

# How much of the productivity surge of 2007-2009 was real?

Michael Mandel  
Progressive Policy Institute.  
mmandel@visibleeconomy.com

Posted to <http://innovationandgrowth.wordpress.com> on March 28, 2011

This is a PDF of a post on my blog. So even though it is lengthy, it should be read at the level of informality of a blog post.

## Summary

In the 2007-2009 period, the U.S. economy experienced its worst recession since the Great Depression. Nevertheless, despite this deep downturn, the near-collapse of the financial system and unprecedented global economic turmoil, U.S. productivity growth actually seemed to accelerate in the 2007-2009 period, or at least maintain its previous pace.

The 2007-2009 productivity gain had a major impact on both economic policy and political discourse. First, it gave the Fed a free hand to feed mammoth amounts of liquidity into the system without worrying about inflation. Second, it convinced the economists of the Obama Administration that the economy was basically sound, and that the big problem was a demand shortfall. That's why they expected things to get back to normal after the fiscal stimulus.

However, I'm going to show in this post that the productivity gain of 2007-2009 is highly suspect. Using BEA statistics, I identify the industries that contributed the most to the apparent productivity gain, including primary metals, mining, agriculture, and computers and electronic products. Then I analyze these high-productivity growth industries in detail using physical measures such as barrels of oil and tons of steel. I conclude these 'high-productivity' industries did not deliver the gains that the official numbers show.

Based on my analysis, I estimate that the actual productivity gains in 2007-2009 may have been very close to zero. In addition, the drop in real GDP in this period was probably significantly larger than the numbers showed. I then explore some implications for economic policy.

## Overview

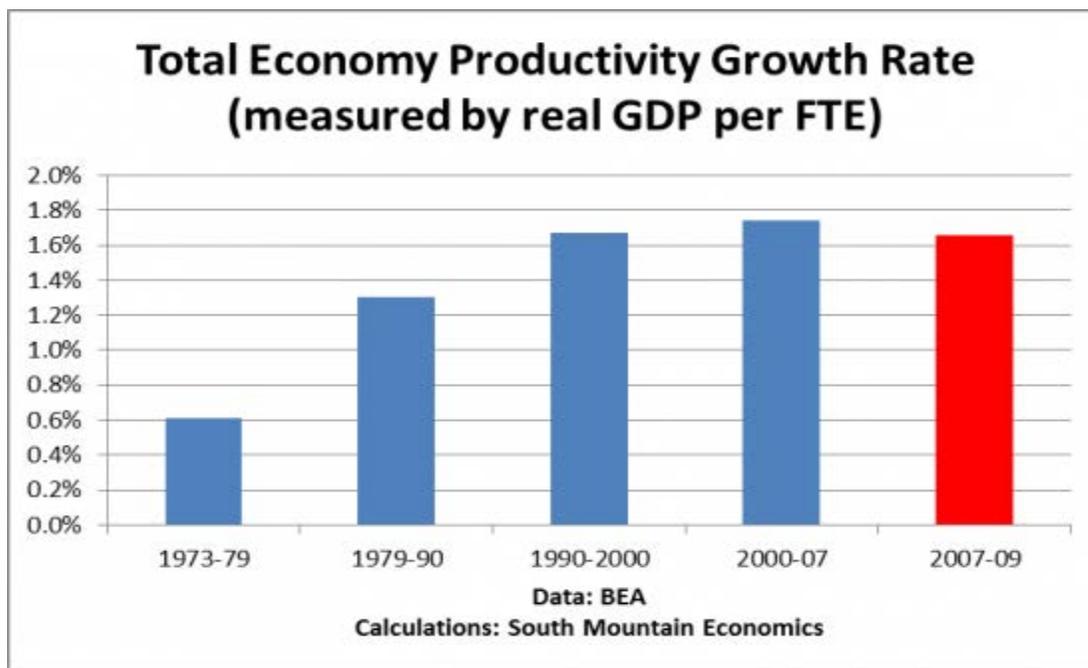
I start by giving several data points.

\*From 2007 to 2009, business productivity rose at a 2.4% rate, according to the official data. By comparison, the productivity growth rate was only 1.2% over the previous two years (2005-2007).

\*If you like your numbers quarterly, business productivity grew at a 1.8% annual rate from the peak in 07IV to the trough in 09II. That's somewhat faster than the 1.6% growth rate of the previous 3 years.

\*Looking at total economy productivity—real GDP divided by full-time equivalents (FTE)—we see that productivity growth in 2007-2009, at 1.6% per year, was double that of the previous two years (0.8% per year).

This chart breaks down productivity growth by business cycle, using the official statistics. You can see that the productivity growth of the 2007-2009 period—the worst recession in 70 years—appeared to be basically a continuation of the productivity gains during the boom. You can't tell from this chart that anything bad happened to the economy.



The strong measured productivity growth during the crisis years reflects a steep drop in employment (-5.7%) combined with an apparently mild two-year decline in real GDP (-2.6%). Of course, 'apparently mild' is still nastier than any post-war U.S. recession, but a 2.6% decline in real GDP translates into a 4.4% decline in per capita GDP, which puts the U.S. at the low end of financial crises described by Reinhart and Rogoff.

What's more, the 2007-2009 productivity gain had a major impact on both economic policy and political discourse. First, it gave the Fed a free hand to feed mammoth amounts of liquidity into the system without worrying about inflation.

Second, the productivity gains convinced the economists of the Obama Administration that the economy was basically sound and there was nothing wrong with the 'supply-side of the economy. Instead, the Obama Administration concluded that the big problem was a demand shortfall, which is why they expected things to get back to normal after the fiscal stimulus.

Consider this April 2010 [speech](#) from Christina Romer, then head of the CEA. She said:

The high unemployment that the United States is experiencing reflects a severe shortfall of aggregate demand. Despite three quarters of growth, real GDP is approximately 6 percent below its trend path. Unemployment is high fundamentally because the economy is producing dramatically below its capacity. That is, far from being “the new normal,” it is “the old cyclical.”

Perhaps more important from a political perspective, the productivity surge helped convince the Obama economists that the job loss was ‘normal’ in some sense. That is, the rise in productivity suggests that the financial crisis apparently just accelerated the normal process of making U.S. businesses leaner and more competitive, and there was no structural problem. In fact, that’s what Romer said in her speech:

In short, in my view the overwhelming weight of the evidence is that the current very high — and very disturbing — levels of overall and long-term unemployment are not a separate, structural problem, but largely a cyclical one. It reflects the fact that we are still feeling the effects of the collapse of demand caused by the crisis. Indeed, at one point I had tentatively titled my talk “It’s Aggregate Demand, Stupid”; but my chief of staff suggested that I find something a tad more dignified.

If productivity was rising, then the job loss was due to a demand shortfall and could be dealt with by stimulating aggregate demand. That, in turn, helps explain why the “job problem” didn’t seem so urgent to the Obama administration, and why they spent more time on other policy issues such as healthcare and regulation.

### Analysis

But when we take a step back, the productivity surge of 2007-2009 is really quite a surprise. Historically, productivity is pro-cyclical. That means in the early stages of recession demand and output typically drop faster than employment, which dampens productivity growth and sometimes even turns it negative. This makes sense, since it typically takes time for businesses to figure out that the bottom has fallen out and that they need to cut workers.

Financial crises, because they are so sudden and so turbulent, should be the classic case of procyclical productivity. During the ‘acute’ stage of the Great Depression (1929-1933), output per hour fell by a remarkable 8%-10%, a sharp break from the rapid productivity increases in the run-up to the Depression. (These are historical numbers, so the exact figure depends on which series we pick).

Why was this financial crisis different? To put it more precisely, in the third quarter of 2008, the world went to hell—does it really make sense that business productivity rose by 4.8% over the next year?

To understand the 2007-2009 productivity surge better, I looked at the productivity performance of individual industries, using the data on industry real value-added and full-time equivalent

employment (FTE) put together by the Bureau of Economic Analysis (see [here](#)). The productivity growth of an industry, then, is the growth of real value-added per FTE.

First I wanted to figure out which industries were driving the overall good productivity performance. The following chart sorts industries by their productivity growth rate from 2007 to 2009 (the top purple bar) . In addition, the blue, lower bar shows the growth rate of real value-added for each industry (blue bar). For example, the computer and electronic products industry had annual productivity growth of about 16% per year between 2007 and 2009, and annual real value-added growth of 9% a year. (Important note: The label for the sector 'Educational services... got cut off. The full label is 'Educational services, health care, and social assistance').

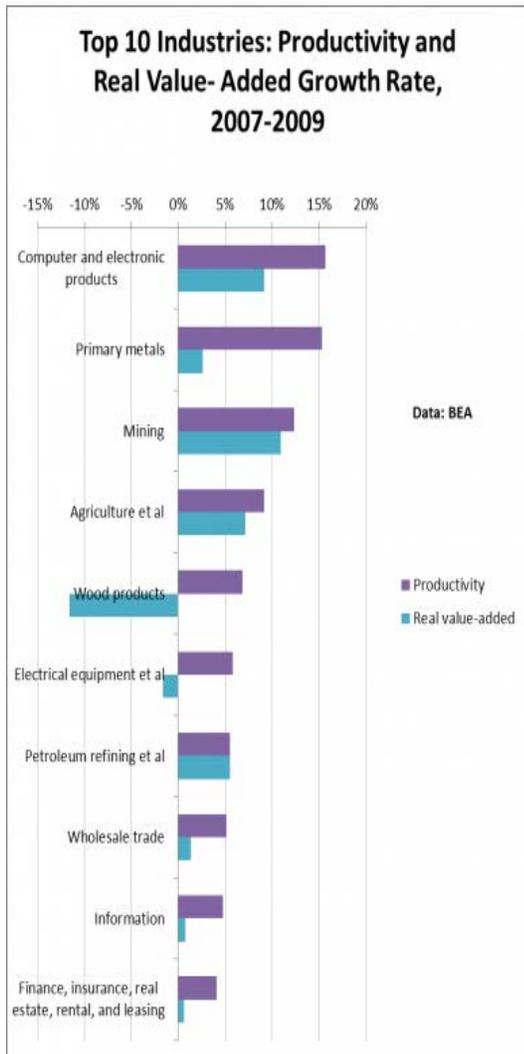
## Productivity and Real Value Added Growth Rate, 2007-2009



We can see a couple of things. Most industries had negative or very low productivity growth in this period, as you would expect during a financial crisis. But there were a few industries where productivity growth was quite strong.

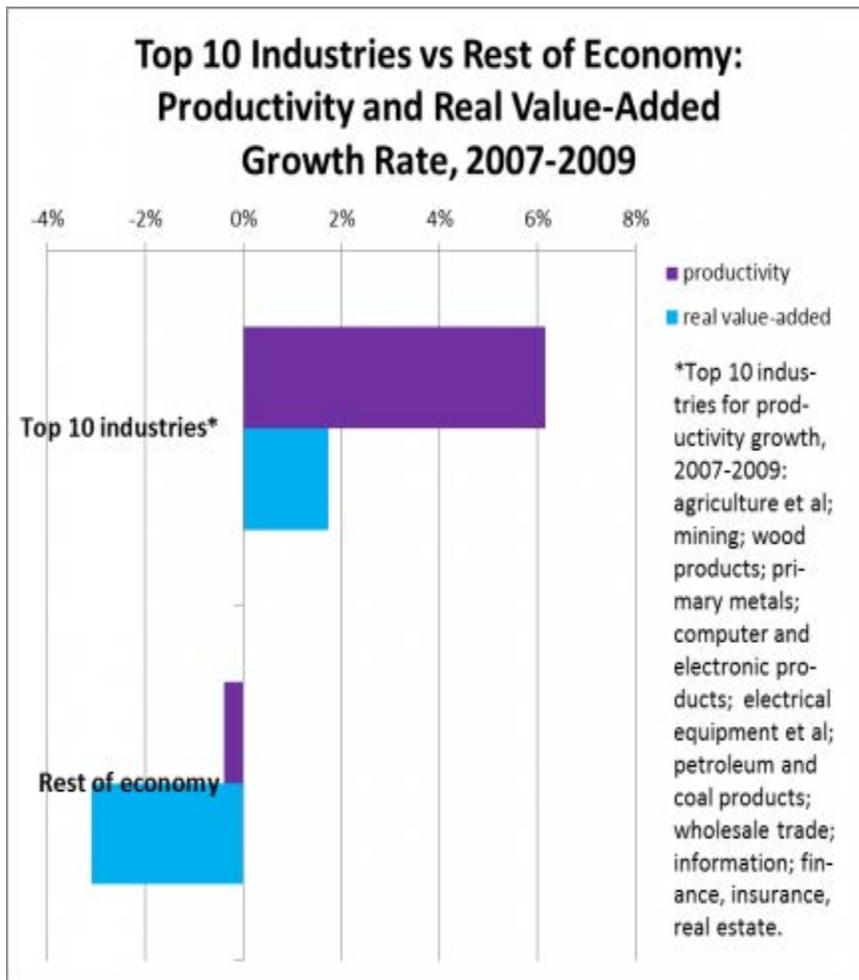
This next chart is a closer-in look at the top 10 industries, ranked by productivity growth in the 2007-2009 period. These ‘high-performing’ industries are computers and electronic products; primary metals; mining (including oil and gas); agriculture and related products; wood products; electrical machinery and related products; petroleum refining; wholesale trade; information; and finance, insurance, and real estate .

This chart already raises some questions. Remember that the blue bar is the growth rate of real-value added. According to the BEA’s data, the productivity gains in 9 out of 10 of these industries is actually accompanied by a rise in output (real value-added) from 2007 to 2009. That means during the biggest financial crisis in recent history, the BEA is estimating that domestic value-added actually increased in industries such as primary metals, mining, agriculture, and financial services (yes, that’s right, financial services). That’s a bit odd.



With a little bit of figuring, I can calculate the performance of the top 10 industries versus the rest of the economy. So the top 10 industries had a productivity growth rate of more than 6% per year during the financial crisis, while the rest of the economy had negative productivity growth. So clearly if we want to understand the mystery of the productivity boom, we need to understand how these industries apparently did so well while the economy was tanking.

Luckily, in many of these industries, we actually have physical measures— tons of steel, barrels of oil, bushels of corn—that gives us a check on the BEA numbers (The one exception is computers et al, but I've got a different approach there). And the physical measures tell a very different story than the BEA data.



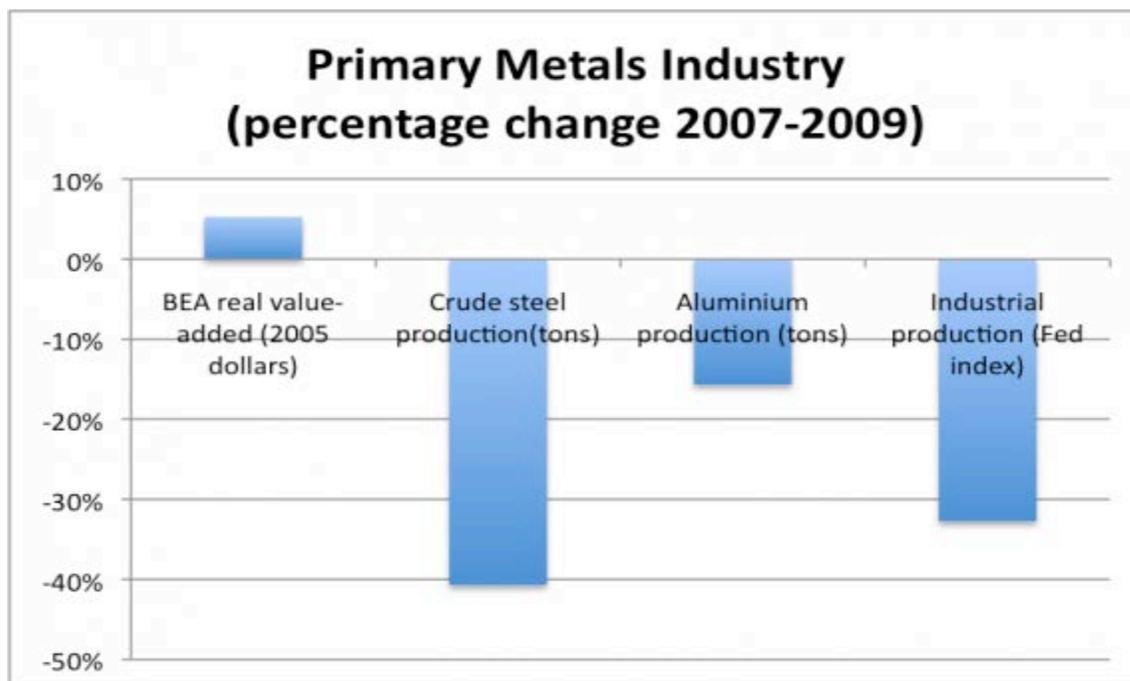
### Primary Metals

Let's start with primary metals, which includes steel, aluminum, and other metals. According to the BEA's official statistics, real value-added in the primary metals industry—steel, aluminum and the like—rose by a total of 5.3% from 2007 to 2009 (this is a cumulative figure, rather than an annual rate). If this statistic is correct, this is a truly amazing performance by an industry which went through tough times over the past twenty years. The performance is especially inspiring

given that the auto industry—one of the biggest customers for steel—was completely flattened by the crisis.

But a look at physical output—steel and aluminum production, measured in tons—tells a much different story. In fact, over this two year period, steel production cratered by more than 40%. Aluminum production wasn't hit as badly, but it too fell. And the Federal Reserve's industrial production index for primary metals fell sharply as well.

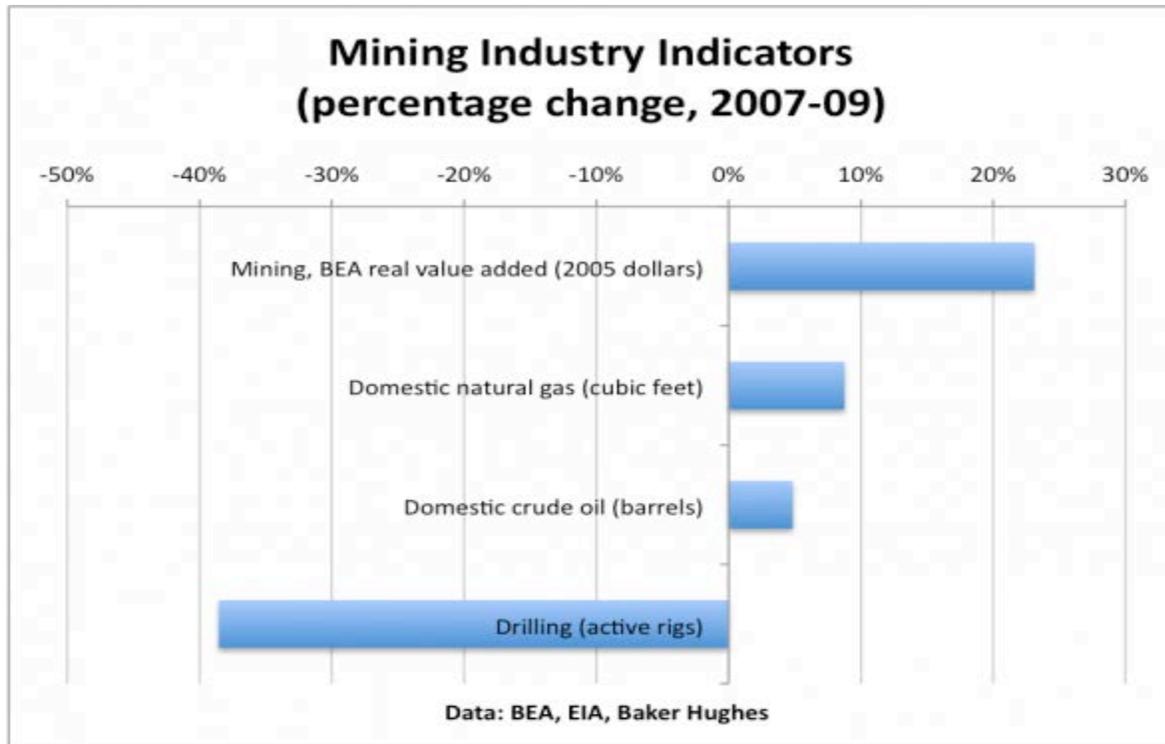
I'd say that the BEA's gain in real value-added and productivity in primary metals seems implausible, at best (a bit later on I'll explain how the BEA did its calculation, and where the problem is)



### Mining and Petroleum Products

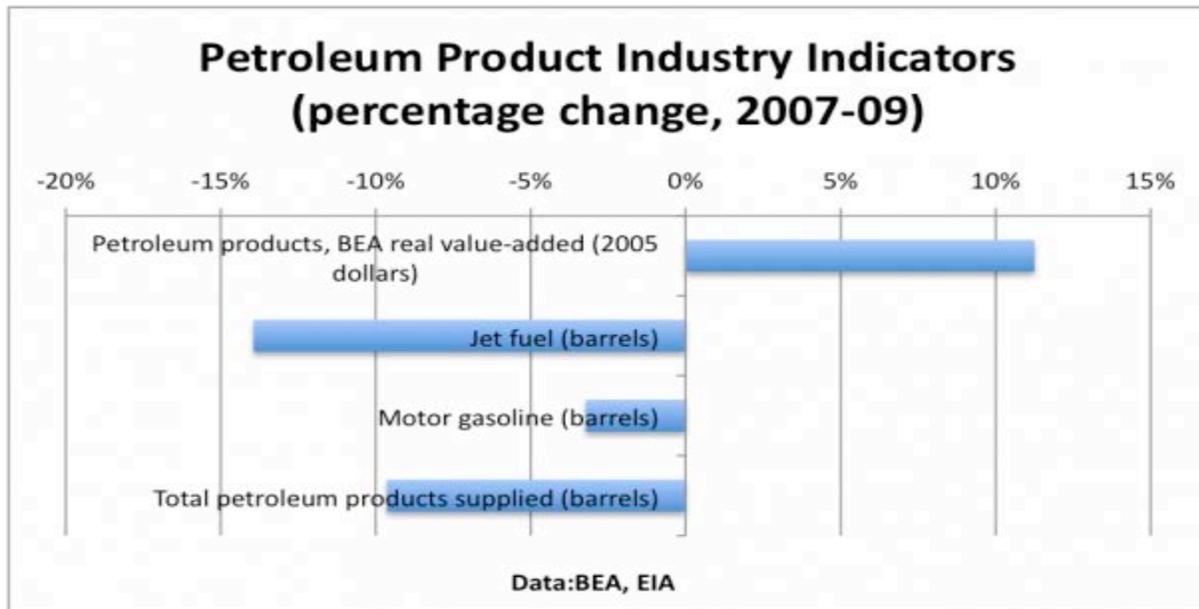
We don't usually think of mining as a fountain of productivity. But in the 2007-2009 period, domestic mining—which is mostly oil and gas extraction—showed a 26% two-year productivity gain according to the official statistics, and a 23% two-year real value-added gain. Similarly, the petroleum products industry—refining and the like—apparently recorded 11% productivity and real value-added gains from 2007 to 2009.

In theory, these sorts of gains are not impossible, even during a recession. The oil industry has always marched to its own beat, and it is technologically sophisticated. However, such gains should be reflected in physical measures such as barrels of oil and natural gas, drilling rigs, and petroleum products produced. Did we suddenly see a great spurt of production? No, no, and no.



Domestic production of oil and natural gas ticked up modestly during this period, just as you'd expect—but far short of the industry growth estimated by the BEA. In other words, if we judged output and productivity growth in the mining industry by physical measures, we'd get a much lower number than the BEA reported.

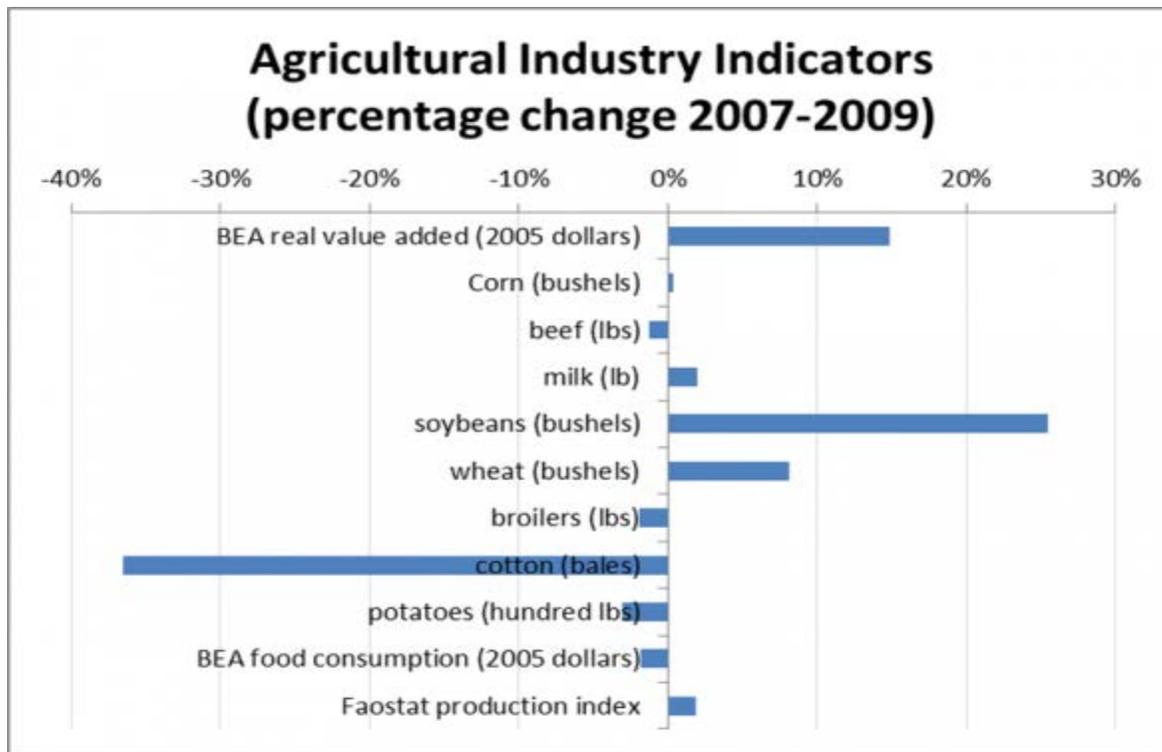
The same thing is true in the petroleum products industry—refining and the like. Here too we have physical measures—the production of things like jet fuel, gasoline, and other products. And here too we can see that physical production fell, even while reported real value rose—another industry where we should be suspicious of the reported productivity growth.



### Agriculture

I'm sure you are getting a bit bored here, but I could do this all day. Agriculture is ranked number #4 on the list of top productivity growth industries during the financial crisis, with an average annual productivity gain of 9% per year in 2008 and 2009, and an average value-added gain of 7% per year, or roughly 15% over two years. Not bad, if true, for an often-overlooked sector.

But once again we have physical measures, and once again they tell a much more pessimistic story than the BEA figures. The agriculture industry has a very diverse set of products, so I can only offer you a sample of them in the chart below.



We see one major crop—soybeans—which outperformed expectations. The rest ranged from doing okay (wheat) to weak positive growth (corn, milk) to moderate shrinkage (beef, broilers, potatoes) to collapse (cotton). (‘Broilers’ are chickens raised to be eaten, in case you didn’t know). The FAOstat production index for U.S. agriculture, as calculated by the Food and Agriculture Organization of the United Nations ([FAO](#)), shows a very tepid gain of 1.9%—much lower than the 15% gain shown by the BEA figures.

### Computer and Electronic Products

And now we come to #1, the top productivity gainer in the 2007-2009 period, the computer and electronics products industry. Whooooee! Despite the economic disaster, the shining light of the U.S. economy apparently produced huge productivity and output gains in this period.

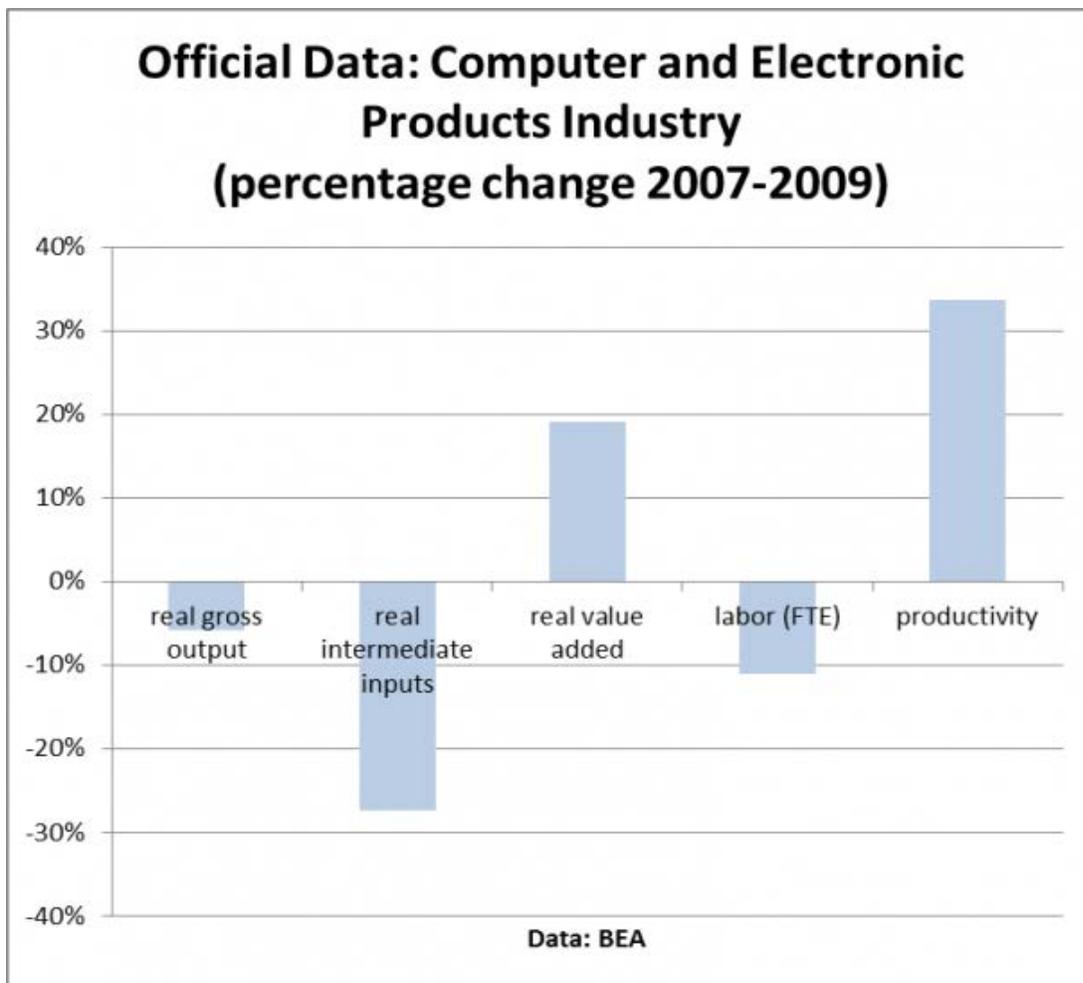
But wait! There’s something tricky going on here. I’m going to have to pause for a second and hit you all with a few definitions. *Gross output* for an industry is basically its sales. *Intermediate inputs*, in the words of the BEA, are “the goods and services (including energy, raw materials, semi-finished goods, and services that are purchased from all sources) that are used in the production process to produce other goods or services rather than for final consumption.” And *value-added* is gross output minus intermediate inputs. GDP is a value-added measure, and productivity is based on value-added as well (important!)

So for a U.S.-made computer, for example, the gross output would be the value of sales based on the producer price of the computer (the price as it leaves the factory). The intermediate inputs would be the value of all the purchased components that go into the computer—chips,

software, case etc—plus the energy to run the factory, plus all the purchased services (legal, financial, computer programming, and so forth).

If a company outsources some of its activities—say, the internal customer help desk is outsourced to a firm in North Dakota, or routine computer programming is outsourced to Bangalore—then the amount of intermediate inputs can go up, even if the gross output stays the same. So if the ratio of intermediate inputs to gross output in an industry increases, that in some sense tells us that the amount of outsourcing/offshoring is increasing, because companies are buying more intermediate inputs to produce a given amount of output.

So how did the computer and electronics industry achieve its big productivity gains in the financial crisis? Did IT manufacturers accelerate their offshoring/outsourcing to take even more advantage of cheap Asian producers? Um, no...not according to the official data. According to the official data, IT manufacturers cut way back on their offshoring/outsourcing in the 2007-2009 period, even while they were cutting back on workers.

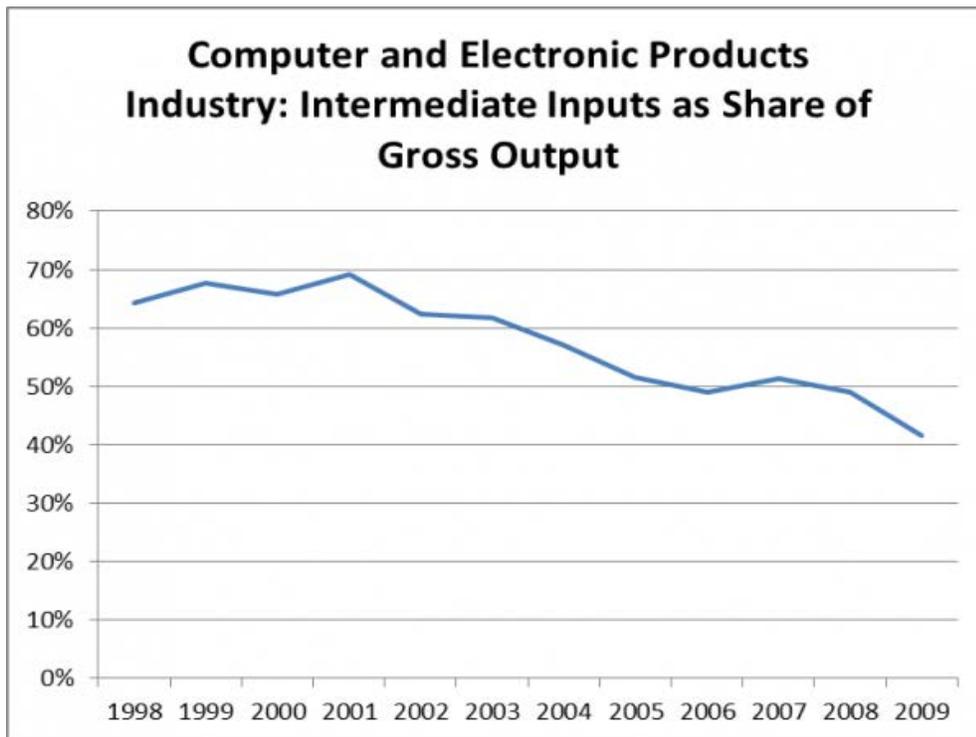


Take a close look at this chart, because it's very important. We're going to read from left to right here. (All calculations are based on official BEA data for the 2007-2009 period). First, real gross output fell by about 6%. If you think of that as sales of domestically produced electronics equipment, adjusted for inflation, that sounds about right. Second—and this is the weird one—the official data apparently shows that intermediate inputs fell by 27%. In other words, to produce 6% less gross output over the course of two years, manufacturers managed to figure out a way to use 27% less inputs, both imported and domestic, in only two years.

This apparent fall in intermediate inputs is what drives the entire stellar productivity performance of the industry. Value-added is the difference between gross output and intermediate inputs. Because real intermediate inputs fell so much faster than real gross output, pure arithmetic means that real value-added grew very quickly. And since productivity is the ratio of value-added to labor, with value-added rising and labor falling, that generates an enormous productivity increase.

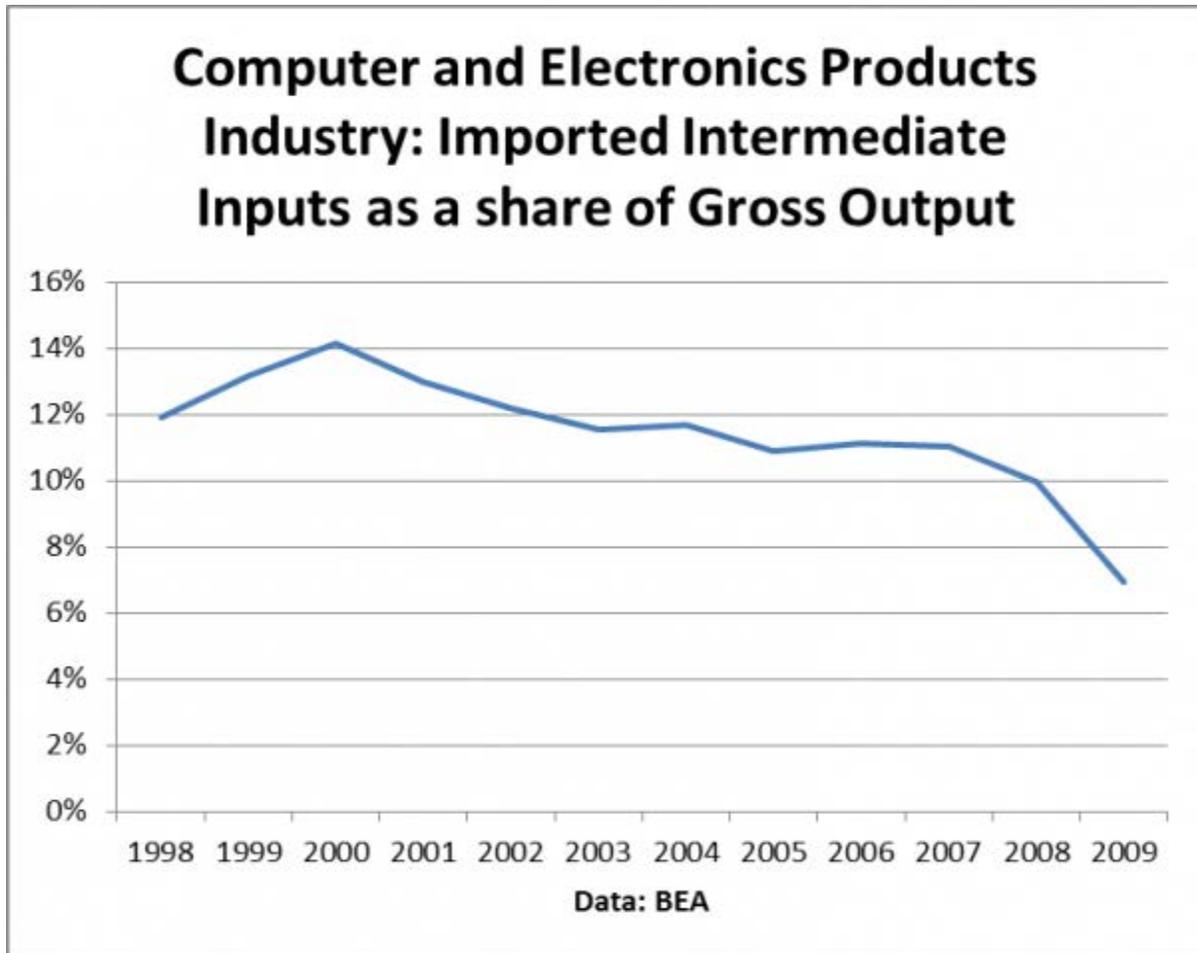
So now we must ask ourselves the question: Is it reasonable that domestic electronics manufacturers could achieve such efficiencies in only two years that they could produce almost as much output with so much less input? In theory, I guess it can't be ruled out.

However, a more persuasive explanation, in my opinion, is systematic mismeasurement of intermediate inputs. To see the problem, let's take a step back and look at the last decade. We would expect that as U.S. IT hardware companies outsourced more and more of their manufacturing and services to Asia, the ratio of intermediate inputs to output would have gone up. After all, the IT industry is the world champion in outsourcing. But when we look at the data, we see something weird.



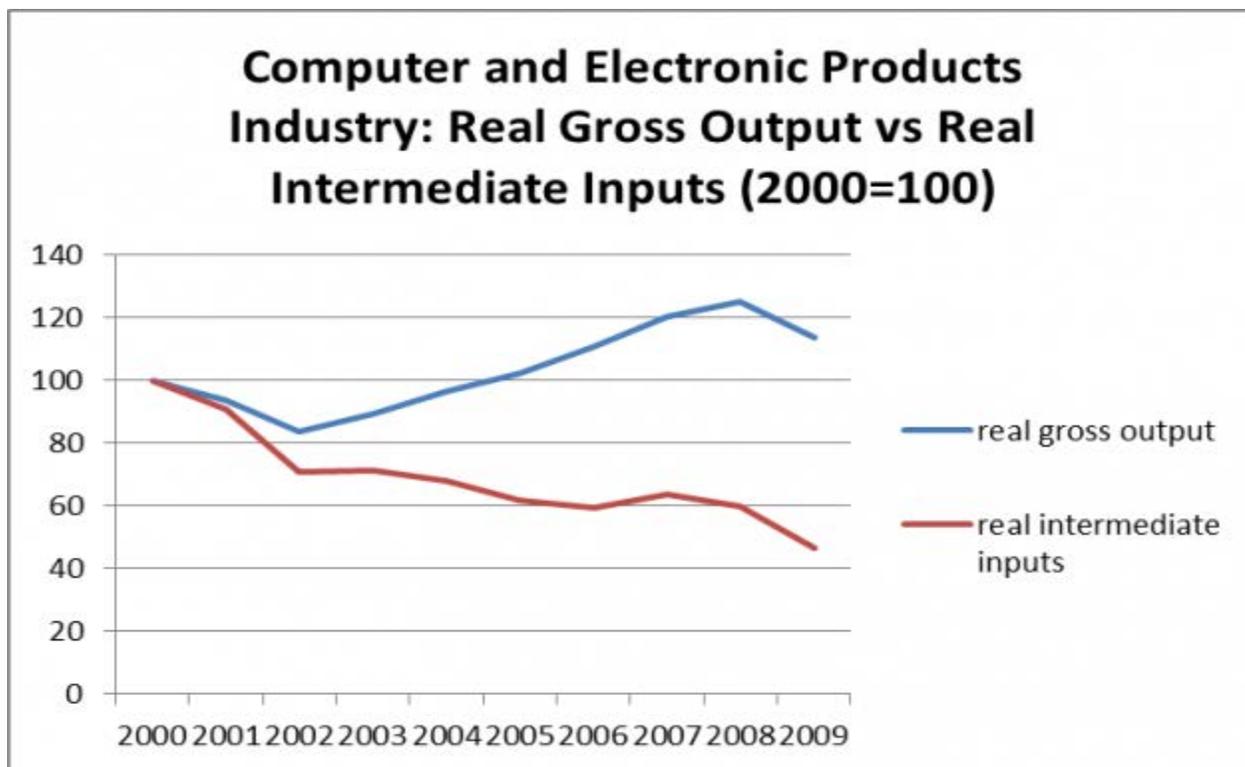
Whoa! It looks like the decline in intermediate inputs actually started back in 2000, if you believe the official data. That would imply a long-term trend of the computer and electronics industry getting more integrated, and the degree of outsourcing/offshoring going down, rather than increasing.

The official data are even harder to believe when it comes to imports. We all know that anything electronic that you buy is packed full of imported components. But according to the official numbers, imported intermediate inputs have been falling as a share of gross output, to well under 10%. Odd.



So far we've been measuring intermediate inputs in nominal dollars, not adjusted for inflation. But when we add in the adjustment for changing prices, the problem gets worse rather than better. According to the official data, real intermediate inputs for the computer and electronics products have actually been declining for the past ten years, even as output has been rising! In other words, the computer and electronics industry has been getting less dependent on offshoring/outsourcing, rather than more.

This makes no sense at all. It simply doesn't fit with the evidence of our eyes.



#### Explanation

The best explanation is that the statistical system is not set up to properly capture the impact of offshoring/outsourcing, especially to other countries (for one technical summary, see Susan Houseman's [article](#) in the February 2011 Survey of Current Business). I'm not going to go into detail here, but there are three interlocking statistical problems.

\*First, outsourcing to low-cost foreign producers is not picked up correctly in the data. If a part that costs \$10 to make in the U.S. is offshored to a foreign producer who charges \$1 per piece, there is no statistical agency that tracks or measures that decline in price. So all those price drops from offshoring? It's as if they didn't happen.

\*Second, when a product is already offshored, and switched to an even lower-cost foreign supplier, that decline too is probably not measured.

\*Three, in product categories with declining prices and rapid model changes—such as cell phones, computers, consumer electronics—the official import price indices underestimate the size of the price decline for product categories with rapid model changes (I call this the 'Nakamura-Steinsson effect,' after the two [economists](#) who discovered it). The reason is simple—when a new model of an imported good is introduced, the BLS typically treats it as a new good, and misses all the price decline from one model to its successor.

Taken together, these three problems have the implication that the price drop of imported intermediate inputs is being severely underestimated, which means that the real growth of intermediate inputs is being severely underestimated as well (see Houseman's paper for the links

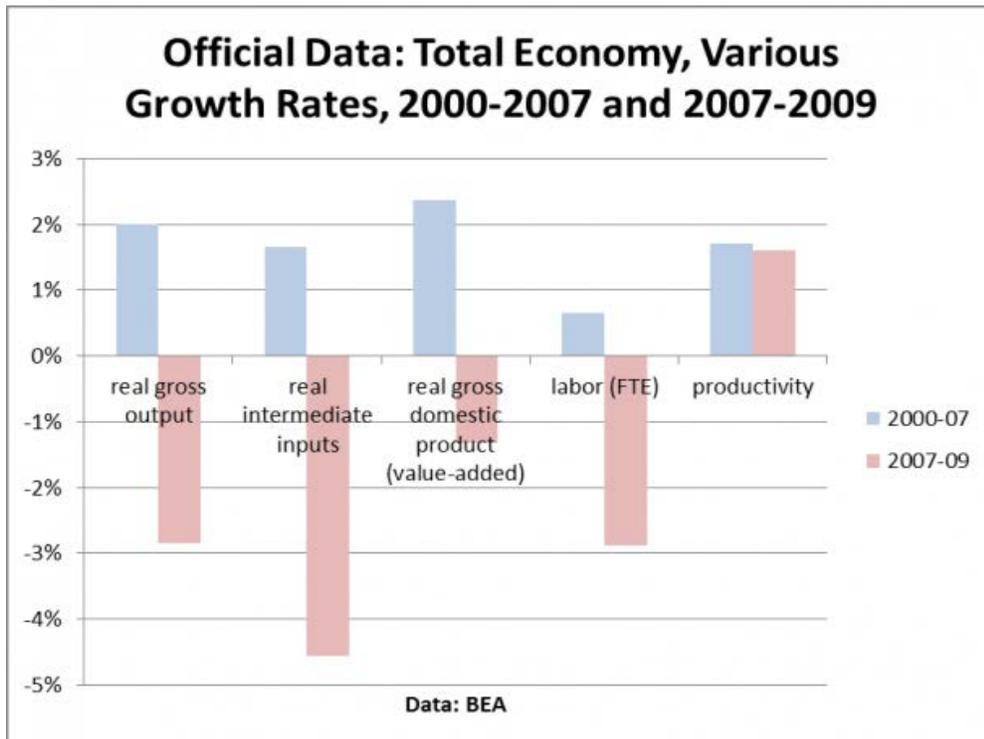
to the relevant literature). It simply does not make sense that intermediate inputs for the computer and electronic product industry should be getting less important over the past ten years.

Given that the long-term decline in intermediate inputs is not believable, I think it's reasonable to be suspicious of the apparent short-term collapse of intermediate inputs as well. It makes more sense to believe that use of intermediate inputs was cut proportionally to output, rather than far faster. In that case, the two-year productivity gain in the computer and electronics industry would look more like 6%, rather than 34%.

The mismeasurement problem for intermediate inputs, though, is not just restricted to the computer and electronics industry—it extends to the overall economy as well. Given the central role of offshoring/outsourcing in the economy, we would expect to see intermediate inputs for the whole increasing in importance over time. Instead, intermediate inputs have been growing less quickly than gross output, implying they are becoming less important.

From 2000 to 2007, the BEA measure of real intermediate inputs for the whole economy grew at a 1.6% growth rate, compared to 2.0% for gross output. That doesn't seem like much of a difference, but it's enough to push up the growth rate of real GDP from 2% to 2.4%, with a comparable effect on productivity growth.

The growth gap between gross output and intermediate inputs was much bigger in the 2007-2009 period. In the official stats, real gross output fell at a 2.9% rate, compared to a 4.6% decline for intermediate inputs. Ordinarily we might say, well, that's okay, we had a collapse in financial services. But in fact the official data says that financial services barely shrunk in this period. So that's not it.

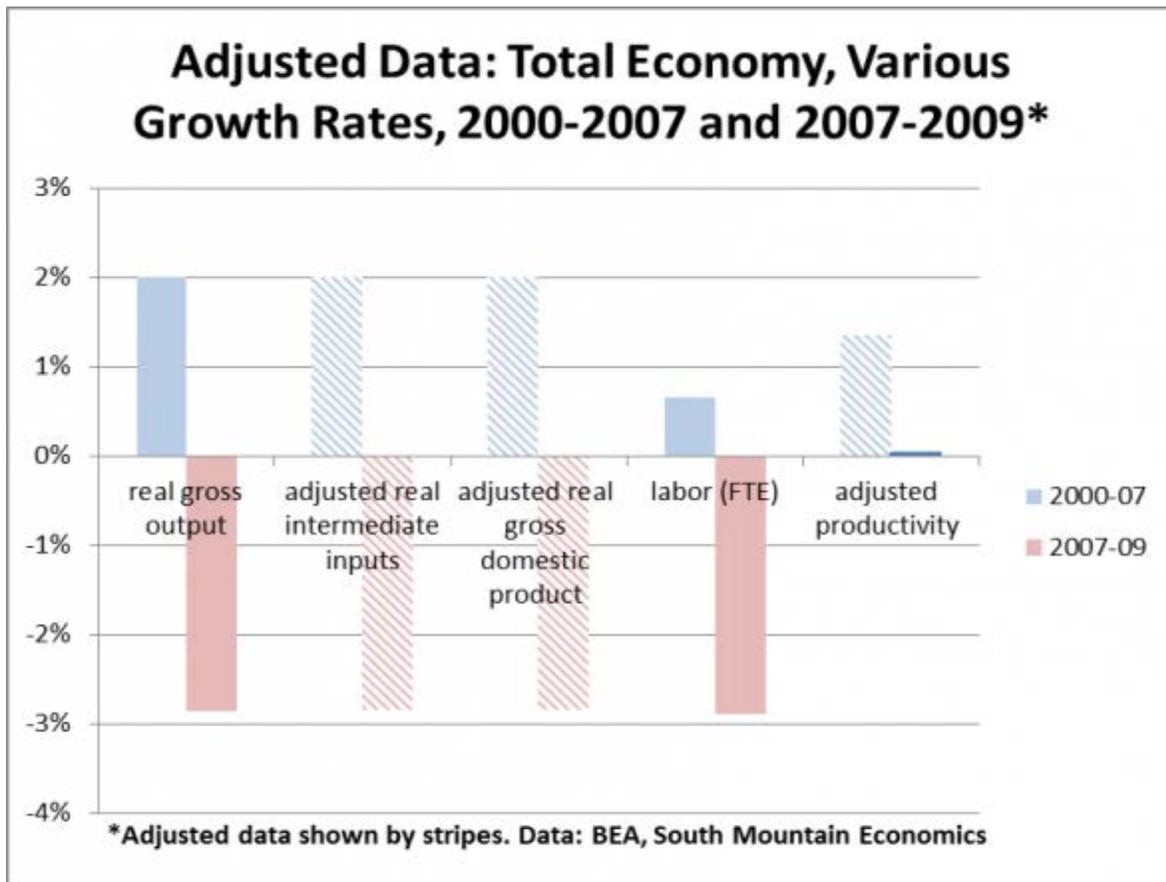


In fact, the ‘high-productivity’ industries that we discussed earlier—primary metals, mining, and agriculture—show the same phenomenon. Each of them appears to have a drop in intermediate inputs exceed the decline in gross output. This difference, in turn, drives the apparent productivity growth. Mining, for example, combines a 10% drop in real gross output with an apparent 46% drop in real intermediate inputs, leading to a reported 23% gain in real value-added and a 26% gain in productivity. It’s very hard to understand how intermediate inputs decline four times as fast as output!

Adjusted statistics

Given the clear internationalization of the economy, we would expect that real intermediate inputs would grow faster than gross output over the long run, not slower. And during the financial crisis, with companies doing everything they could to survive, we would expect that financially-pressed businesses would rely relatively more on cheap suppliers, not less.

So let’s split the difference—let’s see what happens if we allow intermediate inputs to grow (and shrink) at the same rate as gross output. The resulting ‘adjusted data’ won’t show a relative increase in offshoring/outsourcing, but at least they won’t show a decline. The chart below shows the new adjusted numbers. The striped bars are the ones that were changed, the solid bars remained the same.



What's the difference between the original data and the adjusted data? In the period 2000-2007, making the adjustment basically takes about four-tenths of a percentage point off GDP and productivity growth. That is, making the adjustment reduces total economy productivity growth from an average 1.7% per year to an average 1.3% per year over the seven year stretch from 2000 to 2007. That's significant but not enormous.

However, the effect of the adjustment on the 2007-2009 period is spectacular. Productivity growth, which had been 1.6% annually in the original data, basically disappears. The decline in real GDP is twice as large (-1.3% per year in the original data, -2.9% in the adjusted data). And economists are no longer presented with the confounding puzzle of why unemployment rose so much with such a modest decrease in GDP—it's because the decrease in GDP was not so modest. (see a [piece](#) here on Okun's Law, which links GDP changes with unemployment changes).

This picture looks a lot more like we'd expect a financial crisis to look—a sharp drop in GDP, and little or no productivity growth.

Before going on to the implications for policy, let me say two important caveats. First, I still can't connect all the dots between the undeniable problems in the data, and the final macro outcomes. I'm getting closer, but it's still not there yet. Second, remember that we just adjusted intermediate input growth to show a constant level of offshoring/outsourcing. If we assumed growing offshoring/outsourcing, then the change would have been bigger.

### Implications

First, the measured rapid productivity growth allowed the Obama Administration to treat the jobs crisis as purely one of a demand shortfall rather than worrying about structural problems in the economy. Moreover, the relatively small size of the reported real GDP drop probably convinced the Obama economists that their stimulus package had been effective, and that it was only a matter of time before the economy recovered.

A more accurate reading on the economy would have—perhaps—cause the Obama Administration to spend more time and political capital on the jobs crisis, rather than on health care. In some sense, the results of the election of 2010 may reflect this mismatch between the optimistic Obama rhetoric and the facts on the ground.

But there's a broader issue as well. As we saw above, the mismeasurement problem obscures the growing globalization of the U.S. economy, which may in fact be the key trend over the past ten years. Policymakers look at strong productivity growth, and think they are seeing a positive indicator about the domestic economy. In fact, the mismeasurement problem means that the reported strong productivity growth includes some combination of domestic productivity growth, productivity growth at foreign suppliers, and productivity growth "in the supply chain". That is, if U.S. companies were able to intensify the efficiency of their offshoring during the crisis, that would show up as a gain in domestic productivity. (The best case is probably Apple, which has done a great job in managing its supply chain for the iPod, iPhone and iPad and extracting rents).

From an economic and policy point of view, there's a big difference between purely domestic productivity gains, productivity gains at foreign suppliers, and productivity gains 'in the supply chain'. The benefits of domestic productivity gains will like accrue to the broad array of production and nonproduction workers in the U.S. The benefits of productivity gains 'in the supply chain' will likely go to the executives and professionals, both in the U.S. and outside, who set up, maintain, improve, and control supply chains. That's a much smaller, globally mobile group. And the benefits of productivity gains at foreign suppliers? Well, that depends on how much power U.S. buyers have vis-a-vis their suppliers...that is, competitiveness.

From this perspective, the Obama Administration needs to pay much more attention to trade and competitiveness than it has managed up to now. And thinking about measured productivity gains as a combination of purely domestic, purely foreign, and supply chain productivity offers a different way of thinking about globalization. If we want to bring jobs back to the U.S., we need to be able to improve purely domestic productivity. If we want to be able to make sure that we keep getting the benefits of gains in foreign supplier productivity, that means we have to be able to offer something unique to the suppliers—new products or advanced R&D or better information about the U.S. markets. And if we want to get the benefits of supply chain productivity gains, then we have to make sure more Americans are globally sophisticated and able to participate.

I'm going to stop here, since this post is already far far longer than it should have been. The next paper, on new measures of competitiveness, will come out from the Progressive Policy Institute soon. Stay tuned!