The Future of the Professions[1]

DANIEL SUSSKIND
Fellow in Economics
Balliol College, Oxford University

RICHARD SUSSKIND
Visiting Professor
Oxford Internet Institute

Two Futures

There are two possible futures for the professions. Both of these rest on technology. The first is reassuringly familiar to most professionals—it is simply a more efficient version of what we have today. In this future, professionals of many different types use technology, but largely to streamline and optimize their traditional ways of working. In the language of economists, technologies "complement" them in these activities. The second future is a different proposition. Here, increasingly capable systems and machines, either operating alone or designed and operated by people who look quite unlike doctors and lawyers, teachers and accountants, and others, gradually take on more of the tasks that we associate with those traditional professionals. New technologies instead, in the words of economists, "substitute" for professionals in these activities.

For now, and in the medium term, we anticipate that these two futures will be realized in parallel. As we do today, we will continue to see examples of both uses of technology. In the long run, however, we expect that the second future will dominate. Through technological progress, we will find new and more efficient ways to solve the sorts of important problems that, traditionally, only very particular types of professionals have been able to tackle. This presents an existential challenge to traditional professionals, which is one central theme of our book *The Future of the Professions*.

1 Read 27 April 2017. This paper is based on Richard Susskind and Daniel Susskind, *The Future of the Professions* (Oxford: Oxford University Press, 2015). All references in this paper are taken from that work, unless cited otherwise.

Why Do We Have the Professions?

We begin *The Future of the Professions* by asking a fundamental question: *Why* do we have the professions at all? Various theorists have tried to make sense of the professions and their dominance in many walks of life. Some of these are "functionalists," concerned with the different roles that the professions perform—for instance, correcting imbalances of knowledge, strengthening the moral character of society, or maintaining social order. Others are "traitists," more interested in the particular features of the professions than the functions they perform. Like zoologists, they try to identify and classify different species of occupations, drawing up exhaustive checklists of important features and organizing their specimens in careful taxonomies. Most of these theorists, from a variety of perspectives, are fascinated with the "exclusivity" of the professions, their ability to ring-fence, isolate, and effectively exclude others from large expanses of knowledge. There are sometimes traces of a more conspiratorial mind-set at work here. George Bernard Shaw's line, that professions are "conspiracies against the laity," is the rallying call of the deeply suspicious.

Our view of the professions draws on all these ways of thinking. Although the professions from a distance can look quite different, we argue that, in analogous ways, they are all solutions to the same underlying problem—that no one can know everything. Human beings have what Herbert Hart, the legal philosopher, called "limited understanding" of the world around them. Most people, acting alone, are unable to resolve all the important challenges that they might face in life. And so they turn to professionals because they have the *practical expertise*—our term for their knowledge, wisdom, experience, skills, and know-how—that is required to solve these important challenges. In what we call a Print-based Industrial Society, we built the professions to help us create, manage, and apply these great bodies of practical expertise. We (citizens and the state) established a *grand bargain* with the professions, an arrangement that entitles them, often to the exclusion of others, to provide certain types of services, though the nature of the bargain can vary significantly across professions and jurisdictions. In return, the professions are entrusted to act as *gatekeepers*, each responsible for their own unique body of practical expertise—doctors look after medical practical expertise, lawyers after legal practical expertise, and so on across the professions.

Yet we are no longer in a Print-based Industrial Society. We are now in what we call a Technology-based Internet Society. And those traditional professions, such as health and education, working under the grand bargain, are creaking. They are unaffordable, in that most people and institutions do not have access to first-rate professionals—or indeed any professionals. They are antiquated, by and large relying on old-fashioned ways of producing and sharing practical expertise, despite the existence of feasible alternatives. They are opaque, sometimes because the work they do is genuinely too complicated for a layperson to understand. A visitor, for example, to a British Court—witnessing the oak paneling and wigs, and the arcane language and procedure—would be forgiven for thinking there was also some intentional obfuscation at work in the professions. Finally, the professions underperform. Given the way we organize practice expertise in society using the professions as gatekeepers, the finest practical expertise is a very scarce resource. Only a very privileged and lucky few recipients have access to it.

And so, we ask the following question—as we move from a Print-based Society to an Internet Society, might there be new ways of organizing professional work? Might there be new ways to solve the sorts of problems that, in the past, the professions alone have solved? Do we still need the traditional gatekeepers?

The Change

In our work, we draw on hundreds of case studies. The following selection provides a sense of the transformation that is already taking place in the professions.

In education, for instance, more people signed up for Harvard's online courses in a single year than had attended the actual university in its entire existence up until that point. In tax accounting, almost 48 million Americans now use online tax preparation systems to help them file their tax returns, rather than a traditional tax adviser. In medicine, a team of Stanford researchers developed a system that can diagnose whether or not a lesion is cancerous, from a photo, as accurately as leading dermatologists.[2] In journalism, the Associated Press uses algorithms to write earnings reports and sports results, producing 15 times the number of the former than it did when it relied upon traditional financial journalists alone. In law, J. P. Morgan uses a program called COIN to scan commercial loan agreements, doing in a matter of seconds what would have taken, it is claimed, about 360,000 hours of

2 Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan Swetter, Helen Blau, and Sebastian Thrun, "Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks," *Nature* 542 (2017): 155–88.

a traditional lawyer's time.[3] In architecture, the Elbphilharmonie, the new concert hall in Hamburg, contains a strikingly beautiful auditorium composed of 10,000 interlocking acoustic panels, yet this space was designed by an algorithm, not inspired by the refined aesthetic sensibility of a human being.[4] And in the world of divinity, in 2011, the Catholic Church issued the first *digital* imprimatur—the official license granted by the Church to religious texts—to an app called Confession, to help people prepare for confession. (It has tools for tracking sin, and drop-down panels of options for contrition.)

From these examples, and many other case studies like them, we were able to identify eight high-level patterns, and 30 more granular trends to describe the changes that are taking place in the professions. Again, to get a broad sense of these, consider the following five trends.

First, the *more-for-less challenge*. Across the professions, institutions and individuals are being asked to deliver more service, with fewer resources at their disposal. The second trend is the existence of *new competition*. What is notable about many of these cases of change are that they are being driven by people and institutions outside the boundaries of the traditional professions (often tech start-ups), with very different training and experience to traditional professionals. The third is a *move away from bespoke service*. Many professionals think of their work as a form of craft, like an artist starting each project afresh with a blank sheet of paper, or akin to a tailor stitching a suit to fit the particular bodily contours of their clients. We are seeing a move away from that view, recognizing that professional work does not have to be handled in this bespoke way. The fourth is the increasing *decomposition* of professional work. Many professionals think of their work as monolithic, indivisible lumps of endeavor that must all be handled by particular types of professionals, working in certain ways, organized in specific forms of institution. Increasingly, though, we are instead seeing professional work being broken down into composite tasks and activities. Once this is done, it often becomes clear that this work can either be performed by nonprofessionals, or can be automated. And this relates, finally, to the fifth trend, the increasing *routinization* of professional work. When professional work is broken down in this way, it transpires that many of the tasks involved in professional work are not particular complex, they are relatively "routine," and can be automated accordingly.

---

3 Debra Cassens Weiss, "JP Morgan Chase Uses Tech to Save 360,000 Hours of Annual Work by Lawyers and Loan Officers," *ABA Journal,* March 2, 2017.

4 Elizabeth Stinson, "What Happens When Algorithms Design a Concert Hall? The Stunning Elbphilharmonie," *Wired*, December 1, 2017.

Our View of Technology

These developments are driven by technological advances. In our work, we have developed a particular way of thinking about what is taking place in technology, describing these advances in four dimensions—the exponential growth in the underlying technologies, the increasing capability of these systems and machines, their pervasiveness in economic and social life, and our increasing connectedness as human beings. We say a little about each below.

Exponential Growth

The first is exponential growth in the underlying technologies. The law to note here is not a law of the land, but Moore's Law. This story is now a familiar one. In 1965, Gordon Moore, three years before he co-founded Intel, made the prediction that roughly every two years, engineers would be able to double the number of transistors they could fit on a silicon chip. In practice, this would mean that processing power would double every two years. Since then, Moore's Law has broadly held. If this growth in processing power were to continue, it would have profound effects. It means that by 2020, the average desktop computer will have the processing power of the human brain. More remarkable yet, it means that by 2050, if the trend continues, the average desktop computer will have the processing power . . . of all of humanity combined. Skeptics have noted that Moore's Law may now be slowing. Yet even if its pace were to half, with doubling in processing power taking place every four years, that simply means it would be our grandchildren, rather than our children, whose personal machines would have more power than all of humanity combined.

Increasing Capability

The second feature of technological advance is the increasing capability of these systems and machines. They are not simply more powerful in computational terms, as noted before, but we also can use them to perform a wider range of tasks and activities than was possible in the past. In other work, we have called this phenomenon "task encroachment."[5] And this increasing capability manifests itself in four ways.

The first is what is sometimes called "Big Data." As a consequence of our lives becoming increasingly digitized, data can be and now are collected and stored that are derived from our online decisions and behavior. This so-called "data exhaust," which trails behind us in

5 Daniel Susskind, "Rethinking the Capabilities of Technology in Economics," Oxford University Working Paper No. 825, May 2017.

everyday life, can yield patterns, insights, and correlations that human beings acting alone could not perceive. In the legal profession, for instance, there is now a system that can predict the outcome of U.S. Supreme Court decisions as accurately as leading legal scholars. It "knows" or "understands" nothing about the law. Instead, it makes a prediction based on 200 years of case data, each one described by up to 240 variables (the nature of the case, the justices involved, and so on).[6]

The second is the ability of the machines to solve problems and answer questions. The canonical example of this is Watson, a computer system developed by IBM that appeared on the U.S. quiz show *Jeopardy!* in 2011 and beat the two best-ever human champions. This was a system that could answer questions, in a particular format, on anything under the sun, more accurately than the leading human experts. It is noticeable that today's largest technology companies are trying to develop similar problem-solving and question-answering systems—Siri at Apple, Go Google at Google, Alexa at Amazon, Cortana at Microsoft. Each time a question is put to these systems they are designed to provide an answer.

The third dimension of machines' increasing capability originates in the field of "affective computing." This is a largely neglected but critically important field, dedicated to designing systems that can detect and respond to human emotions. There are now systems that, it is said, can distinguish between a smile of genuine joy and one of social conformity more accurately than a human being. Likewise, there are systems that can distinguish between a face showing fake pain and genuine pain.[7] For professionals, who imagine that their interpersonal skills are a distinctive part of what they do, these achievements raise a variety of challenging questions.

Finally, there is the field of robotics. Advances in these capabilities are of immediate relevance to professionals where manual dexterity is thought to be important. For surgeons, for instance, advances in robotic surgery have transformed the way in which many procedures are now performed. The most compelling case of robotic advance, though, is the driverless car—and while this may not be of direct relevance to most professionals, it is of indirect relevance as a caution against underestimating the future capabilities of machines. Until recently, many leading economists who study the labor market thought that the task of driving a car could not be readily automated. Yet, today, it clearly can be. All

6 Daniel Marin Katz, Michael J. Bommarito II, and Josh Blackman, "A General Approach for Predicting the Behaviour of the Supreme Court of the United States," *PLOS ONE*, April 12, 2017.

7 Raffi Khatchadourian, "We Know How You Feel," *New Yorker,* January 19, 2015.

major car manufacturers have driverless programs, and commercial vehicles are anticipated in the next few years.[8]

Increasing Pervasiveness

The third feature of technological advance is the increasing pervasiveness of these new technologies. It is not simply that more people have smartphones in their pockets or tablets in their briefcases (the installed base of smartphones is about 2.8 billion).[9] We are referring here largely to the Internet of Things—where we embed processors, sensors, and Internet connectivity into everyday objects. There are, for instance, alarm clocks linked to train timetables that can let their owners sleep longer if there are delays, and umbrellas that check weather forecasts and light up at the front door to advise their owners when rain is expected. It is estimated that, by 2020, there will be 40 to 50 billion devices connected to the Internet in this way.

Increasingly Connected

The final feature of technological advance is that human beings are becoming increasingly connected to one another, too. Today, about 3.4 billion people are connected to one network—the Internet.[10] And this connectivity expresses itself in many different dimensions. People can *communicate* in new ways. About 28 emails per human being are sent each day. We can also *research* in new ways—libraries and encyclopedias, the dominant information sources of professionals in the past, are largely been superseded by engines like Google. Humans can *socialize* in new ways—more than a quarter of the world's population is on Facebook. We can *share* in new ways—every minute, 300 hours of video are uploaded onto YouTube. We can *cooperate* in new ways—Wikipedia is written collaboratively by 69,000 main contributors, with over 35 million articles in more than 280 languages. We can *compete* in new ways—on Kaggle, data is supplied to a network of statisticians who vie with one another to provide the best analysis. And we can trade in new ways. In 1919, John Maynard Keynes wrote how remarkable it was that "[t]he inhabitant of London could order by telephone, sipping his morning tea in bed, the various products of the whole earth, in such quantity as he might see fit, and reasonably expect their early

8 Susskind, "Rethinking the Capabilities."

9 Mary Meeker, "Internet Trends 2017–Code Conference," *Kleiner Perkins Caufield Byers,* May 31, 2017.

10 Ibid.

delivery upon his doorstep."[11] At that time, "early delivery" meant a few weeks; today, it can mean as little as an hour or less.

Artificial Intelligence

Alongside these general trends, one technology has captured the popular imagination in the last few years—artificial intelligence (AI). We have a particular way of understanding what has happened in the field of AI and why it is significant for the professions in particular.

Our account begins in what we call the *first wave of AI*, a period of research that began in the late 1950s and continued through the 1980s. In the eighties, we were at the heart of the research community working on AI and the law. Our project at the University of Oxford was focused on "expert systems"—a particular type of AI application that could solve legal problems and offer legal advice. In 1988, we moved out of the research lab and into the marketplace, and co-developed the first commercially available AI system in the law. This was in response to a piece of legislation that was passed in the United Kingdom in 1986, known as the Latent Damage Act. The leading expert in the world on this legislation was Phillip Capper, who also happened to be the chair of the law school at Oxford at the time. And together, we worked to build a system, based on Phillip's expertise, for non-experts to use. This system, known as the Latent Damage System, told a user whether this particular piece of legislation applied to them and advised when a possible legal action could no longer be raised because too much time had elapsed. Users no longer had to consult a human expert like Phillip to get an answer. However, soon after that particular development, general progress in AI began to stall. A period known as the "AI winter" began.

A turning point came in 1997 when Deep Blue, another system developed by IBM, beat Garry Kasparov at chess. In the 1980s, when we were working on AI, we wrongly imagined this sort of victory was impossible. The reason for our mistake is important. At the time, we thought the only way to build a system to outperform human experts at given tasks was to sit down with experts (in our case, with Phillip Capper), get them to explain how they solved particular problems (in our case, a legal problem), and then build a system based on that explanation for others to use. These systems contained explicit representations of the knowledge of their "domain expertise." This was the approach of developers of the first wave of AI. But here was the problem—if you sat down with a chess player like Kasparov and asked

11 John Maynard Keynes, *The Economic Consequences of the Peace* (New York: Harcourt, Brace, and Howe, Inc., 1920).

him to explain how he played the game so extraordinarily well, he would struggle. He would say you could not reduce his expertise to a set of instructions for a machine to follow. He might be able to give us a few tips or tricks, but he would ultimately appeal to faculties like "gut reaction" and "intuition," honed through "experience." And so, we thought, if human beings cannot articulate how they perform given tasks, then these tasks cannot be automated—because it was not clear where to begin in composing a set of instructions for a machine to follow.

In retrospect, what we had not anticipated in the 1980s was the exponential growth in processing power that was to follow. By the time Garry Kasparov played against Deep Blue, the supercomputer could consider up to 330 million possible moves a second—whereas a great human player, with a following wind, can manage to explore about 100 moves in any one turn.[12] Deep Blue was not trying to replicate the reasoning or thinking processes that Kasparov followed. It was playing the game in a different way. This observation led us to one of the central ideas in our work, the *AI Fallacy*—"the mistaken supposition that the only way to develop systems that perform tasks at the level of experts or higher is to replicate the thinking processes of human specialists."[13] Advances in processing power, data storage capability, and algorithm design mean that there are now ways to build machines that perform tasks as well as the best humans, but which do not rely on copying the reasoning of human beings. In the case of Deep Blue, it no longer mattered that Kasparov could not explain how he played—the machine was able to win in a different way. As Patrick Winston, one of the earliest AI researchers, put it to us in private correspondence, "there are lots of ways of being smart that aren't smart like us."

Many professionals, when considering whether their work is at risk of automation, commit the AI Fallacy. They believe, mistakenly, that the only way to carry out a task that they perform is to copy them—and given the complexity of their thinking and reasoning, they imagine that most of their tasks cannot be automated. As an example, take a faculty like judgment. Professionals will often argue that exercising judgment is a critical part of their role. And while they may concede that the technological advances taking place are remarkable, they argue that a faculty like judgment is out of the reach of even the most capable machines. Like Kasparov, they would struggle to explain what

12 Murray Campbell, A. Joseph Hoane Jr., and Feng-hsiung Hsu, "Deep Blue," *Artificial Intelligence* 134 (2002): 57–83; Hubert Dreyfus, *What Computers Can't Do* (New York: Harper Colophon, 1979), 102.

13 Susskind and Susskind, *Future of the Professions*, 45.

judgment actually involves in practice. They would again likely resort to explanations in terms of their ineffable "intuition" and "experience."

Yet in light of the AI Fallacy, the question "How could a machine ever exercise judgment?"—the one that tends to be raised by skeptical professionals—is the wrong question to be asking. There are, in fact, two more important questions to be addressed. The first is this: "To what problem is judgment the solution?" Why do clients and patients, students and recipients, approach their professionals and ask them to exercise their "judgment"? And the common answer to that question, our research suggests, is "because of uncertainty." When people's circumstances are unclear; when information is ambiguous; when they lack the knowledge to classify and handle their situations; when humans, to use Herbert Hart's terminology again, have a "limited understanding" of what to do, they turn to their professionals and ask them for help—"in your experience, in your judgment, what should I do?" They want someone knowledgeable to help them make sense of their uncertainty. So, the second, more important question, is not, "How could a machine ever exercise judgment?" but, "Could a machine handle uncertainty more effectively than a human being?" And the answer to that question, in many cases, is that of course it can. In fact, that is precisely what these latest technologies are good at doing. They can handle far larger bodies of data than human beings, and make sense of them in ways that unaided human beings cannot.

The system developed at Stanford University to detect whether or not a lesion is cancerous, mentioned earlier, is a good example of this. It does not try to replicate the judgment of a doctor. Instead, it runs a pattern recognition algorithm through a collection of 129,450 past cases, hunting for similarities between those cases and a photo of the particular lesion in question. It performs a task that might have required judgment when performed by a human dermatologist, but it does so in a fundamentally different way, searching through more cases than a human doctor could expect to have seen in his or her lifetime.

The Future of Professional Work

One of the unhelpful things we do when we talk about the future of work is that we tend to talk about the different jobs that people do. In the professions, for instance, we tend to talk about lawyers and doctors, teachers and accountants, and so on. This is unhelpful because it encourages us to think of the work that the professions do, to use a phrase from before, as monolithic indivisible lumps of endeavor. Yet that is not what professionals do in their daily work. It will be obvious to all practitioners that their jobs are, in fact, made up of a variety of different tasks and activities.

When we think top-down in terms of "jobs," it encourages us to imagine that the only way that new technologies can affect the work of professionals is by substituting their "job" in its entirety, in a disruptive and sudden way. When *The Economist* reviewed our book shortly after publication, the column was accompanied by an image of a robot dressed in the familiar ornaments of professional life—a lawyer's wig, a doctor's stethoscope, an accountant's ledger.[14] They named this robot "Professor Dr Robot QC." When we think about the future of work in terms of "jobs," images like this encourage us to think that one day professionals will turn up at work and find Professor Dr Robot QC sitting at their desks or walking their wards. Their jobs will have been displaced by robots. But this is not how new technologies affect the work of professionals. It helps in this context to think bottom-up in terms of tasks, rather than top-down in terms of jobs. Entire jobs do not disappear in an instant. Instead, new technologies change the sorts of tasks that professionals do in their work. Some tasks still require traditional professionals while others might require different types of people, or not require people at all.

In the medium run, our expectation is that technological change will not lead to *mass unemployment* in the professions, but instead will lead to *substantial redeployment*. This will mean that while some traditional tasks will no longer require traditional professionals, there will be new opportunities for these professionals to adapt and embrace new skills and capabilities. In that spirit, in our work we identify 13 new roles for tomorrow's professionals—from the data scientist to the knowledge engineer, the process analyst to the digital security guard. In *The Future of the Professions*, we explore each of these roles in depth.[15] For current purposes, two observations are worth making. First, many professionals are so steeped in their old ways of working that they do not recognize that these new roles might be part of their job description in the future. The second is that many of these roles require skills and capabilities quite unlike those that we currently train young professionals to perform. Both these observations raise the question as to whether traditional professionals, immersed in time-honoured habits, will, in the end, be sufficiently adaptable to perform new tasks in the future.

---

14 "Professor Dr Robot QC," *Economist,* October 15, 2015.

15 The "digital security guard" is the 13th role, set out in the paperback edition of Richard Susskind and Daniel Susskind, *The Future of the Professions* (Oxford: Oxford University Press, 2017).

The Fundamental Question

When we began our work in 2010, our main preoccupation was with the work of the traditional professions. We wanted to know what technological change would mean for today's doctors and lawyers, teachers and accountants. The temptation, as the title of our book encourages, was simply to ask, "What is the future of the professions?" Yet this formulation assumes that the professions indeed have a future. We felt, on the contrary, that our study should not commit us, before we had conducted our research, to a future in which traditional professionals would necessarily play a central role. As our thinking progressed, we realized there was a more fundamental question that we had to answer. We came upon this by considering what lies at the heart of today's professional service and we concluded that the central role of professionals is to produce and distribute practical expertise. This led us to argue, then, that the fundamental question for the future should be: *How do we produce and distribute practical expertise in society?*

The traditional answer to that question, in the Print-based Industrial Society, has been "through the professions." However, another way to think about what is unfolding is that, as we move into a Technology-based Internet Society, a different set of models for producing and sharing practical expertise are emerging. We identify six of them, as summarized below.

The first is the *networked experts* model. Others have called this "workers on tap."[16] Here, it is still professionals that are involved in producing the practical expertise. But rather than being employed in a particular bricks-and-mortar institution (a firm, hospital, or school), professionals instead use online platforms to work in a far more flexible, more ad hoc way, in solving professional problems. Doctor on Demand in medicine and Axiom Law in the legal world are two examples.

The second is the *para-professional* model. Here, less expert people, using new technologies, are able to perform tasks that would have required more expert people in the past. Take the medical diagnostic system developed at Stanford. It is entirely conceivable that, in primary care of the future, you will be met not by a doctor, but by a nurse practitioner who, using one of these systems, is able to offer the sort of diagnostic support that might have required a more expert person in the past.

The third is the *knowledge-engineering* model. This is what we were doing in the 1980s—engineering systems, derived from the knowledge of experts, for non-experts to use (in our case, to help solve legal

16 See, for instance, "Workers on Tap," *Economist,* December 30, 2014.

problems). Many readily available online DIY tax preparation software and contract drafting tools rely on this model.

The fourth is the *communities of experience* model. Social networks are now a ubiquitous feature of contemporary life. Also familiar are professional networks, where practitioners gather to share their expertise with one another. Less familiar, though, are communities of experience—where patients, rather than practitioners, meet to share their experience and advice. Take, for instance, PatientsLikeMe,[17] an online network of more than 600,000 patients who come together to share experiences of their symptoms and treatments, receiving support and solving problems that might have required more expert medical professionals in the past.

The fifth is the *embedded knowledge* model. To grasp this, consider the card game Solitaire (known also as Patience). If this game is played with physical playing cards and a player tries to put a red five under a red six, this is possible (even if it is called "cheating"). Putting two cards of the same color on top of one other is, of course, against the rules. Now imagine a player who is enjoying the same game but on a smartphone. If the player tries the same move, this is actually not possible, because the system simply returns the offending card to where it came from. The rules are embedded in the system. Breach is not simply prohibited. It is not possible. Likewise, as more of our lives become digitized, practical expertise will not be invoked through the intervention of human beings but will be embedded in our everyday systems instead.

Finally, there is the *machine-generated* model. Here, increasingly capable systems and machines produce and share practical expertise without any human involvement. Of the six models, this is the most radical—where traditional recipients of professional work instead have access to technologies that obviate the need for human experts altogether. Although this scenario is the most widely discussed in popular debate, it is important to keep in mind that this model is only one of six.

Two Moral Questions

When we speak about technology and the professions in our work, we are conscious that it might sound as though we think the future is mapped out in front of us—that we are "technological determinists." But this is not the case. As we make clear in our work, we take the view that how we choose to use technology in the professions is very largely

17 https://www.patientslikeme.com.

in our own hands. This discretion presents us, in turn, with two moral questions.

The first is whether there are some tasks that machines *ought* not to perform, even if they can? As machines become increasingly capable, they will continue to encroach on tasks that once only could be performed by human beings. But are there uses of these new technologies in the professions that we would like to prohibit? Today, U.S. courts use systems to assist judges in making parole decisions.[18] But would we be comfortable with similar systems passing life sentences? There are systems that can rival leading physicians in making medical diagnoses. But would we be comfortable with similar systems making decisions to turn off life-support machines? Where should we set the moral limits to these technologies in the professions? We have called for greater public debate on this question.

The second moral question is this: Who ought to own and control the new bodies of practical expertise? In Print-based Industrial Societies, the professions have acted as gatekeepers, each responsible for their own bodies of practical expertise, supported by the grand bargain. Yet, as explained before, our traditional professions are creaking, open to clear criticism, and in need of reform. Most people do not have affordable access to a good education or to knowledge of their legal entitlements, to reliable medical advice or to sound guidance on how to manage their financial affairs. In Technology-based Internet Societies, where practical expertise can instead be made available online, it is less obvious that the professions—or indeed any individual or institution—should be entitled to act as the gatekeeper to important bodies of knowledge and assistance. The promise is that new technologies will liberate this practical expertise, providing far more affordable access to those who have traditionally been excluded. But the risk is the rise of a new set of gatekeepers in their place, such as large technology companies or Internet service providers, who might build new barriers between recipients and these bodies of practical expertise. This is a significant threat.

In our age of relentless technological advance, it is time to revisit the grand bargain that we have struck with the traditional professions and to ask if this deal remains fit for purpose. Our sense is that it is not. As we argue in the closing lines of our book, we now have the means to share practical expertise much more widely across our world—we should also have the will.

18 Matt O'Brien and Dake Kang, "AI in the Court: When Algorithms Rule on Jail Time," *Associated Press,* January 31, 2018.