, 1959–98." Paper presented at the meeting of the BEA Advisory Committee, U.S. Department of Commerce, Bureau of Economic Analysis, Washington, May 5. http://www.bea.gov/papers/pdf/software.pdf

Reinsdorf, Marshall. 1993. "The Effect of Outlet Price Differentials on the U.S. Consumer Price Index." In *Price Measurements and Their Uses,* edited by Murray F. Foss, Marilyn E. Manser, and Allan H. Young. University of Chicago Press.

Reinsdorf, Marshall, and Robert Yuskavage. Forthcoming. "Offshoring, Sourcing Substitution Bias and the Measurement of U.S. Import Prices, GDP and Productivity." *Review of Income and Wealth.*

Shapiro, Matthew D., and David W. Wilcox. 1996. "Mismeasurement in the Consumer Price Index: An Evaluation." *NBER Macroeconomics Annual* 11: 93–154.

Schreyer, Paul, and W. Erwin Diewert. 2014. "Household Production, Leisure, and Living Standards." In *Measuring Economic Sustainability and Progress,* edited by Dale W. Jorgenson, J. Steven Landefeld, and Paul Schreyer. University of Chicago Press.

Schreyer, Paul, Nicola Brandt, and Vera Zipperer. 2015. "Productivity Measurement with Natural Capital." Paper presented at the Society for Economic Measurement Annual Conference, Paris, July 24. http://repository.cmu.edu/cgi/viewcontent.cgi?article=1007&context=sem_conf

Schultze, Charles L., and Christopher Mackie, editors. 2002. *At What Price? Conceptualizing and Measuring Cost-of-Living and Price Indexes.* Washington: National Academies Press.

Sichel, Daniel E. 1997. "The Productivity Slowdown: Is a Growing Unmeasurable Sector the Culprit?" *Review of Economics and Statistics* 79, no. 3: 367–70.

Soloveichik, Rachel. 2015a. "Exploring the Boundaries of GDP: Cultivated Assets and Valuing 'Free' Media." Presentation. Washington: U.S. Department of Commerce, Bureau of Economic Analysis. https://bea.gov/about/pdf/acm/2015/november/exploring-the-boundaries-of-gdp.pdf

———. 2015b. "Valuing 'Free' Entertainment in GDP: An Experimental Approach." Working paper. Washington: U.S. Department of Commerce, Bureau of Economic Analysis. http://bea.gov/about/pdf/acm/2015/november/valuing-free-entertainment-in-gdp-for-aea-paper.pdf

Stallman, Richard. 1985. "The GNU Manifesto." http://www.gnu.org/gnu/manifesto.en.html

Stein, Lincoln D. 2010. "The Case for Cloud Computing in Genome Informatics." *Genome Biology* 11: Article 207. http://genomebiology.biomedcentral.com/articles/10.1186/gb-2010-11-5-207

Syverson, Chad. 2016. "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown." Working Paper no. 21974. Cambridge, Mass.: National Bureau of Economic Research.

Triplett, Jack E. 1989. "Price and Technological Change in a Capital Good: A Survey of Research on Computers." In *Technology and Capital Formation,* edited by Dale W. Jorgenson and Ralph Landau. MIT Press.

———. 1999. "Economic Statistics, the New Economy, and the Productivity Slowdown." *Business Economics* 34, no. 2: 13–17.

Triplett, Jack E., and Barry P. Bosworth. 2004. *Productivity in the U.S. Services Sector: New Sources of Economic Growth.* Brookings Institution Press.

United Nations and others. 2009. *System of National Accounts 2008.* New York: European Commission, International Monetary Fund, Organization for Economic Cooperation and Development, United Nations, and World Bank.

Varian, Hal. 2011. "The Economic Value of Google." Presentation at the Web 2.0 Expo, San Francisco, March 29. http://www.web2expo.com/webexsf2011/public/schedule/detail/17666

Wasshausen, Dave, and Brent R. Moulton. 2006. "The Role of Hedonic Methods in Measuring Real GDP in the United States." Paper presented at 31st CEIES Seminar: Are We Measuring Productivity Correctly? in Rome, October 12–13. http://www.bea.gov/papers/pdf/hedonicGDP.pdf

Wetterstrand, Kris A. 2016. "DNA Sequencing Costs: Data from the NHGRI Genome Sequencing Program (GSP)." https://www.genome.gov/27541954/dna-sequencing-costs/

Zheng, Simon, and Harry Bloch. 2014. "Australia's Mining Productivity Decline: Implications for MFP Measurement." *Journal of Productivity Analysis* 41, no. 2: 201–12.

Comments and Discussion

COMMENT BY

MARTIN NEIL BAILY This is a terrific paper that changed my views on the slowdown in productivity growth. I had thought that a significant fraction of the post-2004 growth slowdown could be explained by measurement errors of two types: first, underestimation of the pace of productivity in the computer and semiconductor industry; and second, the fact that the National Income and Product Accounts do not include the contribution of "free stuff" like Google and Facebook to final output because they are paid for by advertising and are, therefore, considered intermediate production. David Byrne, John Fernald, and Marshall Reinsdorf show persuasively that neither type of measurement error is significant enough to change our estimates of the slowing of growth in about 2004. The authors document, along with Chad Syverson (2016), that the productivity growth slowdown is very large—output would be larger by about $3 trillion today if there had been no slowdown. The errors of measurement would need to be very large indeed to explain much of this loss of output, and the authors show that this is not the case.

It is worth backing up a little and pointing out the strange, paradoxical economic times that form the backdrop to this paper. Productivity growth has been very slow since about 2004, and the slowdown appears to be getting worse, with output per hour in the nonfarm business sector in 2015 only 2.6 percent higher than its 2010 value, according to the Bureau of Labor Statistics (BLS). This is bad news for living standards and economic growth. At the same time, other signs seem to point to rapid technological change. For example, in a survey released in June 2015 of the *Fortune* 500 CEOs, they named their greatest challenges; the top challenge, listed by 70 percent of the CEOs, was rapid technological change. And a front-page article in the May 3, 2016, *Financial Times* said, "Surging investment in artificial intelligence is giving the United States an early advantage in the global race to dominate a new era of robotics, according to investors and experts in an industry set to become one of the most strategically important in the coming decades" (Waters and Inagaki 2016). Erik Brynjolfsson and Andrew McAfee (2014) are leading experts on technology, and their 2014 book *The Second Machine Age* has been an important and influential guide to changing technology. In chapter 2 they write, "Most of the innovations described in this chapter have occurred in just the past few years. They've taken place in areas where . . . the best thinking often led to the conclusion that it wouldn't speed up. But then digital progress became sudden after being gradual for so long" (Brynjolfsson and McAfee 2014, p. 37). In short, they suggest that technological change has actually speeded up. Mismeasurement had offered one way to resolve the paradox of slow growth and rapid technological change, and Byrne, Fernald, and Reinsdorf have punctured this balloon.

Although I agree with the bottom line of this paper, there are some points of disagreement or places where I would have put a different emphasis. Most important, a casual reader of their paper might believe that productivity growth is being well measured overall, but that is not the case. In important segments of the economy, there is no serious effort to capture the impact of technological change on the quality of output. Most notable is health care, where advances in screening and diagnostics, surgical procedures, and medical devices are constantly being introduced. Thanks to the research by Ernst Berndt of the Massachusetts Institute of Technology and others,[1] there has been an effort to capture the benefits of new drugs, but improvements in the quality of hospital care are being missed. The authors are aware of this problem, but their tight focus on the post-2004 slowdown means they do not give it much attention.

Unlike health care, the high-tech industry has seen efforts made to capture quality change in output, and the authors do an admirable job of exploring the various price deflators covering this industry. They make a good case that errors of measurement are not likely to explain much of the slowdown; but as they note, much uncertainty remains. The reader is left wishing that the BLS would do a systematic review of its price index methodology. One of the big changes in recent years has been to make computers and tablets much lighter and more user-friendly, and I do not think this quality change is being captured.

1. For a list of Ernst Berndt's publications, many of which relate to price measurement, see http://mitsloan.mit.edu/faculty-and-research/faculty-directory/detail/?id=41392.

The authors point out that much of the manufacturing of high-tech equipment has moved offshore, so this sector is providing a smaller boost to U.S. productivity. They also mention the outlet substitution bias that has likely occurred as component production has moved offshore. These points are certainly correct, but I wonder if changes in industry structure mean that some U.S. productivity growth is being missed. Much of the design work for computer chips, iPhones, and the like is still being done in the United States, so the quality change that accounts for most of the productivity increase in high-tech industries is still attributable to economic activity located in the United States and not in the countries where the products are manufactured. One reason for this is that the United States' marginal corporate tax rate is higher than those of other countries, so multinational companies minimize the U.S. content of products made and sold internationally. Even without tax distortions, in a world where supply chains are global, it is intrinsically very difficult to correctly account for productivity by country. This measurement problem is not huge, but it probably has been getting worse over time.

When I ask both economists and noneconomists whether they think that the free stuff on their phones or computers is significant enough to shift the needle of productivity growth, I get bipolar answers. Some people believe that all the new stuff is fantastic and is changing their lives, while others dismiss it as trivial. The answers to my nonscientific personal survey are somewhat age-related; older people are usually less enthusiastic than younger people, but not always so. Some older people and some with disabilities find that their lives have been greatly enhanced. Byrne, Fernald, and Reinsdorf are convincing in showing that the magnitude of uncounted free services is just not large enough to make much of a dent in the $3 trillion hole in productivity. They evaluate free services using a time-use framework, which is an entirely reasonable decision given that the literature has used this framework. But I do not find this approach to be all that compelling. All consumption involves time, but we do not try to capture this when we measure productivity. Automobiles, for example, require time to drive, but the contribution of automobile production to productivity depends on quality-adjusted output in relation to inputs. There is no need to estimate the time spent driving.

The tricky part with Google and other similar services, as the authors note, is that consumers do not pay for them directly, but only through advertising. Such "free" services did not start with Google; television used to be entirely paid for by advertising, and much of it still is. Like Google, television was an intermediate goods industry. Given the value that most

Figure 1. Contributions to Aggregate Multifactor Productivity Growth Using Domar Weights, by Major Industry Sector, 1987–2013

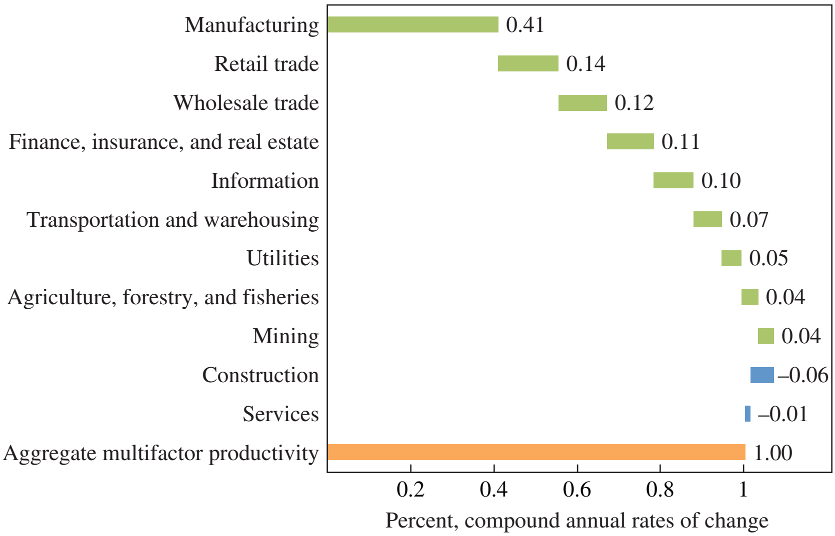

Source: Author's calculations, based on data from the Bureau of Labor Statistics' multifactor productivity database.

consumers place on watching TV, the exclusion of this service from final output in earlier periods meant that productivity growth was understated. Other countries chose to pay for television through license fees, and their TV-related industries were then counted in final output and presumably contributed to measured productivity growth.

If we lived in a world where the budgets of the statistical agencies were greatly expanded, I would urge them to create satellite accounts to estimate the value of free services. If time-use analysis provided the best approach, then I would be a convert. Another alternative would be to use conjoint analyses from survey data, the approach used by market researchers to assess how consumers value different products and product attributes.

The authors have calculated productivity growth by industry, and they use these results in their paper, but there are some additional lessons worth drawing from these data, and I turn now to describe them. My figure 1 shows the Domar-weighted contribution of each of the major private sector industries to multifactor productivity (MFP) growth in the aggregate over the entire period from 1987 to 2013 for which consistent data are available from the BLS. The first striking fact is that manufacturing contributes about 40 percent of aggregate MFP growth. This is despite the fact that manufacturing is only about 10 percent of employment. From a productivity point of view, manufacturing still matters. The second striking fact is that two large sectors of the economy, construction and services, had zero or negative MFP growth contributions from 1987 to 2013. This raises a red flag that measurement problems may actually be quite important for estimated growth over the whole period. Construction productivity has been a mystery for a very long time. Nonresidential output is very hard to price and measure accurately. On the residential side, I participated in a number of cross-country comparisons of residential productivity in the 1990s with the McKinsey Global Institute, and we never found a country where productivity was higher than in the United States.[2] That is not impossible to square with U.S. productivity that declines over time, but it is odd.

The subindustries within services that drag down the total are education and health care, and I discussed health care above. Having no productivity growth in education is perhaps not surprising, given how little the format of education has changed. The content of what is being learned has changed a lot, however, especially in higher education. And the industry structure and the way teaching takes place both seem poised to change, as new technology-related tools are introduced and competition increases. It would be a step forward if measurement methods for this industry were able to keep up with changing educational methods.

My next two figures both illustrate that the industries whose productivity growth increased after 1995 were generally the industries that showed slower growth after 2004. My figure 2 shows that the pattern holds for the major industry sectors, and my figure 3 shows the same result when all of BLS's subindustries are included. One interpretation of this pattern is that a broad productivity opportunity opened up, creating the scope for a rapid productivity increase. Some industries had business models and market conditions that were conducive to taking advantage of this opportunity, but other industries were less well suited and continued along the old path. After 10 years or so, the impact of the innovative surge was over, and the rapid growers fell back to their prior pattern. The obvious candidate for the productivity opportunity was the rapid improvement and dissemination of information technology. There may also have been other contributory

2. Details of the studies are available on the McKinsey Global Institute website, at http://www.mckinsey.com/mgi/our-research/productivity-competitiveness-and-growth.

Figure 2. Productivity Acceleration Post-1995 versus Deceleration Post-2004, by Major Industry Sector

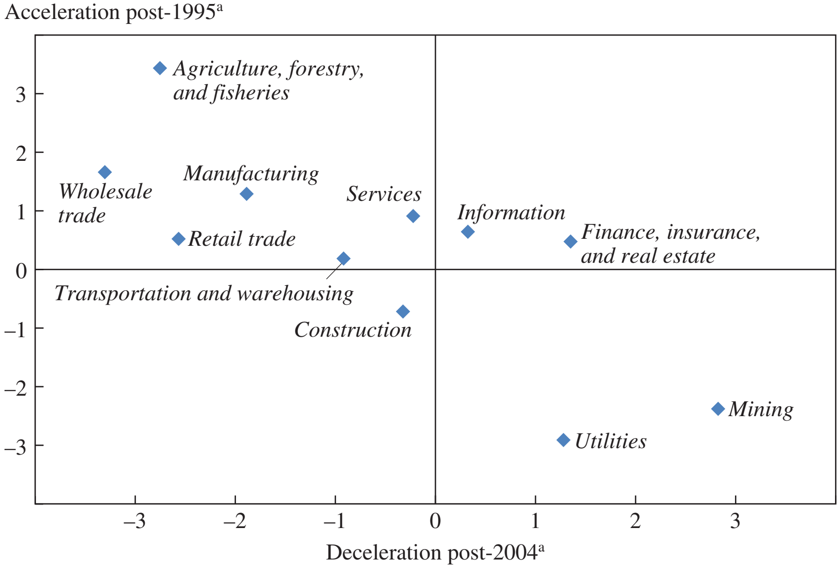

Source: Author's calculations, based on data from the Bureau of Labor Statistics' multifactor productivity database.
a. Percent, compound annual rates of change.

factors, such as a willingness to take risks, intense competition, and strong aggregate demand growth.[3]

The authors of this paper took on a very hard task, asking whether measurement error can explain the 2004 slowdown in productivity growth, and they came up with a convincing answer: It cannot. In his impressive review of U.S. economic history, Robert Gordon (2016) concludes that the information technology revolution was the last in a series of productivity waves dating back to the Industrial Revolution. The authors of this paper are less pessimistic, and I agree with them. For one thing, it seems unlikely that the productivity wave from information technology has run its course. Moore's law must end eventually, but there are many new ways to take advantage of cheap processing power and low-cost communications. The crowdsourcing of design, robots, and the Internet of Things are three such ongoing advances. Innovations in biotechnology and materials science are

3. There may also be some regression toward the mean, although over 10-year periods that effect is likely to be muted.

Figure 3. Productivity Acceleration Post-1995 versus Deceleration Post-2004, by Major Industry Sector and Subindustries

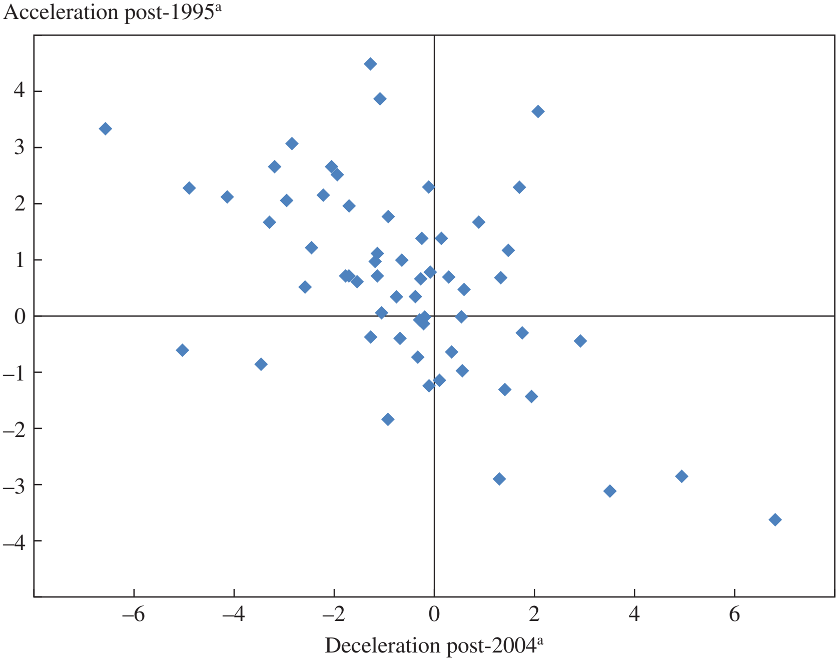

Source: Author's calculations, based on data from the Bureau of Labor Statistics' multifactor productivity database.

a. Percent, compound annual rates of change.

in the works. Productivity growth in wholesale and retail trade has slowed, as brick-and-mortar retailers deal with excess capacity because they face competition from online retailers and slow overall consumption growth, but a competitive shakeout of the industry will eventually result in higher productivity. Gordon (2016) may be correct, however, in saying that the era of 3 percent annual productivity growth over multiple decades is over. Future advances are likely to be lumpier, with surges of productivity from time to time, not all the time.

REFERENCES FOR THE BAILY COMMENT

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies.* New York: W.W. Norton.

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

Syverson, Chad. 2016. "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown." Working Paper no. 21974. Cambridge, Mass.: National Bureau of Economic Research.

Waters, Richard, and Kana Inagaki. 2016. "Investment Surge Gives U.S. the Early Lead in Rise of the Robots." *Financial Times,* May 3.

COMMENT BY

ROBERT J. GORDON The slow productivity growth since 2004, particularly since 2010, represents one of the outstanding economic puzzles of our time. It implies slow current growth of potential real GDP. If it continues, it implies slow future growth of potential GDP and fewer resources to address the nation's problems, including education, infrastructure, and the looming shortfalls in funding for Social Security and Medicare. It would be reassuring for puzzle solving, although disconcerting for the integrity of the nation's statistical system, to learn that the entire post-2004 or post-2010 slowdown in productivity growth was due to well-identified errors in measurement, and that the underlying "true" rate of productivity growth has not decelerated at all.

THE QUESTIONS TO BE ADDRESSED This paper by David Byrne, John Fernald, and Marshall Reinsdorf places its main emphasis on the post-2004 productivity growth slowdown and highlights the $3 trillion in additional business sector real GDP that would have been produced in 2015 if the productivity growth rate of 1995–2004 had continued after 2004. But the post-2004 slowdown is not the only productivity puzzle to be explained. My figure 1 displays the five-year moving average growth rate of quarterly utilization-adjusted total factor productivity (TFP) growth going back to 1952. This data plot identifies four separate eras of TFP growth—consistently rapid, at about 2.0 percent a year, through 1973; then slower and erratic, in the range of 0 to 1.5 percent, from the early 1970s through the mid-1990s; then healthy again for a decade between 1995 and 2004; and finally a sharp slowdown, to about 0.5 percent a year, during most of the past decade.

Average TFP growth rates over selected periods are listed in my table 1. The first two rows contrast the 26 years from 1947 to 1973 with the 42 years since 1973. When the postwar era is divided at 1973, the TFP growth rate slows by more than half, from 2.10 to 0.82 percent a year. This sharp contrast poses the first productivity puzzle: What caused the post-

Figure 1. Five-Year Moving Average of Utilization-Adjusted Total Factor Productivity Growth, 1952–2015

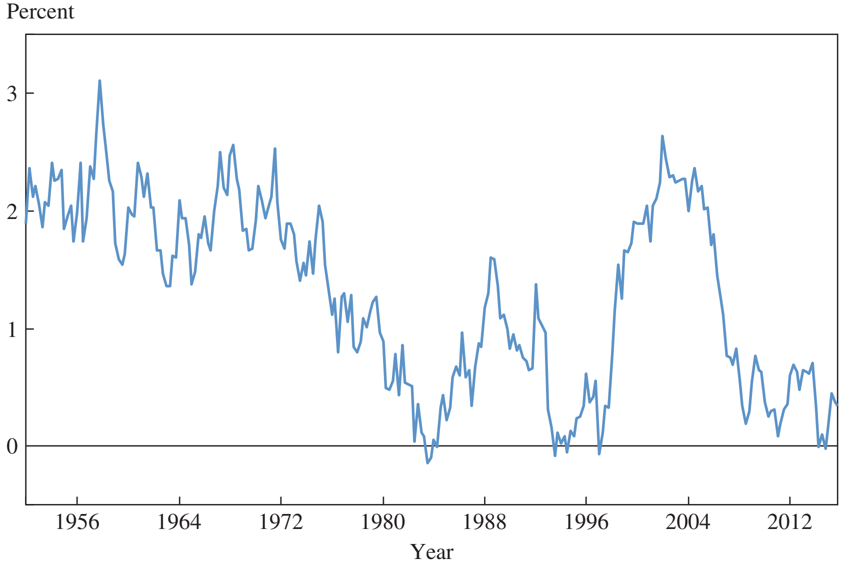

Source: Federal Reserve Bank of San Francisco (http://www.frbsf.org/economic-research/indicators-data/total-factor-productivity-tfp/).

1973 slowdown? The next three rows divide up the post-1973 interval at 1995 and 2004. The TFP growth rates for 1973–95 and for 2004–15 are almost identical, at about 0.5 percent, in sharp contrast to the growth rate of almost 2.0 percent achieved between 1995 and 2004. This leads to the second puzzle: What caused TFP growth to revive? And to the third puzzle: Why was that revival only temporary? The final two rows of my table 1

Table 1. Utilization-Adjusted Total Factor Productivity Growth, 1947–2015

Period	*TFP growth (percent)*
1947:Q2–1973:Q1	2.10
1973:Q1–2015:Q4	0.82
1973:Q1–1995:Q1	0.52
1995:Q1–2004:Q1	1.99
2004:Q1–2015:Q4	0.48
2004:Q1–2010:Q1	0.73
2010:Q1–2015:Q4	0.21

Source: Federal Reserve Bank of San Francisco (http://www.frbsf.org/economic-research/indicators-data/total-factor-productivity-tfp/).

divide the 2004–15 era in early 2010 and show that TFP growth during the 2004–10 period was slightly faster than from 1973 to 1995, but during the 2010–15 period was somewhat slower.

This alternation between relatively fast and relatively slow TFP growth over the four eras of the postwar epoch places a broader perspective on Byrne, Fernald, and Reinsdorf's topic of mismeasurement. Techniques of measurement have been relatively constant since 1947, and thus it is implausible to argue that the slowdown from the first era (1947–73) to the second era (1973–95) happened because measurement became worse by an average of 1.5 percentage points a year. In the same way, it is implausible to argue that the revival to the third era (1995–2004) occurred because measurement became better, at a rate of 1.5 percent a year. Likewise, it is implausible to argue that the slowdown to the fourth era (2004–15) occurred because measurement once again became worse, at a rate of 1.5 percent a year, as had occurred previously, after 1973. These alleged appearances and disappearances of measurement errors of 1.5 percent a year in both directions are implausible, because measurement techniques were relatively constant across the four postwar eras.

MEASUREMENT ERRORS HAVE DIMINISHED IN IMPORTANCE IN THE MARKET ECONOMY Byrne, Fernald, and Reinsdorf place their primary emphasis on the measurement of the private business economy. Their main focus is on mismeasurement in the form of biased deflators for information and communication technology (ICT) equipment in the National Income and Product Accounts (NIPA). They survey recent research on the prices of ICT equipment and conclude that the NIPA deflators systematically understate the rate of decline in the quality-adjusted prices of ICT equipment and thus understate the rate of growth of real ICT investment as well as real GDP and labor productivity. TFP is affected less, because the use of improved deflators not only raises real GDP growth but also raises the growth rate of capital input that is subtracted from output in the calculation of TFP.

However, this price index bias does not help at all in understanding the post-2004 productivity growth slowdown, because the price index bias is roughly constant both before and after 2004. The difference between the rate of change of the authors' liberal ICT price index and the corresponding NIPA index is −5.1 percent a year during the period 1995–2004 and an almost identical −4.5 percent in 2004–14. Both the liberal index and the NIPA index exhibit a sharp deceleration in the rate of the price decline after 2004, which points to a declining rate of technological improvement in ICT equipment as a substantive reason for the productivity growth slowdown.

Figure 2. Import Penetration of Computer Equipment Investment, 2002–13

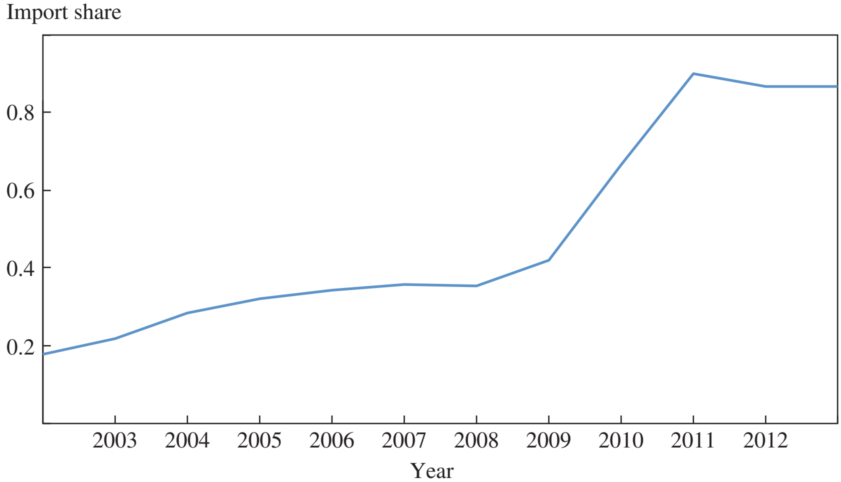

Source: Byrne and Pinto (2015).

The constant post-2004 price index bias is not the end of the story, however, because the relative importance of ICT equipment in the economy has changed in two ways. First, ICT investment is a smaller share of GDP. Comparing the two years 1999–2000 with 2014–15, the GDP share of information processing equipment declined from 2.77 to 1.79 percent, and that of computers and peripherals fell fully by half, from 1.00 to 0.45 percent. When combined with the relatively constant pre- and post-2004 price index corrections, the shrinking ICT shares imply that the price index bias for GDP and labor productivity became smaller after 2004.

The second reason why the price index bias has become less important stems from a sharp shift of ICT investment from domestic production to imports. My figure 2 exhibits the startling shift in computer purchases—from 17.8 percent imported in 2002 to an average of 87.9 percent imported in 2011–13. Consider the implications of the extreme case in which all ICT equipment is imported. An upward price index bias for imported computers would lead to an understatement of the growth rate of both computer investment and computer imports, netting out to zero impact on GDP and labor productivity. The understatement of growth in capital input, however, would lead to an overstatement of TFP growth. Thus the shift to computer imports in the last decade has caused true TFP growth to slow down more since 2004 than in the official NIPA data. This tendency has been exacerbated by the fact that the price indexes for imported computers used in the NIPA decline at a substantially slower rate than the deflators for domestically produced computers, whereas the observed shift of computer purchases to imports suggest the opposite—that the true prices of imported computers have declined faster than prices of domestic production.

These two factors, the declining share of ICT investment in GDP and the shift to imports, together with the substantial upward bias in the price index for imported computers, suggest that since 2004 measurement issues have caused the poor performance of TFP growth to be even worse in reality than in the government's statistics. Byrne, Fernald, and Reinsdorf's treatment concludes, in the third column of their table 2, that measurement errors cause labor productivity growth to be understated by 0.49 percentage point for the period 1995–2004 and by 0.18 point for 2004–14, for a net measurement *improvement* of 0.31 percentage point. Because of offsetting adjustments to output and capital input, the effect on TFP growth is much smaller and goes in the opposite direction, with measurement errors causing TFP growth to be overstated by 0.08 percentage point during 1995–2004 and by a slightly greater 0.12 percentage point during 2004–14. The conclusion is that improved measurement causes the post-2004 slowdown in both labor productivity and TFP growth to become even worse than in the official data.

These important conclusions of Byrne, Fernald, and Reinsdorf's analysis combine an upward bias in the price indexes of computers with a shrinking share of computer investment and of the domestic production share of computer equipment. If this upward price index bias were larger, their conclusions would be magnified. In his important historical study of the price indexes of computers, William Nordhaus (2007, table 10, p. 153) concludes that the price of computer power during the 1990–2002 period decreased at an annual rate of −57.5 percent, as compared with Byrne, Fernald, and Reinsdorf's alternative price index for computers and peripherals in their table 1, which declines at a much slower annual rate of −27.3 percent during the 1995–2004 period. Though there are conceptual differences between Nordhaus's performance-based measure and Byrne, Fernald, and Reinsdorf's hedonic price indexes, Nordhaus's index has the advantage that it includes data for both mainframe and personal computers, whereas Byrne, Fernald, and Reinsdorf's indexes for the 1990s are based only on personal computers. Compared with mainframes, personal computers achieve a much lower price per calculation, and thus the transition from mainframes to personal computers that took place in the 1980s and 1990s reduced the average price per calculation more rapidly than the price decline for personal computers alone.[1] To the extent that Nordhaus's (2007) approach is a better guide to the overall price behavior of computer investment in the 1990s, there has been an even greater tendency for official data to understate that rapid growth of labor productivity during the 1995–2004 period and to understate the true decline in its growth rate since 2004.

THE WELFARE EFFECTS OF FREE INTERNET CONTENT Free Internet information was available both before and after the 2004 transition from fast to slow productivity growth. The proportion of American households connected to the Internet rose from 5 percent in 1995 to 56 percent in 2004, followed by a more gradual increase to 75 percent by 2013 (Gordon 2016, figure 13-4, p. 455). The mismeasurement hypothesis refers to the difference in the welfare benefits from free Internet content available after 2004, as compared with before 2004. An aspect of the post-2004 improvement is the transition from dial-up to broadband access; the proportion of households with broadband increased from 3 percent in 2000 to 29 percent in 2004 and to 65 percent after 2009 (Gordon 2016, figure 13-5, p. 455). Thus, even if the same amount of time were spent on Internet access before and after 2004, there was a quality change in the form of much faster response times made possible by the spread of broadband. Byrne, Fernald, and Reinsdorf recognize that download speed is not considered a quality change in current deflators for Internet access, but they assert that "only a small amount of extra digital service output is missed."

Most of Byrne, Fernald, and Reinsdorf's treatment of free digital services refers to their role in the market economy. The authors consider alternatives to the current national accounts treatment of advertising-supported media as an intermediate good. Even when free Internet services and other forms of entertainment are treated as final consumption, their role in the market economy can be no larger than their advertising revenue. Because total advertising revenue from all sources, including television and print media, amounts to only 1.3 percent of GDP, and because real advertising revenue has grown faster than business sector GDP by only a small margin, the authors find the impact of free Internet services on market GDP to be close to zero. Intuitively, the growing advertising revenue of Google and Facebook is largely canceled out by the decline in advertising revenue from older forms of media, particularly print publications.

1. Nordhaus (2007, p. 155) provides an example of a 2002 IBM supercomputer that had a price-per-performance ratio roughly 34 times higher than a typical Dell personal computer in 2004.

Advocates of the mismeasurement hypothesis have in mind the broader scope of free Internet services as a source of increased consumer welfare, going well beyond market GDP. They point to the rapid increase in Internet usage, particularly of mobile services accessed on smartphones and tablets, as a focus of consumers' leisure-time activity. How important is the increase in consumer welfare resulting from increased Internet use? More than one-third of the U.S. population uses Facebook, and the time each day that its users devote to Facebook reached 50 minutes in 2015 (Stewart 2016). Taken together with other Internet services, most notably YouTube and Google, total daily time devoted to the Internet has been estimated at two hours (Karaian 2015). By comparison, the American Time Use Survey (ATUS) reports that in 2014 Americans on average spent 2.8 hours a day watching television.[2] A problem with the estimation of consumer surplus is that the ATUS does not report on Internet usage as a separate category, with the exception of "household and personal e-mail and messages," which accounted for a trivial 0.03 hour per day. Some mobile phone usage may occur during the ATUS category of time devoted to "socializing and communicating" (0.71 hour per day). Otherwise, the omission of Internet usage could imply that multitasking is the standard mode of behavior, with mobile phones accessed while watching television, eating meals, riding public transit, or standing in line.

To assess the consumer welfare aspects of free Internet services, Byrne, Fernald, and Reinsdorf develop an explicit model based on Gary Becker's (1965) theory of the economics of time, in which total consumption is subject to both a monetary and time budget constraint. However, they do not provide their own estimates of the time value of free Internet services, relying instead on a paper by Erik Brynjolfsson and Joo Hee Oh (2014) that values the incremental consumer surplus from free digital services as "equivalent of adding about 0.3 percentage point per year to business sector output." Note that this addition of 0.3 percentage point exactly cancels out the authors' estimate in their table 2 that the post-2004 decline in business sector labor productivity has been understated by the same 0.3 percentage point a year.

Is this estimate too large or too small? Chad Syverson (2016) provides a survey of different approaches to measuring the consumer surplus of the Internet and compares the resulting surplus estimates to the total missing annual output from the post-2004 productivity slowdown, which he

2. See Bureau of Labor Statistics, "Table A-1: Time Spent in Detailed Primary Activities and Percent of the Civilian Population Engaging in Each Activity, Averages per Day by Sex, 2014 Annual Averages" (http://www.bls.gov/tus/tables/a1_2014.pdf).

estimates, like Byrne, Fernald, and Reinsdorf, to be about $3 trillion in 2015. Most of the literature surveyed by Syverson provides surplus estimates of about $100 billion a year, a trivial fraction of the missing $3 trillion. But one approach developed by Austan Goolsbee and Peter Klenow (2006), when updated by Syverson's numbers, yields a post-2004 incremental surplus estimate of $842 billion, almost one-third of the missing $3 trillion. This one-third translates into 0.6 percentage point of the total post-2004 slowdown of 1.8 percentage points in the growth rate of business sector productivity. How reasonable is this estimate?

Applying $842 billion to the 80 percent of the population with broadband access yields an annual per capita sum of $3,300. If incremental post-2004 Internet usage, compared with the use of the Internet before 2004, is one hour per day, this would imply an Internet value of $9 per hour. This compares with Syverson's (2016) estimate of the average 2015 after-tax wage of $22. The difference between $9 and $22 makes sense, because it reflects the fact that only half the population is employed, and the leisure time spent on the Internet is inframarginal. The use of one hour per day in this calculation, rather than the two hours reported as average daily Internet use, reflects not only the fact that some time was allocated to the Internet before 2004 but also the multitasking implied by the ATUS time allocation. Much if not most Internet use, according to the ATUS, is not occurring during hours of leisure that previously had zero value, but rather as multitasking during hours that previously had value obtained from socializing or watching television.

Therefore, Byrne, Fernald, and Reinsdorf appear to be too dismissive of the consumer surplus contributed by free Internet services. Their estimate, taken from Brynjolfsson and Oh's (2014) findings, values incremental post-2004 Internet services as worth 0.3 percent of business sector output per year, just enough to offset the authors' table 2 estimate that the decline in the growth rate of business sector productivity has been understated by the same 0.3 percent per year. In contrast, Syverson's (2016) approach values the consumer surplus at 0.6 percent per year, twice as much. This is enough to offset the 0.3 percentage point measurement understatement from the authors' table 2, as well as contributing an additional 0.3 percentage point consumer surplus bonus that provides a partial counterweight to the overall 1.8 percentage point productivity growth slowdown. The relatively large size of this Internet valuation naturally leads to the question of how much consumer surplus was contributed by inventions in the past, including the value of free Internet services introduced in the 1995–2004 decade, such as e-mail, search engines, Wikipedia, and the early phase of e-commerce provided by Amazon, iTunes, and airline websites. The technique of multiplying hours of leisure by an hourly value based on the wage rate would yield a particularly large value of incremental consumer surplus in the early 1950s, as free television broadcasts reached almost every home between 1950 and 1955.

History is full of examples of added consumer surplus that was not included as an increase in GDP. During the period from 1900 to 1940, as motor vehicles replaced horses, real GDP did not value the removal of horse droppings and urine from city streets and rural highways. And real GDP did not value the increase in speed and load-carrying capacity made possible by automobiles, nor their flexibility, which gave birth to a new industry called "personal travel." Moreover, real GDP did not value the increase in consumer surplus as clean running water arrived at the in-home tap and replaced the previous need to carry pails of water into the house from nearby wells or streams. Finally, real GDP did not value the replacement of the outhouse and the need to physically dispose of human waste with the silent efficiency of public sewers. Running water, the electric washing machine and refrigerator, the automobile, and all the other inventions of the late 19th and early 20th centuries are linked into the GDP statistics, which means that zero value is placed on their invention. Real GDP did not value the reduction of infant mortality from 22 percent of new births in 1890 to about 1 percent in 1950. By some estimates, this change created more welfare value than all the other sources of increased consumer welfare taken together.

CONCLUSION Because newly invented products and services have always provided consumer surplus that supplements growth in real GDP by unknown amounts, the authors of this paper are correct to focus their main attention on the measurement issues that arise in the private business sector of the market economy. They convincingly demonstrate that the post-2004 slowdown in the growth rate of labor productivity and TFP has not been due to measurement errors. On the contrary, because the most evident source of measurement bias is in the price indexes for ICT equipment, the understatement of business sector output and productivity was greater in the 1995–2004 period than in the 2004–15 period, both because ICT investment was a greater share of GDP and because a much greater share of ICT equipment was manufactured in the domestic economy, whereas after 2004 there was a sharp increase in the share of such equipment that was imported. The authors thus conclude that the post-2004 slowdown is real. I agree with their interpretation that after 2004, productivity growth returned "back to normal," to a rate roughly the same as that achieved between 1973 and 1995, and that the burst of faster productivity growth between 1995 and 2004 reflected a one-time-only conversion of the economy to a modern world whose ICT equipment and software replaced the previous world's typewriters, calculating machines, and file cabinets.

REFERENCES FOR THE GORDON COMMENT

Becker, Gary S. 1965. "A Theory of the Allocation of Time." *Economic Journal* 75, no. 299: 493–517.

Brynjolfsson, Erik, and Joo Hee Oh. 2014. "The Attention Economy: Measuring the Value of Free Digital Services on the Internet." Working paper. https://www.scribd.com/doc/314594157/1-The-Attention-Economy-Measuring-the-Value-of-Free-Goods-on-the-Internet

Byrne, David, and Eugenio Pinto. 2015. "The Recent Slowdown in High-Tech Equipment Price Declines and Some Implications for Business Investment and Labor Productivity." FEDS Notes. Washington: Board of Governors of the Federal Reserve System.

Goolsbee, Austan, and Peter J. Klenow. 2006. "Valuing Consumer Products by the Time Spent Using Them: An Application to the Internet." *American Economic Review* 96, no. 2: 108–13.

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

Karaian, Jason. 2015. "We Now Spend More Than Eight Hours a Day Consuming Media." *Quartz,* June 1.

Nordhaus, William D. 2007. "Two Centuries of Productivity Growth in Computing." *Journal of Economic History* 67, no. 1: 128–59.

Stewart, James B. 2016. "Facebook Has 50 Minutes of Your Time Each Day; It Wants More." *New York Times,* May 5.

Syverson, Chad. 2016. "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown." Working Paper no. 21974. Cambridge, Mass.: National Bureau of Economic Research.

GENERAL DISCUSSION Dan Sichel commended the paper for making a good case that mismeasurement in the technology sector is not the right way to understand the productivity slowdown, as the paper convincingly shows that there has not been either (i) a big shift in shares toward the unmeasured or quality-measured sectors, or (ii) a really big step up in the amount of mismeasurement, what he called the "ingredients" of the mismeasurement story. However, he expressed concern that a casual reader might come away with the wrong impression. In particular, while mismeasurement is not a good explanation for the productivity slowdown, mismeasurement remains a big problem, particularly in the technology, education, and health care sectors. For instance, Sichel cited a paper that he coauthored with David Byrne and Stephen Oliner that convincingly documents big measurement problems in the official price indexes for semiconductors.[1]

In addition, Sichel noted that the present paper highlighted a range of other potential measurement problems for the whole span of the technology sector. Though the authors use rough calibrations, he thought that the calibrations were sensible, and that they highlighted the need for what Matthew Shapiro and David Wilcox called the "house-to-house combat of price measurement."[2] That is, rather than simply applying a plausible calibration, one should actually go and quantify the mismeasurement. Even though mismeasurement is not an explanation of the productivity slowdown, technology is a really dynamic, important sector in the economy, and it remains important to support the Bureau of Economic Analysis (BEA) and the Bureau of Labor Statistics (BLS) in their efforts to try to do a better job at correcting it. More broadly, one would not want the casual reader to conclude that, because mismeasurement does not explain the productivity slowdown, there is no need to worry about mismeasurement anymore. Mismeasurement is still a big problem, he concluded, and it needs to be addressed.

Martin Feldstein had three short comments. First, he believed that one big problem was not just the recent decline observed in productivity but also the volatility in the series. Second, he expressed interest in further examining the long-term bias in output measurement. In a piece written for the *Wall Street Journal*, Feldstein argued that, despite all the best efforts of the BLS, the problems associated with measurement for new and improved products means that the rate of improvement is underestimated.[3] However, he stresses that products are only a small part of the problem, and that the measurement of services is much more important. Roughly 80 percent of private sector employment is in services, and figuring out how to measure improved output in many service sectors, he

1. David M. Byrne, Stephen D. Oliner, and Daniel E. Sichel, "Is the Information Technology Revolution Over?" Finance and Economics Discussion Series no. 2013-36 (Washington: Board of Governors of the Federal Reserve System, 2013).

2. Matthew D. Shapiro, and David W. Wilcox, "Mismeasurement in the Consumer Price Index: An Evaluation," *NBER Macroeconomics Annual* 11 (1996): 93–154.

3. Martin Feldstein, "The U.S. Underestimates Growth," *Wall Street Journal,* May 18, 2015.

believes, is going to be hopelessly difficult. And third, he made a point about the interpretation of these statistics. Some mistakenly tend to treat productivity statistics as indicators of well-being, or of consumer utility, when—as the authors correctly emphasize—they are only about market activities. There has long been a tension in the history of the National Income and Product Accounts and in what the BLS does in focusing on market activities where, when people read statistics about, say, income growth, they interpret that in terms of the value to users, rather than just the market aspect of it.

Jason Furman agreed with Feldstein that the paper was excellent, and that it removed whatever sliver of doubt he had that mismeasurement was an important part of the explanation for the productivity slowdown. He followed with two comments. First, he noted that the less weight put on mismeasurement as the explanation for the slowdown, the more optimistic one should be about the growth measured over the next decade, the reason being that the methods used to measure productivity growth are considerably more persistent than the underlying "true" productivity growth itself. He suggested that a lack of capital investment over the last 5 to 10 years was a big part of the low-productivity story, and that he was optimistic that productivity would eventually rebound. Second, he warned against discounting measurement errors not in the form of persistent biases or incomplete sources. As more data become available in the coming years, past productivity numbers will likely need to be revised, and while one should think of those revisions as being unbiased, given the recent large disconnect between strong employment growth and weak output growth, it could be more likely that output and productivity will be revised up than down. If, for instance, productivity growth is presently measured to be 0.5 percent per year, Furman believed that it was more likely in five years the number would be revised up to 1, rather than revised down to 0.

Robert Hall noted that an important distinction needed to be made between *productivity* and *consumer surplus*, concepts that discussants in the room seemed to him to be conflating. It is really important to understand that productivity cannot be measured by consumer surplus, and how much consumer surplus is associated with output is a totally separate question. He urged discussants to be sure to keep the two concepts apart. Similarly, Hall noted the important distinction between *output per hour* and *total factor productivity* (TFP). Output per hour is just another endogenous variable, he pointed out, whereas TFP is fundamental. He encouraged everyone to keep an eye on what the paper basically does on TFP, and not to confuse it with looking at output per hour, though there is of course a close relationship between the two.

On the question of the future, Hall cited a paper that he recently encountered by Diego Comin and Mark Gertler that stakes out the claim that the reason productivity grew so slowly from 2010 onward was a failure to adopt existing technology, and that as the economy returns to normal—in particular, as the things that have held back investment subside—there will be a closing of the gap, because the creation of new technology will have gone on at normal rates.[4] Hall added that a very interesting fact not mentioned in the discussion at all was that research and development spending did not decline at all during the crisis or afterward.

David Romer echoed Furman and Feldstein in stating how terrific the paper was, and added that it crushed any sliver of hope he had that mismeasurement could explain the productivity slowdown. However, he thought the paper conceded too much on one point, which concerned the question of whether we can "rescue" some of the slowdown—not of productivity growth, but in the growth of consumer welfare—by appealing to the consumer surplus created by new technologies. For growth, the relevant question is not whether the recent innovations have increased consumer surplus but whether they are contributing *more* to consumer surplus than earlier innovations. Since—as discussant Robert Gordon is fond of pointing out—the amount of surplus created by earlier innovations was enormous, he was skeptical that bringing in consumer surplus would eliminate any noticeable part of the growth slowdown. He suggested that in discussing consumer surplus, the authors not just address how much recent innovations have contributed to consumer surplus but also at least mention that what is relevant for growth is *changes* in those contributions over time.

Justin Wolfers made three points, jokingly claiming that one point he did not believe, one was not true, and one was for introductory economics, so it may be true. The first point, which he did not believe, concerned the question of whether or not any of this really mattered. That is, what would economists do differently if productivity growth was measured appropriately? In terms of cyclical policy, productivity growth is already understood to be sufficiently noisy, and one does not typically look to the productivity numbers to figure out whether the latest slowdown was due to a productivity shock. In terms of long-run policy, whether productivity

4. Diego Anzoategui, Diego Comin, Mark Gertler, and Joseba Martinez, "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence," Working Paper no. 22005 (Cambridge, Mass.: National Bureau of Economic Research, 2016).

growth in the past decade was 0.05 percent, 1.5 percent, or 2.5 percent, one would still want to pick up any policy that would add another 1 or 2 points to that. According to Wolfers, most of the things an economist might want to advocate policymakers do would probably not change even if a better job was done at measuring productivity.

The second point Wolfers made related to consumer surplus and economists' apparent obsession with free stuff. There is a view that free services such as Facebook have significantly added to growth and consumer surplus, more so than goods and services that are not free. He pointed out that goods and services do not necessarily have to be free to generate a lot of consumer surplus. As a humorous example, he suggested that the introduction of the discount furniture store IKEA to the United States—which made "Scandinavian clean lines and bright colors" more affordable—arguably generated more of a consumer surplus than Facebook.

Last, Wolfers appealed to introductory economics to make the case to the authors that the idea of creating some sort of "national consumer surplus accounts" would be bad, and that the authors were right to resist. The "diamond–water paradox," an idea often taught in introductory economics, is the apparent contradiction that, although water is on the whole more useful than diamonds in terms of survival, diamonds command a higher price in the market. Suppose six glasses of water are sufficient to sustain life—and that the seventh and eighth only serve to ensure "healthy glowing skin." Pricing out the consumer surplus from the first six glasses, one might pay $200,000 for the first six glasses, because otherwise he would be dead. By that logic, the increment to consumer surplus from free services such as Facebook relative to the increment to consumer surplus from the fact that water exists is going to be very small. Rather than try to measure consumer surplus, one could simply measure well-being, and have a "well-being count." The simple way to do that, Wolfers pointed out, is by asking people how happy they are, and it turns out we already know how to do that.

John Haltiwanger spoke next about the micro productivity evidence for the ideas presented in the paper. The productivity slowdown in the high-tech sector, he noted, seems to be supported by the micro evidence. Likewise, the surge in retail trade productivity over the 1990s is evident in both the macro and micro evidence. What is puzzling, Haltiwanger pointed out, is the apparent collapse in retail trade productivity that the authors report around 2004; according to the micro data, there does not seem to have been this collapse. Haltiwanger suggested two findings from the micro data that raise questions about the macro evidence. First, the restructuring from single-unit establishment firms toward large national chains continues unabated over this period. Second, the productivity gap between single-unit establishment firms and large national chains is just as big as it ever was. The bottom line is that the micro dynamics still suggest that retail trade productivity ought to be doing quite well from these reallocation dynamics. According to the macro data, however, apparently it is not.

Melissa Kearney wondered if it was worth taking seriously the idea that the ways in which people use some free digital services might actually increase the labor supply. Many of the discussants had pointed to technology as being a substitute for leisure time, but perhaps access to technology also frees up time for additional work. For example, an individual might use her smart phone for home production—such as family shopping or scheduling—late at night, freeing up time for work during the day. She cited a paper by Lisa Dettling that found that exogenously determined high-speed home Internet access led to an increase in labor force participation for married women.[5] Kearney wondered if the authors had considered anything like this in their analysis.

Jonathan Pingle commended the authors for contributing to what is really important work suggesting a noticeable deceleration in structural productivity growth. However, adding on the implications of Bruce Fallick and others' cohort component model for aggregate labor supply—presented at the Fall 2014 *Brookings Papers* meeting[6]—and the normal business sector GDP gap, one struggles to talk about potential output growth of 1.5 percent in the United States now, and certainly a deceleration of over 2 percentage points. With many firms, businesses, and policymakers continuing to plan based on backward-looking expectations, that is potentially a big problem.

Joe Beaulieu found, countering Wolfers, that the idea of thinking about consumer surplus is a somewhat interesting but sideline conversation to TFP. He suggested that many of the free digital goods and services talked about essentially take the form of advertising. If there is a huge productivity increase in the advertising industry, that is an intermediate good, which probably means that prices for advertising services have fallen considerably. Therefore, he concluded that productivity properly measured for

5. Lisa J. Dettling, "Broadband in the Labor Market: The Impact of Residential High-Speed Internet on Married Women's Labor Force Participation," Finance and Economics Discussion Series no. 2013-065 (Washington: Board of Governors of the Federal Reserve System, 2013).

6. Stephanie Aaronson, Tomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher Smith, and William Wascher, "Labor Force Participation: Recent Developments and Future Prospects," *Brookings Papers on Economic Activity,* Fall 2014: 197–255.

the rest of the sectors in the economy must be even worse than originally thought. This also raises a second interesting question, which is how one should think about advertising and the implications for the economy. A lot of the advertising industry may involve rent-seeking behavior, which then has interesting implications not only for how one thinks about the economy and what is going on now versus a few years ago but also for how one might actually measure some of these things.

Alan Auerbach raised the question of how the production of multinationals is measured in the U.S. accounts, given that tax-induced profit shifting leads them to understate U.S. profits for tax purposes. Marshall Reinsdorf noted that much of the production could really be taking place in the United States, but that it gets reported overseas. He remarked that a recent paper by the BEA found that roughly 1 percent of GDP is explained by the fact that multinational corporations do most of their production based on where their labor and physical capital are located. He suggested that one might think of this estimate as an upper bound.

David Byrne noted that in addition to failing to reject the hypothesis that there is a productivity slowdown, in order to square the difference between the perspectives of Silicon Valley—that there is in fact a productivity slowdown in the technology sector—with what is actually observed, one should return to a *Fortune* 500 survey that discussant Martin Baily brought up during his remarks. In that survey, many of the companies said that their most important problem at the moment was adapting to new technology. One way to interpret this is that there is a technical frontier that has moved outward, and that these companies have not figured out what to do with it yet, as it is something that requires a tremendous amount of intangible investment. But another way to look at this situation is that it appears to be somewhat harder now to move from the back to the frontier perhaps than it once was, a finding supported by the work of Haltiwanger and others. This is potentially a different way of looking at the results of the paper.

Sichel had earlier cited a paper coauthored with Byrne and Stephen Oliner, which documented big measurement problems in the official price indexes for semiconductors, and noted that the tenors of that paper and the present paper seemed to be a little different. The right way to think about the Sichel–Byrne–Oliner paper, according to Byrne, from a GDP perspective is that microprocessors (MPUs) and semiconductors act as intermediate goods, so they are only going to show up to the extent that their production affects net trade. As it happens for MPUs—which are one of the most important individual goods that the United States trades—they are big in exports, but they are also big in imports. So, if one makes the price index for MPUs fall faster to a first approximation, it has no effect on GDP—not only because MPUs are intermediate goods but also because of the global value chain. The chips may be fabricated in the United States, tested in Costa Rica, and sent back to the United States to be used in data centers, and the transaction price is captured going both ways. The better way to think about the Sichel–Byrne–Oliner paper and the MPUs, according to Byrne, is to look at the MPU price index as a barometer of what is going on in the frontier of the information technology (IT) sector. The huge slowdown observed in the price index for recent years is very alarming, he noted; what the authors find in the present paper is that, for a number of reasons, that MPU price index is just not that accurate.

It had also been mentioned that IT is less important today than it once was, and Byrne noted that that is certainly true as a share of investment. Nevertheless, IT is still an important portion of production. BEA's investment deflator for IT includes things like computers, equipment software, and other special-purpose equipment. In the 1980s, that high-volume production used to be a much bigger share of IT investment than it is now—about 50 percent then, compared with only about 25 percent today. Investment and production of the low-volume, special-purpose equipment is much more difficult to measure. All this is not to say that mismeasurement is bigger than was found in the present paper. Rather, it is that the confidence interval around the estimate is much bigger than it once was. And it is not that the statistical agencies do not know about these problems; it is that they need more funding for measurement. He joked, "We are not supposed to be making price indexes at the Federal Reserve Board," but that, nevertheless, they do because despite it being well known that there are solutions to the problems of mismeasurement, the funding just is not there for the statistical agencies.

On the question of why more was not said about the dire labor productivity outcomes during the past five years, John Fernald noted that, at least when looking at TFP, it does not actually look worse; the past five years for TFP look about the same as the years before the Great Recession. In growth accounting, one observes that while so many people have been newly hired, capital accumulation has not kept up. One common theme in the room, Fernald noted, was that much is still unknown, and the paper only touches narrowly. He conceded that the measurement of health care, services, and new goods more broadly are all really challenging issues.

One final point Fernald made pertained to the value of the Internet. In terms of growth, the Internet brings no more than a few basis points, which he admitted was probably an overstatement; the authors' best estimate was closer to zero. Part of this, he explained, is because the market portion of the Internet is in advertising support. The big benefits are nonmarket, but these are still important, and things at the border of the market are shifting in interesting ways. Perhaps technology has increased the opportunity cost of working for less-educated workers. There is much need to study this, he concluded, adding that it would be great if the statistical agencies could acquire the resources to be able to develop better satellite accounts, and if researchers were working on ways to measure well-being.