

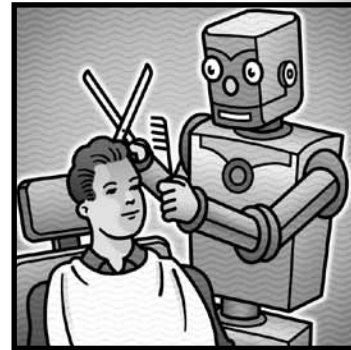
# The Second Machine Age

BY ERIK BRYNJOLFSSON  
AND ANDREW MCAFEE

On first look, *The Second Machine Age* might be mistaken for a slew of other books extolling technological change as the

cure for much of what ails the global economy and polity. And indeed, the authors, Erik Brynjolfsson and Andrew McAfee, are card-carrying techno-boosters who can paint rosy pictures of our digital future with the best of them. ¶ But Brynjolfsson, an economist at MIT's Sloan School of Management, and McAfee, a principal research scientist at MIT's Center for Digital Business, are all too aware that the technology juggernaut has a way of making roadkill of those who don't remain a step ahead. In this jargon-free treatise, they assess the danger that high-speed technical change will leave us with a growing, ill-paid and unemployed underclass. ¶ You've probably read scary stories about what used to be called automation. But dollars to DRAMs, I'll bet you haven't read such a clear-eyed assessment of the risks. Check out the excerpt here – and then buy the book to find out what Brynjolfsson and McAfee propose to do about this looming problem.

— Peter Passell



“The test of our progress is not whether we add more to the abundance of those who have much; it is whether we provide enough for those who have little.”

— Franklin D. Roosevelt

Consider the paradox: GDP has never been higher and innovation has never been faster, yet people are increasingly pessimistic about their children's future living standards. And no wonder; adjusted for inflation, the combined net worth on *Forbes's* billionaire list has more than quintupled since 2000, but the income of the median household in America has fallen.

The economic statistics underscore the dichotomy of what we call the bounty of technology [the increase in volume, variety and quality, and the decrease in the cost of many products and services] and the spread [the ever-greater differences in material success among households]. The economist Jared Bernstein, a senior fellow at the Center on Budget and Policy Priorities, brought our attention to the way productivity and employment have become decoupled. While these two key economic statistics tracked each other for most of the postwar period, they diverged in the late 1990s.

Productivity continued its upward path as employment sagged. Today, the employment-to-population ratio is lower than any time in at least 20 years, and the real income of the median worker is lower than in the 1990s. Meanwhile, like productivity, GDP, corporate investment and after-tax profits are at record highs.

In a place like Silicon Valley or a research university like MIT, the rapid pace of innovation is particularly easy to see. Startups flour-

ish, minting new millionaires and billionaires, while research labs churn out astonishing new technologies. At the same time, however, a growing number of people face financial hardships: students struggle with enormous debt, recent graduates have difficulty finding new jobs and millions have turned to borrowing to temporarily maintain their living standards.

Here, we'll address three important questions about the future of the bounty and the spread. First, will the bounty overwhelm the spread? Second, can technology not only increase inequality but also create structural unemployment? And third, what about globalization, the other great force transforming the economy? Could it explain recent declines in wages and employment?

#### **WHAT'S BIGGER, BOUNTY OR SPREAD?**

Thanks to technology, we are creating a more abundant world – one in which we get more and more output from less raw materials, capital and labor. In the years to come we will continue to benefit from things that are relatively easy to measure, such as higher productivity, and things that are less susceptible to metrics, such as the boost we get from free digital goods.

The previous paragraph describes our current bounty in the dry vocabulary of economics. This is a shame and needs to be corrected – a phenomenon so fundamental and wonderful deserves better language. “Bounty” doesn't simply mean more cheap consumer



goods and empty calories. It also means more choice, greater variety and higher quality in many areas of our lives. It means heart surgeries performed without cracking the sternum and opening the chest cavity. It means constant access to the world's best teachers combined with personalized self-assessments that let students know how well they're mastering the material. It means that households have to spend less of their total budget on groceries, cars, clothing and utilities. It means returning hearing to the deaf and, eventually, sight to the blind. It means less need to work doing boring, repetitive tasks and more opportunity for creative and interactive work.

The manifestations of progress are all based at least in part on digital technologies. When combined with political and economic

systems that offer people choices instead of locking them in, technological advance is an awe-inspiring engine of betterment and bounty. But it is also an engine driving spread, creating larger and larger differences in wealth, income, standards of living and opportunities for advancement. We wish that progress in digital technologies were a rising tide that lifted all boats equally in all seas, but it's not.

Technology is certainly not the only force causing this rise in spread, but it is the main one. Today's information technologies favor more-skilled over less-skilled workers, increase the returns to capital owners over labor, and increase the advantages that superstars have over everybody else. All of these trends increase spread – between those that have jobs and those that don't, between highly

skilled and educated workers and less advanced ones, between superstars and the rest of us. It's clear from everything we've seen and learned recently that, all else equal, future technologies will tend to increase spread, just as they will boost the bounty.

The fact that technology brings both bounty and spread leads to an important question: since there's so much bounty, should we be concerned about the spread? We might consider rising inequality less of a problem if people at the bottom are also seeing their lives improve thanks to technology.

Some observers advance what we will call the "strong bounty" argument, which essentially says that a focus on spread is inappropriate since bounty is the more important phenomenon and exists even at the bottom of the spread. This argument acknowledges that highly skilled workers are pulling away from the rest – and that superstars are pulling so far away as to be out of sight – but then essentially asks, "So what?"

If all people's living standards are getting better, why should we be concerned if some are getting a lot better?" As Harvard economist Greg Mankiw has argued, the enormous income earned by the 1 percent is not necessarily a problem if it reflects the just deserts of people who are creating value for everyone else.

Capitalist economic systems work in part because they provide strong incentives to innovators: if your offering succeeds in the marketplace, you'll reap at least some of the financial rewards. And if your offering succeeds like crazy, the rewards can be huge. When these incentives are working well (and not doing things like providing risk-free rewards to people taking inappropriate risks within the financial system), the benefits can be both large and broad. Everyone benefits, even though not all benefits are distributed equally. As former Treasury Secretary Larry Summers put it, "We do need to recognize that a component of this inequality is the other side of successful entrepreneurship."

We particularly want to encourage entrepreneurship because technological progress typically helps even the poorest people. Innovations like mobile telephones, for example, are improving incomes, health and other measures of well-being in developing countries. As Moore's Law – the rule of thumb that data density in integrated circuits doubles approximately every 18 months – continues to drive down the cost and increase the capability of digital devices, the benefits they bring will continue to add up.

If the strong bounty argument is correct, we have nothing significant to worry about as



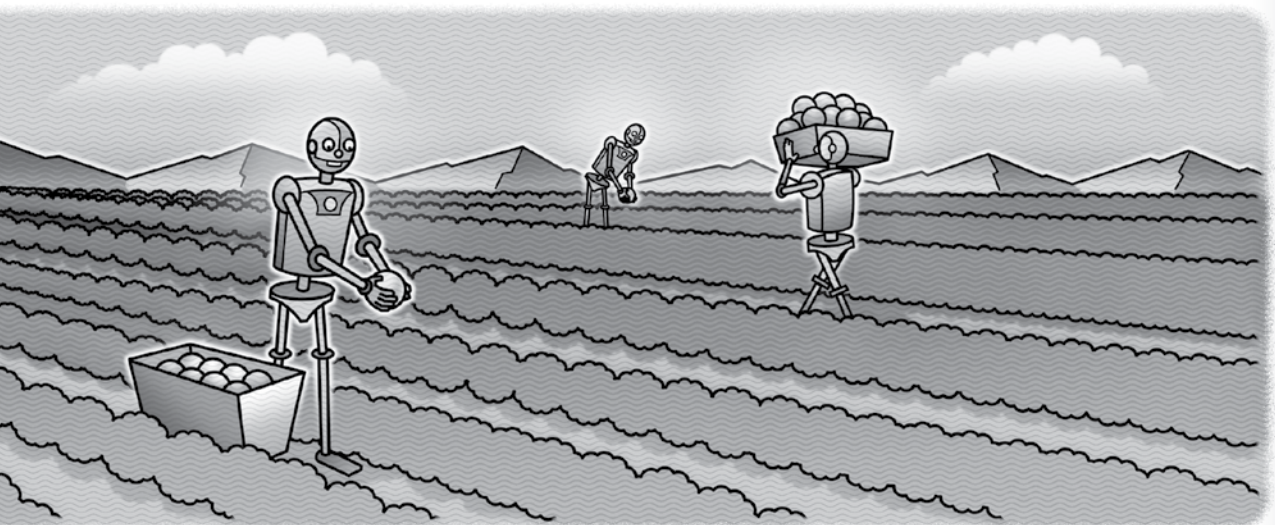
we head deeper into the second machine age. But it isn't. Many people are losing ground, not just relative to others but in absolute terms as well. In America, the income of the median worker is lower in real dollars than it was in 1999. And the story largely repeats itself when we look at households instead of individual workers, or total wealth instead of income.

Some proponents of the strong bounty argument believe that while these declines are real, they're still less important than the unmeasured price decreases, quality improvements and other benefits that we've been experiencing. Economists Donald Boudreaux (George Mason University) and Mark Perry (University of Michigan-Flint) write that:

Spending by households on many of modern life's "basics" – food at home, automobiles, clothing and footwear, household furnishings and equipment, and housing and utilities – fell from 53% of disposable income in 1950 to 44% in 1970 to 32% today ... [and] the quantities and qualities of what ordinary Americans consume are closer to that of rich Americans than they were in decades past. Consider the electronic products that every middle-class teenager can now afford – iPhones, iPads, iPods and laptop computers. They aren't much inferior to the electronic gadgets now used by the top 1% of American income earners, and often they are exactly the same.

These are intriguing arguments. We particularly like the insight that the average worker today is better off in important ways than his or her counterpart in earlier generations precisely because of the bounty brought by innovation and technology. For anything related to information, media, communication and computation, the improvements are so large that they can hardly be believed in retrospect or anticipated in advance. And the bounty doesn't stop there: technological progress also leads to cost reductions and quality improvements in other areas, such as food and power, that may not seem high tech on the surface but actually are when you look under the hood.

Nonetheless, we are not convinced that people at the lower ranges of the spread are doing OK. For one thing, some critical items that they (and everyone else) would like to purchase are getting much more expensive. This phenomenon is well summarized in research by Jared Bernstein, who compared increases in median family income between 1990 and 2008 with changes in the cost of housing, health care and college. He found that while family income grew by around 20 percent during that time, prices for housing and college grew by about 50 percent, and health care by more than 150 percent. Since



Americans' median incomes in real terms have been falling in the years since 2008, these comparisons would be even more unfavorable if extended.

However American households are spending their money, many of them are left without a financial cushion. The economists Annamaria Lusardi (George Washington University), Daniel Schneider (Princeton) and Peter Tufano (Oxford) conducted a study in 2011 asking people about "their capacity to come up with \$2,000 in 30 days. And their findings are troubling. They concluded that:

Approximately one-quarter of Americans report that they would certainly not be able to come up with such funds, and an additional 19 percent would do so by relying at least in part on pawning or selling possessions or taking payday loans.... [A] sizable fraction of seemingly 'middle class' Americans ... judge themselves to be financially fragile.

Other data – about poverty rates, access to health care, the number of people who want full-time jobs but can only find part-time work, and so on – confirm that while the economic bounty from technology is real, it is not sufficient to compensate for huge increases in spread. And those increases are not purely a consequence of the Great Recession, nor a recent or transient phenomenon.

That many Americans face stagnant or falling incomes is bad enough, but it is now combined with decreasing social mobility – an ever lower chance that children born at the bottom end of the spread will escape their circumstances and move upward throughout their lives and careers. Recent research makes it clear that the American dream of upward mobility, which was real in earlier generations, is greatly diminished today. To take just one example, a 2013 study of U.S. tax returns from 1987 to 2009 conducted by economist Jason DeBacker and colleagues found that the 35,000 households in their sample tended to

stay in roughly the same order of richest to poorest year after year, with little reshuffling, even as the differences in household income grew over time.

More recently, the sociologist Robert Putnam has illustrated how for Americans in cities like Port Clinton, Ohio (his hometown), economic conditions and prospects have worsened in recent decades for the children of parents with only high school educations, even as they've improved for college-educated families. This is exactly what we'd expect to see as skill-biased technical change accelerates.

Many Americans believe that they still live in the land of opportunity – the country that offers the greatest chance of economic advancement. But this is no longer the case. As *The Economist* sums up:

Back in its Horatio Alger days, America was more fluid than Europe. Now it is not. Using one-generation measures of social mobility – how much a father's relative income influences that of his adult son – America does half as well as Nordic countries, and about the same as Britain and Italy, Europe's least-mobile places.

So the spread seems to be not only large, but also self-perpetuating. Too often, people at the bottom and middle stay where they are over their careers, and families stay locked-in across generations. This is not healthy for an economy or society.

It would be even healthier if the spread were to diminish the bounty – if inequality and its consequences somehow impeded technological progress, keeping us from enjoying all the potential benefits of the new machine age. Although a common argument is that high levels of inequality can motivate people to work harder, boosting overall economic growth, inequality can also dampen growth.

In their book *Why Nations Fail*, economist Daron Acemoglu and political scientist James

Robinson aimed at uncovering, as the book's subtitle puts it, "the origins of power, prosperity, and poverty." According to Acemoglu and Robinson, the true origins are not geography, natural resources or culture; they're institutions like democracy, property rights and the rule of law (or lack thereof). When they turn their attention to America's current condition, they offer important cautions:

The U.S. generated so much innovation and economic growth for the last two hundred years because, by and large, it rewarded innovation and investment. This did not happen in a vacuum; it was supported by a particular set of political arrangements – inclusive political institutions – which prevented an elite or another narrow group from monopolizing political power and using it for their own benefit and at the expense of society.

So here is the concern: economic inequality will lead to greater political inequality, and those who are further empowered politically will use this to gain greater economic advantage, stacking the cards in their favor and increasing economic inequality still further – a quintessential vicious circle. And we may be in the midst of it.

Their analysis hits on a final reason to worry about the large and growing inequality of recent years: it could lead to the creation of "extractive" institutions that slow our journey into the second machine age. We think this would be something more than a shame; it would be closer to a tragedy.

#### **TECHNOLOGICAL UNEMPLOYMENT**

We've seen that the overall pie of the economy is growing, but some people, even a majority of them, can be made worse off by advances in technology. As demand falls for labor, particularly relatively unskilled labor, wages fall. But can technology actually lead to unemployment?

We're not the first people to ask this question. In fact, it has been debated vigorously,

even violently, for at least 200 years. Between 1811 and 1817, a group of English textile workers whose jobs were threatened by the automated looms of the first Industrial Revolution rallied around a perhaps mythical, Robin Hood-like figure named Ned Ludd, attacking mills and machinery before being suppressed by the British government.

Economists and other scholars saw in the Luddite movement an early example of a broad and important new pattern: large-scale automation entering the workplace and affecting wage and employment prospects. Researchers soon fell into two camps. The first and largest argued that while technological progress and other factors definitely cause some workers to lose their jobs, the fundamentally creative nature of capitalism creates other, usually better, opportunities for them. Unemployment, therefore, is only temporary and not a serious problem.

John Bates Clark (after whom the medal for the best economist under the age of 40 is named) wrote in 1915 that:

In the actual [economy], which is highly dynamic, such a supply of unemployed labor is always at hand, and it is neither possible [nor] normal that it should be altogether absent. The well-being of workers requires that progress should go on, and it cannot do so without causing temporary displacement of laborers.

The following year, the political scientist William Leiserson took this argument further. He described unemployment as something close to a mirage: "the army of the unemployed is no more unemployed than are firemen who wait in firehouses for the alarm to sound, or the reserve police force ready to meet the next call." The creative forces of capitalism, in short, required a supply of ready labor, which came from people displaced by previous instances of technological progress.

John Maynard Keynes was less confident that things would always work out so well for workers. His 1930 essay “Economic Possibilities for our Grandchildren,” while mostly optimistic, nicely articulated the position of the second camp – that automation could, in fact, put people out of work permanently, especially if more and more processes were automated. His essay looked past the immediate hard times of the Great Depression and offered a prediction:

We are being afflicted with a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come – namely, technological unemployment. This means unemployment due to our discovery of means of economizing the use of labor outrunning the pace at which we can find new uses for labor.

The extended joblessness of the Great Depression seemed to confirm Keynes, but it eventually eased. Then came World War II and its insatiable demands for labor, both on the battlefield and the home front, and the threat of technological unemployment receded.

After the war, the debate about technology’s impact on the labor force resumed, and took on new life once computers appeared. A commission of scientists and social theorists sent an open letter to President Lyndon Johnson in 1964 arguing that:

A new era of production has begun. Its principles of organization are as different from those of the industrial era as those of the industrial era were different from the agricultural. The cybernation revolution has been brought about by the combination of the computer and the automated self-regulating machine. This results in a system of almost unlimited productive capacity which requires progressively less human labor.

The Nobel-winning economist Wassily Leontief agreed, writing in 1983 that “the role of humans as the most important factor of production is bound to diminish in the same way

that the role of horses in agricultural production was first diminished and then eliminated by the introduction of tractors.”

Just four years later, however, a panel of economists assembled by the National Academy of Sciences disagreed with Leontief and made a clear, comprehensive and optimistic statement in their report “Technology and Employment”:

By reducing the costs of production and thereby lowering the price of a particular good in a competitive market, technological change frequently leads to increases in output demand: greater output demand results in increased production, which requires more labor, offsetting the employment effects of reductions in labor requirements per unit of output stemming from technological change.... Historically and, we believe, for the foreseeable future, reductions in labor requirements per unit of output resulting from new process technologies have been and will continue to be outweighed by the beneficial employment effects of the expansion in total output that generally occurs.

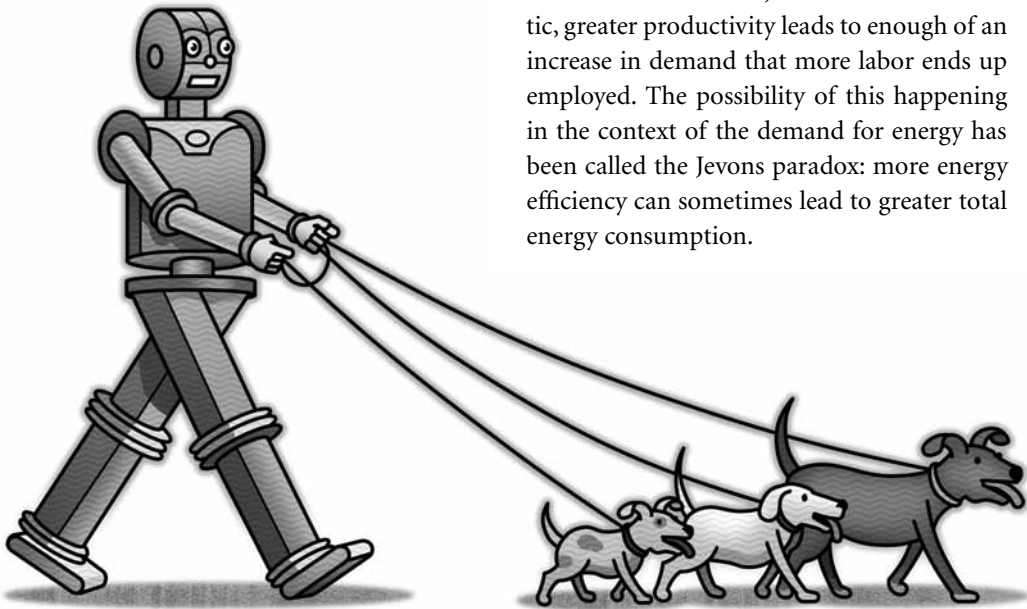
This view – that automation and other forms of technological progress in aggregate create more jobs than they destroy – has come to dominate the discipline of economics. To believe otherwise is to succumb to the “Luddite fallacy.” So in recent years, most of the people arguing that technology is a net job destroyer have not been mainstream economists.

The argument that technology cannot create ongoing structural unemployment, rather than just temporary spells of joblessness during recessions, rests on two pillars: theory and 200 years of historical evidence. But both are less solid than they initially appear.

First, the theory. Three economic mechanisms are candidates for explaining technological unemployment: inelastic demand, rapid change and severe inequality.

If technology leads to more efficient use of labor, then as the economists on the National

Academy of Sciences panel pointed out, technological change does not automatically lead to reduced demand for labor. Lower costs may lead to lower prices for goods, and in turn, lower prices lead to greater demand for the goods, which can ultimately lead to an increase in demand for labor as well. Whether or not this will actually happen depends on the “elasticity of demand,” defined as the percentage increase in the quantity demanded for each percentage decline in price.



For some goods and services, such as automobile tires and household lighting, demand has been relatively inelastic and thus insensitive to price declines. Cutting the price of artificial light in half did not double the amount of illumination that consumers and businesses demanded, so the total revenues for the lighting industry have fallen as lighting became more efficient. In a great piece of historical sleuthing, economist William Nordhaus documented how technology has reduced the price of lumens by over a thousand-fold since the days of candles and whale oil lamps, al-

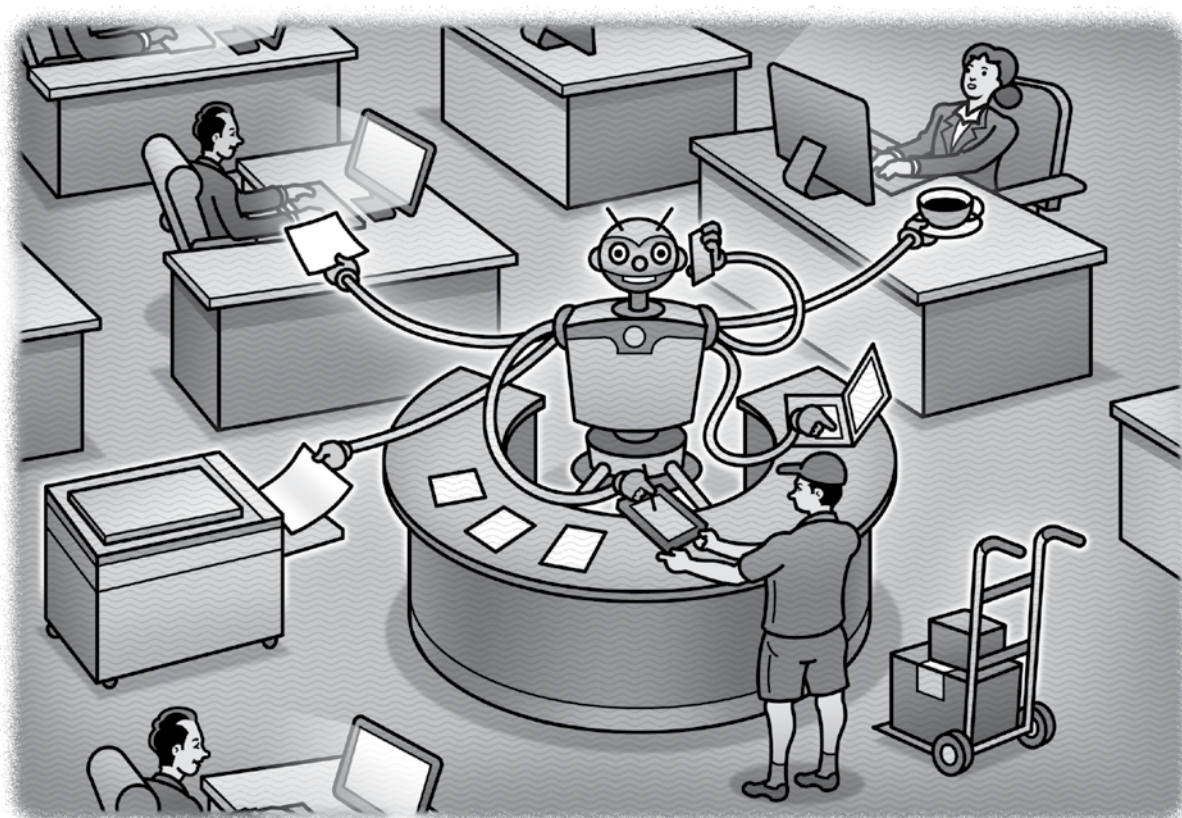
lowing us to expend far less on labor while getting all the light we need.

Whole sectors of the economy, not just product categories, can face relatively inelastic demand for labor. Over the years agriculture and manufacturing have each experienced falling employment as they became more efficient. The lower prices and improved quality of their outputs did not lead to enough increased demand to offset improvements in productivity.

On the other hand, when demand is elastic, greater productivity leads to enough of an increase in demand that more labor ends up employed. The possibility of this happening in the context of the demand for energy has been called the Jevons paradox: more energy efficiency can sometimes lead to greater total energy consumption.

But to economists there is no paradox, just an inevitable implication of elastic demand. This is especially common in new industries like information technology. If elasticity is exactly equal to one (i.e., a one percent decline in price leads to a one percent increase in quantity), then total revenues (price times quantity) will be unchanged. In other words, an increase in productivity will be exactly matched by an identical increase in demand to keep everyone just as busy as they were before.

And while an elasticity of exactly one might seem like a very special case, a good



(if not airtight) argument can be made that, in the long run, this is what happens in the overall economy. For instance, falling food prices might reduce demand for agricultural labor, but they free up just enough money to be spent elsewhere in the economy so that overall employment is maintained. The money is spent not just buying more of the existing goods, but on newly invented products and services. This is the core of the economic argument that technological unemployment is impossible.

Keynes disagreed. He thought that in the long run, demand would not be elastic. That is, ever-lower (quality-adjusted) prices would not necessarily mean we would consume ever-more goods and services. Instead, we would become satiated and choose to consume less. He predicted that this would lead to a dramatic reduction in working hours to as few as

15 per week, as less and less labor was needed to produce all the goods and services that people demanded.

However, it's hard to see this type of technological unemployment as an economic problem. After all, in that scenario people are working less because they are satiated. The "economic problem" of scarcity is replaced by the entirely more appealing problem of what to do with abundant wealth and copious leisure. As the futurist Arthur C. Clarke is purported to have put it, "The goal of the future is full unemployment, so we can play."

Keynes was more concerned with short-term "maladjustments," which brings us to the second, more serious argument for technological unemployment: the inability of our skills, organizations and institutions to keep pace with technical change. When technology

eliminates one type of job, or even the need for a whole category of skills, the affected workers will have to develop new skills and find new jobs. Of course, that can take time, and in the interim they may be unemployed. The optimistic argument maintains that this is temporary. Eventually, the economy will find a new equilibrium and full employment will be restored as entrepreneurs invent new businesses and the workforce adapts its human capital.

But what if this process takes a decade? And what if, by then, technology has changed again? This is the possibility that Wassily Leontief had in mind in a 1983 article in which he speculated that many workers could end up permanently unemployed, like horses unable to adjust to the invention of tractors. Once one concedes that it takes time for workers and organizations to adjust to technical change, then it becomes apparent that accelerating change can lead to widening gaps and increasing possibilities for technological unemployment. Faster technological progress may ultimately bring greater wealth and longer life spans, but it also requires faster adjustments by both people and institutions. With apologies to Keynes, in the long run we may not be dead – but we will still need jobs.

The third argument for technological unemployment may be the most troubling of all. It goes beyond “temporary” maladjustments. Recent advances in technology have created both winners and losers via skill-biased technical change, capital-biased technical change and the proliferation of superstars in winner-take-all markets. This has reduced the demand for some types of work and skills. In a free market, prices adjust to restore equilibrium between supply and demand, and indeed, real wages have fallen for millions in the United States.

In principle, the equilibrium wage could be one dollar an hour for some workers, even as other workers command a wage thousands

of times higher. Most people in advanced countries would not consider one dollar an hour a living wage, and don’t expect society to require people to work at that wage under threat of starvation.

What’s more, in extreme winner-take-all markets, the equilibrium wage might be zero: even if we offered to sing “Satisfaction” for free, people would still prefer to pay for the version sung by Mick Jagger. In the market for music, Mick can now, in effect, make digital copies of himself that compete with us.

A near-zero wage is not a living wage. Rational people would rather look for another gig, and look, and look and look, than depend on a near-zero wage for sustenance. Thus, there is a floor on how low wages for human labor can go. In turn, that floor can lead to unemployment: people who want to work, but are unable to find jobs. If neither the worker nor employers can think of a profitable task that requires that worker’s skills, that worker will go unemployed indefinitely.

Over history, this has happened to many other inputs to production that were once valuable, from whale oil to horse labor. They are no longer needed in today’s economy even at zero price. In other words, just as technology can create inequality, it can also create unemployment. And in theory, this can affect a large number of people, even a majority of the population, and even if the overall economic pie is growing.

So that’s the theory; what about the data? For most of the 200 years since the Luddite rebellion, technology has boosted productivity enormously. But the data show that employment grew alongside productivity up until the end of the 20th century. This demonstrates that productivity doesn’t always lead to job destruction. It’s even tempting to suppose that productivity somehow inevitably leads to job creation, as technology boosters sometimes

argue. However, the data also show that, more recently, job growth decoupled from productivity in the late 1990s.

Which history should we take guidance from: the two centuries ending in the late 1990s, or the 15 years since? We can't know for sure, but our reading of technology tells us that the power of exponential, digital and combinatorial forces, as well as the dawning of machine intelligence and networked intelligence, presage even greater disruptions.

### THE ANDROID EXPERIMENT

Imagine that tomorrow a company introduced androids [the robot sort, not the Google OS] that could do absolutely everything a human worker could do, including building more androids. There's an endless supply of these robots, and they're extremely cheap to buy and virtually free to run. They work all day, every day, without breaking down.

Clearly, the economic implications of such an advance would be profound. First, productivity and output would skyrocket. The androids would operate the farms and factories. Food and manufactures would become much cheaper to produce. In a competitive market, in fact, their prices would fall close to the cost of the raw materials. Around the world, we'd see an amazing increase in the volume, variety and affordability of offerings. The androids, in short, would bring great bounty.

They'd also bring severe dislocations to the labor force. Every economically rational employer would prefer androids, since compared to the status quo they would provide equal capability at lower cost. So they would very quickly replace most, if not all, human workers. Entrepreneurs would continue to develop novel products, create new markets and found companies, but they'd staff these companies with androids instead of people. The owners of the androids and other capital

assets or natural resources would capture all the value in the economy, and do all the consuming. Those with no assets would have only their labor to sell, and their labor would be worthless.

This thought experiment reflects the reality that there is no iron law that technological progress must always be accompanied by broad job creation.

One slight variation on the experiment imagines that the androids can do everything a human worker can do except for one skill – say, cooking. Because there would be so much competition for these jobs, however, companies that employed cooks could offer much lower wages and still fill their open positions. The total number of hours spent cooking in the economy would stay the same (at least as long as people kept eating in restaurants), but the total wages paid to cooks would go down.

The only exception might be superstar chefs with some combination of skill and reputation that could not be duplicated by other people. Superstars would still be able to command high wages; other cooks would not. So in addition to bringing great bounty of output, the androids would also greatly increase the spread in income.

How useful are these thought experiments, which sound more like science fiction than any current reality? Fully functional humanoid robots are not rumbling around at American companies today. And until recently, progress had been slow in making machines that could take the places of human workers in areas like pattern recognition, complex communication, sensing and mobility. But the pace of progress here has greatly accelerated in recent years.

The more readily machines can substitute for human workers, the more likely they'll drive down the wages of humans with similar skills. The lesson from economics and busi-

ness strategy is that you don't want to compete against close substitutes, especially if they have a cost advantage.

But in principle, machines can have very different strengths and weaknesses than humans. When engineers work to amplify these differences, building on the areas where machines are strong and humans are weak, the machines are more likely to complement humans rather than substitute for them. Effective production is more likely to require both human and machine inputs, and the value of the human inputs will grow, not shrink, as the power of machines increases.

A second lesson of economics and business strategy is that it's great to be a complement to something that's increasingly plentiful. Moreover, this approach is more likely to create opportunities to produce goods and services that could never have been created by unaugmented humans – or by machines that simply mimicked people, for that matter. These new goods and services provide a path for productivity growth based on increased output rather than reduced inputs.

Thus in a very real sense, as long as there are unmet needs and wants in the world, un-

employment is a loud warning that we simply aren't thinking hard enough about what needs doing. We aren't being creative enough about solving the problems we have in using the freed-up time and energy of the people whose old jobs were automated away. We can do more to invent technologies and business models that augment the unique capabilities of humans to create new sources of value, instead of automating the ones that already exist. This is the real challenge facing our policymakers, our entrepreneurs and each of us individually.

#### AN ALTERNATIVE EXPLANATION: GLOBALIZATION

Technology isn't the only factor transforming the economy. The other big force of our era is globalization. Could this be the reason that median wages have stagnated in the United States and other advanced economies? A number of thoughtful economists have made exactly that argument. The story is one of *factor price equalization*. This means that in any single market, competition will tend to bid the prices of the factors of production – such as labor or capital – to a single, common price. Over the past few decades, lower transportation and communication costs have helped create one big global market for many products and services.

Businesses can identify and hire workers with skills they need anywhere in the world. If a worker in China can do the same work as an American, then what economists call “the law of one price” demands that they earn essentially the same wages because the market will arbitrage away differences just as it would for other commodities. That's good news for the Chinese worker and for overall economic productivity. But is not good news for the American worker who now faces low-cost competition.



The factor-price equalization story yields a testable prediction: American manufacturers would be expected to shift production overseas, where costs are lower. And indeed, manufacturing employment in the United States has fallen over the past 20 years. Economists David Autor, David Dorn and Gordon Hanson estimate that competition from China can explain about a quarter of the decline in U.S. manufacturing employment.

However, when one looks more closely at the data, the globalization explanation becomes less compelling. Since 1996, manufacturing employment in China has actually fallen as well, coincidentally by an estimated 25 percent. That means 30 million fewer Chinese were employed in the sector, even as output soared by 70 percent. It's not that American workers are being replaced by Chinese workers. It's that both American and Chinese workers are being made more efficient by automation. As a result, both countries are producing more output with fewer workers.

In the long run, the biggest effect of automation likely won't be on workers in America and other developed nations, but on workers in developing nations that currently rely on low-cost labor for their competitive advantage. If you take most of the costs of labor out of the equation by installing robots and other types of automation, the competitive advantage of low wages largely disappears.

This is already beginning to happen. Terry Guo, the founder of Foxconn, the giant China-based manufacturer, has been aggressively installing hundreds of thousands of robots to replace human workers. He says he plans to buy millions more in the coming years. The first wave is going into factories in China and Taiwan, but once an industry becomes largely automated, the case for locating a factory in a low-wage country becomes less compelling.

There may still be logistical advantages if the local business ecosystem is strong, making it easier to get spare parts, supplies and custom components. But inertia may be overcome by the advantages of reducing transit times for finished products and being closer to customers, engineers and designers, educated workers or even regions where the rule of law is strong. This could bring manufacturing back to America.

A similar argument applies outside of manufacturing.

For instance, interactive voice-response systems are automating jobs in call centers. United Airlines, for example, has been successful in making the transition. This can disproportionately affect low-cost workers in places like India and the Philippines. Similarly, many medical doctors have had their dictation sent overseas to be transcribed. But an increasing number are now happy with computer transcription. In more and more domains, intelligent and flexible machines, not humans in other countries, are the most cost-effective source for "labor."

If you look at the types of tasks that have been offshored in the past 20 years, you see that they tend to be relatively routine, well-structured tasks. Interestingly, these are precisely the tasks that are easiest to automate. If you can give precise instructions to someone else on exactly what needs to be done, you can often write a precise computer program to do the same task. In other words, offshoring is often only a way station on the road to automation.

In the long run, low wages will be no match for Moore's Law. Trying to fend off advances in technology by cutting wages is only a temporary protection. It is no more sustainable than asking folk legend John Henry to lift weights to better compete with a steam-powered hammer. **M**