THE GLASS CAGE

AUTOMATION AND US

NICHOLAS CARR

New York Times best-selling author of *THE SHALLOWS*

THE GLASS CAGE

Automation and Us

Nicholas Carr

W. W. NORTON & COMPANY
NEW YORK LONDON

To Ann

CONTENTS

INTRODUCTION

ALERT FOR OPERATORS

CHAPTER ONE

PASSENGERS

CHAPTER TWO

THE ROBOT AT THE GATE

CHAPTER THREE

ON AUTOPILOT

CHAPTER FOUR

THE DEGENERATION EFFECT

INTERLUDE, WITH DANCING MICE

CHAPTER FIVE

WHITE-COLLAR COMPUTER

CHAPTER SIX

WORLD AND SCREEN

CHAPTER SEVEN

AUTOMATION FOR THE PEOPLE

INTERLUDE, WITH GRAVE ROBBER

CHAPTER EIGHT

YOUR INNER DRONE

CHAPTER NINE

THE LOVE THAT LAYS THE SWALE IN ROWS

NOTES

ACKNOWLEDGMENTS

INDEX

No one

to witness

and adjust, no one to drive the car

—*William Carlos Williams*

THE GLASS CAGE

INTRODUCTION

ALERT FOR OPERATORS

ON JANUARY 4, 2013, the first Friday of a new year, a dead day newswise, the Federal Aviation Administration released a one-page notice. It had no title. It was identified only as a “safety alert for operators,” or SAFO. Its wording was terse and cryptic. In addition to being posted on the FAA’s website, it was sent to all U.S. airlines and other commercial air carriers. “This SAFO,” the document read, “encourages operators to promote manual flight operations when appropriate.” The FAA had collected evidence, from crash investigations, incident reports, and cockpit studies, indicating that pilots had become too dependent on autopilots and other computerized systems. Overuse of flight automation, the agency warned, could “lead to degradation of the pilot’s ability to quickly recover the aircraft from an undesired state.” It could, in blunter terms, put a plane and its passengers in jeopardy. The alert concluded with a recommendation that airlines, as a matter of operational policy, instruct pilots to spend less time flying on autopilot and more time flying by hand.[1]

This is a book about automation, about the use of computers and software to do things we used to do ourselves. It’s not about the technology or the economics of automation, nor is it about the future of robots and cyborgs and gadgetry, though all those things enter into the story. It’s about automation’s human consequences. Pilots have been out in front of a wave that is now engulfing us. We’re looking to computers to shoulder more of our work, on the job and off, and to guide us through more of our everyday routines. When we need to get something done today, more often than not we sit down in front of a monitor, or open a laptop, or pull out a smartphone, or strap a net-connected accessory to our forehead or wrist. We run apps. We consult screens. We take advice from digitally simulated voices. We defer to the wisdom of algorithms.

Computer automation makes our lives easier, our chores less burdensome. We’re often able to accomplish more in less time—or to do things we simply couldn’t do before. But automation also has deeper, hidden effects. As aviators have learned, not all of them are beneficial. Automation can take a toll on our work, our talents, and our lives. It can narrow our perspectives and limit our choices. It can open us to surveillance and manipulation. As computers become our constant companions, our familiar, obliging helpmates, it seems wise to take a closer look at exactly how they’re changing what we do and who we are.

CHAPTER ONE

PASSENGERS

AMONG THE HUMILIATIONS OF MY TEENAGE YEARS WAS ONE that might be termed psycho-mechanical: my very public struggle to master a manual transmission. I got my driver's license early in 1975, not long after I turned sixteen. The previous fall, I had taken a driver's ed course with a group of my high-school classmates. The instructor's Oldsmobile, which we used for our on-the-road lessons and then for our driving tests at the dread Department of Motor Vehicles, was an automatic. You pressed the gas pedal, you turned the wheel, you hit the brakes. There were a few tricky maneuvers—making a three-point turn, backing up in a straight line, parallel parking—but with a little practice among pylons in the school parking lot, even they became routine.

License in hand, I was ready to roll. There was just one last roadblock. The only car available to me at home was a Subaru sedan with a stick shift. My dad, not the most hands-on of parents, granted me a single lesson. He led me out to the garage one Saturday morning, plopped himself down behind the wheel, and had me climb into the passenger seat beside him. He placed my left palm over the shift knob and guided my hand through the gears: "That's first." Brief pause. "Second." Brief pause. "Third." Brief pause. "Fourth." Brief pause. "Down over here"—a pain shot through my wrist as it twisted into an unnatural position—"is Reverse." He glanced at me to confirm I had it all down. I nodded helplessly. "And that"—wiggling my hand back and forth—"that's Neutral." He gave me a few tips about the speed ranges of the four forward gears. Then he pointed to the clutch pedal he had pinned beneath his loafer. "Make sure you push that in while you shift."

I proceeded to make a spectacle of myself on the roads of the small New England town where we lived. The car would buck as I tried to find the correct gear, then lurch forward as I mistimed the release of the clutch. I'd stall at every red light, then stall again halfway out into the intersection. Hills were a horror. I'd let the clutch out too quickly, or too slowly, and the car would roll backward until it came to rest against the bumper of the vehicle behind me. Horns were honked, curses cursed, birds flipped. What made the experience all the more excruciating was the Subaru's yellow paint job—the kind of yellow you get with a kid's rain slicker or a randy male goldfinch. The car was an eye magnet, my flailing impossible to miss.

From my putative friends, I received no sympathy. They found my struggles a source of endless, uproarious amusement. "Grind me a pound!" one of them would yell with glee from the backseat whenever I'd muff a shift and set off a metallic gnashing of gear teeth. "Smooth move," another would snigger as the engine rattled to a stall. The word "spaz"—this was well before anyone had heard of political correctness—was frequently lobbed my way. I had a suspicion that my incompetence with the stick was something my buddies laughed about behind my back. The metaphorical implications were not lost on me. My manhood, such as it was at sixteen, felt deflated.

But I persisted—what choice did I have?—and after a week or two I began to get the hang of it. The gearbox loosened up and became more forgiving. My arms and legs stopped working at cross-purposes and started cooperating. Soon, I was shifting without thinking about it. It just happened. The car no longer stalled or bucked or lurched. I no longer had to sweat the hills or the intersections. The transmission and I had become a team. We meshed. I took a quiet pride in my accomplishment.

Still, I coveted an automatic. Although stick shifts were fairly common back then, at least in the econoboxes and junkers that kids drove, they had already taken on a behind-the-times, hand-me-down quality. They seemed fusty, a little yesterday. Who wanted to be "manual" when you could be "automatic"? It was like the difference between scrubbing dishes by hand and sticking them in a dishwasher. As it turned out, I didn't have to wait long for my wish to be granted. Two years after I got my license, I managed to total the Subaru during a late-night misadventure, and not long afterward I took stewardship of a used, cream-colored, two-door Ford Pinto. The car was a piece of crap—some now see the Pinto as marking the nadir of American manufacturing in the twentieth century—but to me it was redeemed by its automatic transmission.

I was a new man. My left foot, freed from the demands of the clutch, became an appendage of leisure. As I tooled around town, it would sometimes tap along jauntily to the thwacks of Charlie Watts or the thuds of John Bonham—the Pinto also had a built-in eight-track deck, another touch of modernity—but more often than not it just stretched out in its little nook under the left side of the dash and napped. My right hand became a beverage holder. I not only felt renewed and up-to-date. I felt liberated.

It didn't last. The pleasures of having less to do were real, but they faded. A new emotion set in: boredom. I didn't admit it to anyone, hardly to myself even, but I began to miss the gear stick and the clutch pedal. I missed the sense of control and involvement they had given me—the ability to rev the engine as high as I wanted, the feel of the clutch releasing and the gears grabbing, the tiny thrill that came with a downshift at speed. The automatic made me feel a little less like a driver and a little more like a passenger. I came to resent it.

■ ■ ■ ■

MOTOR AHEAD thirty-five years, to the morning of October 9, 2010. One of Google's in-house inventors, the German-born roboticist Sebastian Thrun, makes an extraordinary announcement in a blog post. The company has developed "cars that can drive themselves." These aren't some gawky, gearhead prototypes puttering around the Googleplex's parking lot. These are honest-to-goodness street-legal vehicles—Priuses, to be precise—and, Thrun reveals, they've already logged more than a hundred thousand miles on roads and highways in California and Nevada. They've cruised down Hollywood Boulevard and the Pacific Coast Highway, gone back and forth over the Golden Gate Bridge, circled Lake Tahoe. They've merged into freeway traffic, crossed busy intersections, and inched through rush-hour gridlock. They've swerved to avoid collisions. They've done all this by themselves. Without human help. "We think this is a first in robotics research," Thrun writes, with sly humility.[1]

Building a car that can drive itself is no big deal. Engineers and tinkerers have been constructing robotic and remote-controlled automobiles since at least the 1980s. But most of them were crude jalopies. Their use was restricted to test-drives on closed tracks or to races and rallies in deserts and other remote areas, far away from pedestrians and police. The Googlemobile, Thrun's announcement made clear, is different. What makes it such a breakthrough, in the history of both transport and automation, is its ability to navigate the real world in all its chaotic, turbulent complexity. Outfitted with laser range-finders, radar and sonar transmitters, motion detectors, video cameras, and GPS receivers, the car can sense its surroundings in minute detail. It can see where it's going. And by processing all the streams of incoming information instantaneously—in "real time"—its onboard computers are able to work the accelerator, the steering wheel, and the brakes with the speed and sensitivity required to drive on actual roads and respond fluidly to the unexpected events that drivers always encounter. Google's fleet of self-driving cars has now racked up close to a million miles, and the vehicles have caused just one serious accident. That was a five-car pileup near the company's Silicon Valley headquarters in 2011, and it doesn't really count. It happened, as Google was quick to announce, "while a person was manually driving the car."[2]

Autonomous automobiles have a ways to go before they start chauffeuring us to work or ferrying our kids to soccer games. Although Google has said it expects commercial versions of its car to be on sale by the end of the decade, that's probably wishful thinking. The vehicle's sensor systems remain prohibitively expensive, with the roof-mounted laser apparatus alone going for eighty thousand dollars. Many technical challenges remain to be met, such as navigating snowy or leaf-covered roads, dealing with unexpected detours, and interpreting the hand signals of traffic cops and road workers. Even the most powerful computers still have a hard time distinguishing a bit of harmless road debris (a flattened cardboard box, say) from a dangerous obstacle (a nail-studded chunk of plywood). Most daunting of all are the many legal, cultural, and ethical hurdles a driverless car faces. Where, for instance, will culpability and liability reside should a computer-driven automobile cause an accident that kills or injures someone? With the car's owner? With the manufacturer that installed the self-driving system? With the programmers who wrote the software? Until such thorny questions get sorted out, fully automated cars are unlikely to grace dealer showrooms.

Progress will sprint forward nonetheless. Much of the Google test cars' hardware and software will come to be incorporated into future generations of cars and trucks. Since the company went public with its autonomous vehicle program, most of the world's major carmakers have let it be known that they have similar efforts under way. The goal, for the time being, is not so much to create an immaculate robot-on-wheels as to continue to invent and refine automated features that enhance safety and convenience in ways that get people to buy new cars. Since I first turned the key in my Subaru's ignition, the automation of driving has already come a long way. Today's automobiles are stuffed with electronic gadgetry. Microchips and sensors govern the workings of the cruise control, the antilock brakes, the traction and stability mechanisms, and, in higher-end models, the variable-speed transmission, parking-assist system, collision-avoidance system, adaptive headlights, and dashboard displays. Software already provides a buffer between us and the road. We're not so much controlling our cars as sending electronic inputs to the computers that control them.

In coming years, we'll see responsibility for many more aspects of driving shift from people to software. Luxury-car makers like Infiniti, Mercedes, and Volvo are rolling out models that combine radar-assisted adaptive cruise control, which works even in stop-and-go traffic, with computerized steering systems that keep a car centered in its lane and brakes that slam themselves on in emergencies. Other manufacturers are rushing to introduce even more advanced controls. Tesla Motors, the electric car pioneer, is developing an automotive autopilot that "should be able to [handle] 90 percent of miles driven," according to the company's ambitious chief executive, Elon Musk.[3]

The arrival of Google's self-driving car shakes up more than our conception of driving. It forces us to change our thinking about what computers and robots can and can't do. Up until that fateful October day, it was taken for granted that many important skills lay beyond the reach of automation. Computers could do a lot of things, but they couldn't do everything. In an influential 2004 book, *The New Division of Labor: How Computers Are Creating the Next Job Market*, economists Frank Levy and Richard Murnane argued, convincingly, that there were practical limits to the ability of software programmers to replicate human talents, particularly those involving sensory perception, pattern recognition, and conceptual knowledge. They pointed specifically to the example of driving a car on the open road, a talent that requires the instantaneous interpretation of a welter of visual signals and an ability to adapt seamlessly to shifting and often unanticipated situations. We hardly know how we pull off such a feat ourselves, so the idea that programmers could reduce all of driving's intricacies, intangibilities, and contingencies to a set of instructions, to lines of software code, seemed ludicrous. "Executing a left turn across oncoming traffic," Levy and Murnane wrote, "involves so many factors that it is hard to imagine the set of rules that can replicate a driver's behavior." It seemed a sure bet, to them and to pretty much everyone else, that steering wheels would remain firmly in the grip of human hands.[4]

In assessing computers' capabilities, economists and psychologists have long drawn on a basic distinction between two kinds of knowledge: *tacit* and *explicit*. Tacit knowledge, which is also sometimes called procedural knowledge, refers to all the stuff we do without thinking about it: riding a bike, snagging a fly ball, reading a book, driving a car. These aren't innate skills—we have to learn them, and some people are better at them than others—but they can't be expressed as a simple recipe. When you make a turn through a busy intersection in your car, neurological studies show, many areas of your brain are hard at work, processing sensory stimuli, making estimates of time and distance, and coordinating your arms and legs.[5] But if someone asked you to document everything involved in making that turn, you wouldn't be able to, at least not without resorting to generalizations and abstractions. The ability resides deep in your nervous system, outside the ambit of your conscious mind. The mental processing goes on without your awareness.

Much of our ability to size up situations and make quick judgments about them stems from the fuzzy realm of tacit knowledge. Most of our creative and artistic skills reside there too. Explicit knowledge, which is also known as declarative knowledge, is the stuff you can actually write down: how to change a flat tire, how to fold an origami crane, how to solve a quadratic equation. These are processes that can be broken down into well-defined steps. One person can explain them to another person through written or oral instructions: do this, then this, then this.

Because a software program is essentially a set of precise, written instructions—do this, then this, then this—we've assumed that while computers can replicate skills that depend on explicit knowledge, they're not so good when it comes to skills that flow from tacit knowledge. How do you translate the ineffable into lines of code, into the rigid, step-by-step instructions of an algorithm? The boundary between the explicit and the tacit has always been a rough one—a lot of our talents straddle the line—but it seemed to offer a good way to define the limits of automation and, in turn, to mark out the exclusive precincts of the human. The sophisticated jobs Levy and Murnane identified as lying beyond the reach of computers—in addition to driving, they pointed to teaching and medical diagnosis—were a mix of the mental and the manual, but they all drew on tacit knowledge.

Google's car resets the boundary between human and computer, and it does so more dramatically, more decisively, than have earlier breakthroughs in programming. It tells us that our idea of the limits of automation has always been something of a fiction. We're not as special as we think we are. While the distinction between tacit and explicit knowledge remains a useful one in the realm of human psychology, it has lost much of its relevance to discussions of automation.

■ ■ ■ ■

THAT DOESN'T mean that computers now have tacit knowledge, or that they've started to think the way we think, or that they'll soon be able to do everything people can do. They don't, they haven't, and they won't. Artificial intelligence is not human intelligence. People are mindful; computers are mindless. But when it comes to performing demanding tasks, whether with the brain or the body, computers are able to replicate our ends without replicating our means. When a driverless car makes a left turn in traffic, it's not tapping into a well of intuition and skill; it's following a program. But while the strategies are different, the outcomes, for practical purposes, are the same. The superhuman speed with which computers can follow instructions, calculate probabilities, and receive and send data means that they can use explicit knowledge to perform many of the complicated tasks that we do with tacit knowledge. In some cases, the unique strengths of computers allow them to perform what we consider to be tacit skills better than we can perform them ourselves. In a world of computer-controlled cars, you wouldn't need traffic lights or stop signs. Through the continuous, high-speed exchange of data, vehicles would seamlessly coordinate their passage through even the busiest of intersections—just as computers today regulate the flow of inconceivable numbers of data packets along the highways and byways of the internet. What's ineffable in our own minds becomes altogether effable in the circuits of a microchip.

Many of the cognitive talents we've considered uniquely human, it turns out, are anything but. Once computers get quick enough, they can begin to mimic our ability to spot patterns, make judgments, and learn from experience. We were first taught that lesson back in 1997 when IBM's Deep Blue chess-playing supercomputer, which could evaluate a billion possible moves every five seconds, beat the world champion Garry Kasparov. With Google's intelligent car, which can process a million environmental readings a second, we're learning the lesson again. A lot of the very smart things that people do don't actually require a brain. The intellectual talents of highly trained professionals are no more protected from automation than is the driver's left turn. We see the evidence everywhere. Creative and analytical work of all sorts is being mediated by software. Doctors use computers to diagnose diseases. Architects use them to design buildings. Attorneys use them to evaluate evidence. Musicians use them to simulate instruments and correct bum notes. Teachers use them to tutor students and grade papers. Computers aren't taking over these professions entirely, but they are taking over many aspects of them. And they're certainly changing the way the work is performed.

It's not only vocations that are being computerized. Avocations are too. Thanks to the proliferation of smartphones, tablets, and other small, affordable, and even wearable computers, we now depend on software to carry out many of our daily chores and pastimes. We launch apps to aid us in shopping, cooking, exercising, even finding a mate and raising a child. We follow turn-by-turn GPS instructions to get from one place to the next. We use social networks to maintain friendships and express our feelings. We seek advice from recommendation engines on what to watch, read, and listen to. We look to Google, or to Apple's Siri, to answer our questions and solve our problems. The computer is becoming our all-purpose tool for navigating, manipulating, and understanding the world, in both its physical and its social manifestations. Just think what happens these days when people misplace their smartphones or lose their connections to the net. Without their digital assistants, they feel helpless. As Katherine Hayles, a literature professor at Duke University, observed in her 2012 book *How We Think*, "When my computer goes down or my Internet connection fails, I feel lost, disoriented, unable to work—in fact, I feel as if my hands have been amputated." [6]

Our dependency on computers may be disconcerting at times, but in general we welcome it. We're eager to celebrate and show off our whizzy new gadgets and apps—and not only because they're so useful and so stylish. There's something magical about computer automation. To watch an iPhone identify an obscure song playing over the sound system in a bar is to experience something that would have been inconceivable to any previous generation. To see a crew of brightly painted factory robots effortlessly assemble a solar panel or a jet engine is to view an exquisite heavy-metal ballet, each movement choreographed to a fraction of a millimeter and a sliver of a second. The people who have taken rides in Google's car report that the thrill is almost otherworldly; their earth-bound brain has a tough time processing the experience. Today, we really do seem to be entering a brave new world, a Tomorrowland where computers and automatons will be at our service, relieving us of our burdens, granting our wishes, and sometimes just keeping us company. Very soon now, our Silicon Valley wizards assure us, we'll have robot maids as well as robot chauffeurs. Sundries will be fabricated by 3-

D printers and delivered to our doors by drones. The world of the *Jetsons*, or at least of *Knight Rider*, beckons.

It's hard not to feel awestruck. It's also hard not to feel apprehensive. An automatic transmission may seem a paltry thing beside Google's tricked-out, look-ma-no-humans Prius, but the former was a precursor to the latter, a small step along the path to total automation, and I can't help but remember the letdown I felt after the gear stick was taken from my hand—or, to put responsibility where it belongs, after I begged to have the gear stick taken from my hand. If the convenience of an automatic transmission left me feeling a little lacking, a little *underutilized*, as a labor economist might say, how will it feel to become, truly, a passenger in my own car?

■ ■ ■ ■

THE TROUBLE with automation is that it often gives us what we don't need at the cost of what we do. To understand why that's so, and why we're eager to accept the bargain, we need to take a look at how certain cognitive biases—flaws in the way we think—can distort our perceptions. When it comes to assessing the value of labor and leisure, the mind's eye can't see straight.

Mihaly Csikszentmihalyi, a psychology professor and author of the popular 1990 book *Flow*, has described a phenomenon that he calls "the paradox of work." He first observed it in a study he conducted in the 1980s with his University of Chicago colleague Judith LeFevre. They recruited a hundred workers, blue-collar and white-collar, skilled and unskilled, from five businesses around Chicago. They gave each an electronic pager (this was when cell phones were still luxury goods) that they had programmed to beep at seven random moments a day over the course of a week. At each beep, the subjects would fill out a short questionnaire. They'd describe the activity they were engaged in at that moment, the challenges they were facing, the skills they were deploying, and the psychological state they were in, as indicated by their sense of motivation, satisfaction, engagement, creativity, and so forth. The intent of this "experience sampling," as Csikszentmihalyi termed the technique, was to see how people spend their time, on the job and off, and how their activities influence their "quality of experience."

The results were surprising. People were happier, felt more fulfilled by what they were doing, while they were at work than during their leisure hours. In their free time, they tended to feel bored and anxious. And yet they didn't like to be at work. When they were on the job, they expressed a strong desire to be off the job, and when they were off the job, the last thing they wanted was to go back to work. "We have," reported Csikszentmihalyi and LeFevre, "the paradoxical situation of people having many more positive feelings at work than in leisure, yet saying that they 'wish to be doing something else' when they are at work, not when they are in leisure."[7] We're terrible, the experiment revealed, at anticipating which activities will satisfy us and which will leave us discontented. Even when we're in the midst of doing something, we don't seem able to judge its psychic consequences accurately.

Those are symptoms of a more general affliction, on which psychologists have bestowed the poetic name *miswanting*. We're inclined to desire things we don't like and to like things we don't desire. "When the things we want to happen do not improve our happiness, and when the things we want not to happen do," the cognitive psychologists Daniel Gilbert and Timothy Wilson have observed, "it seems fair to say we have wanted badly."[8] And as slews of gloomy studies show, we're forever wanting badly. There's also a social angle to our tendency to misjudge work and leisure. As Csikszentmihalyi and LeFevre discovered in their experiments, and as most of us know from our own experience, people allow themselves to be guided by social conventions—in this case, the deep-seated idea that being "at leisure" is more desirable, and carries more status, than being "at work"—rather than by their true feelings. "Needless to say," the researchers concluded, "such a blindness to the real state of affairs is likely to have unfortunate consequences for both individual well-being and the health of society." As people act on their skewed perceptions, they will "try to do more of those activities that provide the least positive experiences and avoid the activities that are the source of their most positive and intense feelings."[9] That's hardly a recipe for the good life.

It's not that the work we do for pay is intrinsically superior to the activities we engage in for diversion or entertainment. Far from it. Plenty of jobs are dull and even demeaning, and plenty of hobbies and pastimes are stimulating and fulfilling. But a job imposes a structure on our time that we lose when we're left to our own devices. At work, we're pushed to engage in the kinds of activities that human beings find most satisfying. We're happiest when we're absorbed in a difficult task, a task that has clear goals and that challenges us not only to exercise our talents but to stretch them. We become so immersed in the flow of our work, to use Csikszentmihalyi's term, that we tune out distractions and transcend the anxieties and worries that plague our everyday lives. Our usually wayward attention becomes fixed on what we're doing. "Every action, movement, and thought follows inevitably from the previous one," explains Csikszentmihalyi. "Your whole being is involved, and you're using your skills to the utmost."[10] Such states of deep absorption can be produced by all manner of effort, from laying tile to singing in a choir to racing a dirt bike. You don't have to be earning a wage to enjoy the transports of flow.

More often than not, though, our discipline flags and our mind wanders when we're not on the job. We may yearn for the workday to be over so we can start spending our pay and having some fun, but most of us fritter away our leisure hours. We shun hard work and only rarely engage in challenging hobbies. Instead, we watch TV or go to the mall or log on to Facebook. We get lazy. And then we get bored and fretful. Disengaged from any outward focus, our attention turns inward, and we end up locked in what Emerson called the jail of self-consciousness. Jobs, even crummy ones, are "actually easier to enjoy than free time," says Csikszentmihalyi, because they have the "built-in" goals and challenges that "encourage one to become involved in one's work, to concentrate and lose oneself in it."[11] But that's not what our deceiving minds want us to believe. Given the opportunity, we'll eagerly relieve ourselves of the rigors of labor. We'll sentence ourselves to idleness.

■ ■ ■ ■

IS IT any wonder we're enamored of automation? By offering to reduce the amount of work we have to do, by promising to imbue our lives with greater ease, comfort, and convenience, computers and other labor-saving technologies appeal to our eager but misguided desire for release from what we perceive as toil. In the workplace, automation's focus on enhancing speed and efficiency—a focus determined by the profit motive rather than by any particular concern for people's well-being—often has the effect of removing complexity from jobs, diminishing the challenge they present and hence the engagement they promote. Automation can narrow people's responsibilities to the point that their jobs consist largely of monitoring a computer screen or entering data into prescribed fields. Even highly trained analysts and other so-called knowledge workers are seeing their work circumscribed by decision-support systems that turn the making of judgments into a data-processing routine. The apps and other programs we use in our private lives have similar effects. By taking over difficult or time-consuming tasks, or simply rendering those tasks less onerous, the software makes it even less likely that we'll engage in efforts that test our skills and give us a sense of accomplishment and satisfaction. All too often, automation frees us from that which makes us feel free.

The point is not that automation is bad. Automation and its precursor, mechanization, have been marching forward for centuries, and by and large our circumstances have improved greatly as a result. Deployed wisely, automation can relieve of us drudge work and spur us on to more challenging and fulfilling endeavors. The point is that we're not very good at thinking rationally about automation or understanding its implications. We don't know when to say "enough" or even "hold on a second." The deck is stacked, economically and emotionally, in automation's favor. The benefits of transferring work from people to machines and computers are easy to identify and measure. Businesses can run the numbers on capital investments and calculate automation's benefits in hard currency: reduced labor costs, improved productivity, faster throughputs and turnarounds, higher profits. In our personal lives, we can point to all sorts of ways that computers allow us to save time and avoid hassles. And thanks to our bias for leisure over work, ease over effort, we overestimate automation's benefits.

The costs are harder to pin down. We know computers make certain jobs obsolete and put some people out of work, but history suggests, and most economists assume, that any declines in employment will prove temporary and that over the long haul productivity-boosting technology will create attractive new occupations and raise standards of living. The personal costs are even hazier. How do you measure the expense of an erosion of effort and engagement, or a waning of agency and autonomy, or a subtle deterioration of skill? You can't. Those are the kinds of shadowy, intangible things that we rarely appreciate until after they're gone, and even then we may have trouble expressing the losses in concrete terms. But the costs are real. The choices we make, or fail to make, about which tasks we hand off to computers and which we keep for ourselves are not just practical or economic choices. They're ethical choices. They shape the substance of our lives and the place we make for ourselves in the world. Automation confronts us with the most important question of all: What does *human being* mean?

Csikszentmihalyi and LeFevre discovered something else in their study of people's daily routines. Among all the leisure activities reported by their test subjects, the one that generated the greatest sense of flow was driving a car.

CHAPTER TWO

THE ROBOT AT THE GATE

IN THE EARLY 1950S, Leslie Illingworth, a much-admired political cartoonist at the British satirical magazine *Punch*, drew a dark and foreboding sketch. Set at dusk on what appears to be a stormy autumn day, it shows a worker peering anxiously from the doorway of an anonymous manufacturing plant. One of his hands grips a small tool; the other is balled into a fist. He looks out across the muddy factory yard to the plant's main gate. There, standing beside a sign reading "Hands Wanted," looms a giant, broad-shouldered robot. Across its chest, emblazoned in block letters, is the word "Automation."

The illustration was a sign of its times, a reflection of a new anxiety seeping through Western society. In 1956, it was reprinted as the frontispiece of a slender but influential book called *Automation: Friend or Foe?* by Robert Hugh Macmillan, an engineering professor at Cambridge University. On the first page, Macmillan posed an unsettling question: "Are we in danger of being destroyed by our own creations?" He was not, he explained, referring to the well-known "perils of unrestricted 'push-button' warfare." He was talking about a less discussed but more insidious threat: "the rapidly increasing part that automatic devices are playing in the peace-time industrial life of all civilized countries."[1] Just as earlier machines "had replaced man's muscles," these new devices seemed likely to "replace his brains." By taking over many good, well-paying jobs, they threatened to create widespread unemployment, leading to social strife and upheaval—of just the sort Karl Marx had foreseen a century earlier.[2]

But, Macmillan continued, it didn't have to be that way. If "*rightly* applied," automation could bring economic stability, spread prosperity, and relieve the human race of its toils. "My hope is that this new branch of technology may eventually enable us to lift the curse of Adam from the shoulders of man, for machines could indeed become men's slaves rather than their masters, now that practical techniques have been devised for controlling them automatically."[3] Whether technologies of automation ultimately proved boon or bane, Macmillan warned, one thing was certain: they would play an ever greater role in industry and society. The economic imperatives of "a highly competitive world" made that inevitable.[4] If a robot could work faster, cheaper, or better than its human counterpart, the robot would get the job.

■ ■ ■ ■

"WE ARE brothers and sisters of our machines," the technology historian George Dyson once remarked.[5] Sibling relations are notoriously fraught, and so it is with our technological kin. We love our machines—not just because they're useful to us, but because we find them companionable and even beautiful. In a well-built machine, we see some of our deepest aspirations take form: the desire to understand the world and its workings, the desire to turn nature's power to our own purposes, the desire to add something new and of our own fashioning to the cosmos, the desire to be awed and amazed. An ingenious machine is a source of wonder and of pride.

But machines are ugly too, and we sense in them a threat to things we hold dear. Machines may be a conduit of human power, but that power has usually been wielded by the industrialists and financiers who own the contraptions, not the people paid to operate them. Machines are cold and mindless, and in their obedience to scripted routines we see an image of society's darker possibilities. If machines bring something human to the alien cosmos, they also bring something alien to the human world. The mathematician and philosopher Bertrand Russell put it succinctly in a 1924 essay: "Machines are worshipped because they are beautiful and valued because they confer power; they are hated because they are hideous and loathed because they impose slavery." [6]

As Russell's comment suggests, the tension in Macmillan's view of automated machines—they'd either destroy us or redeem us, liberate us or enslave us—has a long history. The same tension has run through popular reactions to factory machinery since the start of the Industrial Revolution more than two centuries ago. While many of our forebears celebrated the arrival of mechanized production, seeing it as a symbol of progress and a guarantor of prosperity, others worried that machines would steal their jobs and even their souls. Ever since, the story of technology has been one of rapid, often disorienting change. Thanks to the ingenuity of our inventors and entrepreneurs, hardly a decade has passed without the arrival of new, more elaborate, and more capable machinery. Yet our ambivalence toward these fabulous creations, creations of our own hands and minds, has remained a constant. It's almost as if in looking at a machine we see, if only dimly, something about ourselves that we don't quite trust.

In his 1776 masterwork *The Wealth of Nations*, the foundational text of free enterprise, Adam Smith praised the great variety of "very pretty machines" that manufacturers were installing to "facilitate and abridge labour." By enabling "one man to do the work of many," he predicted, mechanization would provide a great boost to industrial productivity.[7] Factory owners would earn more profits, which they would then invest in expanding their operations—building more plants, buying more machines, hiring more employees. Each individual machine's abridgment of labor, far from being bad for workers, would actually stimulate demand for labor in the long run.

Other thinkers embraced and extended Smith's assessment. Thanks to the higher productivity made possible by labor-saving equipment, they predicted, jobs would multiply, wages would go up, and prices of goods would come down. Workers would have some extra cash in their pockets, which they would use to purchase products from the manufacturers that employed them. That would provide yet more capital for industrial expansion. In this way, mechanization would help set in motion a virtuous cycle, accelerating a society's economic growth, expanding and spreading its wealth, and bringing to its people what Smith had termed "convenience and luxury."[8] This view of technology as an economic elixir seemed, happily, to be borne out by the early history of industrialization, and it became a fixture of economic theory. The idea wasn't compelling only to early capitalists and their scholarly brethren. Many social reformers applauded mechanization, viewing it as the best hope for raising the urban masses out of poverty and servitude.

Economists, capitalists, and reformers could afford to take the long view. With the workers themselves, that wasn't the case. Even a temporary abridgment of labor could pose a real and immediate threat to their livelihoods. The installation of new factory machines put plenty of people out of jobs, and it forced others to exchange interesting, skilled work for the tedium of pulling levers and pressing foot-pedals. In many parts of Britain during the eighteenth and the early nineteenth century, skilled workers took to sabotaging the new machinery as a way to defend their jobs, their trades, and their communities. "Machine-breaking," as the movement came to be called, was not simply an attack on technological progress. It was a concerted attempt by tradesmen to protect their ways of life, which were very much bound up in the crafts they practiced, and to secure their economic and civic autonomy. "If the workmen disliked certain machines," writes the historian Malcolm Thomis, drawing on contemporary accounts of the uprisings, "it was because of the use to which they were being put, not because they were machines or because they were new."[9]

Machine-breaking culminated in the Luddite rebellion that raged through the industrial counties of the English Midlands from 1811 to 1816. Weavers and knitters, fearing the destruction of their small-scale, locally organized cottage industry, formed guerrilla bands with the intent of stopping big textile mills and factories from installing mechanized looms and stocking frames. The Luddites—the rebels took their now-notorious name from a legendary Leicestershire machine-breaker known as Ned Ludlam—launched nighttime raids against the plants, often wrecking the new equipment. Thousands of British troops had to be called in to battle the rebels, and the soldiers put down the revolt with brutal force, killing many and incarcerating others.

Although the Luddites and other machine-breakers had some scattered success in slowing the pace of mechanization, they certainly didn't stop it. Machines were soon so commonplace in factories, so essential to industrial production and competition, that resisting their use came to be seen as an exercise in futility. Workers acquiesced to the new technological regime, though their distrust of machinery persisted.

■ ■ ■ ■

IT WAS Marx who, a few decades after the Luddites lost their fight, gave the deep divide in society's view of mechanization its most powerful and influential expression. Frequently in his writings, Marx invests factory machinery with a demonic, parasitic will, portraying it as "dead labour" that "dominates, and pumps dry, living labour power." The workman becomes a "mere living appendage" of the "lifeless mechanism."[10] In a darkly prophetic remark during an 1856 speech, he said, "All our invention and progress seem to result in endowing material forces with intellectual life, and stultifying human life into a material force."[11] But Marx didn't just talk about the "infernal effects" of machines. As the media scholar Nick Dyer-Witheford has explained, he also saw and lauded "their emancipatory promise."[12] Modern machinery, Marx observed in that same speech, has "the wonderful power of shortening and fructifying human labour."[13] By freeing workers from the narrow specializations of their trades, machines might allow them to fulfill their potential as "totally developed" individuals, able to shift between "different modes of activity" and hence "different social functions."[14] In the right hands—those of the workers rather than the capitalists—technology would no longer be the yoke of oppression. It would become the uplifting block and tackle of self-fulfillment.

The idea of machines as emancipators took stronger hold in Western culture as the twentieth century approached. In an 1897 article praising the mechanization of American industry, the French economist Émile Levasseur ticked off the benefits that new technology had brought to "the laboring classes." It had raised workers' wages and pushed down the prices they paid for goods, providing them with greater material comfort. It had spurred a redesign of factories, making workplaces cleaner, better lit, and generally more hospitable than the dark satanic mills that characterized the early years of the Industrial Revolution. Most important of all, it had elevated the kind of work that factory hands performed. "Their task has become less onerous, the machine doing everything which requires great strength; the workman, instead of bringing his muscles into play, has become an inspector, using his intelligence." Levasseur acknowledged that laborers still grumbled about having to operate machinery. "They reproach [the machine] with demanding such continued attention that it enervates," he wrote, and they accuse it of "degrading man by transforming him into a machine, which knows how to make but one movement, and that always the same." Yet he dismissed such complaints as blinkered. The workers simply didn't understand how good they had it.[15]

Some artists and intellectuals, believing the imaginative work of the mind to be inherently superior to the productive labor of the body, saw a technological utopia in the making. Oscar Wilde, in an essay published at about the same time as Levasseur's, though aimed at a very different audience, foresaw a day when machines would not just alleviate toil but eliminate it. "All unintellectual labour, all monotonous, dull labour, all labour that deals with dreadful things, and involves unpleasant conditions, must be done by machinery," he wrote. "On mechanical slavery, on the slavery of the machine, the future of the world depends." That machines would assume the role of slaves seemed to Wilde a foregone conclusion: "There is no doubt at all that this is the future of machinery, and just as trees grow while the country gentleman is asleep, so while Humanity will be amusing itself, or enjoying cultivated leisure—which, and not labour, is the aim of man—or making beautiful things, or reading beautiful things, or simply contemplating the world with admiration and delight, machinery will be doing all the necessary and unpleasant work."[16]

The Great Depression of the 1930s curbed such enthusiasm. The economic collapse prompted a bitter outcry against what had, in the Roaring Twenties, come to be known and celebrated as the Machine Age. Labor unions and religious groups, crusading editorial writers and despairing citizens—all railed against the job-destroying machines and the greedy businessmen who owned them. "Machinery did not inaugurate the phenomenon of unemployment," wrote the author of a best-selling book called *Men and Machines*, "but promoted it from a minor irritation to one of the chief plagues of mankind." It appeared, he went on, that "from now on, the better able we are to produce, the worse we shall be off."[17] The mayor of Palo Alto, California, wrote a letter to

President Herbert Hoover imploring him to take action against the "Frankenstein monster" of industrial technology, a scourge that was "devouring our civilization."[18] At times the government itself inflamed the public's fears. One report issued by a federal agency called the factory machine "as dangerous as a wild animal." The uncontrolled acceleration of progress, its author wrote, had left society chronically unprepared to deal with the consequences.[19]

But the Depression did not entirely extinguish the Wildean dream of a machine paradise. In some ways, it rendered the utopian vision of progress more vivid, more necessary. The more we saw machines as our foes, the more we yearned for them to be our friends. "We are being afflicted," wrote the great British economist John Maynard Keynes in 1930, "with a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come—namely, *technological unemployment*." The ability of machines to take over jobs had outpaced the economy's ability to create valuable new work for people to do. But the problem, Keynes assured his readers, was merely a symptom of "a temporary phase of maladjustment." Growth and prosperity would return. Per-capita income would rise. And soon, thanks to the ingenuity and efficiency of our mechanical slaves, we wouldn't have to worry about jobs at all. Keynes thought it entirely possible that in a hundred years, by the year 2030, technological progress would have freed humankind from "the struggle for subsistence" and propelled us to "our destination of economic bliss." Machines would be doing even more of our work for us, but that would no longer be cause for worry or despair. By then, we would have figured out how to spread material wealth to everyone. Our only problem would be to figure out how to put our endless hours of leisure to good use—to teach ourselves "to enjoy" rather than "to strive."[20]

We're still striving, and it seems a safe bet that economic bliss will not have descended upon the planet by 2030. But if Keynes let his hopes get the best of him in the dark days of 1930, he was fundamentally right about the economy's prospects. The Depression did prove temporary. Growth returned, jobs came back, incomes shot up, and companies continued buying more and better machines. Economic equilibrium, imperfect and fragile as always, reestablished itself. Adam Smith's virtuous cycle kept turning.

By 1962, President John F. Kennedy could proclaim, during a speech in West Virginia, "We believe that if men have the talent to invent new machines that put men out of work, they have the talent to put those men back to work."[21] From the opening "we believe," the sentence is ringingly Kennedyesque. The simple words become resonant as they're repeated: *men, talent, men, work, talent, men, work*. The drum-like rhythm marches forward, giving the stirring conclusion—"back to work"—an air of inevitability. To those listening, Kennedy's words must have sounded like the end of the story. But they weren't. They were the end of one chapter, and a new chapter had already begun.

■ ■ ■ ■

WORRIES ABOUT technological unemployment have been on the rise again, particularly in the United States. The recession of the early 1990s, which saw exalted U.S. companies such as General Motors, IBM, and Boeing fire tens of thousands of workers in massive "restructurings," prompted fears that new technologies, particularly cheap computers and clever software, were about to wipe out middle-class jobs. In 1994, the sociologists Stanley Aronowitz and William DiFazio published *The Jobless Future*, a book that implicated "labor-displacing technological change" in "the trend toward more low-paid, temporary, benefit-free blue- and white-collar jobs and fewer decent *permanent* factory and office jobs."[22] The following year, Jeremy Rifkin's unsettling *The End of Work* appeared. The rise of computer automation had inaugurated a "Third Industrial Revolution," declared Rifkin. "In the years ahead, new, more sophisticated software technologies are going to bring civilization ever closer to a near workerless world." Society had reached a turning point, he wrote. Computers could "result in massive unemployment and a potential global depression," but they could also "free us for a life of increasing leisure" if we were willing to rewrite the tenets of contemporary capitalism.[23] The two books, and others like them, caused a stir, but once again fears about technology-induced joblessness passed quickly. The resurgence of economic growth through the middle and late 1990s, culminating in the giddy dot-com boom, turned people's attention away from apocalyptic predictions of mass unemployment.

A decade later, in the wake of the Great Recession of 2008, the anxieties returned, stronger than ever. In mid-2009, the American economy, recovering fitfully from the economic collapse, began to expand again. Corporate profits rebounded. Businesses ratcheted their capital investments up to pre-recession levels. The stock market soared. But hiring refused to bounce back. While it's not unusual for companies to wait until a recovery is well established before recruiting new workers, this time the hiring lag seemed interminable. Job growth remained unusually tepid, the unemployment rate stubbornly high. Seeking an explanation, and a culprit, people looked to the usual suspect: labor-saving technology.

Late in 2011, two respected MIT researchers, Erik Brynjolfsson and Andrew McAfee, published a short electronic book, *Race against the Machine*, in which they gently chided economists and policy makers for dismissing the possibility that workplace technology was substantially reducing companies' need for new employees. The "empirical fact" that machines had bolstered employment for centuries "conceals a dirty secret," they wrote. "There is no economic law that says that everyone, or even most people, automatically benefit from technological progress." Although Brynjolfsson and McAfee were anything but technophobes—they remained "hugely optimistic" about the ability of computers and robots to boost productivity and improve people's lives over the long run—they made a strong case that technological unemployment was real, that it had become pervasive, and that it would likely get much worse. Human beings, they warned, were losing the race against the machine.[24]

Their ebook was like a match thrown onto a dry field. It sparked a vigorous and sometimes caustic debate among economists, a debate that soon drew the attention of journalists. The phrase "technological unemployment," which had faded from use after the Great Depression, took a new grip on the public mind. At the start of 2013, the TV news program *60 Minutes* ran a segment, called "March of the Machines," that examined how businesses were using new technologies in place of workers at warehouses, hospitals, law firms, and manufacturing plants. Correspondent Steve Kroft lamented "a massive high-tech industry that's contributed enormous productivity and wealth to the American economy but surprisingly little in the way of employment."[25] Shortly after the program aired, a team of Associated Press writers published a three-part investigative report on the persistence of high unemployment. Their grim conclusion: jobs are "being obliterated by technology." Noting that science-fiction writers have long "warned of a future when we would be architects of our own obsolescence, replaced by our machines," the AP reporters declared that "the future has arrived."[26] They quoted one analyst who predicted that the unemployment rate would reach 75 percent by the century's end.[27]

Such forecasts are easy to dismiss. Their alarmist tone echoes the refrain heard time and again since the eighteenth century. Out of every economic downturn rises the specter of a job-munching Frankenstein monster. And then, when the economic cycle emerges from its trough and jobs return, the monster goes back in its cage and the worries subside. This time, though, the economy isn't behaving as it normally does. Mounting evidence suggests that a troubling new dynamic may be at work. Joining Brynjolfsson and McAfee, several prominent economists have begun questioning their profession's cherished assumption that technology-fueled productivity gains will bring job and wage growth. They point out that over the last decade U.S. productivity rose at a faster pace than we saw in the preceding thirty years, that corporate profits have hit levels we haven't seen in half a century, and that business investments in new equipment have been rising sharply. That combination should bring robust employment growth. And yet the total number of jobs in the country has barely budged. Growth and employment are "diverging in advanced countries," says economist Michael Spence, a Nobel laureate, and technology is the main reason why: "The replacement of routine manual jobs by machines and robots is a powerful, continuing, and perhaps accelerating trend in manufacturing and logistics, while networks of computers are replacing routine white-collar jobs in information processing."[28]

Some of the heavy spending on robots and other automation technologies in recent years may reflect temporary economic conditions, particularly the ongoing efforts by politicians and central banks to stimulate growth. Low interest rates and aggressive government tax incentives for capital investment have likely encouraged companies to buy labor-saving equipment and software that they might not otherwise have purchased.[29] But deeper and more prolonged trends also seem to be at work. Alan Krueger, the Princeton economist who chaired Barack Obama's Council of Economic Advisers from 2011 to 2013, points out that even before the recession "the U.S. economy was not creating enough jobs, particularly not enough middle-class jobs, and we were losing manufacturing jobs at an alarming rate."[30] Since then, the picture has only darkened. It might be assumed that, at least when it comes to manufacturing, jobs aren't disappearing but simply migrating to countries with low wages. That's not so. The total number of worldwide manufacturing jobs has been falling for years, even in industrial powerhouses like China, while overall manufacturing output has grown sharply.[31] Machines are replacing factory workers faster than economic expansion creates new manufacturing positions. As industrial robots become cheaper and more adept, the gap between lost and added jobs will almost certainly widen. Even the news that companies like GE and Apple are bringing some manufacturing work back to the United States is bittersweet. One of the reasons the work is returning is that most of it can be done without human beings. “Factory floors these days are nearly empty of people because software-driven machines are doing most of the work,” reports economics professor Tyler Cowen.[32] A company doesn’t have to worry about labor costs if it’s not employing laborers.

The industrial economy—the economy of machines—is a recent phenomenon. It has been around for just two and a half centuries, a tick of history’s second hand. Drawing definitive conclusions about the link between technology and employment from such limited experience was probably rash. The logic of capitalism, when combined with the history of scientific and technological progress, would seem to be a recipe for the eventual removal of labor from the processes of production. Machines, unlike workers, don’t demand a share of the returns on capitalists’ investments. They don’t get sick or expect paid vacations or demand yearly raises. For the capitalist, labor is a problem that progress solves. Far from being irrational, the fear that technology will erode employment is fated to come true “in the very long run,” argues the eminent economic historian Robert Skidelsky: “Sooner or later, we will run out of jobs.”[33]

How long is the very long run? We don’t know, though Skidelsky warns that it may be “uncomfortably close” for some countries.[34] In the near term, the impact of modern technology may be felt more in the distribution of jobs than in the overall employment figures. The mechanization of manual labor during the Industrial Revolution destroyed some good jobs, but it led to the creation of vast new categories of middle-class occupations. As companies expanded to serve bigger and more far-flung markets, they hired squads of supervisors and accountants, designers and marketers. Demand grew for teachers, doctors, lawyers, librarians, pilots, and all sorts of other professionals. The makeup of the job market is never static; it changes in response to technological and social trends. But there’s no guarantee that the changes will always benefit workers or expand the middle class. With computers being programmed to take over white-collar work, many professionals are being forced into lower-paying jobs or made to trade full-time posts for part-time ones.

While most of the jobs lost during the recent recession were in well-paying industries, nearly three-fourths of the jobs created since the recession are in low-paying sectors. Having studied the causes of the “incredibly anemic employment growth” in the United States since 2000, MIT economist David Autor concludes that information technology “has really changed the distribution of occupation,” creating a widening disparity in incomes and wealth. “There is an abundance of work to do in food service and there is an abundance of work in finance, but there are fewer middle-wage, middle-income jobs.”[35] As new computer technologies extend automation into even more branches of the economy, we're likely to see an acceleration of this trend, with a further hollowing of the middle class and a growing loss of jobs among even the highest-paid professionals. "Smart machines may make higher GDP possible," notes Paul Krugman, another Nobel Prize–winning economist, "but also reduce the demand for people—including smart people. So we could be looking at a society that grows ever richer, but in which all the gains in wealth accrue to whoever owns the robots."[36]

The news is not all dire. As the U.S. economy gained steam during the second half of 2013, hiring strengthened in several sectors, including construction and health care, and there were encouraging gains in some higher-paying professions. The demand for workers remains tied to the economic cycle, if not quite so tightly as in the past. The increasing use of computers and software has itself created some very attractive new jobs as well as plenty of entrepreneurial opportunities. By historical standards, though, the number of people employed in computing and related fields remains modest. We can't all become software programmers or robotics engineers. We can't all decamp to Silicon Valley and make a killing writing nifty smartphone apps.* With average wages stagnant and corporate profits continuing to surge, the economy's bounties seem likely to go on flowing to the lucky few. And JFK's reassuring words will sound more and more suspect.

Why might this time be different? What exactly has changed that may be severing the old link between new technologies and new jobs? To answer that question we have to look back to that giant robot standing at the gate in Leslie Illingworth's cartoon—the robot named Automation.

■ ■ ■ ■

THE WORD *automation* entered the language fairly recently. As best we can tell, it was first spoken in 1946, when engineers at the Ford Motor Company felt the need to coin a term to describe the latest machinery being installed on the company's assembly lines. "Give us some more of that automatic business," a Ford vice president reportedly said in a meeting. "Some more of that—that—'automation.' "[37] Ford's plants were already famously mechanized, with sophisticated machines streamlining every job on the line. But factory hands still had to lug parts and subassemblies from one machine to the next. The workers still controlled the pace of production. The equipment installed in 1946 changed that. Machines took over the material-handling and conveyance functions, allowing the entire assembly process to proceed automatically. The alteration in work flow may not have seemed momentous to those on the factory floor. But it was. Control over a complex industrial process had shifted from worker to machine.

The new word spread quickly. Two years later, in a report on the Ford machinery, a writer for the magazine *American Machinist* defined automation as "the art of applying mechanical devices to manipulate work pieces … in timed sequence with the production equipment so that the line can be put wholly or partially under push-button control at strategic stations."[38] As automation reached into more industries and production processes, and as it began to take on metaphorical weight in the culture, its definition grew more diffuse. "Few words of recent years have been so twisted to suit a multitude of purposes and phobias as this new word, 'automation,' " grumbled a Harvard business professor in 1958. "It has been used as a technological rallying cry, a manufacturing goal, an engineering challenge, an advertising slogan, a labor campaign banner, and as the symbol of ominous technological progress." He then offered his own, eminently pragmatic definition: "Automation simply means something *significantly more automatic than previously existed in that plant, industry, or location*."[39] Automation wasn't a thing or a technique so much as a force. It was more a manifestation of progress than a particular mode of operation. Any attempt at explaining or predicting its consequences would necessarily be tentative. As with many technological trends, automation would always be both old and new, and it would require a fresh reevaluation at each stage of its advance.

That Ford's automated equipment arrived just after the end of the Second World War was no accident. It was during the war that modern automation technology took shape. When the Nazis began their bombing blitz against Great Britain in 1940, English and American scientists faced a challenge as daunting as it was pressing: How do you knock high-flying, fast-moving bombers out of the sky with heavy missiles fired from unwieldy antiaircraft guns on the ground? The mental calculations and physical adjustments required to aim a gun accurately—not at a plane's current position but at its probable future position—were far too complicated for a soldier to perform with the speed necessary to get a shot off while a plane was still in range. This was no job for mortals. The missile's trajectory, the scientists saw, had to be computed by a calculating machine, using tracking data coming in from radar systems along with statistical projections of a plane's course, and then the calculations had to be fed automatically into the gun's aiming mechanism to guide the firing. The gun's aim, moreover, had to be adjusted continually to account for the success or failure of previous shots.

As for the members of the gunnery crews, their work would have to change to accommodate the new generation of automated weapons. And change it did. Artillerymen soon found themselves sitting in front of screens in darkened trucks, selecting targets from radar displays. Their identities shifted along with their jobs. They were no longer seen "as soldiers," writes one historian, but rather "as technicians reading and manipulating representations of the world." [40]

In the antiaircraft cannons born of the Allied scientists' work, we see all the elements of what now characterizes an automated system. First, at the system's core, is a very fast calculating machine—a computer. Second is a sensing mechanism (radar, in this case) that monitors the external environment, the real world, and communicates essential data about it to the computer. Third is a communication link that allows the computer to control the movements of the physical apparatus that performs the actual work, with or without human assistance. And finally there's a feedback method—a means of returning to the computer information about the results of its instructions so that it can adjust its calculations to correct for errors and account for changes in the environment. Sensory organs, a calculating brain, a stream of messages to control physical movements, and a feedback loop for learning: there you have the essence of automation, the essence of a robot. And there, too, you have the essence of a living being's nervous system. The resemblance is no coincidence. In order to replace a human, an automated system first has to replicate a human, or at least some aspect of a human's ability.

Automated machines existed before World War II. James Watt's steam engine, the original prime mover of the Industrial Revolution, incorporated an ingenious feedback device—the fly-ball governor—that enabled it to regulate its own operation. As the engine sped up, it rotated a pair of metal balls, creating a centrifugal force that pulled a lever to close a steam valve, keeping the engine from running too fast. The Jacquard loom, invented in France around 1800, used steel punch cards to control the movements of spools of different-colored threads, allowing intricate patterns to be woven automatically. In 1866, a British engineer named J. Macfarlane Gray patented a steamship steering mechanism that was able to register the movement of a boat's helm and, through a gear-operated feedback system, adjust the angle of the rudder to maintain a set course.[41] But the development of fast computers, along with other sensitive electronic controls, opened a new chapter in the history of machines. It vastly expanded the possibilities of automation. As the mathematician Norbert Wiener, who helped write the prediction algorithms for the Allies' automated antiaircraft gun, explained in his 1950 book *The Human Use of Human Beings*, the advances of the 1940s enabled inventors and engineers to go beyond "the sporadic design of individual automatic mechanisms." The new technologies, while designed with weaponry in mind, gave rise to "a general policy for the construction of automatic mechanisms of the most varied type." They paved the way for "the new automatic age." [42]

Beyond the pursuit of progress and productivity lay another impetus for the automatic age: politics. The postwar years were characterized by intense labor strife. Managers and unions battled in most American manufacturing sectors, and the tensions were often strongest in industries essential to the federal government's Cold War buildup of military equipment and armaments. Strikes, walkouts, and slowdowns were daily events. In 1950 alone, eighty-eight work stoppages were staged at a single Westinghouse plant in Pittsburgh. In many factories, union stewards held more power over operations than did corporate managers—the workers called the shots. Military and industrial planners saw automation as a way to shift the balance of power back to management. Electronically controlled machinery, declared *Fortune* magazine in a 1946 cover story titled "Machines without Men," would prove "immensely superior to the human mechanism," not least because machines "are always satisfied with working conditions and never demand higher wages." [43] An executive with Arthur D. Little, a leading management and engineering consultancy, wrote that the rise of automation heralded the business world's "emancipation from human workers." [44]

In addition to reducing the need for laborers, particularly skilled ones, automated equipment provided business owners and managers with a technological means to control the speed and flow of production through the electronic programming of individual machines and entire assembly lines. When, at the Ford plants, control over the pace of the line shifted to the new automated equipment, the workers lost a great deal of autonomy. By the mid-1950s, the role of labor unions in charting factory operations was much diminished.[45] The lesson would prove important: in an automated system, power concentrates with those who control the programming.

Wiener foresaw, with uncanny clarity, what would come next. The technologies of automation would advance far more rapidly than anyone had imagined. Computers would get faster and smaller. The speed and capacity of electronic communication and storage systems would increase exponentially. Sensors would see, hear, and feel the world with ever greater sensitivity. Robotic mechanisms would come "to replicate more nearly the functions of the human hand as supplemented by the human eye." The cost to manufacture all the new devices and systems would plummet. The use of automation would become both possible and economical in ever more areas. And since computers could be programmed to carry out logical functions, automation's reach would extend beyond the work of the hand and into the work of the mind—the realm of analysis, judgment, and decision making. A computerized machine didn't have to act by manipulating material things like guns. It could act by manipulating information. "From this stage on, everything may go by machine," Wiener wrote. "The machine plays no favorites between manual labor and white-collar labor." It seemed obvious to him that automation would, sooner or later, create "an unemployment situation" that would make the calamity of the Great Depression "seem a pleasant joke." [46]

The Human Use of Human Beings was a best seller, as was Wiener's earlier and much more technical treatise, *Cybernetics, or Control and Communication in the Animal and the Machine*. The mathematician's unsettling analysis of technology's trajectory became part of the intellectual texture of the 1950s. It inspired or informed many of the books and articles on automation that appeared during the decade, including Robert Hugh Macmillan's slim volume. An aging Bertrand Russell, in a 1951 essay, "Are Human Beings Necessary?," wrote that Wiener's work made it clear that "we shall have to change some of the fundamental assumptions upon which the world has been run ever since civilization began." [47] Wiener even makes a brief appearance as a forgotten prophet in Kurt Vonnegut's first novel, the 1952 dystopian satire *Player Piano*, in which a young engineer's rebellion against a rigidly automated world ends with an epic episode of machine-breaking.

■ ■ ■ ■

THE IDEA of a robot invasion may have seemed threatening, if not apocalyptic, to a public already rattled by the bomb, but automation technologies were still in their infancy during the 1950s. Their ultimate consequences could be imagined, in speculative tracts and science-fiction fantasies, but those consequences were still a long way from being experienced. Through the 1960s, most automated machines continued to resemble the primitive robotic haulers on Ford's postwar assembly lines. They were big, expensive, and none too bright. Most of them could perform only a single, repetitive function, adjusting their movements in response to a few elementary electronic commands: speed up, slow down; move left, move right; grasp, release. The machines were extraordinarily precise, but otherwise their talents were few. Toiling anonymously inside factories, often locked within cages to protect passersby from their mindless twists and jerks, they certainly didn't look like they were about to take over the world.

They seemed little more than very well-behaved and well-coordinated beasts of burden.

But robots and other automated systems had one big advantage over the purely mechanical contraptions that came before them. Because they ran on software, they could hitch a ride on the Moore's Law Express. They could benefit from all the rapid advances—in processor speed, programming algorithms, storage and network capacity, interface design, and miniaturization—that came to characterize the progress of computers themselves. And that, as Wiener predicted, is what happened. Robots' senses grew sharper; their brains, quicker and more supple; their conversations, more fluent; their ability to learn, more capacious. By the early 1970s, they were taking over production work that required flexibility and dexterity—cutting, welding, assembling. By the end of that decade, they were flying planes as well as building them. And then, freed from their physical embodiments and turned into the pure logic of code, they spread out into the business world through a multitude of specialized software applications. They entered the cerebral trades of the white-collar workforce, sometimes as replacements but far more often as assistants.

Robots may have been at the factory gate in the 1950s, but it's only recently that they've marched, on our orders, into offices, shops, and homes. Today, as software of what Wiener termed "the judgment-replacing type" moves from our desks to our pockets, we're at last beginning to experience automation's true potential for changing what we do and how we do it. Everything is being automated. Or, as Netscape founder and Silicon Valley grandee Marc Andreessen puts it, "software is eating the world." [48]

That may be the most important lesson to be gleaned from Wiener's work—and, for that matter, from the long, tumultuous history of labor-saving machinery. Technology changes, and it changes more quickly than human beings change. Where computers sprint forward at the pace of Moore's law, our own innate abilities creep ahead with the tortoise-like tread of Darwin's law. Where robots can be constructed in a myriad of forms, replicating everything from snakes that burrow in the ground to raptors that swoop across the sky to fish that swim through the sea, we're basically stuck with our old, forked bodies. That doesn't mean our machines are about to leave us in the evolutionary dust. Even the most powerful supercomputer evidences no more consciousness than a hammer. It does mean that our software and our robots will, with our guidance, continue to find new ways to outperform us—to work faster, cheaper, better. And, like those antiaircraft gunners during World War II, we'll be compelled to adapt our own work, behavior, and skills to the capabilities and routines of the machines we depend on.

---

* The internet, it's often noted, has opened opportunities for people to make money through their own personal initiative, with little investment of capital. They can sell used goods through eBay or crafts through Etsy. They can rent out a spare room through Airbnb or turn their car into a ghost cab with Lyft. They can find odd jobs through TaskRabbit. But while it's easy to pick up spare change through such modest enterprise, few people are going to be able to earn a middle-class income from the work. The real money goes to the software companies running the online clearinghouses that connect buyer and seller or lessor and lessee—clearinghouses that, being highly automated themselves, need few employees.

CHAPTER THREE

ON AUTOPILOT

ON THE EVENING OF FEBRUARY 12, 2009, a Continental Connection commuter flight made its way through blustery weather between Newark, New Jersey, and Buffalo, New York. As is typical of commercial flights these days, the two pilots didn't have all that much to do during the hour-long trip. The captain, an affable, forty-seven-year-old Floridian named Marvin Renslow, manned the controls briefly during takeoff, guiding the Bombardier Q400 turboprop into the air, then switched on the autopilot. He and his cabin mate, twenty-four-year-old first officer Rebecca Shaw, a newlywed from Seattle, kept an eye on the computer readouts that flickered across the cockpit's five large LCD screens. They exchanged some messages over the radio with air traffic controllers. They went through a few routine checklists. Mostly, though, they passed the time chatting amiably about this and that—families, careers, colleagues, money—as the turboprop cruised along its northwesterly route at sixteen thousand feet.[1]

The Q400 was well into its approach to the Buffalo airport, its landing gear down, its wing flaps out, when the captain's control yoke began to shudder noisily. The plane's "stick shaker" had activated, a signal that the turboprop was losing lift and risked going into an aerodynamic stall.* The autopilot disconnected, as it's programmed to do in the event of a stall warning, and the captain took over the controls. He reacted quickly, but he did precisely the wrong thing. He jerked back on the yoke, lifting the plane's nose and reducing its air speed, instead of pushing the yoke forward to tip the craft down and gain velocity. The plane's automatic stall-avoidance system kicked in and attempted to push the yoke forward, but the captain simply redoubled his effort to pull it back toward him. Rather than prevent a stall, Renslow caused one. The Q400 spun out of control, then plummeted. "We're down," the captain said, just before the plane slammed into a house in a Buffalo suburb.

The crash, which killed all forty-nine people onboard as well as one person on the ground, should not have happened. A National Transportation Safety Board investigation found no evidence of mechanical problems with the Q400. Some ice had accumulated on the plane, but nothing out of the ordinary for a winter flight. The deicing equipment had operated properly, as had the plane's other systems. Renslow had had a fairly demanding flight schedule over the preceding two days, and Shaw had been battling a cold, but both pilots seemed lucid and wakeful while in the cockpit. They were well trained, and though the stick shaker took them by surprise, they had plenty of time and airspace to make the adjustments necessary to avoid a stall. The NTSB concluded that the cause of the accident was pilot error. Neither Renslow nor Shaw had detected "explicit cues" that a stall warning was imminent, an oversight that suggested "a significant breakdown in their monitoring responsibilities." Once the warning sounded, the investigators reported, the captain's response "should have been automatic, but his improper flight control inputs were inconsistent with his training" and instead revealed "startle and confusion."

An executive from the company that operated the flight for Continental, the regional carrier Colgan Air, admitted that the pilots seemed to lack “situational awareness” as the emergency unfolded.[2] Had the crew acted appropriately, the plane would likely have landed safely.

The Buffalo crash was not an isolated incident. An eerily similar disaster, with far more casualties, occurred a few months later. On the night of May 31, an Air France Airbus A330 took off from Rio de Janeiro, bound for Paris.[3] The jet ran into a storm over the Atlantic about three hours after takeoff. Its air-speed sensors, caked with ice, began giving faulty readings, which caused the autopilot to disengage. Bewildered, the copilot flying the plane, Pierre-Cédric Bonin, yanked back on the control stick. The A330 rose and a loud stall warning sounded, but Bonin continued to pull back heedlessly on the stick. As the plane climbed sharply, it lost velocity. The air-speed sensors began working again, providing the crew with accurate numbers. It should have been clear at this point that the jet was going too slow. Yet Bonin persisted in his mistake at the controls, causing a further deceleration. The jet stalled and began to fall. If Bonin had simply let go of the stick, the A330 might well have righted itself. But he didn’t. The flight crew was suffering what French investigators would later term a “total loss of cognitive control of the situation.” [4] After a few more harrowing seconds, another pilot, David Robert, took over the controls. It was too late. The plane dropped more than thirty thousand feet in three minutes.

“This can’t be happening,” said Robert.

“But what *is* happening?” replied the still-bewildered Bonin.

Three seconds later, the jet hit the ocean. All 228 crew and passengers died.

■ ■ ■ ■

IF YOU want to understand the human consequences of automation, the first place to look is up. Airlines and plane manufacturers, as well as government and military aviation agencies, have been particularly aggressive and especially ingenious in finding ways to shift work from people to machines. What car designers are doing with computers today, aircraft designers did decades ago. And because a single mistake in a cockpit can cost scores of lives and many millions of dollars, a great deal of private and public money has gone into funding psychological and behavioral research on automation’s effects. For decades, scientists and engineers have been studying the ways automation influences the skills, perceptions, thoughts, and actions of pilots. Much of what we know about what happens when people work in concert with computers comes out of this research.

The story of flight automation begins a hundred years ago, on June 18, 1914, in Paris. The day was, by all accounts, a sunny and pleasant one, the blue sky a perfect backdrop for spectacle. A large crowd had gathered along the banks of the Seine, near the Argenteuil bridge in the city’s northwestern fringes, to witness the Concours de la Sécurité en Aéroplane, an aviation competition organized to show off the latest advances in flight safety.[5] Nearly sixty planes and pilots took part, demonstrating an impressive assortment of techniques and equipment. Last on the day’s program, flying a Curtiss C-2 biplane, was a handsome American pilot named Lawrence Sperry. Sitting beside him in the C-2's open cockpit was his French mechanic, Emil Cachin. As Sperry flew past the ranks of spectators and approached the judges' stand, he let go of the plane's controls and raised his hands. The crowd roared. The plane was flying itself!

Sperry was just getting started. After swinging the plane around, he took another pass by the reviewing stand, again with his hands in the air. This time, though, he had Cachin climb out of the cockpit and walk along the lower right wing, holding the struts between the wings for support. The plane tilted starboard for a second under the Frenchman's weight, then immediately righted itself, with no help from Sperry. The crowd roared even louder. Sperry circled around once again. By the time his plane approached the stands for its third pass, not only was Cachin out on the right wing, but Sperry himself had climbed out onto the left wing. The C-2 was flying, steady and true, with no one in the cockpit. The crowd and the judges were dumbfounded. Sperry won the grand prize—fifty thousand francs—and the next day his face beamed from the front pages of newspapers across Europe.

Inside the Curtiss C-2 was the world's first automatic pilot. Known as a "gyroscopic stabilizer apparatus," the device had been invented two years earlier by Sperry and his father, the famed American engineer and industrialist Elmer A. Sperry. It consisted of a pair of gyroscopes, one mounted horizontally, the other vertically, installed beneath the pilot's seat and powered by a wind-driven generator behind the propeller. Spinning at thousands of revolutions a minute, the gyroscopes were able to sense, with remarkable precision, a plane's orientation along its three axes of rotation—its lateral pitch, longitudinal roll, and vertical yaw. Whenever the plane diverged from its intended attitude, charged metal brushes attached to the gyroscopes would touch contact points on the craft's frame, completing a circuit. An electric current would flow to the motors operating the plane's main control panels—the ailerons on the wings and the elevators and rudder on the tail—and the panels would automatically adjust their positions to correct the problem. The horizontal gyroscope kept the plane's wings steady and its keel even, while the vertical one handled the steering.

It took nearly twenty years of further testing and refinement, much of it carried out under the auspices of the U.S. military, before the gyroscopic autopilot was ready to make its debut in commercial flight. But when it did, the technology still seemed as miraculous as ever. In 1930, a writer from *Popular Science* gave a breathless account of how an autopilot-equipped plane—"a big tri-motored Ford"—flew "without human aid" during a three-hour trip from Dayton, Ohio, to Washington, D.C. "Four men leaned back at ease in the passenger cabin," the reporter wrote. "Yet the pilot's compartment was empty. A metal airman, scarcely larger than an automobile battery, was holding the stick." [6] When, three years later, the daring American pilot Wiley Post completed the first solo flight around the world, assisted by a Sperry autopilot that he had nicknamed "Mechanical Mike," the press heralded a new era in aviation. "The days when human skill alone and an almost bird-like sense of direction enabled a flier to hold his course for long hours through a starless night or a fog are over," reported the *New York Times*. "Commercial flying in the future will be automatic."[7]

The introduction of the gyroscopic autopilot set the stage for a momentous expansion of aviation's role in warfare and transport. By taking over much of the manual labor required to keep a plane stable and on course, the device relieved pilots of their constant, exhausting struggle with sticks and pedals, cables and pulleys. That not only alleviated the fatigue aviators endured on long flights; it also freed their hands, their eyes, and, most important, their minds for other, more subtle tasks. They could consult more instruments, make more calculations, solve more problems, and in general think more analytically and creatively about their work. They could fly higher and farther, and with less risk of crashing. They could go out in weather that once would have kept them grounded. And they could undertake intricate maneuvers that would have seemed rash or just plain impossible before. Whether ferrying passengers or dropping bombs, pilots became considerably more versatile and valuable once they had autopilots to help them fly. Their planes changed too: they got bigger, faster, and a whole lot more complicated.

Automatic steering and stabilization tools progressed rapidly during the 1930s, as physicists learned more about aerodynamics and engineers incorporated air-pressure gauges, pneumatic controls, shock absorbers, and other refinements into autopilot mechanisms. The biggest breakthrough came in 1940, when the Sperry Corporation introduced its first electronic model, the A-5. Using vacuum tubes to amplify signals from the gyroscopes, the A-5 was able to make speedier, more precise adjustments and corrections. It could also sense and account for changes in a plane's velocity and acceleration. Used in conjunction with the latest bombsight technology, the electronic autopilot proved a particular boon to the Allied air campaign in World War II.

Shortly after the war, on a September evening in 1947, the U.S. Army Air Forces conducted an experimental flight that made clear how far autopilots had come. Captain Thomas J. Wells, a military test pilot, taxied a C-54 Skymaster transport plane with a seven-man crew onto a remote runway in Newfoundland. He then let go of the yoke, pushed a button to activate the autopilot, and, as one of his colleagues in the cockpit later recalled, "sat back and put his hands in his lap."[8] The plane took off by itself, automatically adjusting its flaps and throttles and, once airborne, retracting its landing gear. It then flew itself across the Atlantic, following a series of "sequences" that had earlier been programmed into what the crew called its "mechanical brain." Each sequence was keyed to a particular altitude or mileage reading. The men on the plane hadn't been told of the flight's route or destination; the plane maintained its own course by monitoring signals from radio beacons on the ground and on boats at sea. At dawn the following day, the C-54 reached the English coast. Still under the control of the autopilot, it began its descent, lowered its landing gear, lined itself up with an airstrip at a Royal Air Force base in Oxfordshire, and executed a perfect landing. Captain Wells then lifted his hands from his lap and parked the plane.

A few weeks after the Skymaster's landmark trip, a writer with the British aviation magazine *Flight* contemplated the implications. It seemed inevitable, he wrote, that the new generation of autopilots would "dispose of the necessity for carrying navigators, radio operators, and flight engineers" on planes. The machines would render those jobs redundant. Pilots, he allowed, did not seem quite so dispensable. They would, at least for the foreseeable future, continue to be a necessary presence in cockpits, if only “to watch the various clocks and indicators to see that everything is going satisfactorily.”[9]

■ ■ ■ ■

In 1988, forty years after the C-54’s Atlantic crossing, the European aerospace consortium Airbus Industrie introduced its A320 passenger jet. The 150-seat plane was a smaller version of the company’s original A300 model, but unlike its conventional and rather drab predecessor, the A320 was a marvel. The first commercial aircraft that could truly be called computerized, it was a harbinger of everything to come in aircraft design. The flight deck would have been unrecognizable to Wiley Post or Lawrence Sperry. It dispensed with the battery of analogue dials and gauges that had long been the visual signature of airplane cockpits. In their place were six glowing glass screens, of the cathode-ray-tube variety, arranged neatly beneath the windscreen. The displays presented the pilots with the latest data and readings from the plane’s network of onboard computers.

The A320’s monitor-wrapped flight deck—its “glass cockpit,” as pilots called it—was not its most distinctive feature. Engineers at NASA’s Langley Research Center had pioneered, more than ten years earlier, the use of CRT screens for transmitting flight information, and jet makers had begun installing the screens in passenger planes in the late 1970s.[10] What really set the A320 apart—and made it, in the words of the American writer and pilot William Langewiesche, “the most audacious civil airplane since the Wright brothers’ Flyer”[11]—was its digital fly-by-wire system. Before the A320 arrived, commercial planes still operated mechanically. Their fuselages and wing cavities were rigged with cables, pulleys, and gears, along with a miniature waterworks of hydraulic pipes, pumps, and valves. The controls manipulated by a pilot—the yoke, the throttle levers, the rudder pedals—were linked, by means of the mechanical systems, directly to the moving parts that governed the plane’s orientation, direction, and speed. When the pilot acted, the plane reacted.

To stop a bicycle, you squeeze a lever, which pulls a brake cable, which contracts the arms of a caliper, which presses pads against the tire’s rim. You are, in essence, sending a command—a signal to stop—with your hand, and the brake mechanism carries the manual force of that command all the way to the wheel. Your hand then receives confirmation that your command has been received: you feel, back through the brake lever, the resistance of the caliper, the pressure of the pads against the rim, the skidding of the wheel on the road. That, on a small scale, is what it was like when pilots flew mechanically controlled planes. They became part of the machine, their bodies sensing its workings and feeling its responses, and the machine became a conduit for their will. Such a deep entanglement between human and mechanism was an elemental source of flying’s thrill. It’s what the famous poet-pilot Antoine de Saint-Exupéry must have had in mind when, in recalling his days flying mail planes in the 1920s, he wrote of how “the machine which at first blush seems a means of isolating man from the great problems of nature, actually plunges him more deeply into them.”[12]

The A320’s fly-by-wire system severed the tactile link between pilot and plane. It inserted a digital computer between human command and machine response. When a pilot moved a stick, turned a knob, or pushed a button in the Airbus cockpit, his directive was translated, via a transducer, into an electrical signal that zipped down a wire to a computer, and the computer, following the step-by-step algorithms of its software programs, calculated the various mechanical adjustments required to accomplish the pilot's wish. The computer then sent its own instructions to the digital processors that governed the workings of the plane's moving parts. Along with the replacement of mechanical movements by digital signals came a redesign of cockpit controls. The bulky, two-handed yoke that had pulled cables and compressed hydraulic fluids was replaced in the A320 by a small "sidestick" mounted beside the pilot's seat and gripped by one hand. Along the front console, knobs with small, numerical LED displays allowed the pilot to dial in settings for airspeed, altitude, and heading as inputs to the jet's computers.

After the introduction of the A320, the story of airplanes and the story of computers became one. Every advance in hardware and software, in electronic sensors and controls, in display technologies reverberated through the design of commercial aircraft as manufacturers and airlines pushed the limits of automation. In today's jet-liners, the autopilots that keep planes stable and on course are just one of many computerized systems. Autothrottles control engine power. Flight management systems gather positioning data from GPS receivers and other sensors and use the information to set or refine a flight path. Collision avoidance systems scan the skies for nearby aircraft. Electronic flight bags store digital copies of the charts and other paperwork that pilots used to carry onboard. Still other computers extend and retract the landing gear, apply the brakes, adjust the cabin pressure, and perform various other functions that had once been in the hands of the crew. To program the computers and monitor their outputs, pilots now use large, colorful flat screens that graphically display data generated by electronic flight instrument systems, along with an assortment of keyboards, keypads, scroll wheels, and other input devices. Computer automation has become "all pervasive" on today's aircraft, says Don Harris, an aeronautics professor and ergonomics expert. The flight deck "can be thought of as one huge flying computer interface."[13]

And what of the modern flyboys and flygirls who, nestled in their high-tech glass cockpits, speed through the air alongside the ghosts of Sperry and Post and Saint-Exupéry? Needless to say, the job of the commercial pilot has lost its aura of romance and adventure. The storied stick-and-rudder man, who flew by a sense of feel, now belongs more to legend than to life. On a typical passenger flight these days, the pilot holds the controls for a grand total of three minutes—a minute or two when taking off and another minute or two when landing. What the pilot spends a whole lot of time doing is checking screens and punching in data. "We've gone from a world where automation was a tool to help the pilot control his workload," observes Bill Voss, president of the Flight Safety Foundation, "to a point where the automation is really the primary flight control system in the aircraft."[14] Writes aviation researcher and FAA advisor Hemant Bhana, "As automation has gained in sophistication, the role of the pilot has shifted toward becoming a monitor or supervisor of the automation."[15] The commercial pilot has become a computer operator. And that, many aviation and automation experts have come to believe, is a problem.

■ ■ ■ ■

LAWRENCE SPERRY died in 1923 when his plane crashed into the English Channel. Wiley Post died in 1935 when his plane went down in Alaska. Antoine de Saint-Exupéry died in 1944 when his plane disappeared over the Mediterranean. Premature death was a routine occupational hazard for pilots during aviation's early years; romance and adventure carried a high price. Passengers died with alarming frequency too. As the airline industry took shape in the 1920s, the publisher of a U.S. aviation journal called on the government to improve flight safety, noting that "a great many fatal accidents are daily occurring to people carried in airplanes by inexperienced pilots."[16]

Air travel's lethal days are, mercifully, behind us. Flying is safe now, and pretty much everyone involved in the aviation business believes that advances in automation are one of the reasons why. Together with improvements in aircraft design, airline safety routines, crew training, and air traffic control, the mechanization and computerization of flight have contributed to the sharp and steady decline in accidents and deaths over the decades. In the United States and other Western countries, fatal airliner crashes have become exceedingly rare. Of the more than seven billion people who boarded U.S. commercial flights in the ten years from 2002 through 2011, only 153 ended up dying in a wreck, a rate of two deaths for every million passengers. In the ten years from 1962 through 1971, by contrast, 1.3 billion people took flights, and 1,696 of them died, for a rate of 133 deaths per million.[17]

But this sunny story carries a dark footnote. The overall decline in the number of plane crashes masks the recent arrival of "a spectacularly new type of accident," says Raja Parasuraman, a psychology professor at George Mason University and one of the world's leading authorities on automation.[18] When onboard computer systems fail to work as intended or other unexpected problems arise during a flight, pilots are forced to take manual control of the plane. Thrust abruptly into a now rare role, they too often make mistakes. The consequences, as the Continental Connection and Air France disasters show, can be catastrophic. Over the last thirty years, dozens of psychologists, engineers, and ergonomics, or "human factors," researchers have studied what's gained and lost when pilots share the work of flying with software. They've learned that a heavy reliance on computer automation can erode pilots' expertise, dull their reflexes, and diminish their attentiveness, leading to what Jan Noyes, a human-factors expert at Britain's University of Bristol, calls "a deskilling of the crew."[19]

Concerns about the unintended side effects of flight automation aren't new. They date back at least to the early days of glass cockpits and fly-by-wire controls. A 1989 report from NASA's Ames Research Center noted that as computers had begun to multiply on airplanes during the preceding decade, industry and governmental researchers "developed a growing discomfort that the cockpit may be becoming too automated, and that the steady replacement of human functioning by devices could be a mixed blessing." Despite a general enthusiasm for computerized flight, many in the airline industry worried that “pilots were becoming over-dependent on automation, that manual flying skills may be deteriorating, and that situational awareness might be suffering.”[20]

Studies conducted since then have linked many accidents and near misses to breakdowns of automated systems or to “automation-induced errors” on the part of flight crews.[21] In 2010, the FAA released preliminary results of a major study of airline flights over the preceding ten years which showed that pilot errors had been involved in nearly two-thirds of all crashes. The research further indicated, according to FAA scientist Kathy Abbott, that automation has made such errors more likely. Pilots can be distracted by their interactions with onboard computers, Abbott said, and they can “abdicate too much responsibility to the automated systems.”[22] An extensive 2013 government report on cockpit automation, compiled by an expert panel and drawing on the same FAA data, implicated automation-related problems, such as degraded situational awareness and weakened hand-flying skills, in more than half of recent accidents.[23]

The anecdotal evidence collected through accident reports and surveys gained empirical backing from a rigorous study conducted by Matthew Ebbatson, a young human-factors researcher at Cranfield University, a top U.K. engineering school.[24] Frustrated by the lack of hard, objective data on what he termed “the loss of manual flying skills in pilots of highly automated airliners,” Ebbatson set out to fill the gap. He recruited sixty-six veteran pilots from a British airline and had each of them get into a flight simulator and perform a challenging maneuver—bringing a Boeing 737 with a blown engine in for a landing during bad weather. The simulator disabled the plane’s automated systems, forcing the pilot to fly by hand. Some of the pilots did exceptionally well in the test, Ebbatson reported, but many performed poorly, barely exceeding “the limits of acceptability.” Ebbatson then compared detailed measures of each pilot’s performance in the simulator—the pressure exerted on the yoke, the stability of airspeed, the degree of variation in course—with the pilot’s historical flight record. He found a direct correlation between a pilot’s aptitude at the controls and the amount of time the pilot had spent flying without the aid of automation. The correlation was particularly strong with the amount of manual flying done during the preceding two months. The analysis indicated that “manual flying skills decay quite rapidly towards the fringes of ‘tolerable’ performance without relatively frequent practice.” Particularly “vulnerable to decay,” Ebbatson noted, was a pilot’s ability to maintain “airspeed control”—a skill crucial to recognizing, avoiding, and recovering from stalls and other dangerous situations.

It’s no mystery why automation degrades pilot performance. Like many challenging jobs, flying a plane involves a combination of psychomotor skills and cognitive skills—thoughtful action and active thinking. A pilot needs to manipulate tools and instruments with precision while swiftly and accurately making calculations, forecasts, and assessments in his head. And while he goes through these intricate mental and physical maneuvers, he needs to remain vigilant, alert to what’s going on around him and able to distinguish important signals from unimportant ones. He can’t allow himself either to lose focus or to fall victim to tunnel vision. Mastery of such a multifaceted set of skills comes only with rigorous practice. A beginning pilot tends to be clumsy at the controls, pushing and pulling the yoke with more force than necessary. He often has to pause to remember what he should do next, to walk himself methodically through the steps of a process. He has trouble shifting seamlessly between manual and cognitive tasks. When a stressful situation arises, he can easily become overwhelmed or distracted and end up overlooking a critical change in circumstances.

In time, after much rehearsal, the novice gains confidence. He becomes less halting in his work and more precise in his actions. There's little wasted effort. As his experience continues to deepen, his brain develops so-called mental models—dedicated assemblies of neurons—that allow him to recognize patterns in his surroundings. The models enable him to interpret and react to stimuli intuitively, without getting bogged down in conscious analysis. Eventually, thought and action become seamless. Flying becomes second nature. Years before researchers began to plumb the workings of pilots' brains, Wiley Post described the experience of expert flight in plain, precise terms. He flew, he said in 1935, "without mental effort, letting my actions be wholly controlled by my subconscious mind."[25] He wasn't born with that ability. He developed it through hard work.

When computers enter the picture, the nature and the rigor of the work change, as does the learning the work engenders. As software assumes moment-by-moment control of the craft, the pilot is, as we've seen, relieved of much manual labor. This reallocation of responsibility can provide an important benefit. It can reduce the pilot's workload and allow him to concentrate on the cognitive aspects of flight. But there's a cost. Psychomotor skills get rusty, which can hamper the pilot on those rare but critical occasions when he's required to take back the controls. There's growing evidence that recent expansions in the scope of automation also put cognitive skills at risk. When more advanced computers begin to take over planning and analysis functions, such as setting and adjusting a flight plan, the pilot becomes less engaged not only physically but mentally. Because the precision and speed of pattern recognition appear to depend on regular practice, the pilot's mind may become less agile in interpreting and reacting to fast-changing situations. He may suffer what Ebbatson calls "skill fade" in his mental as well as his motor abilities.

Pilots are not blind to automation's toll. They've always been wary about ceding responsibility to machinery. Airmen in World War I, justifiably proud of their skill in maneuvering their planes during dogfights, wanted nothing to do with the fancy Sperry autopilots.[26] In 1959, the original Mercury astronauts rebelled against NASA's plan to remove manual flight controls from spacecraft.[27] But aviators' concerns are more acute now. Even as they praise the enormous gains in flight technology, and acknowledge the safety and efficiency benefits, they worry about the erosion of their talents. As part of his research, Ebbatson surveyed commercial pilots, asking them whether "they felt their manual flying ability had been influenced by the experience of operating a highly automated aircraft." More than three-fourths reported that "their skills had deteriorated";

just a few felt their skills had improved.[28] A 2012 pilot survey conducted by the European Aviation Safety Agency found similarly widespread concerns, with 95 percent of pilots saying that automation tended to erode "basic manual and cognitive flying skills."[29] Rory Kay, a long-time United Airlines captain who until recently served as the top safety official with the Air Line Pilots Association, fears the aviation industry is suffering from "automation addiction." In a 2011 interview with the Associated Press, he put the problem in stark terms: "We're forgetting how to fly."[30]

■ ■ ■ ■

CYNICS ARE quick to attribute such fears to self-interest. The real reason for the grumbling about automation, they contend, is that pilots are anxious about the loss of their jobs or the squeezing of their paychecks. And the cynics are right, to a degree. As the writer for *Flight* magazine predicted back in 1947, automation technology has whittled down the size of flight crews. Sixty years ago, an airliner's flight deck often had seats for five skilled and well-paid professionals: a navigator, a radio operator, a flight engineer, and a pair of pilots. The radioman lost his chair during the 1950s, as communication systems became more reliable and easier to use. The navigator was pushed off the deck in the 1960s, when inertial navigation systems took over his duties. The flight engineer, whose job involved monitoring a plane's instrument array and relaying important information to the pilots, kept his seat until the advent of the glass cockpit at the end of the 1970s. Seeking to cut costs following the deregulation of air travel in 1978, American airlines made a push to get rid of the engineer and fly with just a captain and copilot. A bitter battle with pilots' unions ensued, as the unions mobilized to save the engineer's job. The fight didn't end until 1981, when a U.S. presidential commission declared that engineers were no longer necessary for the safe operation of passenger flights. Since then, the two-person flight crew has become the norm—at least for the time being. Some experts, pointing to the success of military drones, have begun suggesting that two pilots may in the end be two too many.[31] "A pilotless airliner is going to come," James Albaugh, a top Boeing executive, told an aviation conference in 2011; "it's just a question of when."[32]

The spread of automation has also been accompanied by a steady decline in the compensation of commercial pilots. While veteran jetliner captains can still pull down salaries close to $200,000, novice pilots today are paid as little as $20,000 a year, sometimes even less. The average starting salary for experienced pilots at major airlines is around $36,000, which, as a *Wall Street Journal* reporter notes, is "darn low for mid-career professionals."[33] Despite the modest pay, there's still a popular sense that pilots are overcompensated. An article at the website Salary.com called commercial jet pilots the "most overpaid" professionals in today's economy, arguing that "many of their tasks are automated" and suggesting their work has become "a bit boring."[34]

But pilots' self-interest, when it comes to matters of automation, goes deeper than employment security and pay, or even their own safety. Every technological advance alters the work they do and the role they play, and that in turn changes how they view themselves and how others see them. Their social status and even their sense of self are in play. So when pilots talk about automation, they're speaking not just technically but autobiographically. Am I the master of the machine, or its servant? Am I an actor in the world, or an observer? Am I an agent, or an object? "At heart," MIT technology historian David Mindell writes in his book *Digital Apollo*, "debates about control and automation in aircraft are debates about the relative importance of human and machine." In aviation, as in any field where people work with tools, "technical change and social change are intertwined."[35]

Pilots have always defined themselves by their relationship to their craft. Wilbur Wright, in a 1900 letter to Octave Chanute, another aviation pioneer, said of the pilot's role, "What is chiefly needed is skill rather than machinery."[36] He was not just voicing a platitude. He was referring to what, at the very dawn of human flight, had already become a crucial tension between the capability of the plane and the capability of the pilot. As the first planes were being built, designers debated how inherently stable an aircraft should be—how strong of a tendency it should have to fly straight and level in all conditions. It might seem that more stability would always be better in a flying machine, but that's not so. There's a trade-off between stability and maneuverability. The greater a plane's stability, the harder it becomes for the pilot to exert control over it. As Mindell explains, "The more stable an aircraft is, the more effort will be required to move it off its point of equilibrium. Hence it will be less controllable. The opposite is also true—the more controllable, or maneuverable, an aircraft, the less stable it will be."[37] The author of a 1910 book on aeronautics reported that the question of equilibrium had become "a controversy dividing aviators into two schools." On one side were those who argued that equilibrium should "be made automatic to a very large degree"—that it should be built into the plane. On the other side were those who held that equilibrium should be "a matter for the skill of the aviator."[38]

Wilbur and Orville Wright were in the latter camp. They believed that a plane should be fundamentally unstable, like a bicycle or even, as Wilbur once suggested, "a fractious horse."[39] That way, the pilot would have as much autonomy and freedom as possible. The brothers incorporated their philosophy into the planes they built, which gave precedence to maneuverability over stability. What the Wrights invented at the start of the twentieth century was, Mindell argues, "not simply an airplane that could fly, but also the *very idea* of an airplane as a dynamic machine under the control of a human pilot." [40] Before the engineering decision came an ethical choice: to make the apparatus subservient to the person operating it, an instrument of human talent and volition.

The Wright brothers would lose the equilibrium debate. As planes came to carry passengers and other valuable cargo over long distances, the freedom and virtuosity of the pilot became secondary concerns. Of primary importance were safety and efficiency, and to increase those, it quickly became clear, the pilot's scope of action had to be constrained and the machine itself invested with more authority. The shift in control was gradual, but every time technology assumed a little more power, pilots felt a little more of themselves slip away. In a quixotic 1957 article opposing attempts to further automate flight, a top fighter-jet test pilot named J. O. Roberts fretted about how autopilots were turning the man in the cockpit into little more than “excess baggage except for monitoring duties.” The pilot, Roberts wrote, has to wonder “whether he is paying his way or not.” [41]

But all the gyroscopic, electromechanical, instrumental, and hydraulic innovations only hinted at what digitization would bring. The computer not only changed the character of flight; it changed the character of automation. It circumscribed the pilot’s role to the point where the very idea of “manual control” began to seem anachronistic. If the essence of a pilot’s job consists in sending digital inputs to computers and monitoring computers’ digital outputs—while the computers govern the plane’s moving parts and choose its course—where exactly is the manual control? Even when pilots in computerized planes are pulling yokes or pushing sticks, what they’re often really involved in is a simulation of manual flight. Every action is mediated, filtered through microprocessors. That’s not to say that there aren’t still important skills involved. There are. But the skills have changed, and they’re now applied at a distance, from behind a scrim of software. In many of today’s commercial jets, the flight software can even override the pilots’ inputs during extreme maneuvers. The computer gets the final say. “He didn’t just fly an airplane,” a fellow pilot once said of Wiley Post; “he put it on.” [42] Today’s pilots don’t wear their planes. They wear their planes’ computers—or perhaps the computers wear the pilots.

The transformation that aviation has gone through over the last few decades—the shift from mechanical to digital systems, the proliferation of software and screens, the automation of mental as well as manual work, the blurring of what it means to be a pilot—offers a roadmap for the much broader transformation that society is going through now. The glass cockpit, Don Harris has pointed out, can be thought of as a prototype of a world where “there is computer functionality everywhere.” [43] The experience of pilots also reveals the subtle but often strong connection between the way automated systems are designed and the way the minds and bodies of the people using the systems work. The mounting evidence of an erosion of skills, a dulling of perceptions, and a slowing of reactions should give us all pause. As we begin to live our lives inside glass cockpits, we seem fated to discover what pilots already know: a glass cockpit can also be a glass cage.

---

* A note on terminology: When people talk about a stall, they’re usually referring to a loss of power in an engine. In aviation, a stall refers to a loss of lift in a wing.

CHAPTER FOUR

THE DEGENERATION EFFECT

A HUNDRED YEARS AGO, in his book *An Introduction to Mathematics*, the British philosopher Alfred North Whitehead wrote, "Civilization advances by extending the number of important operations which we can perform without thinking about them." Whitehead wasn't writing about machinery. He was writing about the use of mathematical symbols to represent ideas or logical processes—an early example of how intellectual work can be encapsulated in code. But he intended his observation to be taken generally. The common notion that "we should cultivate the habit of thinking of what we are doing," he wrote, is "profoundly erroneous." The more we can relieve our minds of routine chores, offloading the tasks to technological aids, the more mental power we'll be able to store up for the deepest, most creative kinds of reasoning and conjecture. "Operations of thought are like cavalry charges in battle—they are strictly limited in number, they require fresh horses, and must only be made at decisive moments."[1]

It's hard to imagine a more succinct or confident expression of faith in automation as a cornerstone of progress. Implicit in Whitehead's words is a belief in a hierarchy of human action. Every time we offload a job to a tool or a machine, or to a symbol or a software algorithm, we free ourselves to climb to a higher pursuit, one requiring greater dexterity, richer intelligence, or a broader perspective. We may lose something with each upward step, but what we gain is, in the end, far greater. Taken to an extreme, Whitehead's sense of automation as liberation turns into the techno-utopianism of Wilde and Keynes, or Marx at his sunniest—the dream that machines will free us from our earthly labors and deliver us back to an Eden of leisurely delights. But Whitehead didn't have his head in the clouds. He was making a pragmatic point about how to spend our time and exert our effort. In a publication from the 1970s, the U.S. Department of Labor summed up the job of secretaries by saying that they "relieve their employers of routine duties so they can work on more important matters."[2] Software and other automation technologies, in the Whitehead view, play an analogous role.

History provides plenty of evidence to support Whitehead. People have been handing off chores, both physical and mental, to tools since the invention of the lever, the wheel, and the counting bead. The transfer of work has allowed us to tackle thornier challenges and rise to greater achievements. That's been true on the farm, in the factory, in the laboratory, in the home. But we shouldn't take Whitehead's observation for a universal truth. He was writing when automation was limited to distinct, well-defined, and repetitive tasks—weaving fabric with a steam loom, harvesting grain with a combine, multiplying numbers with a slide rule. Automation is different now. Computers, as we've seen, can be programmed to perform or support complex activities in which a succession of tightly coordinated tasks is carried out through an evaluation of many variables. In automated systems today, the computer often takes on intellectual work—

observing and sensing, analyzing and judging, even making decisions—that until recently was considered the preserve of humans. The person operating the computer is left to play the role of a high-tech clerk, entering data, monitoring outputs, and watching for failures. Rather than opening new frontiers of thought and action to its human collaborators, software narrows our focus. We trade subtle, specialized talents for more routine, less distinctive ones.

Most of us assume, as Whitehead did, that automation is benign, that it raises us to higher callings but doesn't otherwise alter the way we behave or think. That's a fallacy. It's an expression of what scholars of automation have come to call the "substitution myth." A labor-saving device doesn't just provide a substitute for some isolated component of a job. It alters the character of the entire task, including the roles, attitudes, and skills of the people who take part in it. As Raja Parasuraman explained in a 2000 journal article, "Automation does not simply supplant human activity but rather changes it, often in ways unintended and unanticipated by the designers."[3] Automation remakes both work and worker.

■ ■ ■ ■

WHEN PEOPLE tackle a task with the aid of computers, they often fall victim to a pair of cognitive ailments, *automation complacency* and *automation bias*. Both reveal the traps that lie in store when we take the Whitehead route of performing important operations without thinking about them.

Automation complacency takes hold when a computer lulls us into a false sense of security. We become so confident that the machine will work flawlessly, handling any challenge that may arise, that we allow our attention to drift. We disengage from our work, or at least from the part of it that the software is handling, and as a result may miss signals that something is amiss. Most of us have experienced complacency when at a computer. In using email or word-processing software, we become less vigilant proofreaders when the spell checker is on.[4] That's a simple example, which at worst can lead to a moment of embarrassment. But as the sometimes tragic experience of aviators shows, automation complacency can have deadly consequences. In the worst cases, people become so trusting of the technology that their awareness of what's going on around them fades completely. They tune out. If a problem suddenly crops up, they may act bewildered and waste precious moments trying to reorient themselves.

Automation complacency has been documented in many high-risk situations, from battlefields to industrial control rooms to the bridges of ships and submarines. One classic case involved a 1,500-passenger ocean liner named the *Royal Majesty*, which in the spring of 1995 was sailing from Bermuda to Boston on the last leg of a week-long cruise. The ship was outfitted with a state-of-the-art automated navigation system that used GPS signals to keep it on course. An hour into the voyage, the cable for the GPS antenna came loose and the navigation system lost its bearings. It continued to give readings, but they were no longer accurate. For more than thirty hours, as the ship slowly drifted off its appointed route, the captain and crew remained oblivious to the problem, despite clear signs that the system had failed. At one point, a mate on watch was unable to spot an important locational buoy that the ship was due to pass. He failed to report the fact. His trust in the navigation system was so complete that he assumed the buoy was there and he simply didn't see it. Nearly twenty miles off course, the ship finally ran aground on a sandbar near Nantucket Island. No one was hurt, fortunately, though the cruise company suffered millions in damages. Government safety investigators concluded that automation complacency caused the mishap. The ship's officers were "overly reliant" on the automated system, to the point that they ignored other "navigation aids [and] lookout information" that would have told them they were dangerously off course. Automation, the investigators reported, had "the effect of leaving the mariner out of meaningful control or active participation in the operation of the ship."[5]

Complacency can plague people who work in offices as well as those who ply airways and seaways. In an investigation of how design software has influenced the building trades, MIT sociologist Sherry Turkle documented a change in architects' attention to detail. When plans were hand-drawn, architects would painstakingly double-check all the dimensions before handing blueprints over to construction crews. The architects knew that they were fallible, that they could make the occasional goof, and so they followed an old carpentry dictum: measure twice, cut once. With software-generated plans, they're less careful about verifying measurements. The apparent precision of computer renderings and printouts leads them to assume that the figures are accurate. "It seems presumptuous to check," one architect told Turkle; "I mean, how could I do a better job than the computer? It can do things down to hundredths of an inch." Such complacency, which can be shared by engineers and builders, has led to costly mistakes in planning and construction. Computers don't make goofs, we tell ourselves, even though we know that their outputs are only as good as our inputs. "The fancier the computer system," one of Turkle's students observed, "the more you start to assume that it is correcting your errors, the more you start to believe that what comes out of the machine is just how it should be. It is just a visceral thing." [6]

Automation bias is closely related to automation complacency. It creeps in when people give undue weight to the information coming through their monitors. Even when the information is wrong or misleading, they believe it. Their trust in the software becomes so strong that they ignore or discount other sources of information, including their own senses. If you've ever found yourself lost or going around in circles after slavishly following flawed or outdated directions from a GPS device or other digital mapping tool, you've felt the effects of automation bias. Even people who drive for a living can display a startling lack of common sense when relying on satellite navigation. Ignoring road signs and other environmental cues, they'll proceed down hazardous routes and sometimes end up crashing into low overpasses or getting stuck in the narrow streets of small towns. In Seattle in 2008, the driver of a twelve-foot-high bus carrying a high-school sports team ran into a concrete bridge with a nine-foot clearance. The top of the bus was sheared off, and twenty-one injured students had to be taken to the hospital. The driver told police that he had been following GPS instructions and "did not see" signs and flashing lights warning of the low bridge ahead.[7]

Automation bias is a particular risk for people who use decision-support software to guide them through analyses or diagnoses. Since the late 1990s, radiologists have been using computer-aided detection systems that highlight suspicious areas on mammograms and other x-rays. A digital version of an image is scanned into a computer, and pattern-matching software reviews it and adds arrows or other "prompts" to suggest areas for the doctor to inspect more closely. In some cases, the highlights aid in the discovery of disease, helping radiologists identify potential cancers they might otherwise have missed. But studies reveal that the highlights can also have the opposite effect. Biased by the software's suggestions, doctors can end up giving cursory attention to the areas of an image that haven't been highlighted, sometimes overlooking an early-stage tumor or other abnormality. The prompts can also increase the likelihood of false-positives, when a radiologist calls a patient back for an unnecessary biopsy.

A recent review of mammography data, conducted by a team of researchers at City University London, indicates that automation bias has had a greater effect on radiologists and other image readers than was previously thought. The researchers found that while computer-aided detection tends to improve the reliability of "less discriminating readers" in assessing "comparatively easy cases," it can actually degrade the performance of expert readers in evaluating tricky cases. When relying on the software, the experts are more likely to overlook certain cancers.[8] The subtle biases inspired by computerized decision aids may, moreover, be "an inherent part of the human cognitive apparatus for reacting to cues and alarms."[9] By directing the focus of our eyes, the aids distort our vision.

Both complacency and bias seem to stem from limitations in our ability to pay attention. Our tendency toward complacency reveals how easily our concentration and awareness can fade when we're not routinely called on to interact with our surroundings. Our propensity to be biased in evaluating and weighing information shows that our mind's focus is selective and can easily be skewed by misplaced trust or even the appearance of seemingly helpful prompts. Both complacency and bias tend to become more severe as the quality and reliability of an automated system improve.[10] Experiments show that when a system produces errors fairly frequently, we stay on high alert. We maintain awareness of our surroundings and carefully monitor information from a variety of sources. But when a system is more reliable, breaking down or making mistakes only occasionally, we get lazy. We start to assume the system is infallible.

Because automated systems usually work fine even when we lose awareness or objectivity, we are rarely penalized for our complacency or our bias. That ends up compounding the problems, as Parasuraman pointed out in a 2010 paper written with his German colleague Dietrich Manzey. "Given the usually high reliability of automated systems, even highly complacent and biased behavior of operators rarely leads to obvious performance consequences," the scholars wrote. The lack of negative feedback can in time induce "a cognitive process that resembles what has been referred to as 'learned carelessness.' "[11] Think about driving a car when you're sleepy. If you begin to nod off and drift out of your lane, you'll usually go onto a rough shoulder, hit a rumble strip, or earn a honk from another motorist—signals that jolt you back awake. If you're in a car that automatically keeps you within a lane by monitoring the lane markers and adjusting the steering, you won't receive such warnings. You'll drift into a deeper slumber. Then if something unexpected happens—an animal runs into the road, say, or a car stops short in front of you—you'll be much more likely to have an accident. By isolating us from negative feedback, automation makes it harder for us to stay alert and engaged. We tune out even more.

■ ■ ■ ■

OUR SUSCEPTIBILITY to complacency and bias explains how a reliance on automation can lead to errors of both commission and omission. We accept and act on information that turns out to be incorrect or incomplete, or we fail to see things that we should have seen. But the way that a reliance on computers weakens awareness and attentiveness also points to a more insidious problem. Automation tends to turn us from actors into observers. Instead of manipulating the yoke, we watch the screen. That shift may make our lives easier, but it can also inhibit our ability to learn and to develop expertise. Whether automation enhances or degrades our performance in a given task, over the long run it may diminish our existing skills or prevent us from acquiring new ones.

Since the late 1970s, cognitive psychologists have been documenting a phenomenon called the generation effect. It was first observed in studies of vocabulary, which revealed that people remember words much better when they actively call them to mind—when they *generate* them—than when they read them from a page. In one early and famous experiment, conducted by University of Toronto psychologist Norman Slamecka, people used flash cards to memorize pairs of antonyms, like *hot* and *cold*. Some of the test subjects were given cards that had both words printed in full, like this:

HOT : COLD

Others used cards that showed only the first letter of the second word, like this:

HOT : C

The people who used the cards with the missing letters performed much better in a subsequent test measuring how well they remembered the word pairs. Simply forcing their minds to fill in a blank, to act rather than observe, led to stronger retention of information.[12]

The generation effect, it has since become clear, influences memory and learning in many different circumstances. Experiments have revealed evidence of the effect in tasks that involve not only remembering letters and words but also remembering numbers, pictures, and sounds, completing math problems, answering trivia questions, and reading for comprehension. Recent studies have also demonstrated the benefits of the generation effect for higher forms of teaching and learning. A 2011 paper in *Science* showed that students who read a complex science assignment during a study period and then spent a second period recalling as much of it as possible, unaided, learned the material more fully than students who read the assignment repeatedly over the course of four study periods.[13] The mental act of generation improves people's ability to carry out activities that, as education researcher Britte Haugan Cheng has written, "require conceptual reasoning and requisite deeper cognitive processing." Indeed, Cheng says, the generation effect appears to strengthen as the material generated by the mind becomes more complex.[14]

Psychologists and neuroscientists are still trying to figure out what goes on in our minds to give rise to the generation effect. But it's clear that deep cognitive and memory processes are involved. When we work hard at something, when we make it the focus of attention and effort, our mind rewards us with greater understanding. We remember more and we learn more. In time, we gain know-how, a particular talent for acting fluidly, expertly, and purposefully in the world. That's hardly a surprise. Most of us know that the only way to get good at something is by actually doing it. It's easy to gather information quickly from a computer screen—or from a book, for that matter. But true knowledge, particularly the kind that lodges deep in memory and manifests itself in skill, is harder to come by. It requires a vigorous, prolonged struggle with a demanding task.

The Australian psychologists Simon Farrell and Stephan Lewandowsky made the connection between automation and the generation effect in a paper published in 2000. In Slamecka's experiment, they pointed out, supplying the second word of an antonym pair, rather than forcing a person to call the word to mind, "can be considered an instance of automation because a human activity—generation of the word 'COLD' by participants—has been obviated by a printed stimulus." By extension, "the reduction in performance that is observed when generation is replaced by reading can be considered a manifestation of complacency."[15] That helps illuminate the cognitive cost of automation. When we carry out a task or a job on our own, we seem to use different mental processes than when we rely on the aid of a computer. When software reduces our engagement with our work, and in particular when it pushes us into a more passive role as observer or monitor, we circumvent the deep cognitive processing that underpins the generation effect. As a result, we hamper our ability to gain the kind of rich, real-world knowledge that leads to know-how. The generation effect requires precisely the kind of struggle that automation seeks to alleviate.

In 2004, Christof van Nimwegen, a cognitive psychologist at Utrecht University in the Netherlands, began a series of simple but ingenious experiments to investigate software's effects on memory formation and the development of expertise.[16] He recruited two groups of people and had them play a computer game based on a classic logic puzzle called Missionaries and Cannibals. To complete the puzzle, a player has to transport across a hypothetical river five missionaries and five cannibals (or, in van Nimwegen's version, five yellow balls and five blue ones), using a boat that can accommodate no more than three passengers at a time. The tricky part is that there can never be more cannibals than missionaries in one place, either in the boat or on the riverbanks. (If outnumbered, the missionaries become the cannibals' dinner, one assumes.) Figuring out the series of boat trips that can best accomplish the task requires rigorous analysis and careful planning.

One of van Nimwegen's groups worked on the puzzle using software that provided step-by-step guidance, offering, for instance, on-screen prompts to highlight which moves were permissible and which weren't. The other group used a rudimentary program that offered no assistance. As you'd expect, the people using the helpful software made faster progress at the outset. They could follow the prompts rather than having to pause before each move to recall the rules and figure out how they applied to the new situation. But as the game advanced, the players using the rudimentary software began to excel. In the end, they were able to work out the puzzle more efficiently, with significantly fewer wrong moves, than their counterparts who were receiving assistance. In his report on the experiment, van Nimwegen concluded that the subjects using the rudimentary program developed a clearer conceptual understanding of the task. They were better able to think ahead and plot a successful strategy. Those relying on guidance from the software, by contrast, often became confused and would "aimlessly click around."

The cognitive penalty imposed by the software aids became even clearer eight months later, when van Nimwegen had the same people work through the puzzle again. Those who had earlier used the rudimentary software finished the game almost twice as quickly as their counterparts. The subjects using the basic program, he wrote, displayed "more focus" during the task and "better imprinting of knowledge" afterward. They enjoyed the benefits of the generation effect. Van Nimwegen and some of his Utrecht colleagues went on to conduct experiments involving more realistic tasks, such as using calendar software to schedule meetings and event-planning software to assign conference speakers to rooms. The results were the same. People who relied on the help of software prompts displayed less strategic thinking, made more superfluous moves, and ended up with a weaker conceptual understanding of the assignment. Those using unhelpful programs planned better, worked smarter, and learned more.[17]

What van Nimwegen observed in his laboratory—that when we automate cognitive tasks like problem solving, we hamper the mind's ability to translate information into knowledge and knowledge into know-how—is also being documented in the real world. In many businesses, managers and other professionals depend on so-called expert systems to sort and analyze information and suggest courses of action. Accountants, for example, use decision-support software in corporate audits. The applications speed the work, but there are signs that as the software becomes more capable, the accountants become less so. One study, conducted by a group of Australian professors, examined the effects of the expert systems used by three international accounting firms. Two of the companies employed advanced software that, based on an accountant's answers to basic questions about a client, recommended a set of relevant business risks to include in the client's audit file. The third firm used simpler software that provided a list of potential risks but required the accountant to review them and manually select the pertinent ones for the file. The researchers gave accountants from each firm a test measuring their knowledge of risks in industries in which they had performed audits. Those from the firm with the less helpful software displayed a significantly stronger understanding of different forms of risk than did those from the other two firms. The decline in learning associated with advanced software affected even veteran auditors—those with more than five years of experience at their current firm.[18]

Other studies of expert systems reveal similar effects. The research indicates that while decision-support software can help novice analysts make better judgments in the short run, it can also make them mentally lazy. By diminishing the intensity of their thinking, the software retards their ability to encode information in memory, which makes them less likely to develop the rich tacit knowledge essential to true expertise.[19] The drawbacks to automated decision aids can be subtle, but they have real consequences, particularly in fields where analytical errors have far-reaching repercussions. Miscalculations of risk, exacerbated by high-speed computerized trading programs, played a major role in the near meltdown of the world's financial system in 2008. As Tufts University management professor Amar Bhidé has suggested, "robotic methods" of decision making led to a widespread "judgment deficit" among bankers and other Wall Street professionals.[20] While it may be impossible to pin down the precise degree to which automation figured in the disaster, or in subsequent fiascos like the 2010 "flash crash" on U.S. exchanges, it seems prudent to take seriously any indication that a widely used technology may be diminishing the knowledge or clouding the judgment of people in sensitive jobs. In a 2013 paper, computer scientists Gordon Baxter and John Cartlidge warned that a reliance on automation is eroding the skills and knowledge of financial professionals even as computer-trading systems make financial markets more risky.[21]

Some software writers worry that their profession's push to ease the strain of thinking is taking a toll on their own skills. Programmers today often use applications called integrated development environments, or IDEs, to aid them in composing code. The applications automate many tricky and time-consuming chores. They typically incorporate auto-complete, error-correction, and debugging routines, and the more sophisticated of them can evaluate and revise the structure of a program through a process known as refactoring. But as the applications take over the work of coding, programmers lose opportunities to practice their craft and sharpen their talent. "Modern IDEs are getting 'helpful' enough that at times I feel like an IDE operator rather than a programmer," writes Vivek Haldar, a veteran software developer with Google. "The behavior all these tools encourage is not 'think deeply about your code and write it carefully,' but 'just write a crappy first draft of your code, and then the tools will tell you not just what's wrong with it, but also how to make it better.' " His verdict: "Sharp tools, dull minds." [22]

Google acknowledges that it has even seen a dumbing-down effect among the general public as it has made its search engine more responsive and solicitous, better able to predict what people are looking for. Google does more than correct our typos; it suggests search terms as we type, untangles semantic ambiguities in our requests, and anticipates our needs based on where we are and how we've behaved in the past. We might assume that as Google gets better at helping us refine our searching, we would learn from its example. We would become more sophisticated in formulating keywords and otherwise honing our online explorations. But according to the company's top search engineer, Amit Singhal, the opposite is the case. In 2013, a reporter from the *Observer* newspaper in London interviewed Singhal about the many improvements that have been made to

Google's search engine over the years. "Presumably," the journalist remarked, "we have got more precise in our search terms the more we have used Google." Singhal sighed and, "somewhat wearily," corrected the reporter: " 'Actually, it works the other way. The more accurate the machine gets, the lazier the questions become.' "[23]

More than our ability to compose sophisticated queries may be compromised by the ease of search engines. A series of experiments reported in *Science* in 2011 indicates that the ready availability of information online weakens our memory for facts. In one of the experiments, test subjects read a few-dozen simple, true statements—"an ostrich's eye is bigger than its brain," for instance—and then typed them into a computer. Half the subjects were told the computer would save what they typed; the other half were told that the statements would be erased. Afterward, the participants were asked to write down all the statements they could recall. People who believed the information had been stored in the computer remembered significantly fewer of the facts than did those who assumed the statements had not been saved. Just knowing that information will be available in a database appears to reduce the likelihood that our brains will make the effort required to form memories. "Since search engines are continually available to us, we may often be in a state of not feeling we need to encode the information internally," the researchers concluded. "When we need it, we will look it up."[24]

For millennia, people have supplemented their biological memory with storage technologies, from scrolls and books to microfiche and magnetic tape. Tools for recording and distributing information underpin civilization. But external storage and biological memory are not the same thing. Knowledge involves more than looking stuff up; it requires the encoding of facts and experiences in personal memory. To truly know something, you have to weave it into your neural circuitry, and then you have to repeatedly retrieve it from memory and put it to fresh use. With search engines and other online resources, we've automated information storage and retrieval to a degree far beyond anything seen before. The brain's seemingly innate tendency to offload, or externalize, the work of remembering makes us more efficient thinkers in some ways. We can quickly call up facts that have slipped our mind. But that same tendency can become pathological when the automation of mental labor makes it too easy to avoid the work of remembering and understanding.

Google and other software companies are, of course, in the business of making our lives easier. That's what we ask them to do, and it's why we're devoted to them. But as their programs become adept at doing our thinking for us, we naturally come to rely more on the software and less on our own smarts. We're less likely to push our minds to do the work of generation. When that happens, we end up learning less and knowing less. We also become less capable. As the University of Texas computer scientist Mihai Nadin has observed, in regard to modern software, "The more the interface replaces human effort, the lower the adaptivity of the user to new situations."[25] In place of the generation effect, computer automation gives us the reverse: a degeneration effect.

■ ■ ■ ■

BEAR WITH me while I draw your attention back to that ill-fated, slicker-yellow Subaru with the manual transmission. As you'll recall, I went from hapless gear-grinder to reasonably accomplished stick-handler with just a few weeks' practice. The arm and leg movements my dad had taught me, cursorily, now seemed instinctive. I was hardly an expert, but shifting was no longer a struggle. I could do it without thinking. It had become, well, automatic.

My experience provides a model for the way humans gain complicated skills. We often start off with some basic instruction, received directly from a teacher or mentor or indirectly from a book or manual or YouTube video, which transfers to our conscious mind explicit knowledge about how a task is performed: do this, then this, then this. That's what my father did when he showed me the location of the gears and explained when to step on the clutch. As I quickly discovered, explicit knowledge goes only so far, particularly when the task has a psychomotor component as well as a cognitive one. To achieve mastery, you need to develop tacit knowledge, and that comes only through real experience—by rehearsing a skill, over and over again. The more you practice, the less you have to think about what you're doing. Responsibility for the work shifts from your conscious mind, which tends to be slow and halting, to your unconscious mind, which is quick and fluid. As that happens, you free your conscious mind to focus on the more subtle aspects of the skill, and when those, too, become automatic, you proceed up to the next level. Keep going, keep pushing yourself, and ultimately, assuming you have some native aptitude for the task, you're rewarded with expertise.

This skill-building process, through which talent comes to be exercised without conscious thought, goes by the ungainly name *automatization*, or the even more ungainly name *proceduralization*. Automatization involves deep and widespread adaptations in the brain. Certain brain cells, or neurons, become fine-tuned for the task at hand, and they work in concert through the electrochemical connections provided by synapses. The New York University cognitive psychologist Gary Marcus offers a more detailed explanation: "At the neural level, proceduralization consists of a wide array of carefully coordinated processes, including changes to both gray matter (neural cell bodies) and white matter (axons and dendrites that connect between neurons). Existing neural connections (synapses) must be made more efficient, new dendritic spines may be formed, and proteins must be synthesized."[26] Through the neural modifications of automatization, the brain develops *automaticity*, a capacity for rapid, unconscious perception, interpretation, and action that allows mind and body to recognize patterns and respond to changing circumstances instantaneously.

All of us experienced automatization and achieved automaticity when we learned to read. Watch a young child in the early stages of reading instruction, and you'll witness a taxing mental struggle. The child has to identify each letter by studying its shape. She has to sound out how a set of letters combine to form a syllable and how a series of syllables combine to form a word. If she's not already familiar with the word, she has to figure out or be told its meaning. And then, word by word, she has to interpret the meaning of a sentence, often resolving the ambiguities inherent to language. It's a slow, painstaking process, and it requires the full attention of the conscious mind. Eventually, though, letters and then words get encoded in the neurons of the visual cortex—the part of the brain that processes sight—and the young reader begins to recognize them without conscious thought. Through a symphony of brain changes, reading becomes effortless. The greater the automaticity the child achieves, the more fluent and accomplished a reader she becomes.[27]

Whether it's Wiley Post in a cockpit, Serena Williams on a tennis court, or Magnus Carlsen at a chessboard, the otherworldly talent of the virtuoso springs from automaticity. What looks like instinct is hard-won skill. Those changes in the brain don't happen through passive observation. They're generated through repeated confrontations with the unexpected. They require what the philosopher of mind Hubert Dreyfus terms "experience in a variety of situations, all seen from the same perspective but requiring different tactical decisions."[28] Without lots of practice, lots of repetition and rehearsal of a skill in different circumstances, you and your brain will never get really good at anything, at least not anything complicated. And without continuing practice, any talent you do achieve will get rusty.

It's popular now to suggest that practice is all you need. Work at a skill for ten thousand hours or so, and you'll be blessed with expertise—you'll become the next great pastry chef or power forward. That, unhappily, is an exaggeration. Genetic traits, both physical and intellectual, do play an important role in the development of talent, particularly at the highest levels of achievement. Nature matters. Even our desire and aptitude for practice has, as Marcus points out, a genetic component: "How we respond to experience, and even what type of experience we seek, are themselves in part functions of the genes we are born with."[29] But if genes establish, at least roughly, the upper bounds of individual talent, it's only through practice that a person will ever reach those limits and fulfill his or her potential. While innate abilities make a big difference, write psychology professors David Hambrick and Elizabeth Meinz, "research has left no doubt that one of the largest sources of individual differences in performance on complex tasks is simply what and how much people know: declarative, procedural, and strategic knowledge acquired through years of training and practice in a domain."[30]

Automaticity, as its name makes clear, can be thought of as a kind of internalized automation. It's the body's way of making difficult but repetitive work routine. Physical movements and procedures get programmed into muscle memory; interpretations and judgments are made through the instant recognition of environmental patterns apprehended by the senses. The conscious mind, scientists discovered long ago, is surprisingly cramped, its capacity for taking in and processing information limited. Without automaticity, our consciousness would be perpetually overloaded. Even very simple acts, such as reading a sentence in a book or cutting a piece of steak with a knife and fork, would strain our cognitive capabilities. Automaticity gives us more headroom. It increases, to put a different spin on Alfred North Whitehead's observation, "the number of important operations which we can perform without thinking about them."

Tools and other technologies, at their best, do something similar, as Whitehead appreciated. The brain's capacity for automaticity has limits of its own. Our unconscious mind can perform a lot of functions quickly and efficiently, but it can't do everything.

You might be able to memorize the times table up to twelve or even twenty, but you would probably have trouble memorizing it much beyond that. Even if your brain didn't run out of memory, it would probably run out of patience. With a simple pocket calculator, though, you can automate even very complicated mathematical procedures, ones that would tax your unaided brain, and free up your conscious mind to consider what all that math adds up to. But that only works if you've already mastered basic arithmetic through study and practice. If you use the calculator to bypass learning, to carry out procedures that you haven't learned and don't understand, the tool will not open up new horizons. It won't help you gain new mathematical knowledge and skills. It will simply be a black box, a mysterious number-producing mechanism. It will be a barrier to higher thought rather than a spur to it.

That's what computer automation often does today, and it's why Whitehead's observation has become misleading as a guide to technology's consequences. Rather than extending the brain's innate capacity for automaticity, automation too often becomes an impediment to automatization. In relieving us of repetitive mental exercise, it also relieves us of deep learning. Both complacency and bias are symptoms of a mind that is not being challenged, that is not fully engaged in the kind of real-world practice that generates knowledge, enriches memory, and builds skill. The problem is compounded by the way computer systems distance us from direct and immediate feedback about our actions. As the psychologist K. Anders Ericsson, an expert on talent development, points out, regular feedback is essential to skill building. It's what lets us learn from our mistakes and our successes. "In the absence of adequate feedback," Ericsson explains, "efficient learning is impossible and improvement only minimal even for highly motivated subjects."[31]

Automaticity, generation, flow: these mental phenomena are diverse, they're complicated, and their biological underpinnings are understood only fuzzily. But they are all related, and they tell us something important about ourselves. The kinds of effort that give rise to talent—characterized by challenging tasks, clear goals, and direct feedback—are very similar to those that provide us with a sense of flow. They're immersive experiences. They also describe the kinds of work that force us to actively generate knowledge rather than passively take in information. Honing our skills, enlarging our understanding, and achieving personal satisfaction and fulfillment are all of a piece. And they all require tight connections, physical and mental, between the individual and the world. They all require, to quote the American philosopher Robert Talisse, "getting your hands dirty with the world and letting the world kick back in a certain way."[32] Automaticity is the inscription the world leaves on the active mind and the active self. Know-how is the evidence of the richness of that inscription.

From rock climbers to surgeons to pianists, Mihaly Csikszentmihalyi explains, people who "routinely find deep enjoyment in an activity illustrate how an organized set of challenges and a corresponding set of skills result in optimal experience." The jobs or hobbies they engage in "afford rich opportunities for action," while the skills they develop allow them to make the most of those opportunities. The ability to act with aplomb in the world turns all of us into artists. "The effortless absorption experienced by the practiced artist at work on a difficult project always is premised upon earlier mastery of a complex body of skills."[33] When automation distances us from our work, when it gets between us and the world, it erases the artistry from our lives.