Chapter 6

THE HEALTH CARE CHALLENGE

In May 2012, a fifty-five-year-old man checked into a clinic at the University of Marburg in Germany. The patient suffered from fever, an inflamed esophagus, low thyroid hormone levels, and failing vision. He had visited a series of doctors, all of whom were baffled by his condition. By the time he arrived at the Marburg clinic, he was nearly blind and was on the verge of heart failure. Months earlier, and a continent away, a very similar medical mystery had culminated with a fifty-nine-year-old woman receiving a heart transplant at the University of Colorado Medical Center in Denver.

The answer to both mysteries turned out to be the same: cobalt poisoning.[1] Both patients had previously received artificial hips made from metal. The metal implants had abraded over time, releasing cobalt particles and exposing the patients to chronic toxicity. In a remarkable coincidence, papers describing the two cases were published independently in two leading medical journals on nearly the same day in February 2014. The report published by the German doctors came with a fascinating twist: whereas the American team had resorted to surgery, the German team had managed to solve the

mystery not because of their training but because one of the doctors had seen a February 2011 episode of the television show *House*. In the episode, the show's protagonist, Dr. Gregory House, is faced with the same problem and makes an ingenious diagnosis: cobalt poisoning resulting from a metal prosthetic hip replacement.

The fact that two teams of doctors can struggle to make the same diagnosis—and that they can do so even when the answer to the mystery has been broadcast to millions of prime-time television viewers—is a testament to the extent to which medical knowledge and diagnostic skill are compartmentalized in the brains of individual physicians, even in an age when the Internet has enabled an unprecedented degree of collaboration and access to information. As a result, the fundamental process that doctors use to diagnose and treat illnesses has remained, in important ways, relatively unchanged. Upending that traditional approach to problem solving, and unleashing all the information trapped in individual minds or published in obscure medical journals, likely represents one of the most important potential benefits of artificial intelligence and big data as applied to medicine.

In general, the advances in information technology that are disrupting other areas of the economy have so far made relatively few inroads into the health care sector. Especially hard to find is any evidence that technology is resulting in meaningful improvements in overall efficiency. In 1960, health care represented less than 6 percent of the US economy.[2] By 2013 it had nearly tripled, having grown to nearly 18 percent, and per capita health care spending in the United States had soared to a level roughly double that of most other industrialized countries. One of the greatest risks going forward is that technology will continue to impact asymmetrically, driving down wages or creating unemployment across most of the economy, even as the cost of health care continues to climb. The danger, in a sense, is not too many health care robots but too few. If technology fails to rise to the health care challenge, the result is likely to be a soaring,

and ultimately unsustainable, burden on both individual households and the economy as a whole.

Artificial Intelligence in Medicine

The total amount of information that could potentially be useful to a physician attempting to diagnose a particular patient's condition or design an optimal treatment strategy is staggering. Physicians are faced with a continuous torrent of new discoveries, innovative treatments, and clinical study evaluations published in medical and scientific journals throughout the world. For example, MEDLINE, an online database maintained by the US National Library of Medicine, indexes over 5,600 separate journals—each of which might publish anywhere from dozens to hundreds of distinct research papers every year. In addition, there are millions of medical records, patient histories, and case studies that might offer important insights. According to one estimate, the total volume of all this data doubles roughly every five years.[3] It would be impossible for any human being to assimilate more than a tiny fraction of the relevant information even within highly specific areas of medical practice.

As we saw in Chapter 4, medicine is one of the primary areas where IBM foresees its Watson technology having a transformative impact. IBM's system is capable of churning through vast troves of information in disparate formats and then almost instantly constructing inferences that might elude even the most attentive human researcher. It's easy to imagine a near-term future where such a diagnostic tool is considered indispensable, at least for physicians confronting especially challenging cases.

The MD Anderson Cancer Center at the University of Texas handles over 100,000 patients at its Houston hospital each year and is generally regarded as the best cancer treatment facility in the United States. In 2011, IBM's Watson team began working with MD Anderson's doctors to build a customized version of the system geared

toward assisting oncologists working with leukemia cases. The goal is to create an interactive adviser capable of recommending the best evidence-based treatment options, matching patients with clinical drug trials, and highlighting possible dangers or side effects that might threaten specific patients. Initial progress on the project proved to be somewhat slower than the team expected, largely because of the challenges associated with designing algorithms capable of taking on the complexities of cancer diagnosis and treatment. Cancer, it turns out, is tougher than *Jeopardy!* Nonetheless, by January 2014, the *Wall Street Journal* reported that the Watson-based leukemia system at MD Anderson was "back on track" toward becoming operational.[4] Researchers hope to expand the system to handle other kinds of cancer within roughly two years. It's very likely that the lessons IBM takes away from this pilot program will enable the company to streamline future implementations of the Watson technology.

Once the system is operating smoothly, the MD Anderson staff plans to make it available via the Internet so that it can become a powerful resource for doctors everywhere. According to Dr. Courtney DiNardo, a leukemia expert, the Watson technology has the "potential to democratize cancer care" by allowing any physician to "access the latest scientific knowledge and MD Anderson's expertise." "For physicians who aren't leukemia experts," she added, the system "can function as an expert second opinion, allowing them to access the same knowledge and information" relied on by the nation's top cancer treatment center. DiNardo also believes that, beyond offering advice for specific patients, the system "will provide an unparalleled research platform that can be used to generate questions, explore hypotheses and provide answers to critical research questions."[5]

Watson is currently the most ambitious and prominent application of artificial intelligence to medicine, but there are other important success stories as well. In 2009, researchers at the Mayo Clinic in Rochester, Minnesota, built an artificial neural network designed to

diagnose cases of endocarditis—an inflammation of the inner layer of the heart. Endocarditis normally requires that a probe be inserted into the patient's esophagus in order to determine whether or not the inflammation is caused by a potentially deadly infection—a procedure that is uncomfortable, expensive, and itself carries risks for the patient. The Mayo doctors instead trained a neural network to make the diagnosis based on routine tests and observable symptoms alone, without the need for the invasive technique. A study involving 189 patients found that the system was accurate more than 99 percent of the time and successfully saved over half of the patients from having to needlessly undergo the invasive diagnostic procedure.[6]

One of the most important benefits of artificial intelligence in medicine is likely to be the avoidance of potentially fatal errors in both diagnosis and treatment. In November 1994, Betsy Lehman, a thirty-nine-year-old mother of two and a widely read columnist who wrote about health-related issues for the *Boston Globe,* was scheduled to begin her third round of chemotherapy as she continued her battle against breast cancer. Lehman was admitted to the Dana-Farber Cancer Institute in Boston, which, like MD Anderson, is regarded as one of the country's preeminent cancer centers. The treatment plan called for Lehman to be given a powerful dose of cyclophosphamide—a highly toxic drug intended to wipe out her cancer cells. The research fellow who wrote the medication order made a simple numerical error, which meant that the total dosage Lehman received was about four times what the treatment plan actually called for. Lehman died from the overdose on December 3, 1994.[7]

Lehman was just one of as many as 98,000 patients who die in the United States each year as a direct result of preventable medical errors.[8] A 2006 report by the US Institute of Medicine estimated that at least 1.5 million Americans are harmed by medication errors alone, and that such mistakes result in more than $3.5 billion in additional annual treatment costs.[9] An AI system with access to detailed patient

histories, as well as information about medications, including their associated toxicity and side effects, would potentially be able to prevent errors even in very complex situations involving the interaction of multiple drugs. Such a system could act as an interactive adviser to doctors and nurses, offering instantaneous verification of both safety and effectiveness before medication is administered, and—especially in situations where hospital staff are tired or distracted—it would be very likely to save both lives and needless discomfort and expense.

Once medical applications of artificial intelligence evolve to the point where the systems can act as true advisers capable of providing consistently high-quality second opinions, the technology could also help rein in the high costs associated with malpractice liability. Many physicians feel the need to practice "defensive medicine" and order every conceivable test in an attempt to protect themselves against potential lawsuits. A documented second opinion from an AI system versed in best practice standards could offer doctors a "safe harbor" defense against such claims. The result might be less spending on needless medical tests and scans as well as lower malpractice insurance premiums.*

Looking even further ahead, we can easily imagine artificial intelligence having a genuinely transformative impact on the way medical services are delivered. Once machines demonstrate that they can offer accurate diagnosis and effective treatment, perhaps it will not be necessary for a physician to directly oversee every encounter with every patient.

---

* This raises the question of whether the liability would simply migrate to the manufacturer of the AI system. Since such systems might be used to diagnose tens or even hundreds of thousands of patients, the potential liability for errors could be daunting. However, the US Supreme Court ruled in the 2008 case *Riegel v. Medtronic, Inc.*, that medical device manufacturers are protected from some lawsuits if their products have been approved by the FDA. Perhaps similar reasoning would be extended to diagnostic systems. Another issue is that previous attempts to create "safe harbor" laws for doctors have been vigorously opposed by the trial lawyers, who have a great deal of political influence.

In an op-ed I wrote for the *Washington Post,* shortly after Watson's 2011 triumph at playing *Jeopardy!,* I suggested that there may eventually be an opportunity to create a new class of medical professionals: persons educated with perhaps a four-year college or master's degree, and who are trained primarily to interact with and examine patients—and then to convey that information into a standardized diagnostic and treatment system.[10] These new, lower-cost practitioners would be able to take on many routine cases, and could be deployed to help manage the dramatically growing number of patients with chronic conditions such as obesity and diabetes.

Physicians groups would, of course, be likely to oppose the influx of these less-educated competitors.* However, the reality is that the vast majority of medical school graduates are not especially interested in entering family practice, and they are even less excited about serving rural areas of the country. Various studies predict a shortage of up to 200,000 doctors within the next fifteen years as older doctors retire, the Affordable Care Act plan brings as many as 32 million new patients into the health insurance system, and an aging population requires more care.[11] The shortage will be most acute among primary-care physicians as medical school graduates, typically burdened by onerous levels of student debt, choose overwhelmingly to enter more lucrative specialties.

These new practitioners, trained to utilize a standardized AI system that encapsulates much of the knowledge that doctors acquire during the course of nearly a decade of intensive training, could handle routine cases, while referring patients who require more specialized care to physicians. College graduates would benefit significantly from the availability of a compelling new career path, especially as intelligent software increasingly erodes opportunities in other sectors of the job market.

---

* Nurse practitioners with advanced degrees have been able to overcome such political opposition in seventeen US states and are likely to be an important component of primary care in the future.

In some areas of medicine, particularly those that don't require direct interaction with patients, advances in AI are poised to drive dramatic productivity increases and perhaps eventually full automation. Radiologists, for example, are trained to interpret the images that result from various medical scans. Image processing and recognition technology is advancing rapidly and may soon be able to usurp the radiologist's traditional role. Software can already recognize people in photos posted on Facebook and even help identify potential terrorists in airports. In September 2012, the FDA approved an automated ultrasound system for screening women for breast cancer. The device, designed by U-Systems, Inc., is designed to help identify cancer in the roughly 40 percent of women whose dense breast tissue can render standard mammogram technology ineffective. Radiologists still need to interpret the images, but doing so now takes only about three minutes. That compares with twenty to thirty minutes for images produced using standard handheld ultrasound technology.[12]

Automated systems can also provide a viable second opinion. A very effective—but expensive—way to increase cancer detection rates is to have two radiologists read every mammogram image separately and then reach a consensus on any potential anomalies identified by either doctor. This "double reading" strategy results in significantly improved cancer detection and also dramatically reduces the number of patients who have to be recalled for further testing. A 2008 study published in the *New England Journal of Medicine* found that a machine can step into the role of the second doctor. When a radiologist is paired with a computer-aided detection system, the results are just as good as having two doctors separately interpret the images.[13]

Pathology is another area where artificial intelligence is already encroaching. Each year, over a hundred million women throughout the world receive a Pap test to screen for cervical cancer. The test requires that cervical cells be deposited on a glass microscope slide and

then be examined by a technician or doctor for signs of malignancy. It's a labor-intensive process that can cost up to $100 per test. Many diagnostic labs, however, are now turning to a powerful automated imaging system manufactured by BD, a New Jersey–based medical device company. In a 2011 series of articles about job automation for *Slate,* technology columnist Farhad Manjoo called the BD FocalPoint GS Imaging System "a marvel of medical engineering" whose "image-searching software rapidly scans slides in search of more than 100 visual signs of abnormal cells." The system then "ranks the slides according to the likelihood they contain disease" and finally "identifies 10 areas on each slide for a human to scrutinize."[14] The machine does a significantly better job of finding instances of cancer than human analysts alone, even as it roughly doubles the speed at which the tests can be processed.

Hospital and Pharmacy Robotics

The pharmacy at the University of California Medical Center in San Francisco prepares about 10,000 individual doses of medication every day, and yet a pharmacist never touches a pill or a medicine bottle. A massive automated system manages thousands of different drugs and handles everything from storing and retrieving bulk pharmaceutical supplies to dispensing and packaging individual tablets. A robotic arm continuously picks pills from an array of bins and places them in small plastic bags. Every dose goes into a separate bag and is labeled with a barcode that identifies both the medication and the patient who should receive it. The machine then arranges each patient's daily meds in the order that they need to be taken and binds them together. Later, the nurse who administers the medication will scan the barcodes on both the dosage bag and the patient's wrist band. If they don't match, or if the medication is being given at the wrong time, an alarm sounds. Three other specialized robots automate the preparation of injectable medicines; one of these robots deals exclusively with highly toxic

chemotherapy drugs. The system virtually eliminates the possibility of human error by cutting humans almost entirely out of the loop.

UCSF's $7 million automated system is just one of the more spectacular examples of the robotic transformation that's unfolding in the pharmacy industry. Far less expensive robots, not much larger than a vending machine, are invading retail pharmacies located in drug and grocery stores. Pharmacists in the United States require extensive training (a four-year doctoral degree) and have to pass a challenging licensing exam. They are also well paid, earning about $117,000 on average in 2012. Yet, especially in retail settings, much of the work is fundamentally routine and repetitive, and the overriding concern is to avoid a potentially deadly mistake. In other words, much of what pharmacists do is almost ideally suited to automation.

Once a patient's medication is ready to leave a hospital pharmacy, it's increasingly likely that it will do so in the care of a delivery robot. Such machines already cruise the hallways in huge medical complexes delivering drugs, lab samples, patient meals, or fresh linens. The robots can navigate around obstacles and use elevators. In 2010, El Camino Hospital in Mountain View, California, leased nineteen delivery robots from Aethon, Inc., at an annual cost of about $350,000. According to one hospital administrator, paying people to do the same work would have cost over a million dollars per year.[15] In early 2013, General Electric announced plans to develop a mobile robot capable of locating, cleaning, sterilizing, and delivering the thousands of surgical tools used in operating rooms. The tools would be tagged with radio-frequency identification (RFID) locator chips, making it easy for the machine to find them.[16]

Beyond the specific areas of pharmacy and hospital logistics and delivery, autonomous robots have so far made relatively few inroads. Surgical robots are in widespread use, but they are designed to extend the capabilities of surgeons, and robotic surgery actually costs more than traditional methods. There is some preliminary work being done on building more ambitious surgical robots; for example, the I-Sur project is an EU-backed consortium of European researchers

who are attempting to automate basic procedures like puncturing, cutting, and suturing.[17] Still, for the foreseeable future, it seems inconceivable that any patient would be allowed to undergo an invasive procedure without a doctor being present and ready to intervene, so even if such technology materializes, any cost savings would likely be marginal at best.

Elder-Care Robots

The populations of all advanced countries, as well as many developing nations, are aging rapidly. The United States is projected to have over 70 million senior citizens, making up about 19 percent of the population, by 2030. That's up from just 12.4 percent in 2000.[18] In Japan, longevity combined with a low birth rate make the problem even more extreme; by 2025 fully a third of the population will be over sixty-five. The Japanese also have a nearly xenophobic aversion to the increased immigration that might help mitigate the problem. As a result, Japan already has at least 700,000 fewer elder-care workers than it needs—and the shortage is expected to become far more severe in the coming decades.[19]

This surging global demographic imbalance is creating one of the greatest opportunities in the field of robotics: the development of affordable machines that can assist in caring for the elderly. The 2012 movie *Robot & Frank,* a comedy that tells the story of an elderly man and his robotic caretaker, offers a very hopeful take on the kind of progress we're likely to see. The movie opens by announcing to the viewer that it is set in the "near future." The robot then proceeds to exhibit extraordinary dexterity, carry out intelligent conversations, and generally act just like a person. At one point, a glass is knocked off a table, and the robot snatches it out of midair. That, I'm afraid, is not a "near future" scenario.

Indeed, the main problem with elder-care robots as they exist today is that they really don't do a whole lot. Much of the initial progress has been with therapeutic pets like Paro, a robotic baby

seal that provides companionship (at a cost of up to $5,000). Other robots are able to lift and move elderly people, saving a great deal of wear and tear on human caretakers. However, such machines are expensive and heavy—they may weigh ten times as much as the person they are lifting—and will, therefore, probably be deployed primarily in nursing homes or hospitals. Building a low-cost robot with sufficient dexterity to assist with personal hygiene or using the bathroom remains an extraordinary challenge. Experimental machines capable of specific tasks have appeared. For example, researchers at Georgia Tech have built a robot with a soft touch that can give patients a gentle bed bath, but the realization of an affordable, multitasking elder-care robot that can autonomously assist people who are almost completely dependent on others probably remains far in the future.

One of the ramifications of that daunting technical hurdle is that, despite the theoretically huge market opportunity, there are relatively few start-up companies focused on designing elder-care robots and little venture capital flowing into the field. The best hope almost certainly comes from Japan, which is on the brink of a national crisis and which, unlike the United States, has little aversion to direct collaboration between industry and government. In 2013, the Japanese government initiated a program in which it will pay two-thirds of the costs associated with developing inexpensive, single-task robotic devices that can assist the elderly or their caretakers.[20]

Perhaps the most remarkable elder-care innovation developed in Japan so far is the Hybrid Assistive Limb (HAL)—a powered exoskeleton suit straight out of science fiction. Developed by Professor Yoshiyuki Sankai of the University of Tsukuba, the HAL suit is the result of twenty years of research and development. Sensors in the suit are able to detect and interpret signals from the brain. When the person wearing the battery-powered suit thinks about standing up or walking, powerful motors instantly spring into action, providing mechanical assistance. A version is also available for the upper body and could assist caretakers in lifting the elderly.

Wheelchair-bound seniors have been able to stand up and walk with the help of HAL. Sankai's company, Cyberdyne, has also designed a more robust version of the exoskeleton for use by workers cleaning up the Fukushima Daiichi nuclear plant in the wake of the 2011 disaster. The company says the suit will almost completely offset the burden of over 130 pounds of tungsten radiation shielding worn by workers.* HAL is the first elder-care robotic device to be certified by Japan's Ministry of Economy, Trade, and Industry. The suits lease for just under $2,000 per year and are already in use at over three hundred Japanese hospitals and nursing homes.[21]

Other near-term developments will probably include robotic walkers to assist in mobility and inexpensive robots capable of bringing medicine, providing a glass of water, or retrieving commonly misplaced items like eyeglasses. (This would likely be done by attaching RFID tags to the items.) Robots that can help track and monitor people with dementia are also appearing. Telepresence robots that allow doctors or caretakers to interact with patients remotely are already in use in some hospitals and care facilities. Devices of this type are relatively easy to develop because they skirt around the challenge of dexterity. The near-term nursing-care robotics story is primarily going to be about machines that assist, monitor, or enable communication. Affordable robots that can independently perform genuinely useful tasks will be slower to arrive.

Given that truly capable and autonomous elder-care robots are unlikely to emerge in the near future, it might seem reasonable to expect that the looming shortage of nursing home workers and home health aids will, to a significant extent, offset any technology-driven

---

* The names selected by Sankai seem a bit odd for a company focused primarily on elder care. HAL, of course, was the unfriendly computer that wouldn't open the pod bay doors in *2001: A Space Odyssey.* Cyberdyne was the fictional corporation that built Skynet in the *Terminator* movies. Perhaps the company is eying other markets.

job losses that occur in other sectors of the economy. Maybe employment will simply migrate to the health and elder care sector. The US Bureau of Labor Statistics (BLS) projects that by 2022, there will be 580,000 new jobs for personal-care aids and 527,000 for registered nurses (those are the two fastest-growing occupations in the United States), as well as 424,000 home heath aids and 312,000 nursing aids.[22] That adds up to about 1.8 million jobs.

This sounds like a big number. But now consider that the Economic Policy Institute estimates that, as of January 2014, the United States was still short 7.9 million jobs as a result of the Great Recession. That includes 1.3 million jobs that were lost during the downturn and hadn't yet been recovered as well as another 6.6 million jobs that were never created.[23] In other words, if those 1.8 million jobs all appeared today, they would fill only about a quarter of the hole.

Another factor, of course, is that these jobs are low-paying and not particularly suitable for a large fraction of the population. According to the BLS, home health aids and personal aids both provided a medium 2012 income of under $21,000 and require an education level of "less than high school." Large numbers of workers are likely to lack the temperament necessary to thrive in these jobs. If a worker hates his job stamping out widgets, that's one thing. If he despises his job caring for a dependent older person, that's a major problem.

Assuming the BLS's projections are correct and these jobs do materialize in large numbers, there is also the question of who will actually pay for these workers. Decades of stagnant wages, together with the transition from defined benefit pensions to often under-funded 401k plans, will leave a large fraction of Americans in relatively insecure retirement situations. By the time the majority of older people reach the point where they need personal, daily assistance, relatively few are likely to have the private means to hire home health aids, even if the wages for these jobs continue to be very low. As a result, these will probably be quasi-government jobs funded by programs like Medicare or Medicaid and will therefore be viewed as more of a problem than a solution.

Unleashing the Power of Data

As we saw in Chapter 4, the big data revolution offers the promise of new management insights and significantly improved efficiency. In fact, the increasing importance of all this data may be a powerful argument for consolidation in the health insurance sector, or alternatively creating some mechanism for sharing data among insurance companies, hospitals, and other providers. Access to more data could well mean more innovation. Just as Target, Inc., was able to predict pregnancy based on customer purchasing patterns, hospitals or insurance companies with access to large datasets will potentially discover correlations between specific factors that can be controlled and the likelihood of a positive patient outcome. The original AT&T was famous for sponsoring Bell Labs, where many of the twentieth century's most important advances in information technology took place. Perhaps one or more health insurance companies with sufficient scale could play a somewhat similar role—except that the innovations would come not from tinkering in a lab but from continuously analyzing reams of detailed patient and hospital operational data.

Medical sensors either implanted or attached to patients will provide another important source of data. These devices will produce a continuous stream of biometric information that can be used in both diagnosis and in the management of chronic diseases. One of the most promising areas of research is the design of sensors capable of monitoring glucose in people with diabetes. The sensors could communicate with a smart phone or other external device, instantly alerting patients if their glucose level falls outside the safe range and avoiding the need for uncomfortable blood tests. A number of companies already manufacture glucose monitors that can be embedded under a patient's skin. In January 2014, Google announced that it is working on a contact lens that would contain a tiny glucose detector and wireless chip. The lenses would continuously monitor glucose levels by analyzing tears; if the wearer's blood sugar is too high or too low, a tiny LED light would illuminate, providing an instant alert.

Consumer devices like the Apple Watch, formally announced in September 2014, will likewise result in a torrent of health-related data.

Health Care Costs and a Dysfunctional Market

The March 4, 2013, issue of *Time* magazine featured a cover story by Steven Brill entitled "Bitter Pill." The article delved into the forces underlying ever-escalating health care costs in the United States and highlighted case after case of what can only be categorized as outright price gouging—including, for example, a 10,000 percent markup on the same over-the-counter acetaminophen tablets you could buy at your local drug store or Walmart. Routine blood tests for which Medicare would pay about $14 were marked up to $200 and beyond. CT scans that Medicare prices at about $800 were inflated to over $6,500. A feared heart attack that turned out to be a case of heartburn resulted in a $17,000 charge—not including fees for the doctor.[24]

A few months later, Elisabeth Rosenthal of the *New York Times* wrote a series of articles telling essentially the same story: a laceration requiring three simple stitches came in at well over $2,000. A dab of skin glue on a toddler's forehead cost over $1,600. One patient was charged nearly $80 for a small bottle of local anesthetic that can be purchased for $5 on the Internet. Rosenthal noted that the hospital, which buys such supplies in bulk, would likely pay far less.[25]

Both reporters found that these inflated charges generally originate with a massive, obscure—and often secretive—list of prices known as the "chargemaster." The prices listed in the chargemaster seemingly have no rhyme or reason and no meaningful relationship to actual costs. The only thing one can say with consistent certainty about the chargemaster is that its prices are very, very high. Both Brill and Rosenthal found that the most egregious cases of chargemaster abuse occurred with uninsured patients. Hospitals typically expected these people to pay full list price and often were quick to hire bill

collectors or even file lawsuits if patients couldn't or wouldn't pay. Even major health insurance companies, however, are increasingly billed at rates based on a discount from chargemaster prices. In other words, the costs are first inflated—in many cases by a factor of ten or even a hundred—and then a discount of perhaps 30, or even 50, percent is applied, depending on how effectively the insurer negotiates. Imagine buying a gallon of milk for $20 after negotiating a 50 percent discount from the $40 list price. Given this, it should come as no surprise that hospital charges are the most important single driver of consistently soaring health care costs in the United States.

One of the most important lessons of history is that there is a powerful symbiosis between technological progress and a well-functioning market economy. Healthy markets create the incentives that lead to meaningful innovation and ever-increasing productivity, and this has been the driving force behind our prosperity.* Most intelligent people understand this (and are very likely to bring up Steve Jobs and the iPhone when discussing it). The problem is that health care is a broken market and no amount of technology is likely to bring down costs unless the structural problems in the industry are resolved.

There is also, I think, a great deal of confusion about the nature of the health care market and exactly where an effective market pricing mechanism should come into play. Many people would like to believe that health care is a normal consumer market: if only we could get insurance companies, and especially the government, out of the way and instead push decisions and costs onto the consumer (or patient), then we'd get innovations and outcomes similar to what we've seen in other industries (Steve Jobs might be mentioned again here).

---

* Consider, for example, the Soviet Union, which by all accounts had some of the best scientists and engineers in the world. The Soviets were able to achieve solid results in military and space technology, but they were never able to scale the benefits of innovation across the civilian economy. The reason certainly has a lot to do with the absence of working markets.

The reality, however, is that health care is simply not comparable to other markets for consumer products and services, and this has been well understood for over half a century. In 1963, the Nobel laureate economist Kenneth Arrow wrote a paper detailing the ways in which medical care stands apart from other goods and services. Among other things, Arrow's paper highlighted the fact that medical costs are extremely unpredictable and often very high, so that consumers can neither pay for them out of ongoing income nor effectively plan ahead as they might for other major purchases. Medical care can't be tested before you buy it; it's not like visiting the wireless store and trying out all the smart phones. In emergencies, of course, the patient may be unconscious or about to die. And, in any case, the whole business is so complex and requires so much specialized knowledge that a normal person can't reasonably be expected to make such decisions. Health care providers and patients simply don't come to the table as anything approaching equals, and as Arrow pointed out, "both parties are aware of this informational inequality, and their relation is colored by this knowledge."[26] The bottom line is that the high cost, unpredictability, and complexity of major medical and hospitalization services make some kind of insurance model essential for the health care industry.

It is also critical to understand that health care spending is highly concentrated among a tiny number of very sick people. A 2012 report by the National Institute for Health Care Management found that just 1 percent of the population—the very sickest people—accounted for over 20 percent of total national health care spending. Nearly *half* of all spending, about $623 billion in 2009, went to the sickest 5 percent of the population.[27] In fact, health care spending is subject to the same kind of inequality as income in the United States. If you draw a graph, it will look very much like the winner-take-all/long-tail distribution I described in Chapter 3.

The importance of this intense concentration of spending cannot be overemphasized. The small population of very ill people on

whom we are spending all this money are obviously not in a position to negotiate prices with providers; nor would we want to place such a staggering fiscal responsibility in these people's hands. The "market" that we need to make work exists between the providers and the insurance companies—not between providers and patients. The essential lesson of the articles written by Brill and Rosenthal is that this market is dysfunctional because of a fundamental power imbalance between insurers and providers. While individual consumers may rightly perceive health insurance companies as powerful and domineering, the reality is that—relative to providers like hospitals, doctors, and the pharmaceutical industry—they are, in a great many cases, *too weak*. That imbalance is being steadily worsened by an ongoing wave of consolidations among providers. Brill's article notes that as hospitals increasingly snap up "doctor's practices and competing hospitals, their leverage over insurance companies is increasing."[28]

Imagine a near future where a physician wields a powerful tablet computer that allows her to order a range of medical tests and scans with just a few presses on her touch screen. Once a test is completed, the results are instantly routed to her device. If a patient needs a CT scan, or perhaps an MRI, the results are accompanied by a detailed analysis performed by an artificial intelligence application. The software points out any anomalies in the scan and makes recommendations for further care by accessing a massive database of patient records and identifying similar cases. The doctor can see exactly how comparable patients were treated, any issues that arose, and how things ultimately turned out. All this would, of course, be efficient and convenient and ought to lead to a better outcome for the patient. This is the kind of scenario that gets techno-optimists excited about the revolution soon to unfold in the health care arena.

Now assume that the doctor has a financial interest in the diagnostic company that performs the tests or scans. Or, then again, maybe the hospital has acquired the doctor's practice and also owns the testing facility. The prices for the tests and scans bear little

relation to the actual costs of these services—after all, they're listed in the chargemaster—and they are highly profitable. Every time our doctor presses her touch screen, she essentially mints money.

While this example is, at the moment, imaginary, there is an abundance of evidence demonstrating that new health care technologies very often lead to more spending rather than improved productivity. The primary reason is that there is no effective market pricing mechanism to drive increased efficiency. In the absence of market pressure, providers often invest in technologies designed to increase revenue rather than efficiency, or where they do achieve increased productivity they simply retain the profits rather than lowering prices.

The poster child for technology investment as a driver of health care inflation may well be the "proton beam" facilities that are being built to treat prostate cancer. A May 2013 article by Jenny Gold of *Kaiser Health News* noted that "despite efforts to get health care spending under control, hospitals are still racing to build expensive new technology—even when the devices don't necessarily work better than the cheaper kind."[29] The article describes one proton beam facility as "a giant cement-encased building the size of a football field, with a price tag of more than $200 million." The idea behind this expensive new technology is that it delivers less radiation to patients, and yet, studies have found no evidence that protein beam technology results in better patient outcomes than far less expensive approaches.[30] Health care expert Dr. Ezekiel Emanuel says, "We don't have evidence that there's a need for them in terms of medical care. They're simply done to generate profits."[31]

To me, it seems evident that the American people could in principle be made much better off by a massive technological disruption of the health care sector than of, say, the fast food industry. After all, lower prices and improved productivity in health care will likely lead directly to better and longer lives. Cheaper fast food may well do the opposite. Yet, the fast food industry has well-functioning

markets—and the health care sector does not. As long as that situation is allowed to persist, there are few reasons to be optimistic that accelerating technology alone will succeed in reining in soaring health care costs. Given this reality, I'd like to take a brief detour from our technology narrative in order to suggest two alternate strategies that might help to correct the power imbalance between insurers and providers, and hopefully enable the kind of synergy between markets and technology that might bring the transformation we hope for.

Consolidate the Industry and Treat Health Insurance as a Utility

One of the primary messages that leaps out from an analysis of the prices charged by providers is that Medicare—the government-run program for people aged sixty-five and over—is by far the most efficient portion of our health care system. As Brill writes, "Unless you are protected by Medicare, the health care market is not a market at all. It's a crapshoot." The implementation of the Affordable Care Act (Obamacare) will certainly improve the situation as far as individuals who previously lacked insurance are concerned, but it does relatively little to actively rein in hospital costs; instead, the inflated costs will be shifted to insurers and then ultimately to taxpayers in the form of the subsidies that were put in place to make health insurance affordable to people with moderate incomes.

The fact that Medicare is relatively effective at controlling most patient-related costs, while spending far less than private insurers on administration and overhead, underlies the argument for simply expanding the program to include everyone and, in effect, creating a single-payer system. This has been the path followed by a number of other advanced countries—all of which spend far less on health care than the United States and typically have better outcomes according to metrics like life expectancy and infant mortality. While a single-payer system, managed by the government, has both logic

and evidence to support it, there is no escaping the reality that in the United States the whole idea is ideologically toxic to roughly half the population. Putting such a system in place would also presumably result in the demise of nearly the entire private health insurance sector; that does not seem likely given the enormous political influence wielded by the industry.

A single-payer system is, in practice, always assumed to be run by the government, but in theory this does not have to be the case. Another approach might be to merge all private insurance companies into a single national corporation, which would then be heavily regulated. The model would be the original AT&T before it was broken up in the 1980s. The central idea here is that health care is in many ways akin to the telecommunications system: it is, in essence, a utility. Like water and sanitation systems or the nation's electrical infrastructure, the health care system does not stand alone—it is a systemic industry whose efficient operation is critical to both the economy and society. In many cases, the provision of a utility service leads to natural monopoly scenarios. In other words, it is most efficient if only a single firm operates in the market.

An even more effective variation on this theme might be to allow a small number of large competing insurance companies—in effect, a sanctioned oligopoly. This would inject an element of competition into the system. The companies would still be large enough to have significant market power when negotiating with providers, and they would have little choice but to compete on the basis of enabling high-quality care since their reputations would determine their success. Tight regulation of the industry would limit price increases and prevent the companies from engaging in undesirable practices like, for example, designing insurance plans geared specifically toward "cherry-picking" younger, healthier patients or offering plans with substandard protection. Instead, they would have to focus on genuine innovation and efficiency.

Consolidating existing insurance companies into one or more regulated "health care utilities" might provide many of the advantages

of a single-payer system while preserving the industry. Rather than being wiped out, the shareholders of private insurance companies might conceivably see gains as a result of an industry-wide merger. The mechanism by which such a consolidation might be brought about is, of course, far from obvious. Perhaps the government could issue a small number of operating licenses, and it might even hold an auction as it does for the electromagnetic communications spectrum.*

Set "All-Payer" Rates

An alternate, and perhaps more feasible, strategy is the implementation of an "all-payer" system. In this scenario, the government essentially sets the schedule of prices that can be charged by health care providers. Just as Medicare dictates the prices it will pay, an all-payer system would do the same for all patients receiving care from any given provider. An all-payer approach is used in the health care systems of a number of countries, including France, Germany, and Switzerland. In the United States, Maryland also has such a system for hospitals, and the state has seen relatively slow growth in hospitalization costs.[32] All-payer systems vary in the specifics of their implementation; the rates may be set through collective negotiation between providers and payers, or they might be established by a regulating commission after an analysis of actual costs at particular hospitals.

---

* In the United States, the constitutional authority to create a single-payer system—regardless of whether it is run by the government or by private corporations—probably derives from the government's ability to levy a tax on everyone to pay for the system. Therefore, all or a portion of the premiums would be paid by the government. This is already the case with the insurance subsidies associated with the Affordable Care Act. In other words, the federal government can force everyone to pay for a single-payer system through taxes, but it cannot prohibit a parallel private system. So there still would likely be additional services available to those willing and able to pay out of pocket, just as there are private schools. This is different from the system in Canada, where most private health care services are prohibited—leading some Canadians to seek health care services in the United States.

Since an all-payer system enforces the same prices for all patients, it has important implications for the cost shifting that goes on between private patients and those covered by the public systems in the United States (Medicaid for low-income people and Medicare for those over sixty-five). When a single rate is set, the public prices have to rise considerably, putting more of a burden on taxpayers. Privately insured patients, and especially those who are uninsured, will typically benefit from lower prices as they are no longer subsidizing the public programs. This has been the case with Maryland's program.*

It seems to me that a much simpler approach that might produce immediate savings would be to set an all-payer *ceiling* rather than a specific price. For instance, suppose the ceiling were set at the Medicare rate plus 50 percent. In one example from Brill's article, a blood test that Medicare says is worth $14 might then be priced at any amount up to $21—but it could never reach anything like $200. Insurance companies with sufficient market power would still be free to negotiate a price lower than the ceiling. This strategy would immediately eliminate the worst excesses, and as long as the ceiling was set high enough, it would still provide sufficient revenue to providers. A 2010 fact sheet published by the American Hospital Association claims that Medicare paid "90 cents for every dollar spent by hospitals caring for Medicare patients in 2009."[33] If the industry's own lobbying organization says Medicare is covering 90 percent of hospital costs, then a ceiling somewhat higher than the Medicare rate should be sufficient to allow enough cost shifting to make up for that missing 10 percent.** An all-payer ceiling would also be very

* Maryland has a special waiver that has been in place for over thirty years and allows it to pay higher Medicare rates. As of 2014, Maryland has moved to a new experimental system that is allowed under the Affordable Care Act. In addition to setting all-payer rates, the new program will enforce explicit caps on per capita hospital spending. The state expects to save $330 million in Medicare costs over a five-year period.

** The same fact sheet says that Medicaid (the program for the poor) paid 89 percent of actual hospital costs.

easy to implement since it is based directly on the already-published Medicare rates.

One of the most hopeful approaches to controlling health costs, which is gaining some traction in the current environment, is to transition away from a fee-for-service model and toward an "accountable care" system in which doctors and hospitals are paid a set fee to manage the overall health of patients. One of the primary advantages of this approach is that it would reorient the incentives regarding innovation. Rather than simply offering a new way to hoover up even higher fees according to a fixed schedule, emerging technologies would be viewed in terms of their potential to reduce costs and make care more efficient. The key to making that happen, however, is to push more of the financial risk associated with patient care away from insurers (or the government) and onto hospitals, doctors, and other providers. Needless to say, the latter are unlikely to accept that increased risk willingly. In other words, in order to drive a successful transition toward accountable care, we still need to address the market power imbalance that often exists between insurers and providers.

In order to bring relentlessly increasing health care costs in the United States under control, I think it will probably be necessary to pursue one of the two general strategies I've outlined. We will have to move toward a single-payer system where either the government or one or more large private firms exercise more bargaining power in the health insurance market, or alternatively we will need to have regulators exercise direct control over the rates paid to providers. In either scenario, moving aggressively toward an accountable care model might be a vital part of the solution. Both of these approaches, in various combinations, are used successfully by other advanced countries. The bottom line is that a pure "free market" approach in which we cut government out of the loop and expect patients to operate like consumers shopping for groceries or smart phones is never going to work. As Kenneth Arrow pointed out over fifty years ago, health care is simply different.

This is not to say that there are no significant dangers associated with either approach. Both strategies rely on regulators to either control premiums or set the prices paid to providers. There is an obvious risk of regulatory capture; powerful companies or industries may exert influence that bends government policy in their favor. Attempts at such influence have already been successfully directed at Medicare, which is specifically prohibited from using its market power to negotiate drug prices. The United States is virtually the only country in the world where this is the case; every other national government negotiates prices with the drug companies. The result is that Americans, in effect, subsidize lower drug prices in the rest of the world. The three years between 2006 and 2009 saw a 68 percent increase in the rate of "prescription abandonment" in the United States.[34] This happens when patients request that a prescription be filled, but then walk away when they find out the cost. It's something of a mystery to me why this is not more disturbing to Americans, and to grassroots conservatives in particular. The Tea Party, after all, got started after a famous rant by CNBC personality Rick Santelli, who decried the fact that people with mortgages they couldn't afford might be subsidized by taxpayers. Why aren't average Americans more upset about the fact that they are paying the pharmaceutical freight for the rest of the world—including a number of countries that have significantly higher per capita incomes than the United States?

In spite of this problem, Medicare consistently provides high-quality care at a cost significantly lower than in the highly fragmented private insurance sector. In other words, we should not make the perfect the enemy of the good. Nonetheless, Medicare's prohibition against negotiating with the pharmaceutical industry deserves to be subjected to a great deal more public scrutiny. The industry argues that inflated drug prices in the United States are necessary in order to fund further research. However, there are likely more efficient and certainly more equitable ways to ensure that drug research gets

funded.[35*] The potential to reform or streamline the Federal Drug Administration's procedures for testing and approving new drugs also surely exists.

Another issue with Medicare, and one that touches directly on the subject of this book, is that waste can easily be driven by the direct advertisement of products to senior citizens who are told explicitly to pressure their physicians for a prescription and that Medicare will then pick up nearly the entire cost. One government audit found that up to 80 percent of the motorized scooters paid for by Medicare were not really needed by the elderly patients who received them and may actually be harmful to their health. The two largest scooter manufacturers spent over $180 million on advertisements directed at Medicare recipients in 2011.[36] This is another issue that deserves close scrutiny because, as we've seen, there is soon likely to be a profusion of robotic equipment geared toward providing home-based assistance to senior citizens. Such advances have great potential to improve quality of life for the elderly while reducing the cost of their care—but not if we pay for technology in cases where it is unneeded or perhaps even detrimental. The specter of millions of comfortably seated senior citizens watching advertisements telling them that Medicare will happily pay for a robot capable of retrieving their television remote should give us pause.[**]

---

* A related issue has to do with the patents granted to drug manufacturers. These prevent the introduction of cheaper, generic drugs for long periods. Many economists believe that the pharmaceutical patent system is very inefficient. Other countries can also potentially threaten to void drug patents as a price negotiating mechanism—putting a still higher burden on Americans. The Center for Economic and Policy Research published a briefing in 2004 that outlines these issues and presents some more efficient alternatives for funding drug research. Please see the corresponding endnote for details.

** The whole idea behind requiring prescriptions is that patients are not able (or cannot be trusted) to make these decisions for themselves. Why, then, do we allow drug companies or medical equipment manufacturers to advertise directly to patients?

While recent applications of AI and robotics to the health care field are impressive and advancing rapidly, they are, for the most part, just beginning to nibble at the edges of the hospital cost problem. With the exception of pharmacists, and possibly doctors or technicians who specialize in analyzing images or lab specimens, automating even a significant portion of the jobs done by most skilled health care workers remains a daunting challenge. For those seeking a career that is likely to be relatively safe from automation, a skilled health care profession that requires direct interaction with patients remains an excellent bet. That calculus could, of course, change in the more distant future. Twenty or thirty years from now, I think, it's impossible to say with any real confidence what might be technologically possible.

Technology is not the only consideration, of course. Health care, more than any other sector of the economy, is subject to a complex web of rules and regulations imposed by governments, agencies like the FDA, and licensing authorities. Every action and every decision are also colored by the looming threat of litigation if an error—or perhaps just an unlucky outcome—should occur. Even among retail pharmacists, the specific impact of automation on employment isn't easily discernible. The reason is likely regulation. Farhad Manjoo interviewed one pharmacist who said, "Most pharmacists are employed only because the law says that there has to be a pharmacist present to dispense drugs."[37] That, at least for the moment, is probably something of an exaggeration. Job prospects for newly minted pharmacists have worsened significantly over the past decade, and things may well get worse. A 2012 analysis identifies a "looming joblessness crisis for new pharmacy graduates" and suggests that the unemployment rate could reach 20 percent.[38] However, this is likely due largely to an explosion in the number of new graduates entering the job market as pharmacy schools have dramatically increased enrollments.* Relative

* One could also speculate that technology is *indirectly* contributing to diminished prospects for pharmacy graduates by driving more people into the profession. In the first decade of the new millennium, nearly fifty new pharmacy graduate schools opened their doors (a 60 percent increase), and existing programs also dramatically increased enrollments. The number of newly

to most other occupations, there's little doubt that health care professionals enjoy an extraordinary degree of employment security as a result of factors completely unrelated to the technical challenges associated with automating their jobs.

This may be good news for health care workers, but if technology has only a muted impact on health care costs even as it disrupts other employment sectors, the economic risks we face will be amplified. In that scenario, the burden of soaring health care costs will become even more unsustainable as advancing technology continues to produce unemployment and ever-increasing inequality, as well as stagnant, or even falling, incomes for most workers in other industries. This prospect makes it even more critical to introduce meaningful reforms that will correct the market power imbalance between insurers and providers so that advancing technology can be fully leveraged as a mechanism for increased efficiency across the health care sector. Without that, we run the risk that our market economy will eventually come to be dominated by a sector that is inefficient and, indeed, not an especially well-functioning market at all.

Controlling the health care cost burden is especially critical because, as we'll see in Chapter 8, the last thing American households need is an ever-increasing drain on their discretionary income. Indeed, stagnant incomes and growing inequality are already undermining the broad-based consumer demand that is vital to continued economic growth.

graduated pharmacists could hit 15,000 per year by 2016; that's over twice the number of degrees granted in 2000. Something very similar (and perhaps even more extreme) happened with law schools, and the law school enrollment bubble is now famously bursting. Law school has always been a well-traveled path toward monetizing a liberal arts degree. Pharmacy offers similar potential for an undergraduate biology degree. It may be that soaring demand for these professional degrees results, at least in part, from the evaporation of other good opportunities for college graduates. With relatively few other attractive alternatives, college graduates have clamored to get into law or pharmacy school, and the industry has responded by expanding enrollment and ultimately producing far more graduates than the market could absorb. The fact that both pharmacy and law are also impacted by direct automation makes things even more unsustainable. My prediction for the next professional school bubble: MBA degrees.

So far, we have focused primarily on the ways in which technology is likely to transform existing employment sectors. In the next chapter, we'll leap a decade or more ahead in time and imagine how things might look in a future economy populated with entirely new technologies and industries.