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

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

Chapter 3

INFORMATION TECHNOLOGY: AN UNPRECEDENTED FORCE FOR DISRUPTION

Imagine depositing a penny in a bank account. Now, double the account balance every day. On day three you would go from 2 cents to 4 cents. The fifth day would take your balance from 8 to 16 cents. After less than a month, you would have more than a million dollars. If we had deposited that initial penny in 1949, just as Norbert Wiener was writing his essay about the future of computing, and then let Moore's Law run its course—doubling the amount roughly every two years—by 2015, our technological account would contain nearly $86 million. And as things move forward from this point, that balance will continue to double. Future innovations will be able to leverage that enormous accumulated balance, and as a result the rate of progress in the coming years and decades is likely to far exceed what we have become accustomed to in the past.

Moore's Law is the best-known measure of advancing computer power, but information technology is, in fact, accelerating on many different fronts. For example, computer memory capacity and the amount of digital information that can be carried on fiber-optic lines have both experienced consistent exponential increases. Nor is the acceleration confined to computer hardware; the efficiency of some software algorithms has soared at a rate far in excess of what Moore's Law alone would predict.

While exponential acceleration offers valuable insight into the advance of information technology over relatively long periods, the short-term reality is more complex. Progress is generally not always smooth and consistent; instead, it often lurches forward and then pauses while new capabilities are assimilated into organizations and the foundation for the next period of rapid advance is established. There are also intricate interdependencies and feedback loops between different realms of technology. Progress in one area may drive a sudden burst of innovation in another. As information technology marches forward, its tentacles reach ever deeper into organizations and the overall economy, often transforming the way people work in ways that can further its own advance. Consider, for example, how the rise of the Internet and sophisticated collaboration software has enabled the offshoring of software development; this has made a vastly expanded population of skilled programmers available, and all that new talent is helping to drive still more progress.

Acceleration Versus Stagnation

As information and communications technologies have advanced in their decades-long exponential march, innovation in other areas has been largely incremental. Examples include the basic design of cars, homes, aircraft, kitchen appliances, and our overall transportation and energy infrastructures, none of which, for the most part, have changed significantly since the middle of the twentieth century. PayPal co-founder Peter Thiel's famous comment—"We were promised flying cars, and instead what we got was 140 characters"—captures the sentiment of a generation that expected the future to be way cooler than this.

This lack of broad-based progress stands in stark contrast to what a person who lived through the final decades of the nineteenth century and the first half of the twentieth would have experienced. Indoor plumbing, automobiles, airplanes, electricity, home appliances, and public sanitation and utility systems all came into widespread use during this period. In industrialized countries, at least, people at all levels of society received an astonishing upgrade in the quality of their lives, even as the overall wealth of society was propelled to dizzying new heights.

Some economists have taken note of this plodding rate of advance in most spheres of technology and have tied it to the economic trends we looked at in the previous chapter, and in particular to the stagnation of incomes for most ordinary Americans. One of the foundational principles of modern economics is that such technological change is essential to long-term economic growth. Robert Solow, the economist who formalized this idea, received the Nobel Prize for his work in 1987. If innovation is the primary driver of prosperity, then perhaps stagnant incomes imply that the problem is the rate at which new inventions and ideas are being generated, rather than the impact of technology on the working and middle classes. Maybe computers aren't really all that important, and the slow rate of progress on a broader front is what matters most.

Several economists have made this case. Tyler Cowen, an economist at George Mason University, proposed in his 2011 book *The Great Stagnation* that the US economy has run into a temporary plateau after consuming all the low-hanging fruit of accessible innovation, free land, and underutilized human talent. Robert J. Gordon of Northwestern University is even more pessimistic, arguing in a 2012 paper that economic growth in the United States, hampered by a slow pace of innovation and a number of "headwinds"—including excessive debt, an aging population, and shortfalls in our educational system—may essentially be over.[1]

In order to gain some insight into the factors that influence the pace of innovation, we may find it useful to think in terms of the historical path that nearly all technologies follow. Airplanes are a good example. The first controlled, powered flight occurred in December 1903 and lasted about twelve seconds. Progress accelerated from that humble start, but the primitive initial level of the technology meant it would take years before a practical airplane would emerge. By 1905, Wilbur Wright was able to stay aloft for nearly forty minutes while traveling about twenty-four miles. Within a few years, however, things started to really come together; aircraft technology had progressed along its exponential curve, and the rate of absolute progress picked up dramatically. By World War I, airplanes were engaging in high-speed aerial dog fights. Progress continued its acceleration over the next two decades, ultimately producing high-performance fighter aircraft like the Spitfire, the Zero, and the P-51. Sometime around World War II, however, the rate of advance slowed significantly. Aircraft powered by internal combustion engines driving propellers were now very close to their ultimate technical potential, and design improvements beyond that point would be incremental.

This S-shaped path in which accelerating—or exponential—advance ultimately matures into a plateau effectively illustrates the life story of virtually all specific technologies. Of course, we know that as World War II came to a close, an entirely new aircraft technology appeared on the scene. Jet aircraft would soon offer a level of performance far beyond what was possible for any propeller-driven plane. Jets were a disruptive technology: they had an S-curve of their own. Figure 3.1 shows what this might look like.

If we want to dramatically speed up the pace of innovation in aircraft design, we need to find yet another S-curve, and that curve has to represent a technology that is not only superior in terms of performance but also economically viable.* The problem, of course,

---

* The supersonic Concorde, for example, offered a new S-curve in terms of absolute performance, but it did not prove to be an economically sustainable technology and was never able to capture more than a tiny fraction of the airline passenger market. The Concorde was in service from 1976 until 2003.

Figure 3.1. Aircraft Technology S-Curves

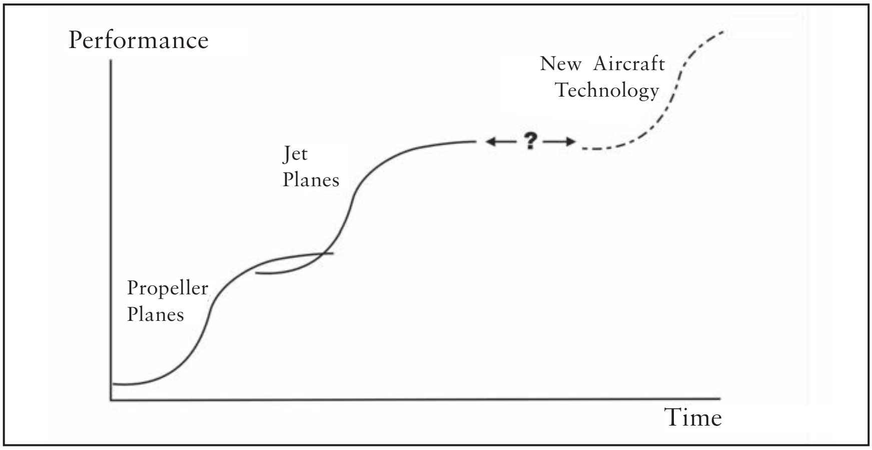

is that so far, that new curve is nowhere to be found. Assuming we can't discover this disruptive new technology simply by hopping the fence at Area 51, it's going to take a giant leap to get to that new S-curve—and this presumes, of course, that the curve even exists.

The critical point here is that while many factors, such as the level of research and development effort and investment, or the presence of a favorable regulatory environment, can certainly have an impact on the relative position of technology S-curves, the most important factor by far is the set of physical laws that govern the sphere of technology in question. We don't yet have a disruptive new aircraft technology and that is primarily due to the laws of physics and the limitations they imply relative to our current scientific and technical knowledge. If we hope to have another period of rapid innovation in a wide range of technological areas—perhaps something comparable to what occurred between approximately 1870 and 1960—we would need to find new S-curves in all these different areas. Obviously, that is likely to represent an enormous challenge.

There is one important reason for optimism, however, and that is the positive impact that accelerating information technology will have on research and development in other fields. Computers have already been transformative in many areas. Sequencing the human genome would certainly have been impossible without advanced computing power. Simulation and computer-based design have greatly expanded the potential for experimentation with new ideas in a variety of research areas.

One information technology success story that has had a dramatic and personal impact on all of us has been the role of advanced computing power in oil and gas exploration. As the global supply of easily accessible oil and gas fields has declined, new techniques such as three-dimensional underground imaging have become indispensable tools for locating new reserves. Aramco, the Saudi national oil company, for example, maintains a massive computing center where powerful supercomputers are instrumental in maintaining the flow of oil. Many people might be surprised to learn that one of the most important ramifications of Moore's Law has been the fact that, at least so far, world energy supplies have kept pace with surging demand.

The advent of the microprocessor has resulted in an astonishing increase in our overall ability to perform computations and manipulate information. Where once computers were massive, slow, expensive, and few in number, today they are cheap, powerful, and ubiquitous. If you were to multiply a single computer's increase in computational power since 1960 by the number of new microprocessors that have appeared since then, the result would be nearly beyond reckoning. It seems impossible to imagine that such an immeasurable increase in our overall computing capacity won't eventually have dramatic consequences in a variety of scientific and technical fields. Nonetheless, the primary determinant of the positions of the technology S-curves we'll need to reach in order to have truly disruptive innovation is still the applicable laws of nature. Computational capability can't change that reality, but it may well help researchers to bridge some of the gaps.

The economists who believe we have hit a technological plateau typically have deep faith in the relationship between the pace of innovation and the realization of broad-based prosperity; the implication is that if we can just jump-start technological progress on a broad front, median incomes will once again begin increasing in real terms. I think there are good reasons to be concerned that this may not necessarily turn out to be the case. In order to understand why, let's look at what makes information technology unique and the ways in which it will intertwine with innovations in other areas.

Why Information Technology Is Different

The relentless acceleration of computer hardware over decades suggests that we've somehow managed to remain on the steep part of the S-curve for far longer than has been possible in other spheres of technology. The reality, however, is that Moore's Law has involved successfully climbing a staircase of cascading S-curves, each representing a specific semiconductor fabrication technology. For example, the lithographic process used to lay out integrated circuits was initially based on optical imaging techniques. When the size of individual device elements shrank to the point where the wavelength of visible light was too long to allow for further progress, the semiconductor industry moved on to X-ray lithography.[2] Figure 3.2 illustrates roughly what climbing a series of S-curves might look like.

One of the defining characteristics of information technology has been the relative accessibility of subsequent S-curves. The key to sustainable acceleration has not been so much that the fruit is low-hanging but, rather, that the tree is climbable. Climbing that tree has been a complex process that has been driven by intensive competition and has required enormous investment. There has also been substantial cooperation and planning. To help coordinate all these efforts, the industry publishes a massive document called the International Technology Roadmap for Semiconductors (ITRS), which essentially offers a detailed fifteen-year preview of how Moore's Law is expected to unfold.

Figure 3.2. Moore's Law as a Staircase of S-Curves

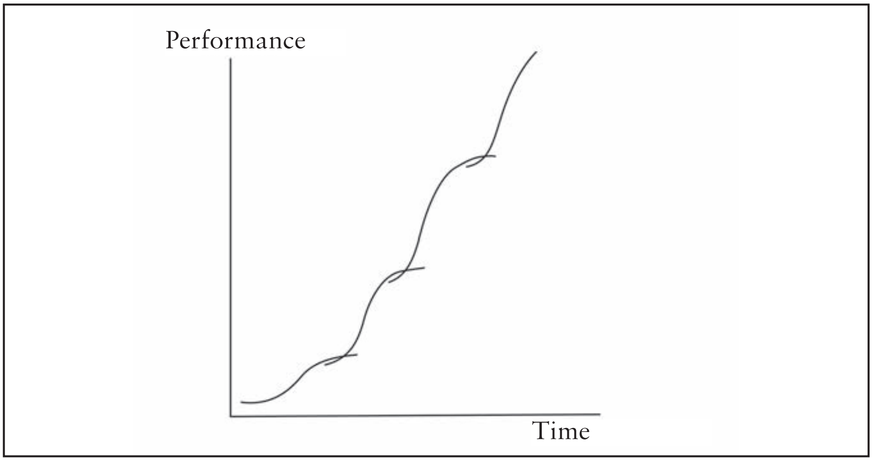

As things stand today, computer hardware may soon run into the same type of challenge that characterizes other areas of technology. In other words, reaching that next S-curve may eventually require a giant—and perhaps even unachievable—leap. The historical path followed by Moore's Law has been to keep shrinking the size of transistors so that more and more circuitry can be packed onto a chip. By the early 2020s, the size of individual design elements on computer chips will be reduced to about five nanometers (billionths of a meter), and that is likely to be very close to the fundamental limit beyond which no further miniaturization is possible. There are, however, a number of alternate strategies that may allow progress to continue unabated, including three-dimensional chip design and exotic carbon-based materials.[3] *

---

* The idea behind 3D chips is to begin stacking circuitry vertically in multiple layers. Samsung Electronics began manufacturing 3D flash memory chips in August 2013. If this technique proves economically viable for the far more sophisticated processor chips designed by companies like Intel and AMD (Advanced Micro Devices), it may represent the future of Moore's Law. Another possibility is to turn to exotic carbon-based materials as an alternative to silicon. Graphene and carbon nanotubes, both of which are the result of recent nanotechnology research, may eventually offer a new medium for very high-performance computing. Researchers at Stanford University have already created a rudimentary carbon nanotube computer, although its performance falls far short of commercial silicon-based processors.

Even if the advance of computer hardware capability were to plateau, there would remain a whole range of paths along which progress could continue. Information technology exists at the intersection of two different realities. Moore's Law has dominated the realm of atoms, where innovation is a struggle to build faster devices and to minimize or find a way to dissipate the heat they generate. In contrast, the realm of bits is an abstract, frictionless place where algorithms, architecture (the conceptual design of computing systems), and applied mathematics govern the rate of progress. In some areas, algorithms have already advanced at a far faster rate than hardware. In a recent analysis, Martin Grötschel of the Zuse Institute in Berlin found that, using the computers and software that existed in 1982, it would have taken a full eighty-two years to solve a particularly complex production planning problem. As of 2003, the same problem could be solved in about a minute—an improvement by a factor of around 43 million. Computer hardware became about 1,000 times faster over the same period, which means that improvements in the algorithms used accounted for approximately a 43,000-fold increase in performance.[4]

Not all software has improved so quickly. This is especially true of areas where software must interact directly with people. In an August 2013 interview with James Fallows of *The Atlantic*, Charles Simonyi, the computer scientist who oversaw the development of Microsoft Word and Excel, expressed the view that software has largely failed to leverage the advances that have occurred in hardware. When asked where the most potential for future improvement lies, Simonyi said: "The basic answer is that nobody would be doing routine, repetitive things anymore."[5]

There is also tremendous room for future progress through finding improved ways to interconnect vast numbers of inexpensive processors in massively parallel systems. Reworking current hardware device technology into entirely new theoretical designs could likewise produce giant leaps in computer power. Clear evidence that a sophisticated architectural design based on deeply complex interconnection can produce astonishing computational capability is provided by what is, by far, the most powerful general computing machine in existence: the human brain. In creating the brain, evolution did not have the luxury of Moore's Law. The "hardware" of a human brain is no faster than that of a mouse and is thousands to millions of times slower than a modern integrated circuit; the difference lies entirely in the sophistication of the design.[6] Indeed, the ultimate in computer capability—and perhaps machine intelligence—might be achieved if someday researchers are able to marry the speed of even today's computer hardware with something approaching the level of design complexity you would find in the brain. Baby steps have already been taken in that direction: IBM released a cognitive computing chip—inspired by the human brain and aptly branded "SyNAPSE"—in 2011 and has since created a new programming language to accompany the hardware.[7]

Beyond the relentless acceleration of hardware, and in many cases software, there are, I think, two other defining characteristics of information technology. The first is that IT has evolved into a true general-purpose technology. There are very few aspects of our daily lives, and especially of the operation of businesses and organizations of all sizes, that are not significantly influenced by or even highly dependent on information technology. Computers, networks, and the Internet are now irretrievably integrated into our economic, social, and financial systems. IT is everywhere, and it's difficult to even imagine life without it.

Many observers have compared information technology to electricity, the other transformative general-purpose technology that came into widespread use in the first half of the twentieth century. Nicholas Carr makes an especially compelling argument for viewing IT as an electricity-like utility in his 2008 book *The Big Switch*. While many of these comparisons are apt, the truth is that electricity is a tough act to follow. Electrification transformed businesses, the overall economy, social institutions, and individual lives to an astonishing degree—and it did so in ways that were overwhelmingly positive. It would probably be very difficult to find a single person in a developed country like the United States who did not eventually receive a major upgrade in his or her standard of living after the advent of electric power. The transformative impact of information technology is likely to be more nuanced and, for many people, less universally positive. The reason has to do with IT's other signature characteristic: cognitive capability.

Information technology, to a degree that is unprecedented in the history of technological progress, encapsulates intelligence. Computers make decisions and solve problems. Computers are machines that can—in a very limited and specialized sense—*think*. No one would argue that today's computers approach anything like human-level general intelligence. But that very often misses the point. Computers are getting dramatically better at performing specialized, routine, and predictable tasks, and it seems very likely that they will soon be poised to outperform many of the people now employed to do these things.

Progress in the human economy has resulted largely from occupational specialization, or as Adam Smith would say, "the division of labour." One of the paradoxes of progress in the computer age is that as work becomes ever more specialized, it may, in many cases, also become more susceptible to automation. Many experts would say that, in terms of *general* intelligence, today's best technology barely outperforms an insect. And yet, insects do not make a habit of landing jet aircraft, booking dinner reservations, or trading on Wall Street. Computers now do all these things, and they will soon begin to aggressively encroach in a great many other areas.

Comparative Advantage and Smart Machines

Economists who reject the idea that machines could someday make a large fraction of our workforce essentially unemployable often base their argument on one of the biggest ideas in economics: the theory of comparative advantage.[8] To see how comparative advantage works, let's consider two people. Jane is truly exceptional. After many years of intensive training and a record of nearly unmatched success, she is considered to be one of the world's leading neurosurgeons. In her gap years between college and medical school, Jane enrolled in one of France's best culinary institutes and is now also a gourmet cook of rarefied talent. Tom is more of an average guy. He is, however, a very good cook, and has been complimented many times on his skills. Still, he can't really come close to matching what Jane can do in the kitchen. And it goes without saying that Tom wouldn't be allowed anywhere near an operating room.

Given that Tom can't compete with Jane as a cook, and certainly not as a surgeon, is there any way that the two could enter into an agreement that would make them both better off? Comparative advantage says "yes" and tells us that Jane could hire Tom as a cook. Why would she do that when she can get a better result by doing the cooking herself? The answer is that it would free up more of Jane's time and energy for the one thing she is truly exceptional at (and the thing that brings in the most income): brain surgery.

The main idea behind comparative advantage is that you should always be able to find a job, provided you specialize in the thing at which you are "least bad" relative to other people. By doing so, you offer others the chance to also specialize and thereby earn a higher income. In Tom's case, least bad meant cooking. Jane is luckier (and a lot richer) because her least bad gig is something she is truly great at, and that talent happens to have a very high market value. Throughout economic history, comparative advantage has been the primary driver of ever more specialization and trade between individuals and nations.

Now let's change the story. Imagine that Jane has the ability to easily and inexpensively clone herself. If you like science fiction movies, think in terms of *Matrix Reloaded,* where Neo battles dozens of copies of the agent, Smith. In that particular struggle, Neo ultimately prevails, but I think you can see that Tom might not be so lucky when it comes to keeping his job working for Jane. Comparative advantage works because of opportunity cost: if a person chooses to do one thing, she must necessarily give up the opportunity to do something else. Time and space are limited; she can't be in two places doing two things at once.

Machines, and particularly software applications, can be easily replicated. In many cases they can be cloned at a cost that is small compared with employing a person. When intelligence can be replicated, the concept of opportunity cost is upended. Jane can now perform brain surgery and cook simultaneously. So why does she need Tom at all? It's a good bet that pretty soon Jane's clones will also start putting less talented brain surgeons out of work. Comparative advantage in the age of smart machines might require something of a rethink.

Imagine the impact of a large corporation being able to train a single employee and then clone him into an army of workers, all of whom instantly possess his knowledge and experience but, from that point on, are also capable of continuing to learn and adapt to new situations. When the intelligence encapsulated in information technology is replicated and scaled across organizations, it has the potential to fundamentally redefine the relationship between people and machines. From the perspective of a great many workers, computers will cease to be tools that enhance their productivity and instead become viable substitutes. This outcome will, of course, dramatically increase the productivity of many businesses and industries—but it will also make them far less labor-intensive.

The Tyranny of the Long Tail

The influence of this distributed machine intelligence is most evident in the information technology industry itself. The Internet has spawned enormously profitable and influential corporations with startlingly diminutive workforces. In 2012, Google, for example, generated a profit of nearly $14 billion while employing fewer than 38,000 people.[9] Contrast that with the automotive industry. At peak employment in 1979, General Motors alone had nearly 840,000 workers but earned only about $11 billion—20 percent less than what Google raked in. And, yes, that's after adjusting for inflation.[10] Ford, Chrysler, and American Motors employed hundreds of thousands more people. Beyond that core workforce, the industry also created millions of peripheral middle-class jobs in areas like driving, repairing, insuring, and renting cars.

Of course, the Internet sector also offers peripheral opportunities. The new information economy is often touted as the great equalizer. After all, anyone can write a blog and run ads on it, publish an ebook, sell stuff on eBay, or develop an iPhone app. While these opportunities do indeed exist, they are dramatically different from all those solid middle-class jobs created by the automotive industry. The evidence shows pretty clearly that the income realized from online activities nearly always tends to follow a winner-take-all distribution. While the Internet may, in theory, equalize opportunity and demolish entry barriers, the actual outcomes it produces are almost invariably highly unequal.

If you graph the traffic coming to websites, advertising revenue generated online, music downloads from the iTunes store, books sold on Amazon, apps downloaded from Apple's AppStore or Google Play, or just about anything else online, you will nearly always end up with something that looks like Figure 3.3. This ubiquitous long-tail distribution is central to the business models of the corporations that dominate the Internet sector. Companies like Google, eBay, and Amazon are able to generate revenue from *every point* on the distribution. If a company controls a large market, then aggregating even tiny sums along the entire curve results in total revenues that can easily reach into the billions.

Figure 3.3. A Winner-Take-All/Long-Tail Distribution

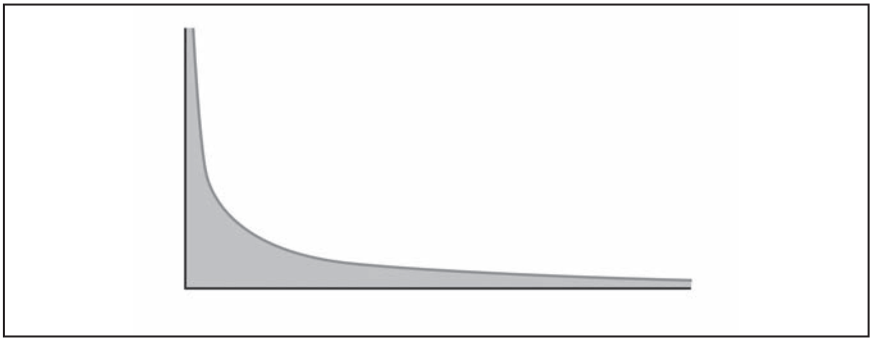

Markets in goods and services that are susceptible to digitalization inevitably evolve into this winner-take-all distribution. Sales of books and music, classified advertising, and movie rentals, for example, are increasingly dominated by a tiny number of online distribution hubs, and one obvious result has been the elimination of vast numbers of jobs for people like journalists and retail store clerks.

The long tail is great if you own it. When, however, you occupy only a single point on the distribution, the story is quite different. Out on the long tail, incomes from most online activities rapidly drop to the pocket-change level. That can work out fine if you have an alternate source of income, or if you happen to be living in your parents' basement. The problem is that as digital technology continues to transform industries, more and more of the jobs that provide that primary-income source are likely to disappear.

As more people lose the dependable income stream that anchors them into the middle class, they are likely to increasingly turn to these long-tail opportunities in the digital economy. A lucky few will provide the anecdotal success stories we will hear about, but the vast majority will struggle to maintain anything approaching a middle-class lifestyle. As techno-visionary Jaron Lanier has pointed out, a great many people are likely to be forced into the type of informal economy that is found in third-world nations.[11] Young adults who find the freedom of the informal economy alluring will quickly discover its drawbacks when they begin to think in terms of maintaining a home, raising children, or planning for retirement. Of course, there have always been people living at the fringes in the United States and other developed economies, but to some extent they free-ride on the wealth generated by a critical mass of middle-class households. The presence of that solid middle is one of the primary factors that differentiates an advanced nation from an impoverished one—and its erosion is becoming increasingly evident, especially in the United States.

Most techno-optimists would likely object to this characterization. They tend to view information technology as universally empowering. It is perhaps not coincidental that they also tend to have been very successful in the new economy. The most prominent digital optimists typically live at the extreme left of the long tail—or, even better, they've perhaps founded a company that owns the entire distribution. In a PBS television special that aired in 2012, inventor and futurist Ray Kurzweil was asked about the possibility of a "digital divide"—meaning that only a small percentage of the population will be able to thrive in the new information economy. Kurzweil dismissed the idea of such a divide and instead pointed to empowering technologies like mobile phones. Anybody with a smart phone, he said, "is carrying around billions of dollars of capability circa 20 or 30 years ago."[12] Left unsaid was how the average person is supposed to leverage that technology into a livable income.

Mobile phones have indeed been shown to improve living standards, but this has been documented primarily in developing countries that lack other communications infrastructure. By far the most celebrated success story involves sardine fishermen in Kerala, a region along the southwest coast of India. In a 2007 research paper, economist Robert Jensen described how mobile phones allowed the fishermen to determine which villages offered the best markets for their fish.[13] Before the advent of wireless technology, targeting a particular village was a guess that often resulted in a mismatch between supply and demand. However, with their new phones, the fishermen knew exactly where the buyers were, and this has resulted in a better functioning market with more stable prices and far less waste.

The sardine fishermen of Kerala have become a kind of standard-bearer for techno-optimism as it relates to developing countries, and their story has been told in numerous books and magazine articles.[14] While mobile phones are unquestionably of great value to third-world fishermen, there is little evidence to suggest that average citizens in developed countries—or, for that matter, even in poor countries—will succeed in deriving a meaningful income from their smart phones. Even skilled software developers find it extremely challenging to generate significant revenue from mobile apps, and the primary reason, needless to say, is that ubiquitous long-tail distribution. Visit almost any online forum populated by Android or iPhone developers and you're likely to find discussions lamenting the winner-take-all nature of the mobile ecosystem and the difficulty in monetizing apps. As a practical matter, for the majority of people who lose middle-class jobs, access to a smart phone may offer little beyond the ability to play Angry Birds while waiting in the unemployment line.

A Moral Question

If we think again in terms of doubling a penny as a proxy for the exponential advance of digital technology, it's clear that today's enormous technological account balance results from the efforts of countless individuals and organizations over the course of decades. Indeed, the arc of progress can be traced back in time at least as far as Charles Babbage's mechanical difference engine in the early seventeenth century.

The innovations that have resulted in fantastic wealth and influence in today's information economy, while certainly significant, do not really compare in importance to the groundbreaking work done by pioneers like Alan Turing or John von Neumann. The difference is that even incremental advances are now able to leverage that extraordinary accumulated account balance. In a sense, the successful innovators of today are a bit like the Boston Marathon runner who in 1980 famously snuck into the race only half a mile from the finish line.

Of course, all innovators stand on the shoulders of those who came before them. This was certainly true when Henry Ford introduced the Model T. However, as we have seen, information technology is fundamentally different. IT's unique ability to scale machine intelligence across organizations in ways that will substitute for workers and its propensity to everywhere create winner-take-all scenarios will have dramatic implications for both the economy and society.

At some point, we may need to ask a fundamental moral question: Should the population at large have some sort of claim on that accumulated technological account balance? The public does, of course, benefit greatly from accelerating digital technology in terms of lower costs, convenience, and free access to information and entertainment. But that brings us back to the problem with Kurzweil's argument about mobile phones: those things won't pay the rent.

It should be kept in mind, as well, that much of the basic research that enabled progress in the IT sector was funded by American taxpayers. The Defense Advanced Research Projects Agency (DARPA) created and funded the computer network that ultimately evolved into the Internet.* Moore's Law has come about, in part, because of university-led research funded by the National Science Foundation. The Semiconductor Industry Association, the industry's political action committee, actively lobbies for increased federal research dollars. Today's computer technology exists in some measure because millions of middle-class taxpayers supported federal funding

---

* DARPA also provided the initial financial backing for the development of Siri (now Apple's virtual assistant technology) and has underwritten the development of IBM's new SyNAPSE cognitive computing chips.

for basic research in the decades following World War II. We can be reasonably certain that those taxpayers offered their support in the expectation that the fruits of that research would create a more prosperous future for their children and grandchildren. Yet, the trends we looked at in the last chapter suggest we are headed toward a very different outcome.

Beyond the basic moral question of whether a tiny elite should be able to, in effect, capture ownership of society's accumulated technological capital, there are also practical issues regarding the overall health of an economy in which income inequality becomes too extreme. Continued progress depends on a vibrant market for future innovations—and that, in turn, requires a reasonable distribution of purchasing power.

In later chapters, we'll look in more detail at some of the overall economic and social implications of digital technology's relentless acceleration. But first, let's look at how these innovations are increasingly threatening the high-skill jobs held by workers with college and even graduate or professional degrees.