PALLADIUM
GOVERNANCE FUTURISM

PATRICK FITZSIMMONS • OCTOBER 3, 2025 • **ARTICLES**

How GDP Hides Industrial Decline

Ed Brown/Frances Perkins Building in Washington, D.C.

For the past few years I have been mulling a paradox: U.S. GDP keeps going up, yet it seems like we make less stuff and that most of the smart people I know work fake jobs. Growing up in the nineties, most of my toys and clothes had tags saying “Made in Hong Kong” or “Made in Vietnam.” But the high-skill, high-tech goods—the washing machine, the car, my computer—were often made in America. Now? From my e-bike to my laptop, from my refrigerator to my mattress, very few goods I own, high-tech or low-tech, were made in the USA.

Meanwhile, I have heard arguments that America is actually making more things than ever. According to *The Economist*, lost jobs are due to automation, not foreign competition, and it is a good thing that machines have liberated us from factory work and enabled more service jobs. Everyone from the American Enterprise Institute to *The Wall Street Journal* to Wikipedia agrees that U.S. manufacturing has not just not significantly fallen—but it has *never been higher*.

Sometimes there are discrepancies between your real-world observations and the data. But this goes far beyond just being a discrepancy: the data is saying the complete opposite of what we see with our own eyes, hear from our acquaintances in the job market, and deduce logically from our knowledge of demographics, technology, industry, and trade. How is this possible? The answer is actually very simple: the data is completely wrong. But you can only figure this out if you go line-by-line into the hundreds of pages of government GDP calculation methodology documentation. Which is exactly what I did.

The most commonly cited graph shared to demonstrate U.S. manufacturing strength is based on the U.S. Bureau of Economic Analysis’s (BEA) manufacturing “real value-added” data, which looks at manufacturing as a subset of total GDP. This graph has been cited by Federal Reserve economists, *Washington Post* columnists, professors—all claiming it refutes the idea that the U.S. economy has been hollowed out. Adjusted for inflation, it shows manufacturing is up 71% since the dataset began in 1997, and up a healthy 37% per capita:

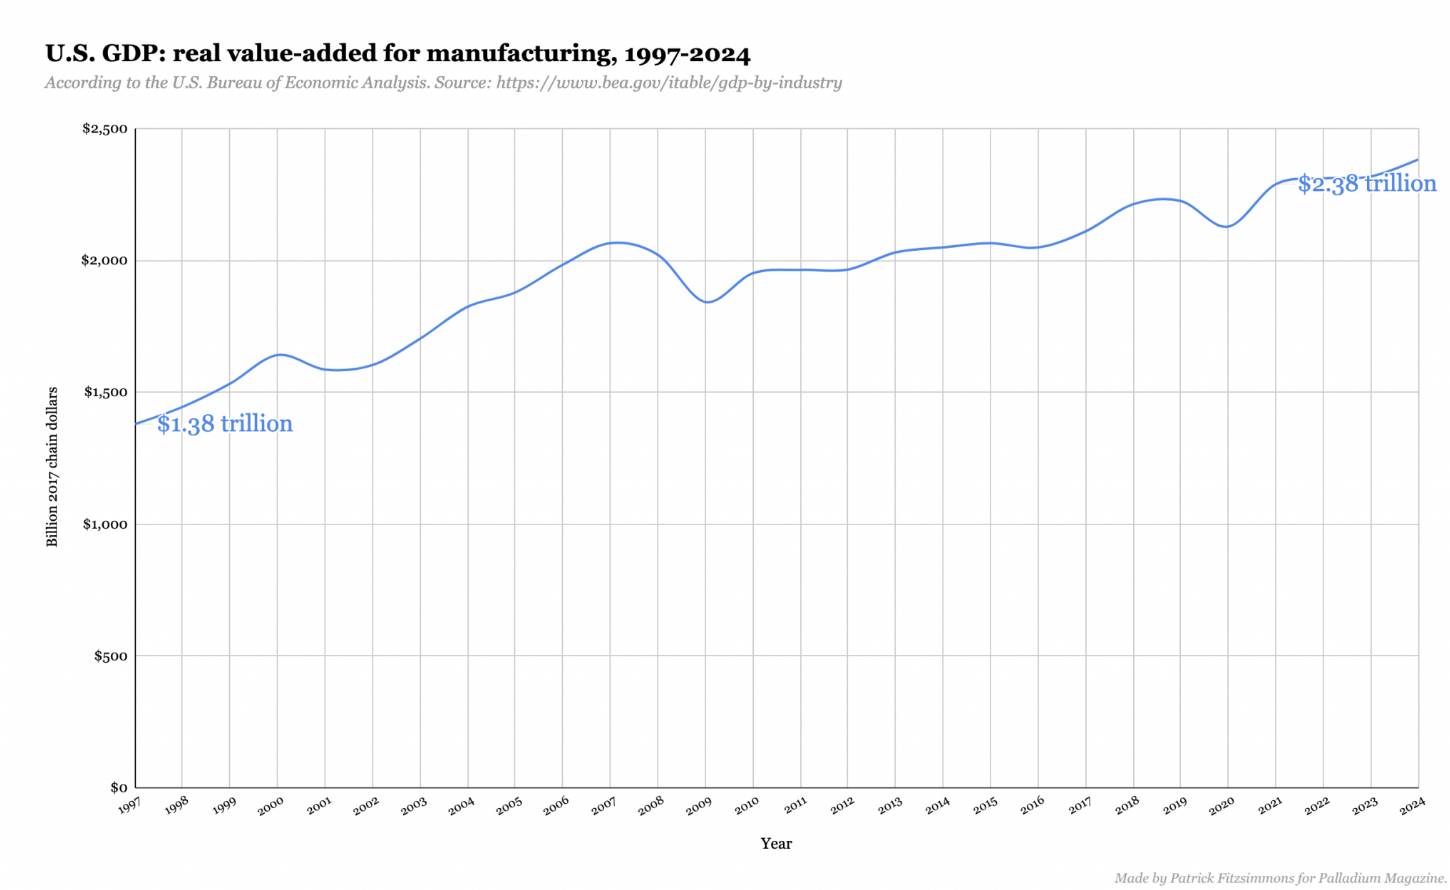

Patrick Fitzsimmons/U.S. real value-added for manufacturing according to the U.S. Bureau of Economic Analysis, 1997-2024

You might think that a measure of manufacturing would in some way measure actual manufactured goods emerging from U.S. factories, like tons of steel rolling out of mills, number of CPUs coming out of chip fabs, and cars rolling off the assembly line. But this is not the case. Despite the name, “real GDP” in practice is the result of hundreds of arbitrary and subjective decisions made by government-employed economists, such as “education administrators are more productive than teachers” or that a 25% increase in automobile “quality” can theoretically show up as a 166% increase in “real GDP value added.”

It is remarkable how there is so much commentary on GDP, yet so few people have truly wrestled with the numbers and where they come from: an advanced MIT macroeconomics textbook will reference GDP over sixty times, yet not once acknowledge the decision-making that goes into making this number. Likewise, pro-market public intellectuals are only too happy to cite GDP to make a point without considering that the metric itself is the antithesis of market capitalism: GDP is a very complicated statistical construct that is made by government bureaucrats behind closed doors without any ability of the public to replicate, audit, or verify assumptions. Sometimes, these kinds of constructs can be useful for accurately representing real-world phenomena, like manufacturing capacity. But a dive into how the sausage is made makes clear that GDP is not one of them.

Where GDP Comes From

Back in the 1930s, U.S. policy-makers and economists were facing two big problems. The first was that the nation was suffering an economic crisis with millions of people and businesses suffering loss of income, yet there was no existing way to sum up incomes across the economy to get a sense of how the nation as a whole was doing. The second problem was that previous economic statistics focused on raw commodity outputs like bushels of wheat grown or tons of steel produced, but now more of the economy was in services, government work, and heterogeneous manufactured products.

The U.S. government commissioned economists to create a comprehensive set of national income statistics. The economists kept working and went on to develop statistics aimed to measure the economic output by category of the economy and then, ultimately, to sum it up into one number that eventually would be known as “gross domestic product,” or GDP.

At first glance, summing up the economy into a single number seems impossible. How do you add up apples and oranges? Or, say, apples, cars, and dentist appointments? One’s first inclination may be to add up total sales receipts in each category, since dollar spending can be compared between products and over time. But this fails because a rise in total spending for a product category may be merely the result of a rise in price, not more production. The clever solution was to combine expenditures or sales receipts in dollar terms with measures of the average price of each item. If spending on cars has doubled, but the price of the average car has also doubled, then there was no real change in production of cars. But if spending on cars doubled while the price only increased 50%, then there was a substantial increase in cars purchased.

Statistical agencies built giant databases of prices across all products, then matched the prices with expenditures for each category. Now it seemed as if they had the ability to do what seemed impossible, and add up changes in the economic sectors into one number. The top economic textbooks have called this "truly among the great inventions of the twentieth century."

Today, GDP numbers are calculated by the Bureau of Economic Analysis, which is staffed by career economists with no political appointees involved. It's under the U.S. Department of Commerce. However, the methodology is also an international effort. U.S. economists join with their foreign counterparts at a United Nations committee to define the "System of National Accounts." These are standards that nations around the world at least try to adhere to—adherence can be a requirement for World Bank loans. The methodology also changes over time. For instance, in 2012 there were major updates to add development of intellectual property (IP) as being its own contributing component to GDP.

Discourse over GDP is frequently confused because there are actually three different calculation approaches: the *income* approach, the *expenditures* approach, and the *value-added* approach. In the textbooks, usually only the expenditures approach is taught. In theory, each approach should sum to the same total number, since everyone's income must come from someone else's spending. However, when comparing sectors, such as government or healthcare, the totals differ for each approach, and this can create a lot of confusion. On top of that, each approach has a *nominal* and a *real* version. Thus when a news report refers to "GDP for healthcare", this could be referencing one of six different numbers!

For the income approach to GDP, the process is to add up every person's compensation, plus corporate retained earnings, plus some adjustments. For the expenditures approach, the formula is to sum the final expenditures of private consumers, the capital expenditures of businesses, the spending of the government, and exports, then subtract imports. "Final expenditures" means that the price of a car bought by a consumer is counted, but the money spent by the car dealership on its electric bill, or the money spent by the car factory on steel, is not counted. Counting non-final expenditures would result in double-counting and would break the number. The expenditures approach is most frequently taught in Economics 101.

Finally, we have the value-added approach. Rather than counting just final sales, this approach counts the sale value minus the input costs at each step. Imagine a gallon of milk sold from a grocery store for $5. Rather than just counting that $5 as a "final expenditure," the value-added approach counts the cow feed sold to the farmer for $1, then the milk milked by the farmer for $4 (thus *adding* $3 in *value*), and then the gallon of milk sold by the store for $5 (adding $1 in value). You add up all the steps and get the same $5.

Each approach has its uses, but you have to be careful with which you use. What percent of GDP is healthcare? You get two different numbers depending on the approach. With the expenditures approach, healthcare is 17% of GDP, but for the value-added approach only 8%. Why? Because the value-added approach only counts expenditures on hospital and clinic workers toward the healthcare category. Money spent on manufacturing medical devices counts as manufacturing; money spent building hospitals counts as construction. For measuring healthcare's share of the economy, it is probably better to use the expenditures approach because it is reasonable to include pharmaceutical production and hospital electricity bills as part of healthcare.

What percent of GDP is government spending? When debating the value of government spending, Elon Musk was fact-checked by his own platform and informed that "government was only 11.3% of GDP." But this is using the value-added approach, which only counts direct government employees. In the value-added approach, the money the government spends on everything from constructing buildings to software licenses does not count. Meanwhile, the expenditures approach counts the government as 17% of GDP. This still excludes interest on debt and transfers like social security, but it includes money the government spends on grants or contracts with private businesses, such as SpaceX.

If you want to know GDP by city or region, the BEA uses the income approach. This is a practical decision. The other approaches are too difficult to tie to specific locations, but income tax data makes it easy to tie income to addresses. Since regional GDP is just measuring income, it does not really tell you if that region's GDP is a result of actual useful market production, or if it is from monopoly profits, rent-seeking, and government deals.

If you want to see what percent of the economy is manufacturing, and how that has changed over time, you can only use the value-added approach. Only the value-added approach separates out each step in the economic chain: from mining the iron ore to transporting it to the factory to manufacturing the product to selling it at the store. The value-added approach categorizes each step, so you can sum together just the increase in price from the manufacturing step across all categories of spending.

There is a second major complexity: for each of these approaches, we have a *real* and a *nominal* version. The nominal version is just based on adding up dollar sales or dollar income. It does not even pretend to be a measure of product or production. To get a measure that purports to measure changes in production over time, economists created a statistic called "real" GDP.

Since the problem with nominal GDP is that year-to-year changes are simply a result of money supply changes, not production, the art of real GDP is to factor out price changes. This is done by using price indexes. These price indexes are a joint effort between the U.S. Department of Labor's Bureau of Labor Statistics (BLS) and the BEA. The BLS actually sends out thousands of agents to look at prices in stores, browse websites, and survey producers. A price index for, say, "cereals and bakery products" is thus created by selecting a representative sample of what consumers actually buy, and then for each item in the sample tracking changes in price and then calculating a weighted average of these changes.

For the expenditures approach, real GDP is calculated by taking every single category of spending, taking the raw amount spent, and then adjusting it by a price index that most closely matches the category of spending. For the value-added approach, real GDP is calculated first by calculating nominal gross output—basically sales—for each establishment in the economy. Then, this nominal gross output is adjusted by a price index that most closely matches what each establishment was actually producing to get real gross output. Then, they find the nominal spending on intermediate inputs—basically anything purchased from a third party, including electricity, components, materials, TV advertising, software licenses, and so on. Then, these inputs must also be adjusted by a matching price index to get real intermediate inputs. Finally, you subtract the inputs from gross output to get *real value-added*.

A Fictitious Measure of Output With No Meaning

While the BEA is in charge of the value-added GDP calculation, they do not collect the data needed to make the calculations. The value-added approach depends on the U.S. Census Bureau and their once-every-five-year economic survey. The Census Bureau surveys all manufacturing establishments to get the total dollar shipments from their factories. The survey also asks the establishments to report the cost of their inputs. This includes imported components, electricity, software subscriptions, materials, basically everything they spend money on except employees.

The sales receipts minus the inputs is the raw “value added.” While all businesses are required to fill out the surveys, their data might not fit for the format of the survey and they are allowed to estimate. Uh oh—as a result, up to 40% of the data ends up being *imputed* and, due to confidentiality promises, neither the general public nor even the BEA can double-check the data or calculations.

The BEA gets this information from the Census Bureau and then combines it with price indexes from the BLS. The trick then is to create a composite price index for the factory output that roughly matches what the factory actually shipped and a composite price index for the inputs that roughly matches the factory's inputs. So if the factory's input spending was 20% on electricity, 50% on iron ore, and 30% on imported machine tools, the BEA might take the BLS electricity price index and weight that 20%, take an iron ore index and weight that 20%, and take a machinery tool index and weight that 30%. At best, this is going to be a rough approximation, because the price index is not going to track the prices of the specific machinery that this factory actually used, or the specific electricity bills of this particular factory. Rather it's all going to be averaged together and it is hoped that it will roughly match.

On the surface, this all seems very smart and very sound. There are many products assembled in the U.S. out of foreign components and the value-added approach avoids the mistake of counting the entire value of the final product toward U.S. manufacturing. Only the final sales price in excess of the input costs counts. By adjusting both the inputs and the outputs for inflation, the GDP calculation avoids counting mere price changes as being changes in output. And by using market prices of goods, and not just tonnage or raw numbers, we can avoid Soviet-style problems of just counting raw cars, even if those cars are unsellable garbage cars made for a quota and that instantly break down.

Unfortunately, real value-added is not a sound number. In fact, it is totally broken and nobody should be citing it.

To illustrate the problem, we need to dig into the GDP-by-industry data tables. We'll be comparing the actual number of things produced, the GDP real gross output, and finally real value-added, between the years 1997 and 2023. Gross output is a calculation adding up all sales from establishments in the industry but without subtracting the costs of the inputs. What we see are bizarre results where "real value-added" is implausibly up even when the production of actual things, or the actual inflation-adjusted value of the things coming out of the factory, is flat or down:

- Motor vehicles, bodies and trailers, and parts: actual light vehicles produced down 11%, real gross output of vehicles up 39%, real value-added up 125%. Inputs as a percent of gross output rose from 74% to

77%.

- Semiconductors: actual computer CPUs shipped is up 94%, gross output of semiconductors up 262%, real value-added up 1698%. Inputs as a percent of gross output fell from 52% to 25%.
- Steel mills & manufacturing from purchased steel: in raw tonnage, steel shipped is down 18%, real gross output is up 5%, real value-added is up 125%. Inputs as a percentage of gross output rose from 73% to 74%.

Actual steel rolling out of the mills is down, the inflation-adjusted value of the steel and steel products rolling out of the mills is flat, inputs as a percentage of output is the same as ever—yet value-added is up 125%!

I challenge anyone who believes in these statistics to tell me what in the real world happened so that raw tonnage of steel was down, real gross output of steel was flat, usage of inputs was up, but “real value-added” was also up, and up hugely. Nobody can explain these numbers. The BEA cannot—I have asked them! If the raw data still exists, nobody has access to it because it was confidential.

The basic problem is that real value-added calculations only work if there are no quality adjustments and there hasn’t been any substitutions in the inputs. If those assumptions do not hold, you can get wild and nonsensical results. Since those assumptions do not actually hold in the real world, those nonsensical results are mixed into the overall calculation in ways that are impossible to account for, thus making the entire number bogus.

My guess is that what happened with steel production is that factories have moved from using raw iron ore to scrap metal as an input. The scrap metal is actually closer to a final good and requires much less energy to turn into steel. But GDP calculations do not know that scrap metal is closer to a final good. What the GDP calculations see is that materials have become more expensive and that energy inputs are less, so it seems like the steel factories are maintaining output with much less input, and thus value-added is greater. The reality, though, is that the United States is not producing any more steel out of factories, the United States is not producing a greater percentage of the steel value chain than in 1997, and the 125% increase in real value-added is a spurious result that represents neither making more stuff nor making better stuff.

Another “quirk” of real value-added is that inflation adjustments and quality adjustments get applied retroactively, which creates wild inflections from small changes. In simplified terms, let’s say that, in 1997, car sales were $100 billion, and were still $100 billion twenty years later in 2017, with no changes due to inflation or input costs. Input costs in both years were $75 billion, meaning $25 billion in value-added in both years. The only thing that changed, let’s say, was that the “quality” of cars got 10% higher thanks to software innovations like Apple CarPlay and design improvements like crumple zones for safety—neither of which add to recurring production input costs. So, let’s say, our economists would adjust the 2017 figure to be $110 billion in “real” terms and show a small 10% increase, right?

Instead, the way it works is that a recent “base year” is taken, in this case 2017, and the base year is never adjusted. So rather than adjusting from $100 billion to $110 billion, the “real” output of 1997 is retroactively adjusted to be *lower*, in this case $91 billion, to get the same 10% increase. But then, our value-added in 1997 has fallen to $16 billion, and the *increase* in “real value-added manufacturing” has jumped from 10% to around 50%! We have created a 50% increase in car manufacturing not by actually producing 50% more cars or “objectively” making cars 50% better, but just by playing around with statistics and definitions.

The effect becomes even more extreme as the quality adjustment gets higher and makes the original value-added shrink to zero or negative. If the quality adjustment is 32%, the value-added increase becomes 652%! And after that it goes infinite and then becomes undefined. Of course, there are further complications. If the inputs are quality-adjusted in the same way as the outputs, the effect might be less, but this probably won’t happen because the methodology is quite different. This is all to demonstrate that value-added is *not* a measure of how much stuff the United States makes. It is a number that produces wild results and thus should not be mixed into aggregate statistics.

I do not know if these scenarios described above are the actual reasons for why value-added is so greatly outpacing gross output. No one else knows either. The examples I just explained are just two of the many possible ways that value-added can give bizarre or counter-intuitive results.

This is not just my critique: a former deputy chief at the BEA, Professor Doug Meade, has sharply criticized real value-added as a metric. In a 2010 conference paper, he wrote, “more than 60 years after it was first introduced, there is still no fundamental agreement on the meaning of real value added, or its price. Most who use it for the study of productivity loosely describe it as a measure of ‘real output’ although strictly speaking it is not that.” He continues to argue that comparing real value-added between years only works under the conditions of no quality adjustments, no input substitutions due to price changes, and no changes to the terms of trade. If those conditions do not hold, then, he says diplomatically, “it would be unclear what [real value-added] is measuring” Or as economist Thomas Rymes, observing the same issues, put it more directly: “a fictitious measure of output with no meaning.”

The Gloomy Reality of U.S. Manufacturing Output

Instead of real value-added, some commentators like to cite *nominal* value-added to show that the U.S. is not falling behind other countries in manufacturing. Nominal value-added cannot be used to track current U.S. manufacturing output versus historical U.S. output, but it is used to compare current U.S. manufacturing output to that of other countries. The nominal value-added tables show that American manufacturing value-added is greater than Germany, Japan, and South Korea combined and still double per-capita that of China. This number is calculated using the GDP value-added approach of adding up manufacturing sales at each step while subtracting out inputs, including imported components. Then, to compare countries, each country’s sum is converted to dollars using current exchange rates. Since nominal value-added is not adjusted by price indexes, it avoids all the problems we discussed with real value-added.

But, once again, the problem with the nominal value-added comparison is that it is not a comparison of actual things—it is a comparison of sales receipts. Thus a given quantity of products that is produced by a bloated cost structure will count as more "GDP" than the same number of products produced by an efficient factory. This is not just a theoretical problem—we know for a fact that the Chinese company BYD produces an equivalent to the Tesla Model 3 for half the price. Thus, $30,000 of manufacturing value-added in the U.S. might represent one car being produced, while for China it might represent two cars, and thus is actually double the output. In general, the China-U.S. dollar exchange rate is not a market rate and thus the conversion does not reflect in any meaningful sense the value of products.

Worse, many U.S. products are more expensive not because they are higher-end and better quality, but because they are protected from competition by tariffs, patents, regulation or national security requirements. For instance, Purism makes an all-in-the-USA phone for $2,000—the phone is no better than a $500 Chinese or South Korean phone, but sells at a premium for the U.S. security market. Others in procurement tell stories of getting quotes for printed circuit boards that cost $5,000 from China but $50,000 in America, thus only government and regulated industries buy American circuit boards. American-made municipal buses can cost three times the price as those made in China, but cities often face rules requiring them to buy American. For a particularly egregious example, thanks to the protections of the Jones Act, American ships cost an astounding *ten times* as much to build as their foreign counterparts.

If Chinese "nominal value-added" is twice that of the U.S. in dollars, but each dollar of sales receipts in China represents two to four times as many cars, chips, buses, and computers, then its actual output may be *four to eight times as great* as American output.

To the extent a country becomes so efficient at manufacturing some high-tech product that the price for it becomes very low, then that product actually counts for little in the nominal value-added numbers, even if the country is producing massive quantities of it. Which is more "output"—one million drones sold for a total of $2 billion dollars, versus one B2 stealth bomber for the same price? A $2,000 custom-made dress for the Met Gala, or one hundred pairs of denim work pants? Nominal value-added comparisons treat them as equivalent.

Nominal value-added cannot tell the difference between a country like 1790s Spain, a manufacturer of luxury goods with inflated nominal prices thanks to New World gold, and 1790s England, a ruthless manufacturer of inexpensive goods that is on its way to world domination. A comparison between countries that simply looks at sales revenues—not at the actual amount of ships, phones, and things produced for that revenue—is simply not a useful comparison.

Besides value-added, there are three other numbers sometimes used to track U.S. manufacturing output. Ultimately, these numbers are also inadequate, but to the extent they tell us anything, their story is more gloomy. The first statistic is the BEA's Manufacturing *real gross output* data, the second is the BLS's *manufacturing real sectoral output index*, and the third is the Federal Reserve's index of industrial production for manufacturing.

When we look at these manufacturing output numbers, the glass half-full view is that "manufacturing output index is up from 1997-2023 and only 7% off its all-time high." The pessimistic spin is that per-capita output is down 16% from 2007 and down 10% from 1997. It seems more reasonable to judge our standard of health as making stuff for the American population, and by that metric manufacturing has been getting worse for as long as this dataset goes back. For the Federal Reserve index, per-capita output is down 16% from 2007 and 2% from 1997.

But actually "down per-capita since 1997" is not the most pessimistic take—both *gross output* and the BLS *manufacturing index* have a crucial problem in that they do not subtract out components imported from abroad. When Dell sells a $5,000 server assembled in North Carolina with a Korean hard-drive, a Chinese motherboard, and a Taiwanese CPU, that server counts for $5,000 of American gross output. A car assembled in America with 60% foreign parts counts 100% for American output. Thus the actual hollowing-out of manufacturing may be even worse than these gross numbers suggest.

Also, all three of these numbers are *quality-adjusted.* "Vehicle" output up 40% does not mean America produced 10 million cars in 1997 and 14 million today. "Semiconductors" up 260% does not mean we make four times as many CPUs. A 40% increase in output means that when you adjust for the changes in quality of today's goods, we produce the equivalent in value of 14 million cars from 1997. Even if we actually produce 8 million vehicles, they are so much better, the "value" produced is much greater than that. The glaring subjectivity of these assessments is rarely pointed out, because it shatters the idea of having an objective, singular metric for economic output.

The Fake Objectivity of Quality Adjustments

While quality adjustments are often justified, they can mislead us in a number of ways. The first problem is honesty and clarity about what the data says. Economists and commentators usually refer to GDP as measuring "output." Colloquially, output and quality are two different things. When we read a headline saying GDP data shows "car output has increased," we think the U.S. has made more cars. We then apply our own views as to whether the quality of the car has changed. When we sneak quality into a measure but still call it "output," we are double-counting and embedding the subjective in the objective, and we lose track of the hard numbers. We are not making more quantity of cars per person like the data says, we are making *fewer* cars, but with Bluetooth and crumple zones.

Worse, almost nobody understands what actually goes into the quality adjustments. For instance, you will hear the claim, "of course the price index for cars should be quality adjusted, because a modern car lasts longer than the clunkers of the 1970s." And in fact, car longevity seems an obvious and objective place to make a quality adjustment—the agencies should adjust the price by the lifespan of the car. But the statistical agencies cannot actually measure how long a 2025 car is likely to last! Thus car longevity is not a quality adjustment that goes into GDP, for better or for worse.

If automakers move the location of a part so it does not burn out so quickly, that is not measured. If a new turbo engine is rife with problems causing early burnout, that is not measured. For automobiles, only quantifiable changes, such as the extra cost of adding a platinum-tipped spark plug, are included. While the BLS provides general information about the quality adjustment process, the specific methodology and the actual decisions are not documented. At the heart of GDP we find this subjective, bureaucratic black box. When we see that "output" of cars has increased since 1997, it is impossible for any commentator to know how that increase in "output" breaks down between actual number of cars, horsepower boosts, safety features, durability improvements, convenience features, blue tooth, power locks, and on and on.

In most cases, there is not actually a single, objective way to measure quality changes. In fact, there are multiple plausible options for quality adjustments, which allows for the final statistic to be virtually any number, from flat-lining to off-the-charts. Consider CPUs. One option would be to adjust by raw instructions-per-second. Since that has increased one-thousand-fold since 1997, if this metric was used to measure quality, then improvements in semiconductors alone would mean U.S. total manufacturing output had increased by 500%!

On the other hand, the time it takes for an ordinary office worker to enter data into a spreadsheet and read the news has barely changed at all. So perhaps there should be no quality adjustment—one CPU is one CPU. In that case, with the computers quality adjustment removed, total manufacturing output for the U.S. would be down decisively per-capita.

Perhaps we could measure CPU improvement by comparing the price of brand new models to existing models when the new model comes out, and then chaining those overlapping improvements together over time. A clever method, but the premium an early adopter pays for the newest CPU is not representative of how the average customer values it. The CPU may be only 10% better but command a 100% premium because they hope to sell a handful of processors to the few early adopters who pay for bragging rights. Or it may be 200% better, but only command a 25% increase in price because of competitive pressures.

We could imagine other methods: comparing how many pixels on a monitor the CPU can drive, or even putting gamers' heads in a scanner and comparing happiness playing a classic video game like *Quake* versus the latest *Overwatch*. How do you adjust for the fact that an average person can now listen to infinite amounts of music for free or almost-free with a cheap smartphone, compared to the millions of dollars in CD, cassette tape, and vinyl costs that would have required thirty years ago? Is it even worth trying to "adjust" for that?

The point is not that any of these methods is right or wrong. The point is that if you have a half-dozen plausible ways of adjusting for quality, none of which from first principles is more objective than another, and you rule out one method for giving ludicrously low results, and one method for ludicrously high results, and just choose a middle route that feels reasonable, then the result of this adjustment is not an objective measure of output. All you have done is launder *vibes* into something that has the appearance of an objective number.

What the BLS actually did in calculating the CPU price index was the price overlap method. But then, in 2010, the improvement in the price index slowed down, as new CPUs were no longer commanding a premium, perhaps because they were burning more heat and energy at a time when people were more energy conscious. Rather than say "well, CPUs are not actually getting better, they just have new trade-offs," the BLS changed their methodology to a custom "hedonic" method that measures all the attributes of the processor.

Again, it is not that the hedonic method is worse, the problem is that if you are changing the methodology after-the-fact because the numbers look wrong, you do not have an objective number, you are converting subjective intuition into something with the appearance of objectivity. The data is no longer telling us how the economy is doing, but instead the consensus feeling of the statistical agencies on how technology has changed is now driving the data.

American Decline Sector After Sector

In computers and semiconductors, the best thing we can say for U.S. performance is that we make almost twice as many CPUs as we did in 1997, and they are 1,000 times faster. But the gloomier take is that we lost the top position to Taiwan. American fabs now only make 12% of chips worldwide, down from 37% in 1990. All other computer components are even worse. The U.S. does not make LCD panels, circuit boards, motherboards, or hard-drives. America now annually imports $400 billion in electronics, compared to exports of $200 billion.

In automobiles, the raw number of vehicles manufactured in the U.S. has stagnated since the 1980s, and declined 39% per-capita since 1987. However, the vehicles produced are heftier, safer, faster, cleaner, and loaded with features. What this means for overall "output" is dependent on your point of view and what question you are trying to answer.

America suffered a catastrophic defeat in solar panels, going from nearly 50% market share in 2008 to practically zero in 2024, with China taking 90% of the market. China has 70% to 90% of the market in lithium batteries. And China dominates the drone industry so completely that they reportedly make 77% of consumer drones worldwide and 90% of those sold in the U.S. In shipbuilding, America is out of the game, going from building 5% of world shipping tonnage in the 1970s to a pathetic 0.2% today.

Boeing, America's last standing big commercial airplane manufacturer, went from averaging 530 planes in the late 1990s to about 440 planes in 2023 and 2024. That's a 33% per-capita decline in airplane production from 25 years ago. And the real decline may be even more, since modern Boeing planes are assembled from parts made around the world. In machine tools, America suffered a bloodbath since the 1970s, with shipments crashing from 75% of domestic use in 1980 to 35% in 2014. The top remaining American machine-tool company, Haas, is reportedly relying heavily on Chinese components.

In pharmaceuticals, it is very difficult to find raw numbers in terms of "doses manufactured of essential medicines." But in dollar terms, the U.S. has transitioned from being a slight net-exporter of pharmaceuticals in the mid-1990s, to a significant net importer, now importing $90 billion while exporting $54 billion. Additionally, the U.S. only manufactures 53% of the active pharmaceutical ingredients that are used in domestic production of medicine.

For other categories, it is more difficult to quantify in terms of raw stuff, but the BEA *real gross output* numbers—manufacturing sales receipts adjusted by the price changes of the goods sold—paint a grim picture. We see significant per-capita declines in: wood products (-9%), mineral products (-31%), machinery (-27%), appliances (-28%), furniture (-45%), textiles (-71%), miscellaneous durables (-41%), paper products (-45%), plastic and rubber products (-24%), printed books and materials (-54%), and chemical products (-10%). America has held steady in packaged food (-3%).

The lone bright spot is medical equipment and supplies, which is up 45%. As noted before, *real gross output* numbers can mislead if vertical integration reduces double-counting, but intermediate inputs as a percentage of gross outputs have not substantially changed since 1997. Finally, we must recognize the one super-high-tech area of manufacturing where America appears to be doing much better than 25 years ago: production of rockets and space-lift capacity. Here, SpaceX single-handedly seized the market from foreign competitors.

Putting it all together, U.S. manufacturing is hurting badly. It is decidedly not the case that American manufacturing is more productive than ever. It is an inexcusable failing that the U.S. does not have comprehensive numbers that track all essential output needed to maintain our quality of life, including not just the number of cars or microchips or refrigerators, but the percent of the components actually made in the USA, and the percent of the machine tools needed to build the components made in the USA. Until we have those numbers and they start going in the right direction, we need less boosterism and more grim resolve to reindustrialize.

Patrick Fitzsimmons is the former founding engineer at HubSpot and now splits his time between independent scholarship and working on Webchiver, an application for creating personal web archives. You can follow him at @pfitzsimmons.