MARTIN FORD

RISE OF THE ROBOTS

TECHNOLOGY AND THE THREAT OF A JOBLESS FUTURE

RISE OF THE ROBOTS

ALSO BY Martin Ford:

The Lights in the Tunnel: Automation, Accelerating Technology and the Economy of the Future

RISE OF THE ROBOTS

TECHNOLOGY AND THE THREAT OF A JOBLESS FUTURE

MARTIN FORD

BASIC BOOKS
A Member of the Perseus Books Group
New York

Published by Basic Books,
A Member of the Perseus Books Group

Books published by Basic Books are available at special discounts for bulk purchases in the United States by corporations, institutions, and other organizations. For more information, please contact the Special Markets Department at the Perseus Books Group, 2300 Chestnut Street, Suite 200, Philadelphia, PA 19103, or call (800) 810-4145, ext. 5000, or email special.markets@perseusbooks.com.

Designed by Pauline Brown

Library of Congress Cataloging-in-Publication Data

Ford, Martin (Martin R.)
Rise of the robots : technology and the threat of a jobless future / Martin Ford.
pages cm
Includes bibliographical references and index.
ISBN 978-0-465-05999-7 (hardback) — ISBN 978-0-465-04067-4 (e-book) 1. Labor supply—Effect of automation on. 2. Labor supply—Effect of technological innovations on. 3. Employment forecasting. 4. Technological innovations—Economic aspects. I. Title.
HD6331.F58 2015
331.13'7042—dc23

2014041327

10 9 8 7 6 5 4 3 2 1

To Tristan, Colin,
Elaine, and Xiaoxiao

CONTENTS

INTRODUCTION

Sometime during the 1960s, the Nobel laureate economist Milton Friedman was consulting with the government of a developing Asian nation. Friedman was taken to a large-scale public works project, where he was surprised to see large numbers of workers wielding shovels, but very few bulldozers, tractors, or other heavy earth-moving equipment. When asked about this, the government official in charge explained that the project was intended as a "jobs program." Friedman's caustic reply has become famous: "So then, why not give the workers spoons instead of shovels?"

Friedman's remark captures the skepticism—and often outright derision—expressed by economists confronting fears about the prospect of machines destroying jobs and creating long-term unemployment. Historically, that skepticism appears to be well-founded. In the United States, especially during the twentieth century, advancing technology has consistently driven us toward a more prosperous society.

There have certainly been hiccups—and indeed major disruptions—along the way. The mechanization of agriculture vaporized millions of jobs and drove crowds of unemployed farmhands into cities in search of factory work. Later, automation and globalization pushed workers out of the manufacturing sector and into new service jobs. Short-term unemployment was often a problem during these transitions, but it never became systemic or permanent. New jobs were created and dispossessed workers found new opportunities.

What's more, those new jobs were often better than earlier counterparts, requiring upgraded skills and offering better wages. At no time was this more true than in the two and a half decades following World War II. This "golden age" of the American economy was characterized by a seemingly perfect symbiosis between rapid technological progress and the welfare of the American workforce. As the machines used in production improved, the productivity of the workers operating those machines likewise increased, making them more valuable and allowing them to demand higher wages. Throughout the postwar period, advancing technology deposited money directly into the pockets of average workers as their wages rose in tandem with soaring productivity. Those workers, in turn, went out and spent their ever-increasing incomes, further driving demand for the products and services they were producing.

As that virtuous feedback loop powered the American economy forward, the profession of economics was enjoying its own golden age. It was during the same period that towering figures like Paul Samuelson worked to transform economics into a science with a strong mathematical foundation. Economics gradually came to be almost completely dominated by sophisticated quantitative and statistical techniques, and economists began to build the complex mathematical models that still constitute the field's intellectual basis. As the postwar economists did their work, it would have been natural for them to look at the thriving economy around them and assume that it was normal: that it was the way an economy was supposed to work—*and would always work*.

In his 2005 book *Collapse: How Societies Choose to Succeed or Fail,* Jared Diamond tells the story of agriculture in Australia. In the nineteenth century, when Europeans first colonized Australia, they found a relatively lush, green landscape. Like American economists in the 1950s, the Australian settlers assumed that what they were seeing was normal, and that the conditions they observed would continue indefinitely. They invested heavily in developing farms and ranches on this seemingly fertile land.

Within a decade or two, however, reality struck. The farmers found that the overall climate was actually far more arid than they were initially led to believe. They had simply had the good fortune (or perhaps misfortune) to arrive during a climactic "Goldilocks period"—a sweet spot when everything happened to be just right for agriculture. Today in Australia, you can find the remnants of those ill-fated early investments: abandoned farm houses in the middle of what is essentially a desert.

There are good reasons to believe that America's economic Goldilocks period has likewise come to an end. That symbiotic relationship between increasing productivity and rising wages began to dissolve in the 1970s. As of 2013, a typical production or nonsupervisory worker earned about 13 percent less than in 1973 (after adjusting for inflation), even as productivity rose by 107 percent and the costs of big-ticket items like housing, education, and health care have soared.[1]

On January 2, 2010, the *Washington Post* reported that the first decade of the twenty-first century resulted in the creation of no new jobs. Zero.[2] This hasn't been true of any decade since the Great Depression; indeed, there has never been a postwar decade that produced less than a 20 percent increase in the number of available jobs. Even the 1970s, a decade associated with stagflation and an energy crisis, generated a 27 percent increase in jobs.[3] The lost decade of the 2000s is especially astonishing when you consider that the US economy needs to create roughly a million jobs per year just to keep up with growth in the size of the workforce. In other words, during those first ten years there were about 10 million missing jobs that should have been created—but never showed up.

Income inequality has since soared to levels not seen since 1929, and it has become clear that the productivity increases that went into workers' pockets back in the 1950s are now being retained almost entirely by business owners and investors. The share of overall national income going to labor, as opposed to capital, has fallen precipitously and appears to be in continuing free fall. Our Goldilocks period has reached its end, and the American economy is moving into a new era.

It is an era that will be defined by a fundamental shift in the relationship between workers and machines. That shift will ultimately challenge one of our most basic assumptions about technology: that *machines are tools* that increase the productivity of workers. Instead, machines themselves are turning into workers, and the line between the capability of labor and capital is blurring as never before.

All this progress is, of course, being driven by the relentless acceleration in computer technology. While most people are by now familiar with Moore's Law—the well-established rule of thumb that says computing power roughly doubles every eighteen to twenty-four months—not everyone has fully assimilated the implications of this extraordinary exponential progress.

Imagine that you get in your car and begin driving at 5 miles per hour. You drive for a minute, accelerate to double your speed to 10 mph, drive for another minute, double your speed again, and so on. The really remarkable thing is not simply the fact of the doubling but the amount of ground you cover after the process has gone on for a while. In the first minute, you would travel about 440 feet. In the third minute at 20 mph, you'd cover 1,760 feet. In the fifth minute, speeding along at 80 mph, you would go well over a mile. To complete the sixth minute, you'd need a faster car—as well as a racetrack.

Now think about how fast you would be traveling—and how much progress you would make in that final minute—if you doubled your speed twenty-seven times. That's roughly the number of times computing power has doubled since the invention of the integrated circuit in 1958. The revolution now under way is happening not just because of the acceleration itself but because *that acceleration has been going on for so long* that the amount of progress we can now expect in any given year is potentially mind-boggling.

The answer to the question about your speed in the car, by the way, is 671 *million* miles per hour. In that final, twenty-eighth minute, you would travel more than 11 million miles. Five minutes or so at that speed would get you to Mars. That, in a nutshell, is where information technology stands today, relative to when the first primitive integrated circuits started plodding along in the late 1950s.

As someone who has worked in software development for more than twenty-five years, I've had a front-row seat when it comes to observing that extraordinary acceleration in computing power. I've also seen at close hand the tremendous progress made in software design, and in the tools that make programmers more productive. And, as a small business owner, I've watched as technology has transformed the way I run my business—in particular, how it has dramatically reduced the need to hire employees to perform many of the routine tasks that have always been essential to the operation of any business.

In 2008, as the global financial crisis unfolded, I began to give serious thought to the implications of that consistent doubling in computational power and, especially, to the likelihood that it would dramatically transform the job market and overall economy in coming years and decades. The result was my first book, *The Lights in the Tunnel: Automation, Accelerating Technology and the Economy of the Future,* published in 2009.

In that book, even as I wrote about the importance of accelerating technology, I underestimated just how rapidly things would in fact move forward. For example, I noted that auto manufacturers were working on collision avoidance systems to help prevent accidents, and I suggested that "over time these systems could evolve into technology capable of driving the car autonomously." Well, it turned out that "over time" wasn't much time at all! Within a year of the book's publication, Google introduced a fully automated car capable of driving in traffic. And since then, three states—Nevada, California, and Florida—have passed laws allowing self-driving vehicles to share the road on a limited basis.

I also wrote about progress being made in the field of artificial intelligence. At the time, the story of IBM's "Deep Blue" computer and how it had defeated world chess champion Garry Kasparov in 1997, was perhaps the most impressive demonstration of AI in action. Once again, I was taken by surprise when IBM introduced Deep Blue's successor, Watson—a machine that took on a far more difficult challenge: the television game show *Jeopardy!* Chess is a game with rigidly defined rules; it is the sort of thing we might expect a computer to be good at. *Jeopardy!* is something else entirely: a game that draws on an almost limitless body of knowledge and requires a sophisticated ability to parse language, including even jokes and puns. Watson's success at *Jeopardy!* is not only impressive, it is highly practical, and in fact, IBM is already positioning Watson to play a significant role in fields like medicine and customer service.

It's a good bet that nearly all of us will be surprised by the progress that occurs in the coming years and decades. Those surprises won't be confined to the nature of the technical advances themselves: the impact that accelerating progress has on the job market and the overall economy is poised to defy much of the conventional wisdom about how technology and economics intertwine.

One widely held belief that is certain to be challenged is the assumption that automation is primarily a threat to workers who have little education and lower-skill levels. That assumption emerges from the fact that such jobs tend to be routine and repetitive. Before you get too comfortable with that idea, however, consider just how fast the frontier is moving. At one time, a "routine" occupation would probably have implied standing on an assembly line. The reality today is far different. While lower-skill occupations will no doubt continue to be affected, a great many college-educated, white-collar workers are going to discover that their jobs, too, are squarely in the sights as software automation and predictive algorithms advance rapidly in capability.

The fact is that "routine" may not be the best word to describe the jobs most likely to be threatened by technology. A more accurate term might be "predictable." Could another person learn to do your job by studying a detailed record of everything you've done in the past? Or could someone become proficient by repeating the tasks you've already completed, in the way that a student might take practice tests to prepare for an exam? If so, then there's a good chance that an algorithm may someday be able to learn to do much, or all, of your job. That's made especially likely as the "big data" phenomenon continues to unfold: organizations are collecting incomprehensible amounts of information about nearly every aspect of their operations, and a great many jobs and tasks are likely to be encapsulated in that data—waiting for the day when a smart machine learning algorithm comes along and begins schooling itself by delving into the record left by its human predecessors.

The upshot of all this is that acquiring more education and skills will not necessarily offer effective protection against job automation in the future. As an example, consider radiologists, medical doctors who specialize in the interpretation of medical images. Radiologists require a tremendous amount of training, typically a minimum of thirteen years beyond high school. Yet, computers are rapidly getting better at analyzing images. It's quite easy to imagine that someday, in the not too distant future, radiology will be a job performed almost exclusively by machines.

In general, computers are becoming very proficient at acquiring skills, especially when a large amount of training data is available. Entry-level jobs, in particular, are likely to be heavily affected, and there is evidence that this may already be occurring. Wages for new college graduates have actually been declining over the past decade, while up to 50 percent of new graduates are forced to take jobs that do not require a college degree. Indeed, as I'll demonstrate in this book, employment for many skilled professionals—including lawyers, journalists, scientists, and pharmacists—is already being significantly eroded by advancing information technology. They are not alone: most jobs are, on some level, fundamentally routine and predictable, with relatively few people paid primarily to engage in truly creative work or "blue-sky" thinking.

As machines take on that routine, predictable work, workers will face an unprecedented challenge as they attempt to adapt. In the past, automation technology has tended to be relatively specialized and to disrupt one employment sector at a time, with workers then switching to a new emerging industry. The situation today is quite different. Information technology is a truly general-purpose technology, and its impact will occur across the board. Virtually every industry in existence is likely to become less labor-intensive as new technology is assimilated into business models—and that transition could happen quite rapidly. At the same time, the new industries that emerge will nearly always incorporate powerful labor-saving technology right from their inception. Companies like Google and Facebook, for example, have succeeded in becoming household names and achieving massive market valuations while hiring only a tiny number of people relative to their size and influence. There's every reason to expect that a similar scenario will play out with respect to nearly all the new industries created in the future.

All of this suggests that we are headed toward a transition that will put enormous stress on both the economy and society. Much of the conventional advice offered to workers and to students who are preparing to enter the workforce is likely to be ineffective. The unfortunate reality is that a great many people will do everything right—at least in terms of pursuing higher education and acquiring skills—and yet will still fail to find a solid foothold in the new economy.

Beyond the potentially devastating impact of long-term unemployment and underemployment on individual lives and on the fabric of society, there will also be a significant economic price. The virtuous feedback loop between productivity, rising wages, and increasing consumer spending will collapse. That positive feedback effect is already seriously diminished: we face soaring inequality not just in income but also in consumption. The top 5 percent of households are currently responsible for nearly 40 percent of spending, and that trend toward increased concentration at the top seems almost certain to continue. Jobs remain the primary mechanism by which purchasing power gets into the hands of consumers. If that mechanism continues to erode, we will face the prospect of having too few viable consumers to continue driving economic growth in our mass-market economy.

As this book will make clear, advancing information technology is pushing us toward a tipping point that is poised to ultimately make the entire economy less labor-intensive. However, that transition won't necessarily unfold in a uniform or predictable way. Two sectors in particular—higher education and health care—have, so far, been highly resistant to the kind of disruption that is already becoming evident in the broader economy. The irony is that the failure of technology to transform these sectors could amplify its negative consequences elsewhere, as the costs of health care and education become ever more burdensome.

Technology, of course, will not shape the future in isolation. Rather, it will intertwine with other major societal and environmental challenges such as an aging population, climate change, and resource depletion. It's often predicted that a shortage of workers will eventually develop as the baby boom generation exits the workforce, effectively counterbalancing—or perhaps even overwhelming—any impact from automation. Rapid innovation is typically framed purely as a countervailing force with the potential to minimize, or even reverse, the stress we put on the environment. However, as we'll see, many of these assumptions rest on uncertain foundations: the story is sure to be far more complicated. Indeed, the frightening reality is that if we don't recognize and adapt to the implications of advancing technology, we may face the prospect of a "perfect storm" where the impacts from soaring inequality, technological unemployment, and climate change unfold roughly in parallel, and in some ways amplify and reinforce each other.

In Silicon Valley the phrase "disruptive technology" is tossed around on a casual basis. No one doubts that technology has the power to devastate entire industries and upend specific sectors of the economy and job market. The question I will ask in this book is bigger: Can accelerating technology disrupt *our entire system* to the point where a fundamental restructuring may be required if prosperity is to continue?

Chapter 1

THE AUTOMATION WAVE

A warehouse worker approaches a stack of boxes. The boxes are of varying shapes, sizes, and colors, and they are stacked in a somewhat haphazard way.

Imagine for a moment that you can see inside the brain of the worker tasked with moving the boxes, and consider the complexity of the problem that needs to be solved.

Many of the boxes are a standard brown color and are pressed tightly against each other, making the edges difficult to perceive. Where precisely does one box end and the next begin? In other cases, there are gaps and misalignments. Some boxes are rotated so that one edge juts out. At the top of the pile, a small box rests at an angle in the space between two larger boxes. Most of the boxes are plain brown or white cardboard, but some are emblazoned with company logos, and a few are full-color retail boxes intended to be displayed on store shelves.

The human brain is, of course, capable of making sense of all this complicated visual information almost instantaneously. The worker easily perceives the dimensions and orientation of each box, and seems to know instinctively that he must begin by moving the boxes at the top of the stack and how to move the boxes in a sequence that won't destabilize the rest of the pile.

This is exactly the type of visual perception challenge that the human brain has evolved to overcome. That the worker succeeds in moving the boxes would be completely unremarkable—were it not for the fact that, in this case, the worker is a robot. To be more precise, it is a snake-like robotic arm, its head consisting of a suction-powered gripper. The robot is slower to comprehend than a human would be. It peers at the boxes, adjusts its gaze slightly, ponders some more, and then finally lunges forward and grapples a box from the top of the pile.* The sluggishness, however, results almost entirely from the staggering complexity of the computation required to perform this seemingly simple task. If there is one thing the history of information technology teaches, it is that this robot is going to very soon get a major speed upgrade.

Indeed, engineers at Industrial Perception, Inc., the Silicon Valley start-up company that designed and built the robot, believe the machine will ultimately be able to move a box every second. That compares with a human worker's maximum rate of a box roughly every six seconds.[1] Needless to say, the robot can work continuously; it will never get tired or suffer a back injury—and it will certainly never file a worker's compensation claim.

Industrial Perception's robot is remarkable because its capability sits at the nexus of visual perception, spatial computation, and dexterity. In other words, it is invading the final frontier of machine automation, where it will compete for the few relatively routine, manual jobs that are still available to human workers.

Robots in factories are, of course, nothing new. They have become indispensable in virtually every sector of manufacturing, from

---

* A video of Industrial Perception's box-moving robot can be seen on the company's website at http://www.industrial-perception.com/technology.html.

automobiles to semiconductors. Electric-car company Tesla's new plant in Fremont, California, uses 160 highly flexible industrial robots to assemble about 400 cars per week. As a new-car chassis arrives at the next position in the assembly line, multiple robots descend on it and operate in coordination. The machines are able to autonomously swap the tools wielded by their robotic arms in order to complete a variety of tasks. The same robot, for example, installs the seats, retools itself, and then applies adhesive and drops the windshield into place.[2] According to the International Federation of Robotics, global shipments of industrial robots increased by more than 60 percent between 2000 and 2012, with total sales of about $28 billion in 2012. By far the fastest-growing market is China, where robot installations grew at about 25 percent per year between 2005 and 2012.[3]

While industrial robots offer an unrivaled combination of speed, precision, and brute strength, they are, for the most part, blind actors in a tightly choreographed performance. They rely primarily on precise timing and positioning. In the minority of cases where robots have machine vision capability, they can typically see in just two dimensions and only in controlled lighting conditions. They might, for example, be able to select parts from a flat surface, but an inability to perceive depth in their field of view results in a low tolerance for environments that are to any meaningful degree unpredictable. The result is that a number of routine factory jobs have been left for people. Very often these are jobs that involve filling the gaps between the machines, or they are at the end points of the production process. Examples might include choosing parts from a bin and then feeding them into the next machine, or loading and unloading the trucks that move products to and from the factory.

The technology that powers the Industrial Perception robot's ability to see in three dimensions offers a case study in the ways that cross-fertilization can drive bursts of innovation in unexpected areas. It might be argued that the robot's eyes can trace their origin to November 2006, when Nintendo introduced its Wii video game console.

Nintendo's machine included an entirely new type of game controller: a wireless wand that incorporated an inexpensive device called an accelerometer. The accelerometer was able to detect motion in three dimensions and then output a data stream that could be interpreted by the game console. Video games could now be controlled through body movements and gestures. The result was a dramatically different game experience. Nintendo's innovation smashed the stereotype of the nerdy kid glued to a monitor and a joystick, and opened a new frontier for games as active exercise.

It also demanded a competitive response from the other major players in the video game industry. Sony Corporation, makers of the PlayStation, elected to essentially copy Nintendo's design and introduced its own motion-detecting wand. Microsoft, however, aimed to leapfrog Nintendo and come up with something entirely new. The Kinect add-on to the Xbox 360 game console eliminated the need for a controller wand entirely. To accomplish this, Microsoft built a webcam-like device that incorporates three-dimensional machine vision capability based in part on imaging technology created at a small Israeli company called PrimeSense. The Kinect sees in three dimensions by using what is, in essence, sonar at the speed of light: it shoots an infrared beam at the people and objects in a room and then calculates their distance by measuring the time required for the reflected light to reach its infrared sensor. Players could now interact with the Xbox game console simply by gesturing and moving in view of the Kinect's camera.

The truly revolutionary thing about the Kinect was its price. Sophisticated machine vision technology—which might previously have cost tens or even hundreds of thousands of dollars and required bulky equipment—was now available in a compact and lightweight consumer device priced at $150. Researchers working in robotics instantly realized the potential for the Kinect technology to transform their field. Within weeks of the product's introduction, both university-based engineering teams and do-it-yourself innovators had hacked into the Kinect and posted YouTube videos of robots that were now able to see in three dimensions.[4] Industrial Perception likewise decided to base its vision system on the technology that powers the Kinect, and the result is an affordable machine that is rapidly approaching a nearly human-level ability to perceive and interact with its environment while dealing with the kind of uncertainty that characterizes the real world.

A Versatile Robotic Worker

Industrial Perception's robot is a highly specialized machine focused specifically on moving boxes with maximum efficiency. Boston-based Rethink Robotics has taken a different track with Baxter, a lightweight humanoid manufacturing robot that can easily be trained to perform a variety of repetitive tasks. Rethink was founded by Rodney Brooks, one of the world's foremost robotics researchers at MIT and a co-founder of iRobot, the company that makes the Roomba automated vacuum cleaner as well as military robots used to defuse bombs in Iraq and Afghanistan. Baxter, which costs significantly less than a year's wages for a typical US manufacturing worker, is essentially a scaled-down industrial robot that is designed to operate safely in close proximity to people.

In contrast to industrial robots, which require complex and expensive programming, Baxter can be trained simply by moving its arms through the required motions. If a facility uses multiple robots, one Baxter can be trained and then the knowledge can be propagated to the others simply by plugging in a USB device. The robot can be adapted to a variety of tasks, including light assembly work, transferring parts between conveyer belts, packing products into retail packaging, or tending machines used in metal fabrication. Baxter is particularly talented at packing finished products into shipping boxes. K'NEX, a toy construction set manufacturer located in Hatfield, Pennsylvania, found that Baxter's ability to pack its products tightly allowed the company to use 20–40 percent fewer boxes.[5] Rethink's robot also has two-dimensional machine vision capability powered by cameras on both wrists and can pick up parts and even perform basic quality-control inspections.

The Coming Explosion in Robotics

While Baxter and Industrial Perception's box-moving robot are dramatically different machines, they are both built on the same fundamental software platform. ROS—or Robot Operating System—was originally conceived at Stanford University's Artificial Intelligence Laboratory and then developed into a full-fledged robotics platform by Willow Garage, Inc., a small company that designs and manufactures programmable robots that are used primarily by researchers at universities. ROS is similar to operating systems like Microsoft Windows, Macintosh OS, or Google's Android but is geared specifically toward making robots easy to program and control. Because ROS is free and also open source—meaning that software developers can easily modify and enhance it—it is rapidly becoming the standard software platform for robotics development.

The history of computing shows pretty clearly that once a standard operating system, together with inexpensive and easy-to-use programming tools, becomes available, an explosion of application software is likely to follow. This has been the case with personal computer software and, more recently, with iPhone, iPad, and Android apps. Indeed, these platforms are now so saturated with application software that it can be genuinely difficult to conceive of an idea that hasn't already been implemented.

It's a good bet that the field of robotics is poised to follow a similar path; we are, in all likelihood, at the leading edge of an explosive wave of innovation that will ultimately produce robots geared toward nearly every conceivable commercial, industrial, and consumer task. That explosion will be powered by the availability of standardized software and hardware building blocks that will make it a relatively simple matter to assemble new designs without the need to reinvent the wheel. Just as the Kinect made machine vision affordable, other hardware components—such as robotic arms—will see their costs driven down as robots begin scaling up to high-volume production. As of 2013, there were already thousands of software components available to work with ROS, and development platforms were cheap enough to allow nearly anyone to start designing new robotics applications. Willow Garage, for example, sells a complete mobile robot kit called TurtleBot that includes Kinect-powered machine vision for about $1,200. After inflation is taken into account, that's far less than what an inexpensive personal computer and monitor cost in the early 1990s, when Microsoft Windows was in the early stages of producing its own software explosion.

When I visited the RoboBusiness conference and tradeshow in Santa Clara, California, in October 2013, it was clear that the robotics industry had already started gearing up for the coming explosion. Companies of all sizes were on hand to showcase robots designed to perform precision manufacturing, transport medical supplies between departments in large hospitals, or autonomously operate heavy equipment for agriculture and mining. There was a personal robot named "Budgee" capable of carrying up to fifty pounds of stuff around the house or at the store. A variety of educational robots focused on everything from encouraging technical creativity to assisting children with autism or learning disabilities. At the Rethink Robotics booth, Baxter had received Halloween training and was grasping small boxes of candy and then dropping them into pumpkin-shaped trick-or-treat buckets. There were also companies marketing components like motors, sensors, vision systems, electronic controllers, and the specialized software used to construct robots. Silicon Valley start-up Grabit Inc. demonstrated an innovative electroadhesion-powered gripper that allows robots to pick up, carry, and place nearly anything simply by employing a controlled electrostatic charge. To round things out, a global law firm with a specialized robotics practice was on hand to help employers navigate the complexities of labor, employment, and safety regulations when robots are brought in to replace, or work in close proximity to, people.

One of the most remarkable sights at the tradeshow was in the aisles—which were populated by a mix of human attendees and dozens of remote-presence robots provided by Suitable Technologies, Inc. These robots, consisting of a flat screen and camera mounted on a mobile pedestal, allowed remote participants to visit tradeshow booths, view demonstrations, ask questions, and otherwise interact normally with other participants. Suitable Technologies offered remote presence at the tradeshow for a minimal fee, allowing visitors from outside the San Francisco Bay area to avoid thousands of dollars in travel costs. After a few minutes, the robots—each with a human face displayed on its screen—did not seem at all out of place as they prowled between booths and engaged other attendees in conversation.

Manufacturing Jobs and Factory Reshoring

In a September 2013 article, Stephanie Clifford of the *New York Times* told the story of Parkdale Mills, a textile factory in Gaffney, South Carolina. The Parkdale plant employs about 140 people. In 1980, the same level of production would have required more than 2,000 factory workers. Within the Parkdale plant, "only infrequently does a person interrupt the automation, mainly because certain tasks are still cheaper if performed by hand—like moving half-finished yarn between machines on forklifts."[6] Completed yarn is conveyed automatically toward packing and shipping machines along pathways attached to the ceiling.

Nonetheless, those 140 factory jobs represent at least a partial reversal of a decades-long decline in manufacturing employment. The

US textile industry was decimated in the 1990s as production moved to low-wage countries, especially China, India, and Mexico. About 1.2 million jobs—more than three-quarters of domestic employment in the textile sector—vanished between 1990 and 2012. The last few years, however, have seen a dramatic rebound in production. Between 2009 and 2012, US textile and apparel exports rose by 37 percent to a total of nearly $23 billion.[7] The turnaround is being driven by automation technology so efficient that it is competitive with even the lowest-wage offshore workers.

Within the manufacturing sector in the United States and other developed countries, the introduction of these sophisticated labor-saving innovations is having a mixed impact on employment. While factories like Parkdale don't directly create large numbers of manufacturing jobs, they do drive increased employment at suppliers and in peripheral areas like driving the trucks that move raw materials and finished products. While a robot like Baxter can certainly eliminate the jobs of some workers who perform routine tasks, it also helps make US manufacturing more competitive with low-wage countries. Indeed, there is now a significant "reshoring" trend under way, and this is being driven both by the availability of new technology and by rising offshore labor costs, especially in China where typical factory workers saw their pay increase by nearly 20 percent per year between 2005 and 2010. In April 2012, the Boston Consulting Group surveyed American manufacturing executives and found that nearly half of companies with sales exceeding $10 billion were either actively pursuing or considering bringing factories back to the United States.[8]

Factory reshoring dramatically decreases transportation costs and also provides many other advantages. Locating factories in close proximity to both consumer markets and product design centers allows companies to cut production lead times and be far more responsive to their customers. As automation becomes ever more flexible and sophisticated, it's likely that manufacturers will trend toward offering more customizable products—perhaps, for example, allowing customers to create unique designs or specify hard-to-find clothing sizes through easy-to-use online interfaces. Domestic automated production could then put a finished product into a customer's hands within days.

There is, however, one important caveat to the reshoring narrative. Even the relatively small number of new factory jobs now being created as a result of reshoring won't necessarily be around over the long term; as robots continue to get more capable and dexterous and as new technologies like 3D printing come into widespread use, it seems likely that many factories will eventually approach full automation. Manufacturing jobs in the United States currently account for well under 10 percent of total employment. As a result, manufacturing robots and reshoring are likely to have a fairly marginal impact on the overall job market.

The story will be very different in developing countries like China, where employment is far more focused in the manufacturing sector. In fact, advancing technology has already had a dramatic impact on Chinese factory jobs; between 1995 and 2002 China lost about 15 percent of its manufacturing workforce, or about 16 million jobs.[9] There is strong evidence to suggest that this trend is poised to accelerate. In 2012, Foxconn—the primary contract manufacturer of Apple devices—announced plans to eventually introduce up to a million robots in its factories. Taiwanese company Delta Electronics, Inc., a producer of power adapters, has recently shifted its strategy to focus on low-cost robots for precision electronics assembly. Delta hopes to offer a one-armed assembly robot for about $10,000—less than half the cost of Rethink's Baxter. European industrial robot manufacturers like ABB Group and Kuka AG are likewise investing heavily in the Chinese market and are currently building local factories to churn out thousands of robots per year.[10]

Increased automation is also likely to be driven by the fact that the interest rates paid by large companies in China are kept artificially low as a result of government policy. Loans are often rolled over continuously, so that the principal is never repaid. This makes capital investment extremely attractive even when labor costs are low and has been one of the primary reasons that investment now accounts for nearly half of China's GDP.[11] Many analysts believe that this artificially low cost of capital has caused a great deal of mal-investment throughout China, perhaps most famously the construction of "ghost cities" that appear to be largely unoccupied. By the same token, low capital costs may create a powerful incentive for big companies to invest in expensive automation, even in those cases where it does not necessarily make good business sense to do so.

One of the biggest challenges for a transition to robotic assembly in the Chinese electronics industry will be designing robots that are flexible enough to keep up with rapid product lifecycles. Foxconn, for example, maintains massive facilities where workers live onsite in dormitories. In order to accommodate aggressive production schedules, thousands of workers can be woken in the middle of the night and set immediately to work. That results in an astonishing ability to rapidly ramp up production or adjust to product design changes, but it also puts extreme pressure on workers—as evidenced by the near epidemic of suicides that occurred at Foxconn facilities in 2010. Robots, of course, have the ability to work continuously, and as they become more flexible and easier to train for new tasks, they will become an increasingly attractive alternative to human workers, even when wages are low.

The trend toward increased factory automation in developing countries is by no means limited to China. Clothing and shoe production, for example, continues to be one of the most labor-intensive sectors of manufacturing, and factories have been transitioning from China to even lower-wage countries like Vietnam and Indonesia. In June 2013, athletic-shoe manufacturer Nike announced that rising wages in Indonesia had negatively impacted its quarterly financial numbers. According to the company's chief financial officer, the long-term solution to that problem is going to be "engineering the labor out of the product."[12] Increased automation is also seen as a way to deflect criticism regarding the sweatshop-like environments that often exist in third-world garment factories.

The Service Sector: Where the Jobs Are

In the United States and other advanced economies, the major disruption will be in the service sector—which is, after all, where the vast majority of workers are now employed. This trend is already evident in areas like ATMs and self-service checkout lanes, but the next decade is likely to see an explosion of new forms of service sector automation, potentially putting millions of relatively low-wage jobs at risk.

San Francisco start-up company Momentum Machines, Inc., has set out to fully automate the production of gourmet-quality hamburgers. Whereas a fast food worker might toss a frozen patty onto the grill, Momentum Machines' device shapes burgers from freshly ground meat and then grills them to order—including even the ability to add just the right amount of char while retaining all the juices. The machine, which is capable of producing about 360 hamburgers per hour, also toasts the bun and then slices and adds fresh ingredients like tomatoes, onions, and pickles only after the order is placed. Burgers arrive assembled and ready to serve on a conveyer belt. While most robotics companies take great care to spin a positive tale when it comes to the potential impact on employment, Momentum Machines co-founder Alexandros Vardakostas is very forthright about the company's objective: "Our device isn't meant to make employees more efficient," he said. "It's meant to completely obviate them."[13] * The company estimates that the average fast food

* The company is not unaware of the potential impact its technology will have on jobs and, according to its website, plans to support a program that will offer discounted technical training to workers who are displaced.

restaurant spends about $135,000 per year on wages for employees who produce hamburgers and that the total labor cost for burger production for the US economy is about $9 billion annually.[14] Momentum Machines believes its device will pay for itself in less than a year, and it plans to target not just restaurants but also convenience stores, food trucks, and perhaps even vending machines. The company argues that eliminating labor costs and reducing the amount of space required in kitchens will allow restaurants to spend more on high-quality ingredients, enabling them to offer gourmet hamburgers at fast food prices.

Those burgers might sound very inviting, but they would come at a considerable cost. Millions of people hold low-wage, often part-time, jobs in the fast food and beverage industries. McDonald's alone employs about 1.8 million workers in 34,000 restaurants worldwide.[15] Historically, low wages, few benefits, and a high turnover rate have helped to make fast food jobs relatively easy to find, and fast food jobs, together with other low-skill positions in retail, have provided a kind of private sector safety net for workers with few other options: these jobs have traditionally offered an income of last resort when no better alternatives are available. In December 2013, the US Bureau of Labor Statistics ranked "combined food preparation and serving workers," a category that excludes waiters and waitresses in full-service restaurants, as one of the top employment sectors in terms of the number of job openings projected over the course of the decade leading up to 2022—with nearly half a million new jobs and another million openings to replace workers who leave the industry.[16]

In the wake of the Great Recession, however, the rules that used to apply to fast food employment are changing rapidly. In 2011, McDonald's launched a high-profile initiative to hire 50,000 new workers in a single day and received over a million applications—a ratio that made landing a McJob more of a statistical long shot than getting accepted at Harvard. While fast food employment was once dominated by young people looking for a part-time income while in school, the industry now employs far more mature workers who rely on the jobs as their primary income. Nearly 90 percent of fast food workers are twenty or older, and the average age is thirty-five.[17] Many of these older workers have to support families—a nearly impossible task at a median wage of just $8.69 per hour.

The industry's low wages and nearly complete lack of benefits have drawn intensive criticism. In October 2013, McDonald's was lambasted after an employee who called the company's financial help line was advised to apply for food stamps and Medicaid.[18] Indeed, an analysis by the Labor Center at the University of California, Berkeley, found that more than half of the families of fast food workers are enrolled in some type of public assistance program and that the resulting cost to US taxpayers is nearly $7 billion per year.[19]

When a spate of protests and ad hoc strikes at fast food restaurants broke out in New York and then spread to more than fifty US cities in the fall of 2013, the Employment Policies Institute, a conservative think tank with close ties to the restaurant and hotel industries, placed a full-page ad in the *Wall Street Journal* warning that "Robots Could Soon Replace Fast Food Workers Demanding a Higher Minimum Wage." While the ad was doubtless intended as a scare tactic, the reality is that—as the Momentum Machines device demonstrates—increased automation in the fast food industry is almost certainly inevitable. Given that companies like Foxconn are introducing robots to perform high-precision electronic assembly in China, there is little reason to believe that machines won't also eventually be serving up burgers, tacos, and lattes across the fast food industry.*

Japan's Kura sushi restaurant chain has already successfully pioneered an automation strategy. In the chain's 262 restaurants, robots

---

* Economists categorize fast food as part of the service sector; however, from a technical standpoint it is really closer to being a form of just-in-time manufacturing.

help make the sushi while conveyor belts replace waiters. To ensure freshness, the system keeps track of how long individual sushi plates have been circulating and automatically removes those that reach their expiration time. Customers order using touch panel screens, and when they are finished dining they place the empty dishes in a slot near their table. The system automatically tabulates the bill and then cleans the plates and whisks them back to the kitchen. Rather than employing store managers at each location, Kura uses centralized facilities where managers are able to remotely monitor nearly every aspect of restaurant operations. Kura's automation-based business model allows it to price sushi plates at just 100 yen (about $1), significantly undercutting its competitors.[20]

It's fairly easy to envision many of the strategies that have worked for Kura, especially automated food production and offsite management, eventually being adopted across the fast food industry. Some significant steps have already been taken in that direction; McDonalds, for example, announced in 2011 that it would install touch screen ordering systems at 7,000 of its European restaurants.[21] Once one of the industry's major players begins to gain significant advantages from increased automation, the others will have little choice but to follow suit. Automation will also offer the ability to compete on dimensions beyond lower labor costs. Robotic production might be viewed as more hygienic since fewer workers would come into contact with the food. Convenience, speed, and order accuracy would increase, as would the ability to customize orders. Once a customer's preferences were recorded at one restaurant, automation would make it a simple matter to consistently produce the same results at other locations.

Given all this, I think it is quite easy to imagine that a typical fast food restaurant may eventually be able to cut its workforce by 50 percent, or perhaps even more. At least in the United States, the fast food market is already so saturated that it seems very unlikely that new restaurants could make up for such a dramatic reduction in the number of workers required at each location. And this, of course, would mean that a great many of the job openings forecast by the Bureau of Labor Statistics might never materialize.

The other major concentration of low-wage service jobs is in the general retail sector. Economists at the Bureau of Labor Statistics rank "retail salesperson" second only to "registered nurse" as the specific occupation that will add the most jobs in the decade ending in 2020 and expect over 700,000 new jobs to be created.[22] Once again, however, technology has the potential to make the government projections seem optimistic. We can probably anticipate that three major forces will shape employment in the retail sector going forward.

The first will be the continuing disruption of the industry by online retailers like Amazon, eBay, and Netflix. The competitive advantage that online suppliers have over brick and mortar stores is already, of course, evident with the demise of major retail chains like Circuit City, Borders, and Blockbuster. Both Amazon and eBay are experimenting with same-day delivery in a number of US cities, with the objective of undermining one of the last major advantages that local retail stores still enjoy: the ability to provide immediate gratification after a purchase.

In theory, the encroachment of online retailers should not necessarily destroy jobs but, rather, would transition them from traditional retail settings to the warehouses and distribution centers used by the online companies. However, the reality is that once jobs move to a warehouse they become far easier to automate. Amazon purchased Kiva Systems, a warehouse robotics company in 2012. Kiva's robots, which look a bit like huge, roving hockey pucks, are designed to move materials within warehouses. Rather than having workers roam the aisles selecting items, a Kiva robot simply zips under an entire pallet or shelving unit, lifts it, and then brings it directly to the worker packing an order. The robots navigate autonomously using a grid laid out by barcodes attached to the floor and are used to automate warehouse operations at a variety of major retailers in addition to

Amazon, including Toys "R" Us, the Gap, Walgreens, and Staples.[23] A year after the acquisition, Amazon had about 1,400 Kiva robots in operation but had only begun the process of integrating the machines into its massive warehouses. One Wall Street analyst estimates that the robots will ultimately allow the company to cut its order fulfillment costs by as much as 40 percent.[24]

The Kroger Company, one of the largest grocery retailers in the United States, has also introduced highly automated distribution centers. Kroger's system is capable of receiving pallets containing large supplies of a single product from vendors and then disassembling them and creating new pallets containing a variety of different products that are ready to ship to stores. It is also able to organize the way that products are stacked on the mixed pallets in order to optimize the stocking of shelves once they arrive at stores. The automated warehouses completely eliminate the need for human intervention, except for loading and unloading the pallets onto trucks.[25] The obvious impact that these automated systems have on jobs has not been lost on organized labor, and the Teamsters Union has repeatedly clashed with Kroger, as well as other grocery retailers, over their introduction. Both the Kiva robots and Kroger's automated system do leave some jobs for people, and these are primarily in areas, such as packing a mixture of items for final shipment to customers, that require visual recognition and dexterity. Of course, these are the very areas in which innovations like Industrial Perception's box-moving robots are rapidly advancing the technical frontier.

The second transformative force is likely to be the explosive growth of the fully automated self-service retail sector—or, in other words, intelligent vending machines and kiosks. One study projects that the value of products and services vended in this market will grow from about $740 billion in 2010 to more than $1.1 trillion by 2015.[26] Vending machines have progressed far beyond dispensing sodas, snacks, and lousy instant coffee, and sophisticated machines that sell consumer electronics products like Apple's iPod and iPad are now common in airports and upscale hotels. AVT, Inc., one of the leading manufacturers of automated retail machines, claims that it can design a custom self-service solution for virtually any product. Vending machines make it possible to dramatically reduce three of the most significant costs incurred in the retail business: real estate, labor, and theft by customers and employees. In addition to providing 24-hour service, many of the machines include video screens and are able to offer targeted point-of-sale advertising that's geared toward enticing customers to purchase related products in much the same way that a human sales clerk might do. They can also collect customer email addresses and send receipts. In essence, the machines offer many of the advantages of online ordering, with the added benefit of instant delivery.

While the proliferation of vending machines and kiosks is certain to eliminate traditional retail sales jobs, these machines will also, of course, create jobs in areas like maintenance, restocking, and repair. The number of those new jobs, however, is likely to be more limited than you might expect. The latest-generation machines are directly connected to the Internet and provide a continuous stream of sales and diagnostic data; they are also specifically designed to minimize the labor costs associated with their operation.

In 2010, David Dunning was the regional operations supervisor responsible for overseeing the maintenance and restocking of 189 Redbox movie rental kiosks in the Chicago area.[27] Redbox has over 42,000 kiosks in the United States and Canada, typically located at convenience stores and supermarkets, and rents about 2 million videos per day.[28] Dunning managed the Chicago-area kiosks with a staff of just seven. Restocking the machines is highly automated; in fact, the most labor-intensive aspect of the job is swapping the translucent movie advertisements displayed on the kiosk—a process that typically takes less than two minutes for each machine. Dunning and his staff divide their time between the warehouse, where new movies arrive, and their cars and homes, where they are able to access and manage the machines via the Internet. The kiosks are designed from the ground up for remote maintenance. For example, if a machine jams it will report this immediately, and a technician can log in with his or her laptop computer, jiggle the mechanism, and fix the problem without the need to visit the site. New movies are typically released on Tuesdays, but the machines can be restocked at any time prior to that; the kiosk will automatically make the movies available for rental at the right time. That allows technicians to schedule restocking visits to avoid traffic.

While the jobs that Dunning and his staff have are certainly interesting and desirable, in number they are a fraction of what a traditional retail chain would create. The now-defunct Blockbuster, for example, once had dozens of stores in greater Chicago, each employing its own sales staff.[29] At its peak, Blockbuster had a total of about 9,000 stores and 60,000 employees. That works out to about seven jobs per store—roughly the same number that Redbox employed in the entire region serviced by Dunning's team.

The third major force likely to disrupt employment in the retail sector will be the introduction of increased automation and robotics into stores as brick and mortar retailers strive to remain competitive. The same innovations that are enabling manufacturing robots to advance the frontier in areas like physical dexterity and visual recognition will eventually allow retail automation to begin moving from warehouses into more challenging and varied environments like stocking shelves in stores. In fact, as far back as 2005, Walmart was already investigating the possibility of using robots that rove store aisles at night and automatically scan barcodes in order to track product inventories.[30]

At the same time, self-service checkout aisles and in-store information kiosks are sure to become easier to use, as well as more common. Mobile devices will also become an ever more important self-service tool. Future shoppers will rely more and more on their phones as a way to shop, pay, and get help and information about products while in traditional retail settings. The mobile disruption of retail is already under way. Walmart, for example, is testing an experimental program that allows shoppers to scan barcodes and then checkout and pay with their phones—completely avoiding long checkout lines.[31] Silvercar, a start-up rental car company, offers the capability to reserve and pick up a car without ever having to interact with a rental clerk; the customer simply scans a barcode to unlock the car and then drives away.[32] As natural language technology like Apple's Siri or even more powerful systems like IBM's Watson continue to advance and become more affordable, it's easy to imagine shoppers soon being able to ask their mobile devices for assistance in much the same way they might ask a store employee. The difference, of course, is that the customer will never have to wait for or hunt down the employee; the virtual assistant will always be instantly available and will rarely, if ever, give an inaccurate answer.

While many retailers may choose to bring automation into traditional retail configurations, others may instead elect to entirely redesign stores—perhaps, in essence, turning them into scaled-up vending machines. Stores of this type might consist of an automated warehouse with an attached showroom where customers could examine product samples and place orders. Orders might then be delivered directly to customers, or perhaps even loaded robotically into vehicles. Regardless of the specific technological path ultimately followed by the retail industry, it's difficult to imagine that the eventual result won't be more robots and machines—and significantly fewer jobs for people.

Cloud Robotics

One of the most important propellants of the robot revolution may turn out to be "cloud robotics"—or the migration of much of the intelligence that animates mobile robots into powerful, centralized computing hubs. Cloud robotics has been enabled by the dramatic acceleration in the rate at which data can be communicated; it is now possible to offload much of the computation required by advanced robotics into huge data centers while also giving individual robots access to network-wide resources. That, of course, makes it possible to build less expensive robots, since less onboard computational power and memory are required, and also allows for instant software upgrades across multiple machines. If one robot employs centralized machine intelligence to learn and adapt to its environment, then that newly acquired knowledge could become instantly available to any other machines accessing the system—making it easy to scale machine learning across large numbers of robots. Google announced support for cloud robotics in 2011 and provides an interface that allows robots to take advantage of all the services designed for Android devices.*

The impact of cloud robotics may be most dramatic in areas like visual recognition that require access to vast databases as well as powerful computational capability. Consider, for example, the enormous technical challenge involved in building a robot capable of performing a variety of housekeeping chores. A robotic maid tasked with clearing up the clutter in a room would need to be able to recognize an almost unlimited number of objects and then decide what to do with them. Each of those items might come in a variety of styles, be oriented in different ways, and perhaps even be somehow entangled with other objects. Compare that challenge to the one taken on by the Industrial Perception box-moving robot we met at the beginning of this chapter. While that robot's ability to discern and grasp individual boxes even when they are stacked in a careless way is an impressive achievement, it is still limited to, well, boxes. That's obviously a very long way from being able to recognize and manipulate virtually any object of any shape and in any configuration.

---

* Google's strong interest in robotics was further demonstrated in 2013, when the company purchased eight robotics start-up companies over a six-month period. Among the companies acquired was Industrial Perception.

Building such comprehensive visual perception and recognition into an affordable robot poses a daunting challenge. Yet, cloud robotics offers at least a glimpse of the path that may eventually lead to a solution. Google introduced its “Goggles” feature for camera-equipped mobile devices in 2010 and has significantly improved the technology since then. This feature allows you to take a photo of things like landmark buildings, books, works of art, and commercial products and then have the system automatically recognize and retrieve information relevant to the photo. While building the ability to recognize nearly any object into a robot’s onboard system would be extraordinarily difficult and expensive, it’s fairly easy to imagine robots of the future recognizing the objects in their environment by accessing a vast centralized database of images similar to the one used by the Goggles system. The cloud-based image library could be updated continuously, and any robots with access to the system would get an instant upgrade to their visual recognition capability.

Cloud robotics is sure to be a significant driver of progress in building more capable robots, but it also raises important concerns, especially in the area of security. Aside from its uncomfortable similarity to “Skynet,” the controlling machine intelligence in the *Terminator* movies starring Arnold Schwarzenegger, there is the much more practical and immediate issue of susceptibility to hacking or cyber attack. This will be an especially significant concern if cloud robotics someday takes on an important role in our transportation infrastructure. For example, if automated trucks and trains eventually move food and other critical supplies under centralized control, such a system might create extreme vulnerabilities. There is already great concern about the vulnerability of industrial machinery, and of vital infrastructure like the electrical grid, to cyber attack. That vulnerability was demonstrated by the Stuxnet worm that was created by the US and Israeli governments in 2010 to attack the centrifuges used in Iran’s nuclear program. If, someday, important infrastructure components are dependent on centralized machine intelligence, those concerns could be raised to an entirely new level.

Robots in Agriculture

Of all the employment sectors that make up the US economy, agriculture stands out as the one that has already undergone the most dramatic transformation as a direct result of technological progress. Most of those new technologies were, of course, mechanical in nature and came long before the advent of advanced information technology. In the late nineteenth century, nearly half of all US workers were employed on farms; by 2000 that fraction had fallen below 2 percent. For crops like wheat, corn, and cotton that can be planted, maintained, and harvested mechanically, the human labor required per bushel of output is now nearly negligible in advanced countries. Many aspects of raising and managing livestock are also mechanized. For example, robotic milking systems are in common use on dairy farms, and in the United States, chickens are grown to standardized sizes so as to make them compatible with automated slaughtering and processing.

The remaining labor-intensive areas of agriculture are primarily geared toward picking delicate, high-value fruits and vegetables, as well as ornamental plants and flowers. As with other relatively routine, manual occupations, these jobs have so far been protected from mechanization primarily because they are highly dependent on visual perception and dexterity. Fruits and vegetables are easily damaged and often need to be selected based on color or softness. For a machine, visual recognition is a significant challenge: lighting conditions can be highly variable, and individual fruits can be in a variety of orientations and may be partly or even completely obscured by leaves.

The same innovations that are advancing the robotics frontier in factory and warehouse settings are finally making many of these remaining agricultural jobs susceptible to automation. Vision

Robotics, a company based in San Diego, California, is developing an octopus-like orange harvesting machine. The robot will use three-dimensional machine vision to make a computer model of an entire orange tree and then store the location of each fruit. That information will then be passed on to the machine's eight robotic arms, which will rapidly harvest the oranges.[33] Boston-area start-up Harvest Automation is initially focused on building robots to automate operations in nurseries and greenhouses; the company estimates that manual labor accounts for over 30 percent of the cost of growing ornamental plants. In the longer run, the company believes that its robots will be able to perform up to 40 percent of the manual agricultural labor now required in the United States and Europe.[34] Experimental robots are already pruning grapevines in France using machine vision technology combined with algorithms that decide which stems should be cut.[35] In Japan, a new machine is able to select ripe strawberries based on subtle color variations and then pick a strawberry every eight seconds—working continuously and doing most of the work at night.[36]

Advanced agricultural robots are especially attractive in countries that do not have access to low-wage, migrant labor. Australia and Japan, for example, are both island nations with rapidly aging workforces. Security considerations likewise make Israel a virtual island in terms of labor mobility. Many fruits and vegetables need to be harvested within a very small time window, so that a lack of available workers at just the right time can easily turn out to be a catastrophic problem.

Beyond reducing the need for labor, agricultural automation has enormous potential to make farming more efficient and far less resource-intensive. Computers have the ability to track and manage crops at a level of granularity that would be inconceivable for human workers. The Australian Centre for Field Robotics (ACFR) at the University of Sydney is focused on employing advanced agricultural robotics to help position Australia as a primary supplier of food for Asia's exploding population—in spite of the country's relative paucity of arable land and fresh water. ACFR envisions robots that continuously prowl fields taking soil samples around individual plants and then injecting just the right amount of water or fertilizer.[37] Precision application of fertilizer or pesticides to individual plants, or even to specific fruits growing on a tree, could potentially reduce the use of these chemicals by up to 80 percent, thereby dramatically decreasing the amount of toxic runoff that ultimately ends up fouling rivers, streams, and other bodies of water.[38] *

Agriculture in most developing countries is notoriously inefficient. The plots of land worked by families are often tiny, capital investment is minimal, and modern technology is unavailable. Even though farming techniques are labor-intensive, the land often has to support more people than are really necessary to cultivate it. As global population grows to 9 billion and beyond in the coming decades, there will be ever-increasing pressure to transition any and all available arable land into larger and more efficient farms that are capable of producing higher crop yields. Advancing agricultural technology will have a significant role to play, especially in countries where water is scarce and ecosystems have been damaged by overuse of chemicals. Increased mechanization, however, will also mean that the land will provide livelihoods for far fewer people. The historical norm has been for those excess workers to migrate to cities and industrial centers in search of factory work—but as we have seen, those factories are themselves going to be transformed by accelerating automation technology. In fact, it seems somewhat difficult to imagine how many developing countries will succeed in navigating these technological disruptions without running into significant unemployment crises.

---

* Precision agriculture—or the ability to keep track of and manage individual plants or even fruits—is part of the "big data" phenomenon, a subject that we'll examine in more depth in Chapter 4.

In the United States, agricultural robotics has the potential to eventually throw a wrench into many of the fundamental assumptions that underlie immigration policy—an area that is already subject to intensely polarized politics. The impact is already evident in some areas that used to employ large numbers of farmworkers. In California, machines skirt around the daunting visual challenge of picking individual almonds by simply grasping the entire tree and violently shaking it. The almonds fall to the ground where they'll be harvested by a different machine. Many California farmers have transitioned from delicate crops like tomatoes to more robust nuts because they can be harvested mechanically. Overall agricultural employment in California fell by about 11 percent in the first decade of the twenty-first century, even as the total production of crops like almonds, which are compatible with automated farming techniques, has exploded.[39]

As robotics and advanced self-service technologies are increasingly deployed across nearly every sector of the economy, they will primarily threaten lower-wage jobs that require modest levels of education and training. These jobs, however, currently make up the vast majority of the new positions being generated by the economy—and the US economy needs to create something on the order of a million jobs per year just to tread water in the face of population growth. Even if we set aside the possibility of an actual reduction in the number of these jobs as new technologies emerge, any decline in the rate at which they are created will have dire, cumulative consequences for employment over the long run.

Many economists and politicians might be inclined to dismiss this as a problem. After all, routine, low-wage, low-skill jobs—at least in advanced economies—tend to be viewed as inherently undesirable, and when economists discuss the impact of technology on these kinds of jobs, you are very likely to encounter the phrase "freed up"—as in, workers who lose their low-skill jobs will be freed up to pursue more training and better opportunities. The fundamental assumption, of course, is that a dynamic economy like the United States will always be capable of generating sufficient higher-wage, higher-skill jobs to absorb all those newly freed up workers—given that they succeed in acquiring the necessary training.

That assumption rests on increasingly shaky ground. In the next two chapters we'll look at the impact that automation has already had on jobs and incomes in the United States and consider the characteristics that set information technology apart as a uniquely disruptive force. That discussion will provide a jumping-off point from which to delve into an unfolding story that is poised to upend the conventional wisdom about the types of jobs most likely to be automated and the viability of ever more education and training as a solution: the machines are coming for the high-wage, high-skill jobs as well.

Chapter 2

IS THIS TIME DIFFERENT?

On the morning of Sunday, March 31, 1968, the Reverend Martin Luther King, Jr., stood in the elaborately carved limestone pulpit at Washington National Cathedral. The building—one of the largest churches in the world and over twice the size of London's Westminster abbey—was filled to capacity with thousands of people packed into the nave and transept, looking down from the choir loft, and squeezed into doorways. At least another thousand people gathered outside on the steps or at nearby St. Alban's Episcopal Church to hear the sermon over loudspeakers.

It would be Dr. King's final Sunday sermon. Just five days later the cathedral would again be overflowing with a far more somber crowd—including President Lyndon Johnson, senior cabinet officials, all nine Supreme Court justices, and leading members of Congress—gathered to honor King at a memorial service the day following his assassination in Memphis, Tennessee.[1]

The title of Dr. King's sermon that day was "Remaining Awake Through a Great Revolution." Civil and human rights were, as might be expected, a major component of his address, but he had in mind revolutionary change on a much broader front. As he explained a short way into his sermon:

> There can be no gainsaying of the fact that a great revolution is taking place in the world today. In a sense it is a triple revolution: that is, a technological revolution, with the impact of automation and cybernation; then there is a revolution in weaponry, with the emergence of atomic and nuclear weapons of warfare; then there is a human rights revolution, with the freedom explosion that is taking place all over the world. Yes, we do live in a period where changes are taking place. And there is still the voice crying through the vista of time saying, "Behold, I make all things new; former things are passed away."[2]

The phrase "triple revolution" referred to a report written by a group of prominent academics, journalists, and technologists that called itself the Ad Hoc Committee on the Triple Revolution. The group included Nobel laureate chemist Linus Pauling as well as economist Gunnar Myrdal, who would be awarded the Nobel Prize in economics, along with Friedrich Hayek, in 1974. Two of the revolutionary forces identified in the report—nuclear weapons and the civil rights movement—are indelibly woven into the historical narrative of the 1960s. The third revolution, which comprised the bulk of the document's text, has largely been forgotten. The report predicted that "cybernation" (or automation) would soon result in an economy where "potentially unlimited output can be achieved by systems of machines which will require little cooperation from human beings."[3] The result would be massive unemployment, soaring inequality, and, ultimately, falling demand for goods and services as consumers increasingly lacked the purchasing power necessary to continue driving economic growth. The Ad Hoc Committee went on to propose a radical solution: the eventual implementation of a guaranteed minimum income made possible by the "economy of abundance" such widespread automation could create, and which would "take the place of the patchwork of welfare measures" that were then in place to address poverty.*

The Triple Revolution report was released to the media and sent to President Johnson, the secretary of labor, and congressional leaders in March 1964. An accompanying cover letter warned ominously that if something akin to the report's proposed solutions was not implemented, "the nation will be thrown into unprecedented economic and social disorder." A front-page story with extensive quotations from the report appeared in the next day's *New York Times,* and numerous other newspapers and magazines ran stories and editorials (most of which were critical), in some cases even printing the entire text of the report.[4]

The Triple Revolution marked what was perhaps the crest of a wave of worry about the impact of automation that had arisen following World War II. The specter of mass joblessness as machines displaced workers had incited fear many times in the past—going all the way back to Britain's Luddite uprising in 1812—but in the 1950s and '60s, the concern was especially acute and was articulated by some of the United States' most prominent and intellectually capable individuals.

In 1949, at the request of the *New York Times,* Norbert Wiener, an internationally renowned mathematician at the Massachusetts Institute of Technology, wrote an article describing his vision for the future of computers and automation.[5] Wiener had been a child prodigy who entered college at age eleven and completed his PhD

---

* The Committee on the Triple Revolution did not advocate the immediate implementation of a guaranteed income. Instead, it proposed a list of nine transitional policies. Many of these were quite conventional, and included things such as greatly increased investment in education, public works projects to create jobs, and the construction of low-cost housing. The report also argued for a greatly expanded role for unions and suggested that organized labor should become an advocate for the unemployed as well as those who held jobs.

when he was just seventeen; he went on to establish the field of cybernetics and made substantial contributions in applied mathematics and to the foundations of computer science, robotics, and computer-controlled automation. In his article—written just three years after the first true general-purpose electronic computer was built at the University of Pennsylvania*—Wiener argued that "if we can do anything in a clear and intelligible way, we can do it by machine" and warned that that this could ultimately lead to "an industrial revolution of unmitigated cruelty" powered by machines capable of "reducing the economic value of the routine factory employee to a point at which he is not worth hiring at any price."**

Three years later, a dystopian future much like the one Wiener had imagined was brought to life in the pages of Kurt Vonnegut's first novel. *Player Piano* described an automated economy in which industrial machines managed by a tiny technical elite did virtually all the work, while the vast majority of the population faced a meaningless existence and a hopeless future. Vonnegut, who went on to achieve legendary status as an author, continued to believe in the relevance of his 1952 novel throughout his life, writing decades later that it was becoming "more timely with each passing day."[6]

Four months after the Johnson administration received the Triple Revolution report, the president signed a bill creating the National Commission on Technology, Automation, and Economic Progress.[7] In his remarks at the bill's signing ceremony, Johnson said that "automation can be the ally of our prosperity if we will just

---

* ENIAC (Electronic Numerical Integrator and Computer) was built at the University of Pennsylvania in 1946. A true programmable computer, it was financed by the US Army and intended primarily for calculating firing tables used to aim artillery.

** Due to a miscommunication, Wiener's article was never published in 1949. A draft copy was discovered by a researcher working with documents in the MIT library archives in 2012, and substantial excerpts were finally published in a May 2013 article by *New York Times* science reporter John Markoff.

look ahead, if we will understand what is to come, and if we will set our course wisely after proper planning for the future." The newly formed commission then—as is almost universally the case with such commissions—quickly faded into obscurity, leaving behind at least three book-length reports of its own.[8]

The irony of all the automation worries in the postwar period was that the economy offered very little in the way of evidence to support such concerns. When the Triple Revolution report was released in 1964, the unemployment rate was just over 5 percent, and it would fall to a low of 3.5 percent by 1969. Even during the four recessions that occurred between 1948 and 1969, unemployment never reached 7 percent, and then it fell rapidly once recovery was under way.[9] The introduction of new technologies did drive substantial increases in productivity, but the lion's share of that growth was captured by workers in the form of higher wages.

By the early 1970s, focus had shifted to the OPEC oil embargo, and then to the subsequent years of stagflation. The potential for machines and computers to cause unemployment was pushed further and further out of the mainstream. Among professional economists in particular, the idea became virtually untouchable. Those who did dare to entertain such thoughts risked being labeled a "neo-Luddite."

Given that the dire circumstances predicted by the Triple Revolution report did not come to pass, we can ask an obvious question: Were the authors of the report definitively wrong? Or did they—like many others before them—simply sound the alarm far too soon?

Norbert Wiener, as one of the early pioneers of information technology, perceived the digital computer as being fundamentally different from the mechanical technologies that preceded it. It was a game changer: a new kind of machine with the potential to usher in a new age—and, ultimately, perhaps rend the very fabric of society. Yet, Wiener's views were expressed at a time when computers were room-sized monstrosities whose calculations were powered by tens of thousands of searingly hot radio vacuum tubes, some number of which could be expected to fail on a near daily basis.[10] It would be decades before the exponential arc of progress would drive digital technology to a level where such views might reasonably be justified.

Those decades are now behind us, and the time is ripe for an open-minded reassessment of the impact of technology on the economy. The data shows that even as concerns about the impact of labor-saving technology receded to the fringes of economic thought, something that had been fundamental to the postwar era of prosperity gradually began to change in the American economy. The nearly perfect historical correlation between increasing productivity and rising incomes broke down: wages for most Americans stagnated and, for many workers, even declined; income inequality soared to levels not seen since the eve of the 1929 stock market crash; and a new phrase—"jobless recovery"—found a prominent place in our vocabulary. In all, we can enumerate at least seven economic trends that, taken together, suggest a transformative role for advancing information technology.

Seven Deadly Trends

Stagnant Wages

The year 1973 was an eventful one in the history of the United States. The Nixon administration was embroiled in the Watergate scandal, and in October, OPEC initiated an oil embargo that would soon result in long lines of angry motorists at gas stations across the country. Even as Nixon descended into his death spiral, however, there was another story unfolding. This story began with an event that went completely unheralded and yet marked the beginning of a trend that would arguably dwarf both Watergate and the oil crisis in importance. For that was the year a typical American worker's pay reached its peak. Measured in 2013 dollars, a typical worker—that is, production and nonsupervisory workers in the private sector, representing well over half the American workforce—earned about $767 per week in

1973. The following year, real average wages began a precipitous decline from which they would never fully recover. A full four decades later, a similar worker earns just $664, a decline of about 13 percent.[11]

The story is modestly better if we look at median household incomes. Between 1949 and 1973, US median household incomes roughly doubled, from about $25,000 to $50,000. Growth in median incomes during this period tracked nearly perfectly with per capita GDP. Three decades later, median household income had increased to about $61,000, an increase of just 22 percent. That growth, however, was driven largely by the entry of women into the workforce. If incomes had moved in lockstep with economic growth—as was the case prior to 1973—the median household would today be earning well in excess of $90,000, over 50 percent more than the $61,000 they do earn.[12]

Figure 2.1 shows the relationship between labor productivity* (which measures the value of workers' hourly output) and compensation (which includes wages and benefits) paid to ordinary private sector workers from 1948 onward. The first segment of the graph (from 1948 to 1973) shows the way economists expect things to work. Growth in productivity moves in almost perfect lockstep with compensation. Prosperity marches upward and is shared broadly by all those who contribute to the economy. Beyond the mid-1970s, the widening gap between the two lines is a graphic illustration of the extent to which the fruits of innovation throughout the economy are now accruing almost entirely to business owners and investors, rather than to workers.

---

* Labor productivity measures the value of the output (either goods or services) produced by workers per hour. It is a critically important gauge of the general efficiency of an economy; to a significant extent it determines the wealth of a nation. Advanced, industrialized countries have high productivity because their workers have access to more and better technology, enjoy better nutrition as well as safer and more healthful environments, and are generally better educated and trained. Poor countries lack these things and are, therefore, less productive; their people must work longer and harder to produce the same level of output.

Figure 2.1. Growth of Real Hourly Compensation for Production and Nonsupervisory Workers Versus Productivity (1948–2011)

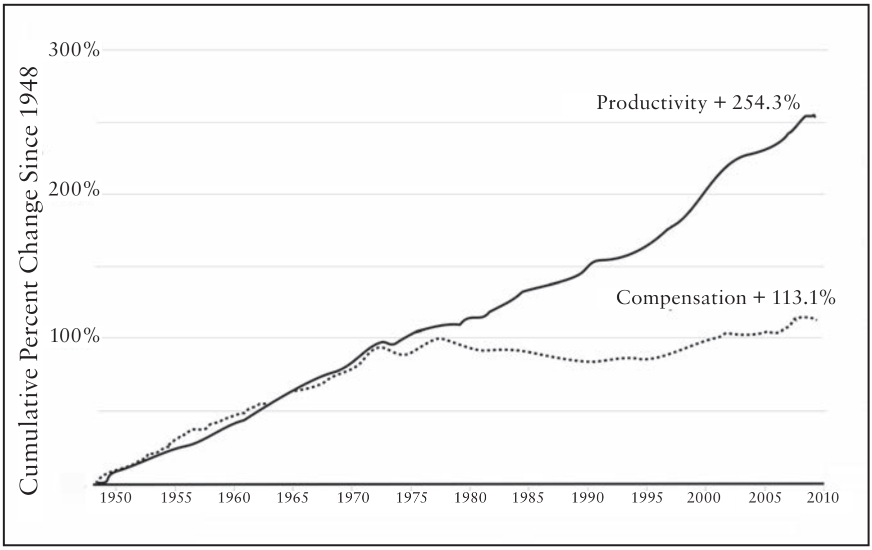

Source: Lawrence Mishel, Economic Policy Institute, based on an analysis of unpublished total economy data from the Bureau of Labor Statistics, the Labor Productivity and Costs program, and the Bureau of Economic Analysis's National Income and Product Accounts public data series.[13]

Despite the clarity of this graph, many economists have still not fully acknowledged the divergence between wage and productivity growth. Figure 2.2 shows how growth rates for compensation and productivity compare during different periods going back to 1947. Productivity has significantly outstripped compensation in every decade from 1980 on. The difference is especially dramatic from 2000 to 2009; although productivity growth nearly matches the 1947–1973 period—the golden era of postwar prosperity—compensation lags far behind. It's difficult to look at this graph and not come away with the impression that productivity growth is pretty clearly blowing the doors off the raises that most workers are getting.

The authors of most college economics textbooks have been especially slow to acknowledge this picture. Consider, for example,

Figure 2.2. Productivity Growth Versus Compensation Growth

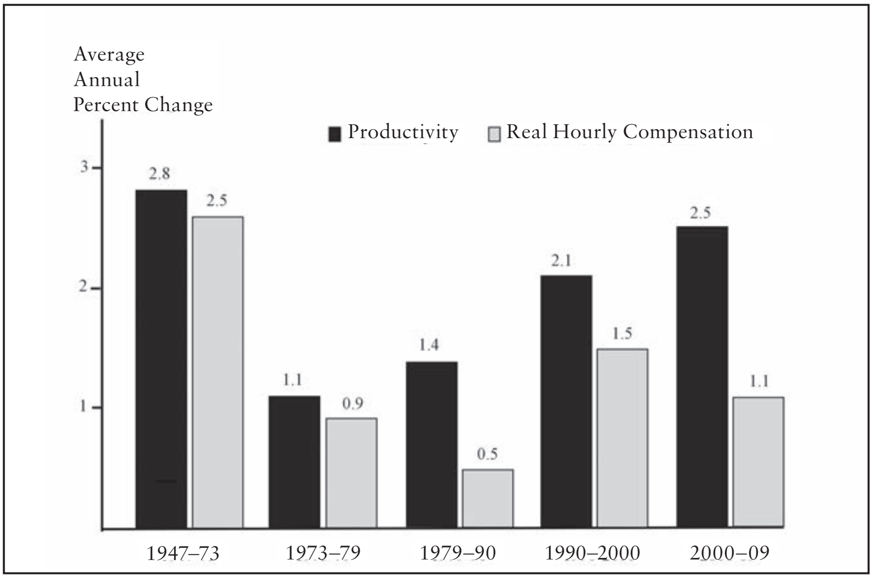

SOURCE: US Bureau of Labor Statistics.[14]

Principles of Economics, an introductory textbook authored by John B. Taylor and Akila Weerapana,[15] the required text for Professor Taylor's wildly popular introductory economics class at Stanford University. It includes a bar chart very similar to Figure 2.2, but still argues for a tight relationship between wages and productivity. What about the fact that productivity leaps away from wages beginning in the 1980s? Taylor and Weerapana note that "the relationship is not perfect." That appears to be something of an understatement. The 2007 edition of another textbook, also titled *Principles of Economics*[16] and co-authored by Princeton professor—and former Federal Reserve chairman—Ben Bernanke, suggests that slow wage growth from 2000 on may have resulted from "the weak labor market that followed the recession of 2001" and that wages ought to "catch up to productivity growth as the labor market returns to normal"—a view that seems to ignore the fact that the tight correlation between wage and productivity growth began to deteriorate long before today's college students were born.*

A Bear Market for Labor's Share, and a Raging Bull for Corporations

Early in the twentieth century, the British economist and statistician Arthur Bowley delved into decades of national income data for the United Kingdom and showed that the fraction of national income going to labor and capital respectively remained relatively constant, at least over long periods. This apparently fixed relationship ultimately became an accepted economic principle known as "Bowley's Law." John Maynard Keynes, perhaps the most famous economist of all time, would later say that Bowley's Law was "one of the most surprising, yet best established facts in the whole range of economic statistics."[17]

As Figure 2.3 shows, during the postwar period, the share of US national income going to labor moved in a fairly tight range, just as

---

* There is also a technical issue that comes into play when discussing the gap between wage growth and productivity growth. Both the wage (or, more broadly, compensation) and productivity numbers must be adjusted for inflation. The standard way to do this, and the method used by the US Bureau of Labor Statistics (BLS), is to use two different measures of inflation. Wages are adjusted using the Consumer Price Index (CPI) because this reflects the prices of products and services that workers actually spend their money on. The productivity figures are adjusted using the GDP deflator (or implicit price deflator), which is a broader measure of inflation in the entire economy. In other words, the GDP deflator incorporates prices for a lot of things that consumers don't actually purchase. One especially important difference is that computers and information technology—which have seen substantial price deflation due to Moore's Law—are much more important in the GDP deflator than in the CPI (computers are not a big component of most household budgets, but are purchased in volume by businesses). Some economists—particularly those who are more conservative—argue that the GDP deflator should be used for *both* wages and productivity. When this method is used, the gap between wage growth and productivity growth narrows significantly. However, this approach almost certainly understates the level of inflation that impacts wage earners.

Figure 2.3. US Labor's Share of National Income (1947–2014)

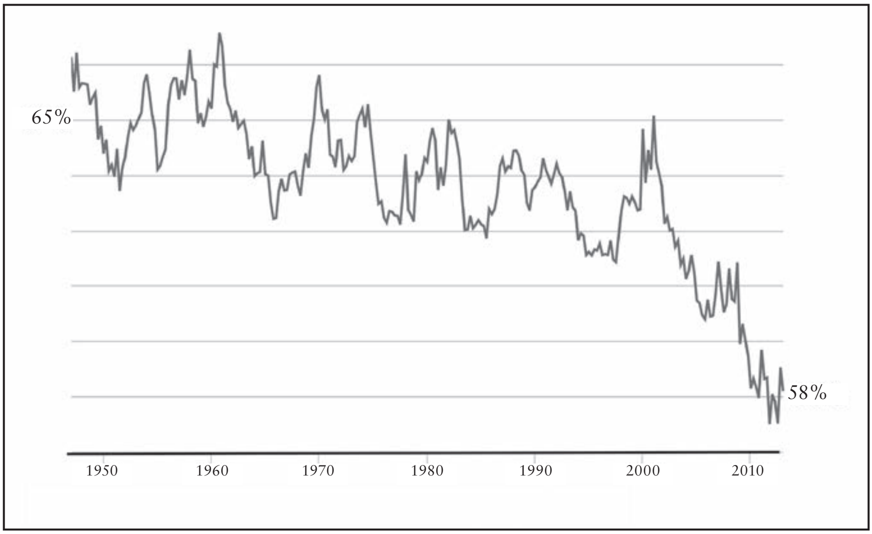

SOURCE: US Bureau of Labor Statistics and Federal Reserve Bank of St. Louis (FRED).[18]

Bowley's Law would have predicted. From the mid-1970s on, however, Bowley's Law began to fall apart as labor's share went first into a gradual decline and then into a seeming free fall just after the turn of the century. The decline is all the more remarkable when we consider that labor's share includes anyone who draws a paycheck. In other words, the enormous salaries of CEOs, Wall Street executives, superstar athletes, and movie stars are all considered labor, and those, of course, haven't been declining at all: they've been skyrocketing. A graph showing the share of national income accruing to ordinary workers—or, more broadly, the bottom 99 percent of the income distribution—would certainly show an even more precipitous plunge.

While labor's share of income plummeted, the story was very different for corporate profits. In April 2012, the *Wall Street Journal* ran a story entitled "For Big Companies, Life Is Good" that documented the astonishing speed at which corporations recovered from the most severe economic crisis since the Great Depression. While millions of workers remained unemployed or accepted jobs at lower pay or with fewer hours, the corporate sector emerged from the downturn "more productive, more profitable, flush with cash and less burdened by debt."[19] Over the course of the Great Recession, corporations had become adept at producing more with fewer workers. In 2011, big companies generated an average of $420,000 in revenue for each employee, an increase of more than 11 percent over the 2007 figure of $378,000.[20] Spending on new plants and equipment, including information technology, by S&P 500 companies had doubled from the year before, bringing capital investment as a percentage of revenue back to pre-crisis levels.

Corporate profits as a percentage of the total economy (GDP) also skyrocketed after the Great Recession (see Figure 2.4). Notice that despite the precipitous plunge in profits during the 2008–2009 economic crisis, the speed at which profitability recovered was unprecedented compared with previous recessions.

Figure 2.4. Corporate Profits as a Percentage of GDP

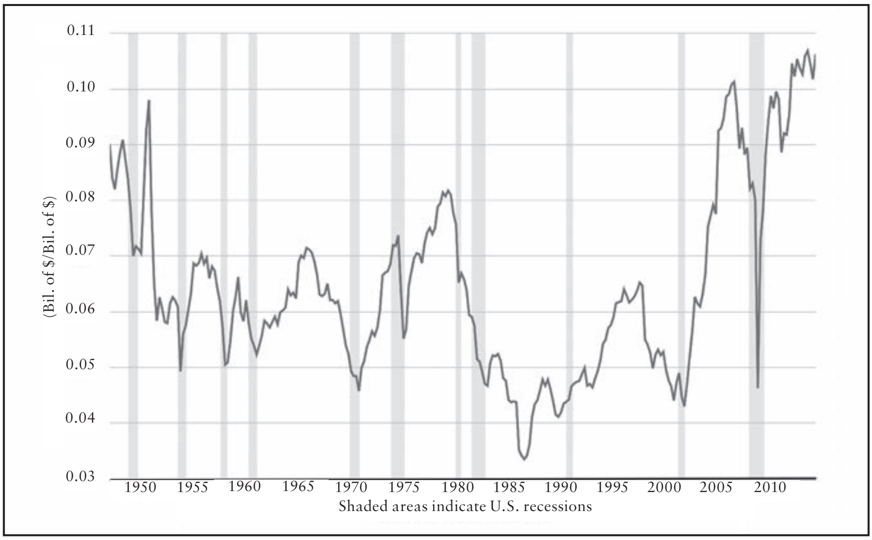

SOURCE: Federal Reserve Bank of St. Louis (FRED).[21]

The decline in labor's share of national income is by no means limited to the United States. In a June 2013 research paper,[22] economists Loukas Karabarbounis and Brent Neiman, both of the University of Chicago's Booth School of Business, analyzed data from fifty-six different countries and found that thirty-eight demonstrated a significant decline in labor's share. In fact, the authors' research showed that Japan, Canada, France, Italy, Germany, and China all had larger declines than the United States over a ten-year period. The decline in labor's share in China—the country that most of us assume is hoovering up all the work—was especially precipitous, falling at three times the rate in the United States.

Karabarbounis and Neiman concluded that these global declines in labor's share resulted from "efficiency gains in capital producing sectors, often attributed to advances in information technology and the computer age."[23] The authors also noted that a stable labor share of income continues to be "a fundamental feature of macroeconomic models."[24] In other words, just as economists do not seem to have fully assimilated the implications of the circa-1973 divergence of productivity and wage growth, they are apparently still quite happy to build Bowley's Law into the equations they use to model the economy.

Declining Labor Force Participation

A separate trend has been the decline in labor force participation. In the wake of the 2008–9 economic crisis, it was often the case that the unemployment rate fell not because large numbers of new jobs were being created, but because discouraged workers exited the workforce. Unlike the unemployment rate, which counts only those people actively seeking jobs, labor-force participation offers a graphic illustration that captures workers who have given up.

As Figure 2.5 shows, the labor force participation rate rose sharply between 1970 and 1990 as women flooded into the workforce. The overall trend disguises the crucial fact that the percentage

Figure 2.5. Labor Force Participation Rate

SOURCE: US Bureau of Labor Statistics and Federal Reserve Bank of St. Louis (FRED).[25]

of men in the labor force has been in consistent decline since 1950, falling from a high of about 86 percent to 70 percent as of 2013. The participation rate for women peaked at 60 percent in 2000; the overall labor force participation rate peaked at about 67 percent that same year.[26]

Labor force participation has been falling ever since, and although this is due in part to the retirement of the baby boom generation, and in part because younger workers are pursuing more education, those demographic trends do not fully explain the decline. The labor force participation rate for adults between the ages of twenty-five and fifty-four—those old enough to have completed college and even graduate school, yet too young to retire—has declined from about 84.5 percent in 2000 to just over 81 percent in 2013.[27] In other words, both the overall labor force participation rate and the participation rate for prime working-age adults have fallen by about three percentage points since 2000—and about half of that decline came before the onset of the 2008 financial crisis.

The decline in labor force participation has been accompanied by an explosion in applications for the Social Security disability program, which is intended to provide a safety net for workers who suffer debilitating injuries. Between 2000 and 2011, the number of applications more than doubled, from about 1.2 million per year to nearly 3 million per year.[28] As there is no evidence of an epidemic of workplace injuries beginning around the turn of the century, many analysts suspect that the disability program is being misused as a kind of last-resort—and permanent—unemployment insurance program. Given all this, it seems clear that something beyond simple demographics or cyclical economic factors is driving people out of the labor force.

Diminishing Job Creation, Lengthening Jobless Recoveries, and Soaring Long-Term Unemployment

Over the past half-century, the US economy has become progressively less effective at creating new jobs. Only the 1990s managed to—just barely—keep up with the previous decade's job growth, and that was largely due to the technology boom that occurred in the second half of the decade. The recession that began in December 2007 and the ensuing financial crisis were a total disaster for job creation in the 2000s; the decade ended with virtually the same number of jobs that had existed in December 1999. Even before the Great Recession hit, however, the new century's first decade was already on track to produce by far the worst percentage growth in employment since World War II.

As Figure 2.6 shows, the number of jobs in the economy had increased by only about 5.8 percent through the end of 2007. Prorating that number for the entire decade suggests that, if the economic crisis had not occurred, the 2000s would likely have finished with a roughly 8 percent job creation rate—less than half of the percentage increase seen in the 1980s and '90s.

That miserable job creation performance is especially disturbing in light of the fact that the economy needs to generate large numbers

Figure 2.6. US Job Creation by Decade

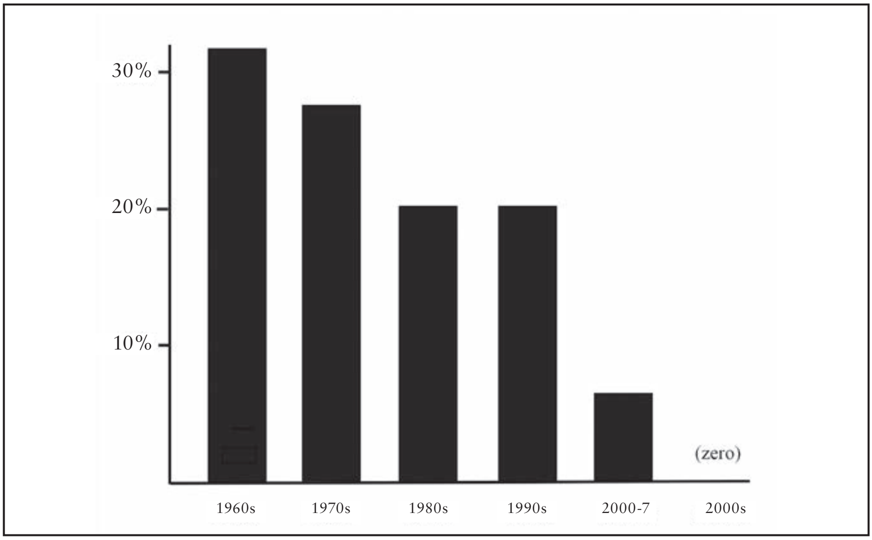

SOURCE: US Bureau of Labor Statistics and Federal Reserve Bank of St. Louis (FRED).[29]

of new jobs—between 75,000 and 150,000 per month, depending on one's assumptions—just to keep up with population growth.[30] Even when the lower estimate is employed, the 2000s still resulted in a deficit of about 9 million jobs over the course of the decade.

Clear evidence also shows that when a recession knocks the wind out of the economy, it is taking longer and longer for the job market to recover. Temporary layoffs have given way to jobless recoveries. A 2010 research report by the Federal Reserve Bank of Cleveland found that recent recessions have seen a dramatic decline in the rate at which unemployed workers are able to land new jobs. In other words, the problem is not that more jobs are being destroyed in downturns; it is that fewer are being created during recoveries. After the onset of the Great Recession in December 2007, the unemployment rate continued to rise for nearly two years, ultimately increasing by a full five percentage points and peaking at 10.1 percent. The Cleveland Fed's analysis found that the increased difficulty faced by workers finding new jobs accounted for over 95 percent of that 5 percent jump in the unemployment rate.[31] This, in turn, has led to a huge jump in the long-term unemployment rate, which peaked in 2010, when about 45 percent of workers had been out of work for more than six months.[32] Figure 2.7 shows the number of months it took for the labor market to recover from recent recessions. The Great Recession resulted in a monstrous jobless recovery; it took until May 2014—a full six and a half years after the start of the downturn—for employment to return to its pre-recession level.

Extended unemployment is a debilitating problem. Job skills erode over time; the risk that workers will become discouraged increases, and many employers seem to actively discriminate against the long-term unemployed, often refusing even to consider their résumés. Indeed, a field experiment conducted by Rand Ghayad, a PhD candidate in economics at Northeastern University, showed

Figure 2.7. US Recessions: Months for Employment to Recover (Measured from Start of Recession)

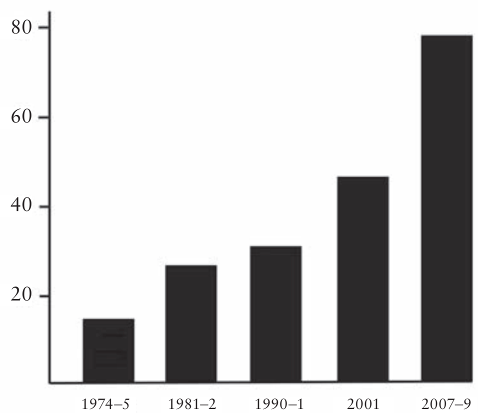

Source: US Bureau of Labor Statistics and Federal Reserve Bank of St. Louis (FRED).[33]

that a recently unemployed applicant with no industry experience was actually more likely to be called in for a job interview than someone with directly applicable experience who had been out of work for more than six months.[34] A separate report by the Urban Institute found that the long-term unemployed are not appreciably different from other workers, suggesting that becoming one of the long-term unemployed—and suffering the stigma that attaches to that category—may largely be a matter of bad luck.[35] If you happen to lose your job at an especially unfavorable time and then fail to find a new position before the dreaded six-month mark (a real possibility if the economy is in free fall), your prospects diminish dramatically from that point on—regardless of how qualified you may be.

Soaring Inequality

The divide between the rich and everyone else has been growing steadily since the 1970s. Between 1993 and 2010 over half of the increase in US national income went to households in the top 1 percent of the income distribution.[36] Since then, things have only gotten worse. In an analysis published in September 2013, economist Emmanuel Saez of the University of California, Berkeley, found that an astonishing 95 percent of total income gains during the years 2009 to 2012 were hoovered up by the wealthiest 1 percent.[37] Even as the Occupy Wall Street movement has faded from the scene, the evidence shows pretty clearly that income inequality in the United States is not just high—it may well be accelerating.

While inequality has been increasing in nearly all industrialized countries, the United States remains a clear outlier. According to the Central Intelligence Agency's analysis, income inequality in America is roughly on a par with that of the Philippines and significantly exceeds that of Egypt, Yemen, and Tunisia.[38] Studies have also found that economic mobility, a measure of the likelihood that the children of the poor will succeed in moving up the income scale, is significantly lower in the United States than in nearly all European nations.

In other words, one of the most fundamental ideas woven into the American ethos—the belief that anyone can get ahead through hard work and perseverance—really has little basis in statistical reality.

From the perspective of any one individual, inequality can be very difficult to perceive. Most people tend to focus their attention locally. They worry about how they are doing relative to the guy next door as opposed to the hedge fund manager they will, in all likelihood, never encounter. Surveys have shown that most Americans vastly underestimate the existing extent of inequality, and when asked to select an "ideal" national distribution of income, they make a choice that, in the real world, exists only in Scandinavian social democracies.[39] *

Nonetheless, inequality has real implications that go far beyond simple frustration about your inability to keep up with the Joneses. Foremost is the fact that the overwhelming success of those at the extreme top seems to be correlated with diminishing prospects for nearly everyone else. The old adage that a rising tide lifts all boats gets pretty tired when you haven't had a meaningful raise since the Nixon administration.

There is also an obvious risk of political capture by the financial elite. In the United States, to a greater degree than in any other advanced democracy, politics is driven almost entirely by money. Wealthy individuals and the organizations they control can mold government policy through political contributions and lobbying, often producing outcomes that are clearly at odds with what the public actually wants. As those at the apex of the income distribution become increasingly detached—living in a kind of bubble that insulates them almost entirely from the realities faced by typical Americans—there

---

* This is true regardless of political party. In one study conducted by Dan Ariely of Duke University, over 90 percent of Republicans and 93 percent of Democrats preferred an income distribution similar to that of Sweden over that of the United States.

is a real risk that they will be unwilling to support investment in the public goods and infrastructure upon which everyone else depends.

The soaring fortunes of those at the very top may ultimately represent a threat to democratic governance. However, the most immediate problem for most middle- and working-class people is that job market opportunities are broadly deteriorating.

Declining Incomes and Underemployment for Recent College Graduates

A four-year college degree has come to be almost universally viewed as an essential credential for entry into the middle class. As of 2012, average hourly wages for college graduates were more than 80 percent higher than the wages of high school graduates.[40] The college wage premium is a reflection of what economists call "skill biased technological change" (SBTC).* The general idea behind SBTC is that information technology has automated or deskilled much of the work handled by less educated workers, while simultaneously increasing the relative value of the more cognitively complex tasks typically performed by college graduates.

Graduate and professional degrees convey still higher incomes, and in fact, since the turn of the century, things are looking quite a bit less rosy for young college graduates who don't also have an advanced degree. According to one analysis, incomes for young workers with only a bachelor's degree declined nearly 15 percent between 2000 and 2010, and the plunge began well before the onset of the 2008 financial crisis.

---

* SBTC and the college wage premium offer a partial explanation for increasing income inequality. However, since nearly a third of the adult US population has a college degree, if this were the only thing going on, it would imply a much tamer form of inequality than actually exists. The real action is at the very top—and things become more extreme the higher you go. The outsized fortunes of the top 1 (or .01) percent cannot reasonably be attributed to better education or training.

Recent college graduates are also underemployed. By some accounts, fully half of new graduates are unable to find jobs that utilize their education and offer access to the crucial initial rung on the career ladder. Many of these unlucky graduates will probably find it very difficult to move up into solid middle-class trajectories.

To be sure, college graduates have, on average, maintained their income premium over workers with only a high school education, but this is largely because the prospects for these less educated workers have become genuinely dismal. As of July 2013, fewer than half of American workers who were between the ages of twenty and twenty-four and not enrolled in school had full-time jobs. Among non-students aged sixteen to nineteen only about 15 percent were working full-time.[41] The return on investment for a college education may be falling, but it still nearly always beats the alternative.

Polarization and Part-Time Jobs

A further new problem is that the jobs being created during economic recoveries are generally worse than those destroyed by recessions. In a 2012 study, economists Nir Jaimovich and Henry E. Siu analyzed data from recent US recessions and found that the jobs mostly likely to permanently disappear are the good middle-class jobs, while the jobs that tend to get created during recoveries are largely concentrated in low-wage sectors like retail, hospitality, and food preparation and, to a lesser extent, in high-skill professions that require extensive training.[42] This has been especially true over the course of the recovery that began in 2009.[43]

Many of these new low-wage jobs are also part-time. Between the start of the Great Recession in December 2007 and August 2013, about 5 million full-time jobs were vaporized, but the number of part-time jobs actually increased by approximately 3 million.[44] That increase in part-time work has occurred entirely among workers who have had their hours cut or who would like a full-time job but are unable to find one.

The propensity for the economy to wipe out solid middle-skill, middle-class jobs, and then to replace them with a combination of low-wage service jobs and high-skill, professional jobs that are generally unattainable for most of the workforce, has been dubbed "job market polarization." Occupational polarization has resulted in an hourglass-shaped job market where workers who are unable to land one of the desirable jobs at the top end up at the bottom.

This polarization phenomenon has been studied extensively by David Autor, an economist at the Massachusetts Institute of Technology. In a 2010 paper, Autor identifies four specific mid-range occupational categories that have been especially hard-hit as polarization has unfolded: sales, office/administrative, production/craft/repair, and operators/fabricators/laborers. Over the thirty years between 1979 and 2009, the percentage of the US workforce employed in these four areas declined from 57.3 percent to 45.7 percent, and there was a noticeable acceleration in the rate of job destruction between 2007 and 2009.[45] Autor's paper also makes it clear that polarization is not limited to the United States, but has been documented in most advanced, industrial economies; in particular, sixteen countries within the European Union have seen a significant decline in the percentage of the workforce engaged in mid-range occupations over the thirteen years between 1993 and 2006.[46]

Autor concludes that the primary driving forces behind job market polarization are "the automation of routine work and, to a smaller extent, the international integration of labor markets through trade and, more recently, offshoring."[47] In their more recent paper showing the relationship between polarization and jobless recoveries, Jaimovich and Siu point out that fully 92 percent of the job losses in mid-range occupations have occurred within a year of a recession.[48] In other words, polarization is not necessarily something that happens according to a grand plan, nor is it a gradual and continuous evolution. Rather, it is an organic process that is deeply intertwined with the business cycle; routine jobs are eliminated for economic reasons during a recession, but organizations then discover that ever-advancing information technology allows them to operate successfully without rehiring the workers once a recovery gets under way. Chrystia Freeland of Reuters puts it especially aptly, writing that "the middle-class frog isn't being gradually boiled; it is being periodically grilled at a very high heat."[49]

A Technology Narrative

It's fairly easy to piece together a hypothetical narrative that puts advancing technology—and the resulting automation of routine work—front and center as the explanation for these seven deadly economic trends. The golden era from 1947 to 1973 was characterized by significant technological progress and strong productivity growth. This was before the age of information technology; the innovations during this period were primarily in areas like mechanical, chemical, and aerospace engineering. Think, for example, of how airplanes evolved from employing internal combustion engines driving propellers to much more reliable and better-performing jet engines. This period exemplified what is written in all those economics textbooks: innovation and soaring productivity made workers more valuable—and allowed them to command higher wages.

In the 1970s, the economy received a major shock from the oil crisis and entered an unprecedented period of high unemployment combined with high inflation. Productivity fell dramatically. The rate of innovation also plateaued as continued technological progress in many areas became more difficult. Jet aircraft changed very little. Both Apple and Microsoft were founded during this period, but the full impact of information technology was still far in the future.

The 1980s saw increased innovation, but it became more focused in the information technology sector. This type of innovation had a different impact on workers; for those with the right skill set, computers increased their value, just as the innovations in the postwar era had done for nearly everyone. For many other workers, however, computers had a less positive effect. Some types of jobs began to be either destroyed entirely or deskilled, making workers less valuable—at least until they were able to retrain for jobs that leveraged computer technology. As information technology gained in importance, labor's share of income gradually began to decline. Jet aircraft remained largely unchanged from the 1970s but increasingly used computers in their instrumentation and controls.

The 1990s saw IT innovation accelerate even more, and the Internet took off in the second half of the decade. The trends that began in the 1980s continued, but the decade also saw the tech bubble and the creation of millions of new jobs, especially in the IT sector. These were good jobs that often involved administering the computers and networks that were rapidly becoming critical to businesses of all sizes. As a result, wages did better in this period, but still fell well short of productivity growth. Innovation was centered even more on IT. The recession of 1990–1991 was followed by a jobless recovery as workers, many of whom had lost good mid-range jobs, struggled to find new positions. The job market gradually became more polarized. Jet aircraft were still essentially similar to the designs of the 1970s; however, they now had "fly by wire" systems, in which computers moved the control surfaces in response to the pilots' inputs, as well as increased flight automation.

In the years following 2000, information technology continued its acceleration and productivity rose as businesses got better at taking full advantage of all the new innovations. Many of those good jobs created in the 1990s began to disappear as corporations automated or offshored jobs, or began to outsource their IT departments to centralized "cloud" computing services. Throughout the economy, computers and machines were increasingly replacing workers rather than making them more valuable, and wage increases fell far short of growth in productivity. Both the share of national income going to labor and the labor force participation rate declined dramatically.

The job market continued to polarize, and jobless recoveries became the norm. Jet aircraft still used the same basic designs and propulsion systems as in the 1970s, but computer-aided design and simulation had resulted in many incremental improvements in areas such as fuel efficiency. The information technology incorporated into aircraft became even more sophisticated and routinely included full-flight automation, which allowed the planes to take off, fly to a destination, and then land—all without human intervention.

Now, you may quite rightfully object to that story as being overly simplistic—or perhaps even completely wrong. After all, wasn't it really globalization, or maybe Reaganomics, that led to all our problems? As I said, this was intended to be a hypothetical narrative: a simple story to help clarify the argument for the importance of technology in these seven documented economic trends. Each of these trends has been studied by teams of economists and others who have attempted to discover the underlying causes, and technology has often been implicated as a contributing, if not always the primary, factor. However, it is when all seven trends are considered together that the argument for advancing information technology as a disruptive economic force is most compelling.

Aside from advancing information technology, there are three other primary possibilities that might conceivably have contributed to all, or at least most, of our seven economic trends: globalization, the growth of the financial sector, and politics (in which I include factors like deregulation and the decline of organized labor).

Globalization

That globalization has had a dramatic impact on certain industries and regions is undeniable—just look at America's rustbelt. But globalization, and in particular trade with China, alone could not have caused wages for most American workers to stagnate over four decades.

First, global trade directly impacts workers who are employed in the tradable sector—in other words, in industries that produce goods or services that can be transported to other locations. The vast majority of American workers now work in nontradable areas like government, education, health care, food services, and retail. For the most part, these people are not directly competing with overseas workers, so globalization is not driving down their wages.

Second, although it may appear that virtually everything sold at Walmart is made in China, most American consumer spending stays in the United States. A 2011 analysis by Galina Hale and Bart Hobijn, two economists at the Federal Reserve Bank of San Francisco, found that 82 percent of the goods and services Americans purchase are produced entirely in the United States; this is largely because we spend the vast majority of our money on nontradable services. The total value of imports from China amounted to less than 3 percent of US consumer spending.[50]

It is undoubtedly true that, as Figure 2.8 shows, the fraction of American workers employed in manufacturing has fallen

Figure 2.8. Percentage of US Workers in Manufacturing

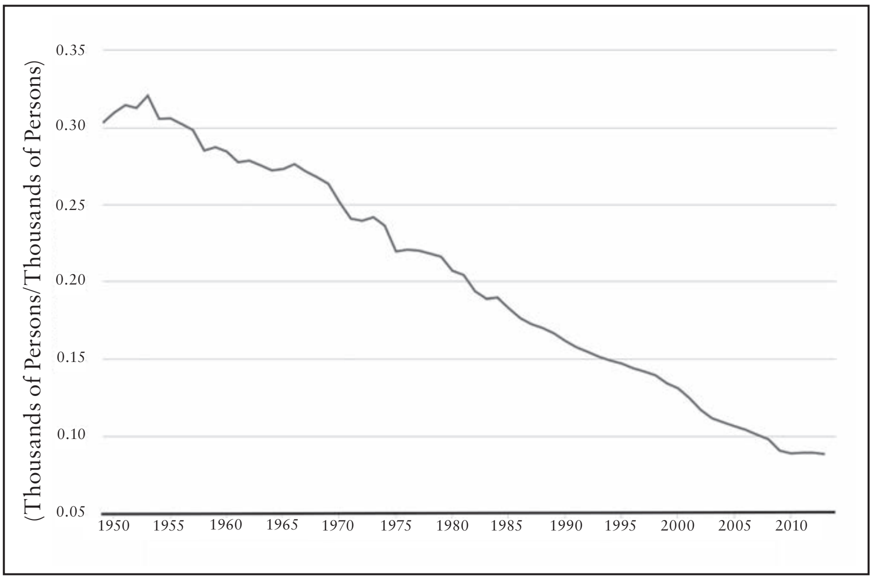

SOURCE: US Bureau of Labor Statistics and Federal Reserve Bank of St. Louis (FRED).[51]

dramatically since the early 1950s. This trend began decades before enactment of the North American Free Trade Agreement (NAFTA) in the 1990s and the rise of China in the 2000s. In fact, the decline seems to have halted at the end of the Great Recession as manufacturing employment has actually outperformed the job market as a whole.

A potent force has been very consistently eliminating jobs in the manufacturing sector. That force is advancing technology. Even as the number of manufacturing jobs has been steadily declining as a percentage of total employment, the inflation-adjusted value of the goods manufactured in the United States has dramatically increased over time. We are making more stuff, but doing so with fewer and fewer workers.

Financialization

In 1950, the US financial sector represented about 2.8 percent of the overall economy. By 2011 finance-related activity had grown more than threefold to about 8.7 percent of GDP. The compensation paid to workers in the financial sector has also exploded over the past three decades, and is now about 70 percent more than the average for other industries.[52] The assets held by banks have ballooned from about 55 percent of GDP in 1980 to 95 percent in 2000, while the profits generated in the financial sector have more than doubled from an average of about 13 percent of all corporate profits in the 1978–1997 timeframe to 30 percent in the period between 1998 and 2007.[53] No matter how you choose to measure it, finance has grown dramatically as a share of economic activity in the United States and, to a somewhat less spectacular degree, in nearly all industrialized countries.

The primary complaint leveled against the financialization of the economy is that much of this activity is geared toward rent seeking. In other words, the financial sector is not creating real value or adding to the overall welfare of society; it is simply finding ever more creative ways to siphon profits and wealth from elsewhere in the economy. Perhaps the most colorful articulation of this accusation came from *Rolling Stone*'s Matt Taibbi in his July 2009 takedown of Goldman Sachs that famously labeled the Wall Street firm "a great vampire squid wrapped around the face of humanity, relentlessly jamming its blood funnel into anything that smells like money."[54]

Economists who have studied financialization have found a strong correlation between the growth of the financial sector and inequality as well as the decline in labor's share of national income.[55] Since the financial sector is, in effect, imposing a kind of tax on the rest of the economy and then reallocating the proceeds to the top of the income distribution, it's reasonable to conclude that it has played a role in a number of the trends we've looked at. Still, it seems hard to make a strong case for financialization as the primary cause of, say, polarization and the elimination of routine jobs.

It's also important to realize that growth in the financial sector has been highly dependent on advancing information technology. Virtually all of the financial innovations that have arisen in recent decades—including, for example, collateralized debt obligations (CDOs) and exotic financial derivatives—would not have been possible without access to powerful computers. Likewise, automated trading algorithms are now responsible for nearly two-thirds of stock market trades, and Wall Street firms have built huge computing centers in close physical proximity to exchanges in order to gain trading advantages measured in tiny fractions of a second. Between 2005 and 2012, the average time to execute a trade dropped from about 10 seconds to just 0.0008 seconds,[56] and robotic, high-speed trading was heavily implicated in the May 2010 "flash crash" in which the Dow Jones Industrial Average plunged nearly a thousand points and then recovered for a net gain, all within the space of just a few minutes.

Viewed from this perspective, financialization is not so much a competing explanation for our seven economic trends; it is rather—at least to some extent—one of the ramifications of accelerating information technology. In this, there is a strong cautionary note as we look to the future: as IT continues its relentless progress, we can be certain that financial innovators, in the absence of regulations that constrain them, will find ways to leverage all those new capabilities—and, if history is any guide, it won't necessarily be in ways that benefit society as a whole.

Politics

In the 1950s, more than a third of the US private sector workforce was unionized. By 2010, that number had declined to about 7 percent.[57] At the height of its power, organized labor was a powerful advocate for the middle class as a whole. The fact that workers were able to consistently capture the lion's share of productivity growth in the 1950s and '60s can likely be attributed at least in part to the negotiating power of unions during that period. The situation today is very different; unions now struggle simply to maintain their existing membership.

The precipitous decline in the power of organized labor is one of the most visible developments associated with the rightward drift that has characterized American economic policy over the past three decades. In their 2010 book *Winner Take All Politics,* political scientists Jacob S. Hacker and Paul Pierson make a compelling case for politics as the primary driver of inequality in the United States. Hacker and Pierson point to 1978 as the pivotal year when the American political landscape began to shift under a sustained and organized assault from conservative business interests. In the decades that followed, industries were deregulated, top marginal tax rates on the wealthy and on corporations were cut to historic lows, and workplaces were made increasingly inhospitable to union organization. Much of this was driven not by electoral politics but, rather, by continuous lobbying on the part of business interests. As the power of organized labor withered, and as the number of lobbyists in Washington exploded, the day-to-day political warfare in the capital became increasingly asymmetric.

While the political situation in the United States seems uniquely detrimental to the middle class, evidence for the impact of advancing technology can be found in a wide range of developed and developing nations. Inequality is increasing in nearly all industrialized countries, while the share of national income claimed by labor is generally falling. Job market polarization has been observed in a majority of European nations. And in Canada—where organized labor remains a powerful national force—inequality is rising, median household incomes have fallen in real terms since 1980, and private sector union membership has declined as manufacturing jobs have disappeared.[58]

To some extent, the question here is one of categorization: if a nation fails to implement policies designed to mitigate the impact of structural changes brought on by advancing technology, should we label that as a problem caused by technology, or politics? Regardless, there is little question that the United States stands alone in terms of the political decisions it has made; rather than simply failing to enact policies that might have slowed the forces driving the country toward higher levels of inequality, America very often has made choices that have effectively put a wind at the back of those forces.

Looking to the Future

The debate over the primary causes of the soaring inequality and decades-long wage stagnation that have developed in the United States is likely to continue unabated, and because it touches on intensely polarizing issues—organized labor, tax rates on the wealthy, free trade, the proper role of government—the dialogue is sure to be colored by ideology. To my mind, the evidence I've presented here demonstrates that information technology has played a significant—though not necessarily dominant—role over the past few decades. Beyond that, I'm content to leave it to economic historians to delve into the data and perhaps someday shine a more definitive light on the precise forces involved in getting us to this point. The real question—and the primary subject of this book—is, What will be most important in the future? Many of the forces that heavily impacted the economy and political environment over the past half-century have largely played out. Unions outside the public sector have been decimated. Women who want careers have entered the workforce or enrolled in colleges and professional schools. There is evidence that the drive toward factory offshoring has slowed significantly, and in some cases, manufacturing is returning to the United States.

Among the forces poised to shape the future, information technology stands alone in terms of its exponential progress. Even in nations whose political environments are far more responsive to the welfare of average workers, the changes wrought by technology are becoming increasingly evident. As the technological frontier advances, many jobs that we would today consider nonroutine, and therefore protected from automation, will eventually be pulled into the routine and predictable category. The hollowed-out middle of the already polarized job market is likely to expand as robots and self-service technologies eat away at low-wage jobs, while increasingly intelligent algorithms threaten higher-skill occupations. Indeed, a 2013 study by Carl Benedikt Frey and Michael A. Osborne at the University of Oxford concluded that occupations amounting to nearly half of US total employment may be vulnerable to automation within roughly the next two decades.[59]

While accelerating information technology is nearly certain to have an outsized impact on the future economy and job market, it will remain deeply intertwined with other powerful forces. The line between technology and globalization will blur as higher-skill jobs become more vulnerable to electronic offshoring. If, as seems likely, advancing technology continues to drive the United States and other industrialized countries toward ever higher inequality, then the political influence wielded by the financial elite can only increase. This may make it even more difficult to enact policies that might serve to counteract the structural shifts occurring in the economy and improve the prospects for those in the middle and bottom of the income distribution.

In my 2009 book *The Lights in the Tunnel,* I wrote that "while technologists are actively thinking about, and writing books about, intelligent machines, the idea that technology will ever truly replace a large fraction of the human workforce and lead to permanent, structural unemployment is, for the majority of economists, almost unthinkable." To their credit, some economists have since begun to take the potential for widespread automation more seriously. In their 2011 ebook *Race Against the Machine,* Erik Brynjolfsson and Andrew McAfee of the Massachusetts Institute of Technology helped bring these ideas into the economic mainstream. Prominent economists including Paul Krugman and Jeffrey Sachs have likewise written about the possible impact of machine intelligence.[60] Nonetheless, the idea that technology might someday truly transform the job market and ultimately demand fundamental changes to both our economic system and the social contract remains either completely unacknowledged or at the very fringes of public discourse.

Indeed, among practitioners of economics and finance there is often an almost reflexive tendency to dismiss anyone who argues that this time might be different. This is very likely the correct instinct when one is discussing those aspects of the economy that are primarily driven by human behavior and market psychology. The psychological underpinnings of the recent housing bubble and bust were almost certainly little different from those that have characterized financial crises throughout history. Many of the political machinations of the early Roman republic could probably be dropped seamlessly onto the front page of today's *Politico*. These things never really change.

It would be a mistake, however, to apply that same reasoning to the impact of advancing technology. Up until the moment the first aircraft achieved sustained powered flight at Kitty Hawk, North Carolina, it was an incontrovertible fact—supported by data stretching back to the beginning of time—that human beings, strapped into heavier-than-air contraptions, *do not fly*. Just as that reality shifted in an instant, a similar phenomenon plays out continuously in nearly every sphere of technology. This time is always different where technology is concerned: that, after all, is the entire point of innovation. Ultimately, the question of whether smart machines will someday eclipse the capability of average people to perform much of the work demanded by the economy will be answered by the nature of the technology that arrives in the future—not by lessons gleaned from economic history.

In the next chapter, we'll examine the nature of information technology and its relentless acceleration, the characteristics that set it apart, and the ways in which it is already transforming important spheres of the economy.