wealth, which translated into dramatically lower demand and GDP. While the recession technically ended in June 2009, as we write this in 2013 the U.S. economy is still operating well below its potential, with unemployment at 7.6 percent and capacity utilization at 78 percent. During such a slump, any metric that includes output in the numerator, such as labor productivity, will often be at least temporarily depressed. In fact, when you look at history, you see that in the early years of the Great Depression, in the 1930s, productivity didn't just slow but actually fell for two years in a row—something it never did in the recent slump. Growth pessimists had even more company in the 1930s than they do today, but the following three decades proved to be the best ones of the twentieth century. Go back to figure 7.2 and look most closely at the dashed line charting the years following the dip in productivity in the early 1930s. You'll see the biggest wave of growth and bounty that the first machine age ever delivered.

The explanation for this productivity surge is in the lags that we always see when GPTs are installed. The benefits of electrification stretched for nearly a century as more and more complementary innovations were implemented. The digital GPTs of the second machine age are no less profound. Even if Moore's Law ground to a halt today, we could expect decades of complementary innovations to unfold and continue to boost productivity. However, unlike the steam engine or electricity, second machine age technologies continue to improve at a remarkably rapid exponential pace, replicating their power with digital perfection and creating even more opportunities for combinatorial innovation. The path won't be smooth—for one thing, we haven't banished the business cycle—but the fundamentals are in place for bounty that vastly exceeds anything we've ever seen before.

---

* The Rule of 70 (or, more precisely, the rule of 69.3 percent) is based on the following equation: $(1 + x)y = 2$ where x is the rate of growth and y is the number of years. Taking the natural logarithm of both sides gives $y \ln (1 + x) = \ln 2$. The ln (2) is 0.693 and for small x, ln (1 + x) is roughly equal to x, so the equation simplifies to $xy = 70$ percent.

* One can also measure capital productivity, which is output per unit of capital input; or multifactor productivity, which is output divided by a weighted average of both capital and labor inputs. Economists sometimes use another term for multifactor productivity, the "Solow Residual," which better reflects the fact that we don't necessarily know its origins. Robert Solow himself noted that it was less a concrete measure of technological progress than a "measure of our ignorance."

† That's a good thing, because there are natural limits to how much we can increase inputs, especially labor. They're subject to diminishing returns—no one is going to work more than twenty-four hours a day, or employ more than 100 percent of the labor force. In contrast, productivity growth reflects ability to innovate—it's limited only by our imaginations.

* Output divided by labor and physical capital inputs is often more ambitiously called 'total factor productivity.' However, that term can be a bit misleading, because there are other inputs to production. For instance, companies can make major investments in intangible organizational capital. The more kinds of inputs we are able to measure, the better we can account for overall output growth. As a result, the residual that we label "productivity" (not explained by

When President Hoover was trying to understand what was happening during the Great Depression and design a program to fight it, a comprehensive system of national accounts did not exist. He had to rely on scattered data like freight car loadings, commodity prices, and stock price indexes that gave only an incomplete and often unreliable view of economic activity. The first set of national accounts was presented to Congress in 1937 based on the pioneering work of Nobel Prize winner Simon Kuznets, who worked with researchers at the National Bureau of Economic Research and a team at the U.S. Department of Commerce. The resulting set of metrics have served as beacons that helped illuminate many of the dramatic changes that transformed the economy throughout the twentieth century.

But as the economy has changed so, too, must our metrics. More and more what we care about in the second machine age are ideas, not things—mind, not matter; bits, not atoms; and interactions, not transactions. The great irony of this information age is that, in many ways, we actually know less about the sources of value in the economy than we did fifty years ago. In fact, much of the change has been invisible for a long time simply because we did not know what to look for. There's a huge layer of the economy unseen in the official data and, for that matter, unaccounted for on the income statements and balance sheets of most companies. Free digital goods, the sharing economy, intangibles and changes in our relationships have already had big effects on our well-being. They also call for new organizational structures, new skills, new institutions, and perhaps even a reassessment of some of our values.

Music to Your Ears

The story of music's move from physical media to computer files has been told often and well, but one of that transition's most interesting aspects is less discussed. Music is hiding itself from our traditional economic statistics. Sales of music on physical media declined from 800 million units in 2004 to less than 400 million units in 2008. Yet over the same time period total units of music purchased still grew, reflecting an even faster increase in the purchases of digital downloads. Digital streams such as iTunes, Spotify, or Pandora also came to prominence, and, of course, the purchase data don't reflect the even larger number of songs that were shared, streamed, or downloaded for free, often via piracy. Before the rise of the MP3, even the most fanatical music fan, with a basement stacked high with LPs, tapes, and CDs, wouldn't have had a fraction of the twenty million songs available on a child's smartphone via services like Spotify or Rhapsody. What's more, clever research by Joel Waldfogel at the University of Minnesota finds quantitative evidence that the overall quality of music has not declined over the past decade and is, if anything, higher than ever.[1] If you're like most people, you are listening to more and better music than ever before.

So how did music disappear? The value of music has not changed, only the price. From 2004 to 2008, the combined revenue from sales of music dropped from $12.3 billion to $7.4 billion—that's a decline of 40 percent. Even when we include all digital sales, throwing in ringtones on mobile phones for good measure, the total revenues to the record companies are still down 30 percent.

Similar economics apply when you read the New York Times, Bloomberg Businessweek, or MIT Sloan Management Review online at a reduced price or for free instead of buying a physical copy at the newsstand, or when you use Craigslist instead of the classified ads, or when you share photos via Facebook instead of mailing prints around to friends and relatives. Analog dollars are becoming digital pennies.

By now, the number of pages of digital text and images on the Web is estimated to exceed one trillion.[2] As discussed in chapter 4, bits are created at virtually zero cost and transmitted almost instantaneously worldwide. What's more, a copy of a digital good is exactly identical to the original. This leads to some very different economics and some special measurement problems. When a business traveler calls home to talk to her children via Skype, that may add zero to GDP, but it's hardly worthless. Even the wealthiest robber baron would have been unable to buy this service. How do we measure the benefits of free goods or services that were unavailable at any price in previous eras?

What GDP Leaves Out

Despite all the attention it gets from economists, pundits, journalist, and politicians, GDP, even if it were perfectly measured, does not quantify our welfare. The trends in GDP growth and productivity growth covered in chapter 7 are important, but they are not sufficient measures of our overall well-being, or even our economic well-being. Robert Kennedy put this poetically in his quote at the beginning of this chapter.

While it would be unrealistic to put a dollar value on stirring oratory like RFK's, we can do a better job of understanding our basic economic progress by considering some of the changes in the goods and services that we are able to consume. It soon becomes clear that the trends in the official statistics not only underestimate our bounty, but in the second machine age they have also become increasingly misleading.

In addition to their vast library of music, children with smartphones today have access to more information in real time via the mobile web than the president of the United States had twenty years ago. Wikipedia alone claims to have over fifty times as much information as Encyclopaedia Britannica, the premier compilation of knowledge for most of the twentieth century.[3] Like Wikipedia but unlike Britannica, much of the information and entertainment available today is free, as are over one million apps on smartphones.[4]

Because they have zero price, these services are virtually invisible in the official statistics. They add value to the economy, but not dollars to GDP. And because our productivity data are, in turn, based on GDP metrics, the burgeoning availability of free goods does not move the productivity dial. There's little doubt, however, that they have real value. When a girl clicks on a YouTube video instead of going to the movies, she's saying that she gets more net value from YouTube than traditional cinema. When her brother downloads a free gaming app on his iPad instead of buying a new video game, he's making a similar statement.

Free: Good for Well-Being, Bad for GDP

In some ways, the proliferation of free products even pushes GDP downward. If the cost of creating and delivering an encyclopedia to your desktop is a few pennies instead of thousands of dollars, then you're certainly better off. But this decrease in costs lowers GDP even as our personal well-being increases, leaving GDP to travel in the opposite direction of our true well-being. A simple switch to using a free texting service like Apple's iChat instead of SMS, free classifieds like Craigslist instead of newspaper ads, or free calls like Skype instead of a traditional telephone service can make billions of dollars disappear from companies' revenues and the GDP statistics.[5]

As these examples show, our economic welfare is only loosely related to GDP. Unfortunately many economists, journalists, and much of the general public still use "GDP growth" as a synonym for "economic growth." For much of the twentieth century, this was a fair comparison. If one assumes that each additional unit of production created a similar increment in well-being, then counting up how many units were produced, as GDP does, would be a fine approximation of welfare. A nation that sells more cars, more bushels of wheat, and more tons of steel probably corresponds to a nation whose people are better off.

With a greater volume of digital goods introduced each year that do not have a dollar price, this traditional GDP heuristic is becoming less useful. As we discussed in chapter 4, the second machine age is often described as an "information economy," and with good reason. More people than ever are using Wikipedia, Facebook, Craigslist, Pandora, Hulu, and Google, with thousands of new digital goods introduced each year.

The U.S. Bureau of Economic Analysis defines the information sector's contribution to the economy as the sum of the sales of software, publishing, motion pictures, sound recording, broadcasting, telecommunications, and information and data processing services. According to the official measures, these account for just 4 percent of our GDP today, almost precisely the same share of GDP as in the late 1980s, before the World Wide Web was even invented. But clearly this isn't right. The official statistics are missing a growing share of the real value created in our economy.

Measuring Growth with a Time Machine: Would You Rather . . . ?

Can we improve on GDP as a measure of well-being? Economists sometimes use an alternate approach that resembles the children's game "Would you rather . . . ?" The 1912 Sears shopping catalog had thousands of items for sale, from a "Sears Motor Car" for $335 (page 1,213) to dozens of pairs of women's shoes, some available for as little as $1.50 (pages 371–79). Suppose I gave you an expanded version of this catalog that listed all the goods and services available in 1912, not just from Sears, but from any seller in the economy of 1912, and all the same prices as 1912.[6] Would you rather shop exclusively in that old catalog, with no other choices, or would you rather pay today's prices for a full selection of today's goods and services?

Or to make the comparison less difficult, pick two more recent catalogs, like 1993 versus 2013. If you had fifty thousand dollars to spend, would you rather be able to buy any 1993-model car (it would be brand-new) and pay 1993 prices, or a 2013 car and pay 2013 prices? Would you rather be able to buy the bananas, contact lenses, chicken wings, shirts, chairs, banking services, airline tickets, movies, telephone service, health care, housing services, light bulbs, computers, gasoline, and other goods and services that were available in 1993 at 1993 prices? Or would you rather buy the equivalent 2013 basket of services at 2013 prices?

Bananas or a gallon of gasoline have not really changed qualitatively since 1993, so the only difference to consider is their price. If that were the only difference, inflation would be easy to calculate, and the "would you rather" comparison would be a lot easier, too. For other goods, though, especially second machine age goods like online information and mobile phone capabilities, there have been big changes in quality, so the real quality-adjusted price may have fallen even if the nominal sticker price has increased. What's more, there are a lot of new goods that didn't exist before, especially digital goods. There are also some older goods and services that have been discontinued or degraded. It's hard to find a good horsehide razor strop these days,[7] or a 1993 vintage personal computer, or a gas station where the attendants routinely wash your windshield for no charge, like they once did.

Once you pick which catalog you like better, the next step asks how much money I would have to pay you to make you indifferent between the two catalogs. If I have to pay you 20 percent more to make you just as happy shopping from the new catalog as you would be shopping from the old catalog, then the overall price index has increased by 20 percent. And if your income has not changed, then that erosion of purchasing power translates to an equivalent fall in your standard of living. Similarly, if your income increases faster than the price index, then your standard of living is increasing.

This approach makes sense conceptually, and it's the basis for the way most modern governments calculate changes in the standard of living. For instance, the cost of living

adjustments used to index Social Security payments are based on this kind of analysis.[8] But the data used for these calculations are almost always drawn, understandably, from market transactions where money changes hands. The free economy is not factored in.

Consumer Surplus: How Much Would You Pay If You Had To?

An alternative approach measures the consumer surplus generated by goods and services. Consumer surplus compares the amount a consumer would have been willing to pay for something to the amount they actually have to pay. If you would happily pay one dollar to read the morning newspaper but instead you get it for free, then you've just gained one dollar of consumer surplus. However, as noted above, replacing a paid newspaper with an equivalent free new service would decrease GDP even though it increased consumer surplus.[9] In this case, consumer surplus would be a better measure of our economic well-being. Yet as appealing as consumer surplus is as a concept, it is also extremely difficult to measure.

The difficulty in measuring the consumer surplus, however, has not stopped a number of researchers from trying to eke out some estimates. In 1993, Erik wrote a paper calculating that the rapidly growing consumer surplus from price declines in computers increased economic welfare by about $50 billion each year.*[10]

Of course, when the product being studied is already free, looking at price declines doesn't work. Recent research that Erik did with Joo Hee Oh, a postdoctoral student at MIT, took a different approach. They started with the observation that even when people don't pay with money, they still give up something valuable whenever they use their Internet: their time.[11] No matter how rich or poor we are, each of us gets twenty-four hours in a day. In order to consume YouTube, Facebook, or e-mail, we must 'pay' attention. In fact, Americans nearly doubled the amount of leisure time they spent on Internet between 2000 and 2011. This implies that they valued it more than the other ways they could spend their time. By considering the value of users' time and comparing leisure time spent on the Internet to time spent in other ways, Erik and Joo Hee estimated that the Internet created about $2,600 of value per user each year. None of this showed up in the GDP statistics but if it had, GDP growth—and thus productivity growth—would have been about 0.3 percent higher each year. In other words, instead of the reported 1.2 percent productivity growth for 2012, it would have been 1.5 percent.

In contrast to leisure, where more time is a good thing, value at work is created by saving time. Hal Varian, the chief economist at Google, looked specifically at time savings gained from Google searches.[12] He and his team gathered a random sample of Google queries, such as: "In making cookies, does the use of butter or margarine affect the size of the cookie?" The team then did their best job to answer the questions without using Google—by looking answers up in the library, for instance. On average it took