PART II

ACCELERATING

TWO

What the Hell Happened in 2007?

John Doerr, the legendary venture capitalist who backed Netscape, Google, and Amazon, doesn't remember the exact day anymore; all he remembers is that it was shortly before Steve Jobs took the stage at the Moscone Center in San Francisco on January 9, 2007, to announce that Apple had reinvented the mobile phone. Doerr will never forget, though, the moment he first laid eyes on that phone. He and Jobs, his friend and neighbor, were watching a soccer match that Jobs's daughter was playing in at a school near their homes in Palo Alto. As play dragged on, Jobs told Doerr that he wanted to show him something.

"Steve reached into the top pocket of his jeans and pulled out the first iPhone," Doerr recalled for me, "and he said, 'John, this device nearly broke the company. It is the hardest thing we've ever done.' So I asked for the specs. Steve said that it had five radios in different bands, it had so much processing power, so much RAM [random access memory], and so many gigabits of flash memory. I had never heard of so much flash memory in such a small device. He also said it had no buttons—it would use software to do everything—and that in one device 'we will have the world's best media player, world's best telephone, and world's best way to get to the Web—all three in one.'"

Doerr immediately volunteered to start a fund that would support creation of applications for this device by third-party developers, but Jobs wasn't interested at the time. He didn't want outsiders messing with his elegant phone. Apple would do the apps. A year later, though, he changed his mind; that fund was launched, and the mobile phone app industry exploded. The moment that Steve Jobs introduced the iPhone turns out to have been a pivotal junction in the history of technology—and the world.

Vint Cerf, who was one of the founding fathers of the Internet, remarked to me: "The spread of the smartphone made the Internet more valuable, because it became all the more accessible." But the Internet "also made the smartphone more valuable," because you could use that phone to tap into all the content and computing power of the Internet. The synthesis of these two technologies in one year would have been historically noteworthy all on its own. But it wasn't alone.

There are vintage years in wine and vintage years in history, and 2007 was definitely one of the latter.

Because not just the iPhone emerged in 2007—a whole group of companies emerged in and around that year. Together, these new companies and innovations have reshaped how people and machines communicate, create, collaborate, and think.

Just go down the list. In 2007, a company called VMware went public. VMware's translation software—a kind of digital Rosetta stone—made it possible for the same computer to run multiple operating systems (and applications) all at the same time, rather than requiring a dedicated computer for each different operating system. It was a crucial technology enabling the growth of cloud computing. Meanwhile, storage capacity for computing exploded thanks to the emergence that year of a framework for software called Hadoop, making "big data" possible for all. In 2007, development began on an open-source platform for writing and collaborating on software, called GitHub, that would vastly expand the ability of software to start, as Netscape founder Marc Andreessen once put it, "eating the world." On September 26, 2006, Facebook, a social networking site that had been confined to users on college campuses and at high schools, was opened to everyone at least thirteen years old with a valid e-mail address, and started to scale globally. In 2007, a microblogging company called Twitter, which had been part of a broader start-up, was spun off as its own separate platform and also started to scale globally. Change.org, the most popular social mobilization website, emerged in 2007.

In late 2006, Google bought YouTube, and in 2007 it launched Android, an open-standards platform for devices that would help smartphones scale globally with an alternative operating system to Apple's iOS. In 2007, AT&T, the iPhone's exclusive connectivity provider, invested in something called "software-enabled networks"—thus rapidly expanding its capacity to handle all the cellular traffic created by this smartphone revolution. According to AT&T, mobile data traffic on its national wireless network increased by *more than 100,000 percent* from January 2007 through December 2014.

In 2007, "Satoshi Nakamoto"—the name used by an unknown person or persons—began working on a digital currency and payment system called "Bitcoin." Nakamoto released the concept on October, 31, 2008, in a research paper entitled "Bitcoin: A Peer-to-Peer Electronic Cash System." The paper proposed that "a purely peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution." A decade later, it appears that Bitcoin's digital currency could well become the backbone of the global banking system in the twenty-first century. According to Wikipedia, "Nakamoto claimed that work on the Bitcoin writing of the code began in 2007."

Also in 2007, Amazon released something called the Kindle, onto which, thanks to Qualcomm's 3G technology, you could download thousands of books anywhere in the blink of an eye, launching the e-book revolution. In 2007, Airbnb was conceived in an apartment in San Francisco. In late 2006, the Internet crossed one billion users worldwide, which seems to have been a tipping point. In 2007, Palantir Technologies, the leading company using big data analytics and augmented intelligence to, among other things, help the intelligence community find needles in haystacks, launched its first platform. "Computing power and storage reached a level that made it possible for us to create an algorithm that could make a lot of sense out of things we could not make sense of before," explained Palantir's cofounder Alexander Karp. In 2005, Michael Dell decided to relinquish his job as CEO of Dell and step back from the hectic pace and just be its chairman. Two years later he realized that was bad timing. "I could see that the pace of change had really accelerated. I realized we could do all this different stuff. So I came back to run the company in . . . 2007."

It was also in 2007 that David Ferrucci, who led the Semantic Analysis and Integration Department at IBM's Watson Research Center in Yorktown Heights, New York, and his team began building a cognitive computer called Watson—"a special-purpose computer system designed to push the envelope on deep question and answering, deep analytics, and the computer's understanding of natural language," noted the website HistoryofInformation.com. "'Watson' became the first cognitive computer, combining machine learning and artificial intelligence."

In 2007, Intel introduced non-silicon materials—known as high-k/metal gates (the term refers to the transistor gate electrode and transistor gate dielectric)—into microchips for the first time. This very technical fix was hugely important. Although non-silicon materials were already used in other parts of the microprocessor, their introduction into the transistor helped Moore's law—the expectation that the power of microchips would double roughly every two years—continue on its path of delivering exponential growth in computing power. At that time there was real concern that Moore's law was hitting a wall with traditional silicon transistors.

"By opening the way to non-silicon materials it gave Moore's law another shot in the arm at a time when many people were thinking it was coming to an end," said Sadasivan Shankar, who worked on Intel's material design team at the time and now teaches materials and computational sciences at the Harvard School of Engineering and Applied Sciences. Commenting on the breakthrough, the *New York Times* Silicon Valley reporter John Markoff wrote on January 27, 2007: "Intel, the world's largest chip maker, has overhauled the basic building block of the information age, paving the way for a new generation of faster and more energy-efficient processors. Company researchers said the advance represented the most significant change in the materials used to manufacture silicon chips since Intel pioneered the modern integrated-circuit transistor more than four decades ago."

For all of the above reasons, 2007 was also "the beginning of the clean power revolution," said Andy Karsner, the U.S. assistant secretary of energy for efficiency and renewable energy from 2006 to 2008. "If anyone in 2005 or 2006 told you their predictive models captured where clean tech and renewable energy went in 2007 they are lying. Because what happened in 2007 was the beginning of an exponential rise in solar energy, wind, biofuels, LED lighting, energy efficient buildings, and the electrification of vehicles. It was a hockey stick moment."

And it wasn't just for renewables. It was in 2007 that the shale revolution took off, propelled by big data, advances in GPS and better software that made horizontal drilling to extract natural gas from shale deposits so much more efficient. "Fracking had been used for decades," noted

John Bringardner, in a May 2, 2014, essay in *The New Yorker*, "but in the late nineties, after years of experimenting with variants of the process, a Texas wildcatter named George Mitchell discovered that a combination of sand and water (known as 'high-volume slick water hydrofracturing') was cheaper and more effective than previous methods, which relied on other, heavier fluids." Around 2007, though, he added, "Devon Energy, an Oklahoma oil-and-gas producer that had bought Mitchell's company in 2002, had found that it could extract even more gas from each well by pairing [Mitchell's] method with horizontal drilling. Other companies began copying the combined technique, setting off a natural-gas boom in 2008."

Estimated reserves of natural gas in the United States in 2008 were 35 percent higher than in 2006—a phenomenal boost in reserves in basically two years, according to the Potential Gas Committee, the authority on U.S. gas supplies. "The jump is the largest increase in the 44-year history of reports from the committee," *The New York Times* reported on June 17, 2009.

Last but certainly not least, in 2007 the cost of DNA sequencing began to fall dramatically as the biotech industry shifted to new sequencing techniques and platforms, leveraging all the computing and storage power that was just exploding. This change in instruments was a turning point for genetic engineering and led to the "rapid evolution of DNA sequencing technologies that has occurred in recent years," according to Genome.gov. In 2001, it cost $100 million to sequence just one person's genome. On September 30, 2015, *Popular Science* reported: "Yesterday, personal genetics company Veritas Genetics announced that it had reached a milestone: participants in its limited, but steadily expanding Personal Genetics Program can get their entire genome sequenced for just $1,000."

As the graphs on the next two pages display, 2007 was clearly a turning point for many technologies.

Technology has always moved up in step changes. All the elements of computing power—processing chips, software, storage chips, networking, and sensors—tend to move forward roughly as a group. As their improving capacities reach a certain point, they tend to meld together into a platform, and that platform scales a new set of capabilities, which becomes the new normal. As we went from mainframes to desktops to laptops to smartphones with mobile applications, each generation of technology

Cost of DNA Sequencing, per Genome

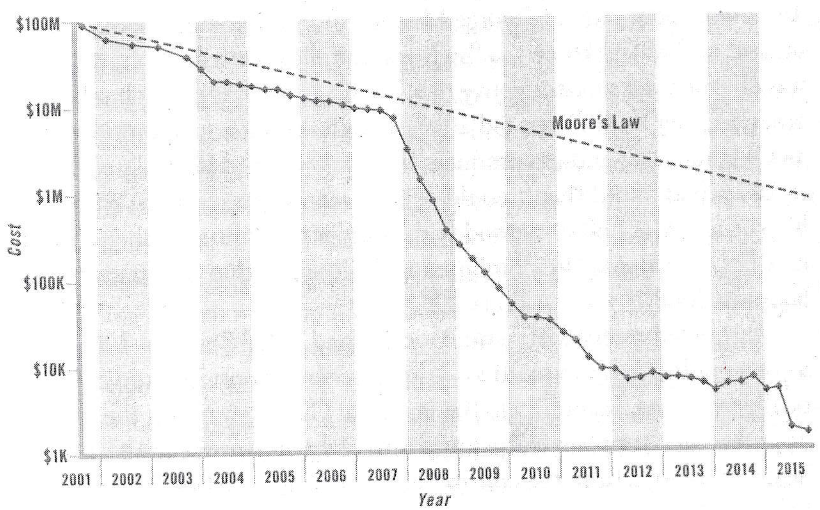

Source: National Human Genome Research Institute

Utility Patent Grants in Biotech, 1963–2014

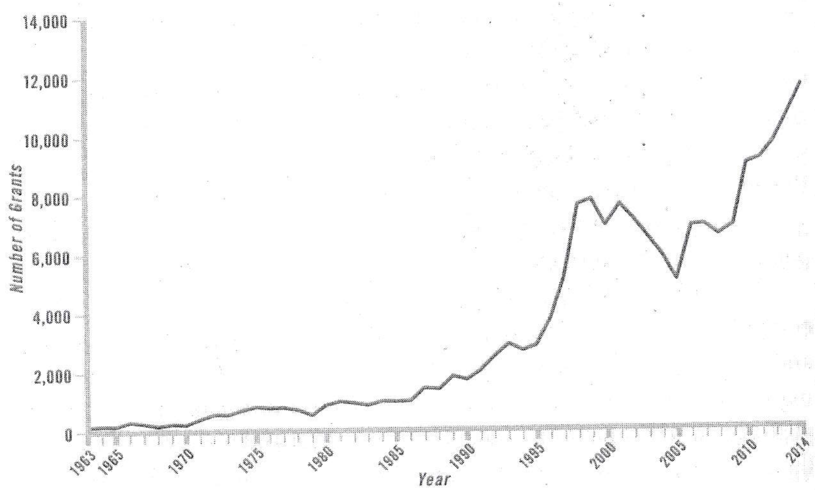

Source: U.S. Patent and Trademark Office

Growth of Solar Power

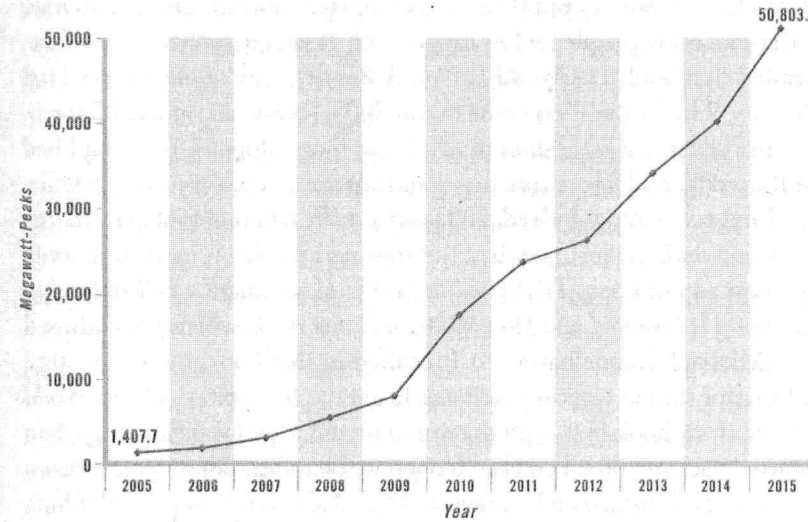

Courtesy of Paula Mints, SVP Market Research

got easier and more natural for people to use than the one before. When the first mainframe computers came out, you needed to have a computer science degree to use them. Today's smartphone can be accessed by young children and the illiterate.

As step changes in technology go, though, the platform birthed around the year 2007 surely constituted one of the greatest leaps forward in history. It suffused a new set of capabilities to connect, collaborate, and create throughout every aspect of life, commerce, and government. Suddenly there were so many more things that could be digitized, so much more storage to hold all that digital data, so many faster computers and so much more innovative software that could process that data for insights, and so many more organizations and people (from the biggest multinationals to the smallest Indian farmers) who could access those insights, or contribute to them, anywhere in the world through their handheld computers—their smartphones.

This is the central technology engine driving the Machine today. It snuck up on us very fast. In 2004, I started writing a book about what I thought then was the biggest force driving the Machine—namely, how the world was getting wired to such a degree that more people in more places had an equal opportunity to compete, connect, and collaborate with more other people for less money with greater ease than ever before. I called that book *The World Is Flat: A Brief History of the Twenty-First Century*. The first edition came out in 2005. I wrote an updated 2.0 edition in 2006 and a 3.0 edition in 2007. And then I stopped, thinking I had built a pretty solid framework that could last me as a columnist for a while.

I was very wrong! Indeed, 2007 was a really bad year to stop thinking.

I first realized just how bad the minute I sat down in 2010 to write my most recent book, *That Used to Be Us: How America Fell Behind in the World It Invented and How We Can Come Back*, which I coauthored with Michael Mandelbaum. As I recalled in that book, the first thing I did when I started working on it was to get the first edition of *The World Is Flat* off my bookshelf—just to remind myself what I was thinking when I started back in 2004. I cracked it open to the index, ran my finger down the page, and immediately discovered that Facebook wasn't in it! That's right—when I was running around in 2004 declaring that the world was flat, Facebook didn't even exist yet, Twitter was still a sound, the cloud was still in the sky, 4G was a parking space, "applications" were what you sent to college, LinkedIn was barely known and most people thought it was a prison, Big Data was a good name for a rap star, and Skype, for most people, was a typographical error. All of those technologies blossomed *after* I wrote *The World Is Flat*—most of them around 2007.

So a few years later, I began updating in earnest my view of how the Machine worked. A crucial impetus was a book I read in 2014 by two MIT business school professors—Erik Brynjolfsson and Andrew McAfee—entitled *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. The first machine age, they argued, was the Industrial Revolution, which accompanied the invention of the steam engine in the 1700s. This period was "all about power systems to augment human muscle," explained McAfee in an interview, "and each successive invention in that age delivered more and more power. But they all required humans to make decisions about them." Therefore, the inventions of that era actually made human control and labor "more valuable and important."

Labor and machines were, broadly speaking, complementary, he added. In the second machine age, though, noted Brynjolfsson, "we

are beginning to automate a lot more cognitive tasks, a lot more of the control systems that determine what to use that power for. In many cases today artificially intelligent machines can make better decisions than humans." So humans and software-driven machines may increasingly be substitutes, not complements.

The key, but not the only, driving force making this possible, they argued, was the exponential growth in computing power as represented by Moore's law: the theory first postulated by Intel cofounder Gordon Moore in 1965 that the speed and power of microchips—that is, computational processing power—would double roughly every year, which he later updated to every two years, for only slightly more money with each new generation. Moore's law has held up close to that pattern for fifty years.

To illustrate this kind of exponential growth, Brynjolfsson and McAfee recalled the famous legend of the king who was so impressed with the man who invented the game of chess that he offered him any reward. The inventor of chess said that all he wanted was enough rice to feed his family. The king said, "Of course, it shall be done. How much would you like?" The man asked the king to simply place a single grain of rice on the first square of a chessboard, then two on the next, then four on the next, with each subsequent square receiving twice as many grains as the previous one. The king agreed, noted Brynjolfsson and McAfee—without realizing that sixty-three instances of doubling yields a fantastically big number: something like eighteen quintillion grains of rice. That is the power of exponential change. When you keep doubling something for fifty years you start to get to some very big numbers, and eventually you start to see some very funky things that you have never seen before.

The authors argued that Moore's law just entered the "second half of the chessboard," where the doubling has gotten so big and fast we're starting to see stuff that is fundamentally different in power and capability from anything we have seen before—self-driving cars, computers that can think on their own and beat any human in chess or *Jeopardy!* or even Go, a 2,500-year-old board game considered vastly more complicated than chess. That is what happens "when the rate of change and the acceleration of the rate of change both increase at the same time," said McAfee, and "we haven't seen anything yet!"

So, at one level, my view of the Machine today is built on the shoulders of Brynjolfsson and McAfee's fundamental insight into how the steady acceleration in Moore's law has affected technology—but I think the Machine today is even more complicated. That's because it's not just pure technological change that has hit the second half of the chessboard. It is also two other giant forces: accelerations in the Market and in Mother Nature.

"The Market" is my shorthand for the acceleration of globalization. That is, global flows of commerce, finance, credit, social networks, and connectivity generally are weaving markets, media, central banks, companies, schools, communities, and individuals more tightly together than ever. The resulting flows of information and knowledge are making the world not only interconnected and hyperconnected but interdependent—everyone everywhere is now more vulnerable to the actions of anyone anywhere.

And "Mother Nature" is my shorthand for climate change, population growth, and biodiversity loss—all of which have also been accelerating, as they, too, enter the second halves of their chessboards.

Here again, I am standing on the shoulders of others. I derive the term "the age of accelerations" from a series of graphs first assembled by a team of scientists led by Will Steffen, a climate change expert and researcher at the Australian National University, Canberra. The graphs, which originally appeared in a 2004 book entitled *Global Change and the Earth System: A Planet Under Pressure,* looked at how technological, social, and environmental impacts were accelerating and feeding off one another from 1750 to 2000, and particularly since 1950. The term "Great Acceleration" was coined in 2005 by these same scientists to capture the holistic, comprehensive, and interlinked nature of all these changes simultaneously sweeping across the globe and reshaping the human and biophysical landscapes of the Earth system. An updated version of those graphs was published in the *Anthropocene Review* on March 2, 2015; they appear on pages 166–167 of this book.

"When we started the project it was ten years since the first accelerations had been published, which ran from 1750 to 2000," explained Owen Gaffney, director of strategy for the Stockholm Resilience Centre, and part of the Great Acceleration team. "We wanted to update the graphs to 2010 to see if the trajectory had altered any"—and indeed it had, he said: it had accelerated.

It is the core argument of this book that these simultaneous accelerations in the Market, Mother Nature, and Moore's law together constitute the "age of accelerations," in which we now find ourselves. These are the central gears driving the Machine today. These three accelerations are impacting one another—more Moore's law is driving more globalization and more globalization is driving more climate change, and more Moore's law is also driving more potential solutions to climate change and a host of other challenges—and at the same time transforming almost every aspect of modern life.

Craig Mundie, a supercomputer designer and former chief of strategy and research at Microsoft, defines this moment in simple physics terms: "The mathematical definition of velocity is the first derivative, and acceleration is the second derivative. So velocity grows or shrinks as a function of acceleration. In the world we are in now, acceleration seems to be increasing. [That means] you don't just move to a higher speed of change. The rate of change also gets faster . . . And when the rate of change eventually exceeds the ability to adapt you get 'dislocation.' 'Disruption' is what happens when someone does something clever that makes you or your company look obsolete. 'Dislocation' is when the whole environment is being altered so quickly that everyone starts to feel they can't keep up."

That is what is happening now. "The world is not just rapidly changing," adds Dov Seidman, "it is being dramatically reshaped—it is starting to operate differently" in many realms all at once. "And this reshaping is happening faster than we have yet been able to reshape ourselves, our leadership, our institutions, our societies, and our ethical choices."

Indeed, there is a mismatch between the change in the pace of change and our ability to develop the learning systems, training systems, management systems, social safety nets, and government regulations that would enable citizens to get the most out of these accelerations and cushion their worst impacts. This mismatch, as we will see, is at the center of much of the turmoil roiling politics and society in both developed and developing countries today. It now constitutes probably the most important governance challenge across the globe.

Astro Teller's Graph

The most illuminating illustration of this phenomenon was sketched out for me by Eric "Astro" Teller, the CEO of Google's X research and development lab, which produced Google's self-driving car, among other innovations. Appropriately enough, Teller's formal title at X is "Captain of Moonshots." Imagine someone whose whole mandate is to come to the office every day and, with his colleagues, produce moonshots—turning what others would consider science fiction into products and services that could transform how we live and work. His paternal grandfather was the physicist Edward Teller, designer of the hydrogen bomb, and his maternal grandfather was Gérard Debreu, a Nobel Prize–winning economist. Good genes, as they say. We were in a conference room at X headquarters, which is a converted shopping mall. Teller arrived at our interview on Rollerblades, which is how he keeps up with his daily crush of meetings.

He wasted no time before launching into an explanation of how the accelerations in Moore's law and in the flow of ideas are together causing an increase in the pace of change that is challenging the ability of human beings to adapt.

Teller began by taking out a small yellow 3M notepad and saying: "Imagine two curves on a graph." He then drew a graph with the Y axis labeled "rate of change" and the X axis labeled "time." Then he drew the first curve—a swooping exponential line that started very flat and escalated slowly before soaring to the upper outer corner of the graph, like a hockey stick: "This line represents scientific progress," he said. At first it moves up very gradually, then it starts to slope higher as innovations build on innovations that have come before, and then it starts to soar straight to the sky.

What would be on that line? Think of the introduction of the printing press, the telegraph, the manual typewriter, the Telex, the mainframe computer, the first word processors, the PC, the Internet, the laptop, the mobile phone, search, mobile apps, big data, virtual reality, human-genome sequencing, artificial intelligence, and the self-driving car.

A thousand years ago, Teller explained, that curve representing scientific and technological progress rose so gradually that it could take one hundred years for the world to look and feel dramatically different. For instance, it took centuries for the longbow to go from development into military use in Europe in the late thirteenth century. If you lived in the twelfth century, your basic life was not all that different than if you lived in the eleventh century. And whatever changes were being introduced in major towns in Europe or Asia took forever to reach the countryside, let alone the far reaches of Africa or South America. Nothing scaled globally all at once.

But by 1900, Teller noted, this process of technological and scientific change "started to speed up" and the curve started to accelerate upward. "That's because technology stands on its own shoulders—each generation of invention stands on the inventions that have come before," said Teller. "So by 1900, it was taking twenty to thirty years for technology to take one step big enough that the world became uncomfortably different. Think of the introduction of the car and the airplane."

Then the slope of the curve started to go almost straight up and off the graph with the convergence of mobile devices, broadband connectivity, and cloud computing (which we will discuss shortly). These developments diffused the tools of innovation to many more people on the planet, enabling them to drive change farther, faster, and more cheaply.

"Now, in 2016," he added, "that time window—having continued to shrink as each technology stood on the shoulders of past technologies—has become so short that it's on the order of five to seven years from the time something is introduced to being ubiquitous and the world being uncomfortably changed."

What does this process feel like? In my first book about globalization, *The Lexus and the Olive Tree*, I included a story Lawrence Summers told me that captured the essence of where we'd come from and where we were heading. It was 1988, Summers recalled, and he was working on the Michael Dukakis presidential campaign, which sent him to Chicago to give a speech. A car picked him up at the airport to take him to the event, and when he slipped into the car he discovered a telephone fixed into the backseat. "I thought it was sufficiently neat to have a cell phone in my car in 1988 that I used it to call my wife to tell her that I was in a car with a phone," Summers told me. He also used it to call everyone else he could think of, and they were just as excited.

Just nine years later Summers was deputy treasury secretary. On a trip to the Ivory Coast in West Africa, he had to inaugurate an American-funded health care project in a village upriver from the main city, Abidjan, that was opening its first potable water well. What he remembered most, though, he told me, was that on his way back from the village, as he stepped into a dugout canoe to return downriver, an Ivory Coast official handed him a cell phone and said: "Washington has a question for you." In nine years Summers went from bragging that he was in a car with a mobile phone in Chicago to nonchalantly using one in the backseat of his dugout canoe in Abidjan. The pace of change had not only quickened but was now happening at a global scale.

That Other Line

So that is what is going on with scientific and technological progress—but Teller wasn't done drawing his graph for me. He'd promised two lines, and he now drew the second, a straight line that began many years ago above the scientific progress line but since then had climbed far more incrementally, so incrementally you could barely detect its positive slope.

"The good news is that there is a competing curve," Teller explained. "This is the rate at which humanity—individuals and society—adapts to changes in its environment." These, he added, can be technological changes (mobile connectivity), geophysical changes (such as the Earth warming and cooling), or social changes (there was a time when we weren't okay with mixed-race marriages, at least here in the United States). "Many of those major changes were driven by society, and we have adapted. Some were more or less uncomfortable. But we adapted."

Indeed, the good news is that we've gotten a little bit faster at adapting over the centuries, thanks to greater literacy and knowledge diffusion. "The rate at which we can adapt is increasing," said Teller. "A thousand years ago, it probably would have taken two or three generations to adapt to something new." By 1900, the time it took to adapt got down to one generation. "We might be so adaptable now," said Teller, "that it only takes ten to fifteen years to get used to something new."

Alas, though, that may not good enough. Today, said Teller, the accelerating speed of scientific and technological innovations (and, I would add, new ideas, such as gay marriage) can outpace the capacity of the average human being and our societal structures to adapt and absorb them. With that thought in mind, Teller added one more thing to the graph—a big dot. He drew that dot on the rapidly sloping technology curve just above the place where it intersected with the adaptability line.

He labeled it: "We are here." The graph, as redrawn for this book, can be seen on the next page.

That dot, Teller explained, illustrates an important fact: even though human beings and societies have steadily adapted to change, on average, the rate of technological change is now accelerating so fast that it has risen above the average rate at which most people can absorb all these changes. Many of us cannot keep pace anymore.

"And that is causing us cultural angst," said Teller. "It's also preventing us from fully benefiting from all of the new technology that is coming along every day . . . In the decades following the invention of the internal combustion engine—before the streets were flooded with mass-produced cars—traffic laws and conventions were gradually put into place. Many of those laws and conventions continue to serve us well today, and over the course of a century, we had plenty of time to adapt our laws to new inventions, such as freeways. Today, however, scientific advances are bringing seismic shifts to the ways in which we use our roads; legislatures and municipalities are scrambling to keep up, tech companies are chafing under outdated and sometimes nonsensical rules, and the public is not sure what to think. Smartphone technology gave rise

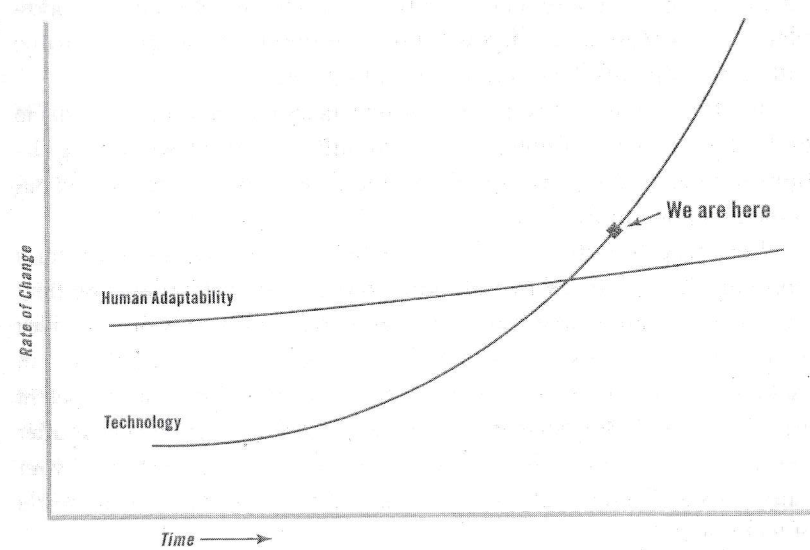

to Uber, but before the world figures out how to regulate ride-sharing, self-driving cars will have made those regulations obsolete."

This is a real problem. When fast gets really fast, being slower to adapt makes you really slow—and disoriented. It is as if we were all on one of those airport moving sidewalks that was going around five miles an hour and suddenly it sped up to twenty-five miles an hour—even as everything else around it stayed roughly the same. That is really disorienting for a lot of people.

If the technology platform for society can now turn over in five to seven years, but it takes ten to fifteen years to adapt to it, Teller explained, "we will all feel out of control, because we can't adapt to the world as fast as it's changing. By the time we get used to the change, that won't even be the prevailing change anymore—we'll be on to some new change."

That is dizzying for many people, because they hear about advances such as robotic surgery, gene editing, cloning, or artificial intelligence, but have no idea where these developments will take us.

"None of us have the capacity to deeply comprehend more than one of these fields—the sum of human knowledge has far outstripped any single individual's capacity to learn—and even the experts in these fields can't predict what will happen in the next decade or century," said Teller. "Without clear knowledge of the future potential or future unintended negative consequences of new technologies, it is nearly impossible to draft regulations that will promote important advances—while still protecting ourselves from every bad side effect."

In other words, if it is true that it now takes us ten to fifteen years to understand a new technology and then build out new laws and regulations to safeguard society, how do we regulate when the technology has come and gone in five to seven years? This is a problem.

Let's take patents as one example of a system that was built for a world in which changes arrived more slowly, explained Teller. The standard patent arrangement was: "We'll give you a monopoly on your idea for twenty years"—usually minus time to issue the actual patent—"in exchange for which people will get to know the information in the patent after it expires." But what if most new technologies are obsolete after four to five years, asked Teller, "and it takes four to five years to get your patents issued? That makes patents increasingly irrelevant in the world of technology."

Another big challenge is the way we educate our population. We go to school for twelve or more years during our childhoods and early adulthoods, and then we're done. But when the pace of change gets this fast, the only way to retain a lifelong working capacity is to engage in lifelong learning. There is a whole group of people—judging from the 2016 U.S. election—who "did not join the labor market at age twenty thinking they were going to have to do lifelong learning," added Teller, and they are not happy about it.

All of these are signs "that our societal structures are failing to keep pace with the rate of change," he said. Everything feels like it's in constant catch-up mode. What to do? We certainly don't want to slow down technological progress or abandon regulation. The only adequate response, said Teller, "is that we try to increase our society's ability to adapt." That is the only way to release us from the society-wide anxiety around tech. "We can either push back against technological advances," argued Teller, "or we can acknowledge that humanity has a new challenge: we must rewire our societal tools and institutions so that they will enable us to keep pace. The first option—trying to slow technology—may seem like the easiest solution to our discomfort with change, but humanity is facing some catastrophic environmental problems of its own making, and burying our heads in the sand won't end well. Most of the solutions to the big problems in the world will come from scientific progress."

If we could "enhance our ability to adapt even slightly," he continued, "it would make a significant difference." He then returned to our graph and drew a dotted line that rose up alongside the adaptability line but faster. This line simulated our learning faster as well as governing smarter, and therefore intersected with the technology/science change line at a higher point.

Enhancing humanity's adaptability, argued Teller, is 90 percent about "optimizing for learning"—applying features that drive technological innovation to our culture and social structures. Every institution, whether it is the patent office, which has improved a lot in recent years, or any other major government regulatory body, has to keep getting more agile—it has to be willing to experiment quickly and learn from mistakes. Rather than expecting new regulations to last for decades, it should continuously reevaluate the ways in which they serve society. Universities are now experimenting with turning over their curriculum much faster and more

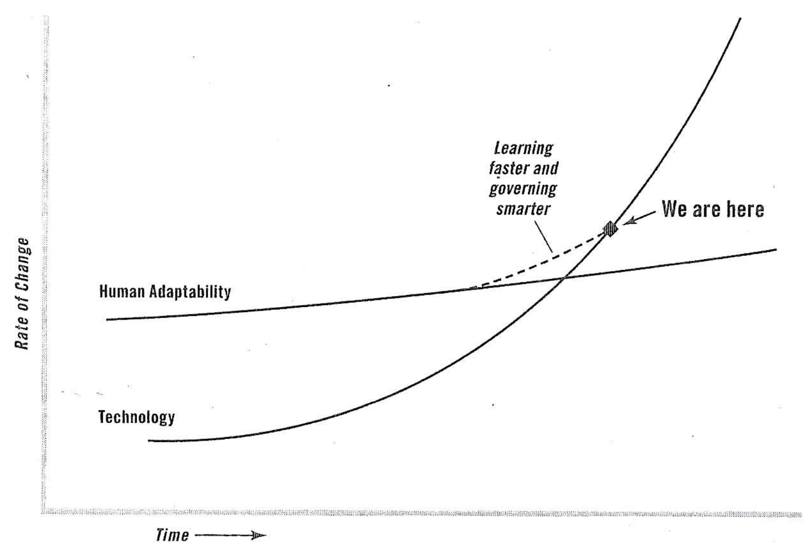

often to keep up with the change in the pace of change—putting a "use-by date" on certain courses. Government regulators need to take a similar approach. They need to be as innovative as the innovators. They need to operate at the speed of Moore's law.

"Innovation," Teller said, "is a cycle of experimenting, learning, applying knowledge, and then assessing success or failure. And when the outcome is failure, that's just a reason to start the cycle over again." One of X's mottos is "Fail fast." Teller tells his teams: "I don't care how much progress you make this month; my job is to cause your rate of improvement to increase—how do we make the same mistake in half the time for half the money?"

In sum, said Teller, what we are experiencing today, with shorter and shorter innovation cycles, and less and less time to learn to adapt, "is the difference between a constant state of destabilization versus occasional destabilization." The time of static stability has passed us by, he added. That does not mean we can't have a new kind of stability, "but the new kind of stability has to be dynamic stability. There are some ways of being, like riding a bicycle, where you cannot stand still, but once you are moving it is actually easier. It is not our natural state. But humanity has to learn to exist in this state."

We're all going to have to learn that bicycle trick.

When that happens, said Teller, "in a weird way we will be calm again, but it will take substantial relearning. We definitely don't train our children for dynamic stability."

We will need to do that, though, more and more, if we want future generations to thrive and find their own equilibrium. The next four chapters are about the underlying accelerations in Moore's law, the Market, and Mother Nature that define how the Machine works today. If we are going to achieve the dynamic stability Teller speaks of, we must understand how these forces are reshaping the world, and why they became particularly dynamic—beginning around 2007.