Azeem Azhar

Creator of *Exponential View*

Exponential

How Accelerating Technology Is Leaving Us Behind and What to Do About It

Contents

About the Author

Azeem Azhar is the creator of *Exponential View*, Britain's leading platform for in-depth tech analysis. His weekly newsletter is read by 200,000 people from around the world, and his hit podcast has featured guests including Yuval Noah Harari, Tony Blair and Kate Raworth. A member of the World Economic Forum's Global Futures Council, Azhar contributes to publications including the *Financial Times*, *Wired* and the *MIT Technology Review*.

To Salman, Sophie and Jasmine
and the Exponential View community

Preface: The Great Transition

My home lies between the neighbourhoods of Cricklewood and Golders Green in north-west London. It is a suburban house in a suburban street of a type familiar across Europe and the United States. And it is a relatively recent addition to the landscape. Look at a map of the area from 1920, and all you'll see is farmland. The plot of my semi-detached house is right in the middle of a field. A bridleway is shown where an access road now runs, and a few gates and hedgerows demarcate what is now my neighbourhood. A couple of hundred metres to the north lies a blacksmith's.

Just a few years later, the area had transformed. Pick up a map of the same area from 1936 and you'll see the farmland has become the streets that I walk through daily. The blacksmith's has disappeared, replaced by a mechanical workshop. The brick-built interwar-era homes are arranged on the same plots they occupy now, perhaps lacking the odd glass extension. It is a remarkable metamorphosis, which reflects the emergence of a recognisably modern way of life.

As late as the 1880s, life in London resembled that of a much earlier era – horses plied the roads, leaving piles of manure in the streets as they did so; most domestic tasks were powered by hand; much of the population inhabited crowded, centuries-old slum buildings. But beginning in the 1890s, and in many cases completed by the 1920s, the key technologies of the twentieth century took hold. Pictures of central London streets in 1925 show them free of horses, replaced by cars and buses. A network of cables would have carried electricity from coal-fired power stations to offices and homes. Telephone lines ran into many houses and allowed people to talk to distant friends.

These changes in turn brought social upheaval. As modern systems of production developed, so too did full-time employment contracts with benefits; new forms of transport brought with them the commute; the electrification of factories helped the rise of large companies with recognisable brand names. Someone living in the 1980s who stepped into a time machine and went back to the 1880s would have seen little they were familiar with. If they travelled back to just the 1930s they would have recognised much more.

This two-decade transformation reflects the sudden, dramatic changes that technology can bring. Since the days of flint axes and wooden digging sticks, humans have been technologists. We seek to make life easier for ourselves; and to do so, we build tools – technologies – that help us achieve our goals. These technologies have long allowed humans to redefine the world around us. They let us farm and then build; travel on land, then through air, then into space; move from nomadic life to villages to cities.

But, as my predecessors in what is now north-west London learnt, the technologies we build can take society in unexpected directions. When a technology takes off, its effects can be enormous, stretching across all the areas of human life:

our jobs, the wars we fight, the nature of our politics, even our manners and habits. To borrow a word from economics, technology is not 'exogenous' to the other forces that define our lives – it combines with political, cultural and social systems, often in dramatic and unforeseen ways.

The unpredictable ways that technology combines with wider forces – sometimes moving slowly, sometimes causing rapid and seismic transformations – are what makes it so difficult to analyse. The emerging discipline of complexity science tries to make sense of the ways in which the different elements of a complicated system interact – how different species relate to each other to make up an ecosystem, for example. Human society is the ultimate 'complex system'; it is made up of countless, constantly interacting elements – individuals, households, governments, companies, beliefs, technologies.

According to complexity science, the connections between different elements mean that small changes in one area of a system can ripple across the whole. And these changes can be chaotic, sudden and profound.[1] Even if we have a significant degree of knowledge about the component parts of the system, establishing where these ripples might end up is rarely straightforward.[2] A new technology might at first cause a small social change – but one that eventually spirals into major repercussions for the whole of society.

When these ripples – or 'feedback loops', in the jargon of complexity science – start to spread, they can feel uncomfortable. One need only glance at the pages of a newspaper from the turn of the twentieth century to realise that sudden change is anxiety-inducing. A quick survey of *New York Times* articles from a century ago reveals that Americans were apprehensive about elevators, the telephone, the television and more.[3]

Of course, jitters in the elevator were rarely the real issue. Rather, these innovations came to symbolise people's fears about the pace of change. We know intuitively that technological changes rarely remain enclosed within one sphere. By allowing us to build ever-taller buildings, elevators revolutionised the layout and economies of cities. By making contact between people easier, the telephone drastically altered how humans interacted with colleagues and friends. After a technology has taken off, its effects are felt everywhere.

Today, we are undergoing another period of dramatic transformation. The clearest sign of this is the way people talk about technology. The PR company Edelman runs a renowned annual survey on trust in the public sphere. One of their key questions – put to 30,000 people in 20 countries – is whether they feel comfortable with how quickly technology was moving. In 2020, more than 60 per cent of respondents felt the 'pace of change was too fast', a number that had been creeping upwards for several years.[4]

It's tempting to assume that people always feel technological and social change is too fast. They thought so a century ago, and they think so again now. But the argument of this book is that we are indeed living through a time of unusually fast change – and this change is being brought about by sudden technological advances. In the early twenty-first century, the defining technologies of the industrial age are metamorphosing. Our society is being propelled forward by several new innovations – computing and artificial intelligence, renewable electricity and energy storage, breakthroughs in biology and manufacturing.

These innovations are improving in ways that we don't yet fully understand. What makes them unique is the fact they are developing: at an exponential pace, getting faster and faster with each passing month. As in previous periods of rapid technological change, their impact is felt across society – not only leading to new services and products, but also altering the relationship between old and new companies, employers and workers, city and country, citizens and the market.

Complexity scientists refer to moments of radical change within a system as a 'phase transition'.[5] When liquid water turns into steam, it is the same chemical, yet its behaviour is radically different. Societies too can undergo phase changes. Some moments feel abrupt, discontinuous, world-changing. Think of the arrival of Columbus in the Americas, or the fall of the Berlin Wall.

The rapid reorganisation of our society today is just such a moment. A phase transition has been reached, and we are witnessing our systems transforming before our very eyes. Water is becoming steam.

The transformation of society in the early twenty-first century is the focus of this book. It is a book about how new technology is getting faster. And it tries to explain the effect this acceleration is having on our politics, our economies and our ways of life.

But it is not a pessimistic book. There is nothing inevitably harmful about the technologies I will describe. The elements of society that are most important to us – our companies, cultures and laws – emerged in response to the changes brought by earlier technologies. One of the defining features of human history is our adaptability. When rapid technological change arrives, it first brings turmoil, then people adapt, and then eventually, we learn to thrive.

Yet I have chosen to write this book because we presently lack the vocabulary to make sense of technological change. When you watch the news, or read the blogs coming out of the longstanding capital of tech, Silicon Valley, it becomes apparent that our public conversation around technology is limited. New technology is changing the world, and yet misunderstandings about what this tech is, why it matters and how we should respond are everywhere.

In my view, there are two main problems with our conversation about technology – problems that this book hopes to address. First, there is a misconception of how humans relate to technology. We often assume that tech is somehow independent of humanity – that it is a force that brought itself into being and doesn't reflect the biases and power structures of the humans who created it. In this rendering, technology is value-free – it is made neutral – and it is the consumers of the technology who determine whether it is used for good or for evil.

This view is particularly common in Silicon Valley. In 2013, Google's executive chairman Eric Schmidt wrote, 'The central truth of the technology industry – that technology is neutral but people are not – will periodically be lost amid all the noise.'[6] Peter Diamandis, an engineer and physician – as well as the founder of a company that offers tech courses, Singularity University – wrote that while the computer 'is clearly the greatest tool for self-empowerment we've yet seen, it's still only a tool, and, like all tools, is fundamentally neutral'.[7]

This is a convenient notion for those who create technology. If technology is neutral, its inventors can concentrate on building their gizmos. If the tech starts to have any insidious effects, society – rather than its inventors – is to blame. But if technology isn't neutral – that is, if it has encoded some form of ideology, or system of power – that might mean its makers need to be more careful. Society might want to manage or regulate the technologists and their creations more carefully. And those regulations might become a hassle.

Sadly for these engineers, their view of technology is a fiction. Technologies are not just neutral tools to be applied (or misapplied) by their users. They are artefacts built by people. And these people direct and design their inventions according to their own preferences. Just as some religious texts say that humans are fashioned in the image of God, so tools are made in the image of the humans who design them. And that means our technologies often recreate the systems of power that exist in the rest of society. Our phones are designed to fit in men's hands rather than women's. Many medicines are less effective on Black and Asian people, because the pharmaceutical industry often develops its treatments for white customers. When we build technology, we might make these systems of power more durable – by encoding them into infrastructure that is more inscrutable and less accountable than humans are.

And so this book doesn't analyse technology as some abstract force, separate from the rest of society. It views tech as something that is built by humans and reflects human desires, even if it can also transform human society in radical and unexpected ways. *Exponential* is as much about the way technology interacts with our forms of social, political and economic organisation as it is about the technology itself.

The second problem with the way we talk about technology is even more insidious. Many people outside of the world of technology make no effort to understand it, nor to develop the right response to it. Politicians frequently demonstrate a fundamental ignorance about even the most basic workings of mainstream technologies.[8] They are like people trying to fuel a car by filling its trunk with hay. The Brexit trade deal, agreed between the UK and the European Union in December 2020, describes the Netscape Communicator as a 'modern e-mail software package'. The software has been defunct since 1997.

Admittedly, understanding new technology is hard. It takes knowledge of a wide range of new innovations. And it also takes an understanding of society's existing rules, norms, institutions and conventions. In other words, effective analysis of technology involves straddling two worlds. It's reminiscent of a famous 1959 lecture by the British scientist and novelist C. P. Snow. He feared that intellectual life was being split between the domains of literature and science, specifically within the context of British public life. These 'two cultures' did not intersect, and those who understood one rarely understood the other – there was a 'gulf of mutual incomprehension', generated by a 'backward looking intelligentsia' made up of artsy Oxbridge graduates who looked down on technological and scientific progress. This, according to Snow, had disastrous implications: 'When those two senses have grown apart, then no society is going to be able to think with wisdom.'[9]

Today, the gap between the two cultures is wider than ever before. Except now it is most pronounced between technologists – whether software engineers, product developers or Silicon Valley executives – and everyone else. The culture of technology is constantly developing in new, dangerous and unexpected directions. The other culture – the world of humanities and social science, inhabited by most commentators and policymakers – cannot keep track of what is happening. In the absence of a dialogue between the two cultures, our leading thinkers on both sides will struggle to offer the right solutions.

This book is my attempt to bring these two worlds together. On the one hand, I will try to help technologists view their efforts in a wider social context. On the other, I'll aim to help non-technologists get a better understanding of the technologies underpinning this period of rapid social change.

This mix of disciplines suits me well. I am a child of the microchip, born the year after the first commercially produced computer processor was released; a young adult of the internet, who discovered the web while at university; and a professional of the tech industry, having launched my first website – for Britain's *Guardian* newspaper – in 1995. I have founded four tech companies and invested in more than 30 start-ups since 1998. I even survived the dot-com frenzy at the turn of the millennium. Later, at Reuters, I ran an innovation group where our teams built wacky, sometimes brilliant products for hedge fund managers and Indian farmers alike. For several years, I worked with venture capitalists in Europe, backing the most ambitious technology founders we could find – and I still invest actively in young technology companies. As a start-up investor, I have spoken with hundreds of technology founders in fields as diverse as artificial intelligence, advanced biology, sustainability, quantum computing, electric vehicles and space flight.

But my academic training is in the social sciences. At university I focused on politics, philosophy and economics – though, unusually, I also took a programming course with a group of physicists who were much smarter than me. And for much of my career, my focus has been on how technology is transforming business and society. In my career as a journalist, first at *The Guardian* and then *The Economist*, I found myself having to explain complicated topics from the world of software engineering to a mainstream audience. And I've taken a particular interest in the political implications of new forms of technology. For a time, I was a non-executive member of Ofcom, the regulator that looks at the telecom, internet and media industries in the UK. In 2018, I became a board member of the Ada Lovelace Institute, where we have been looking at the ethical implications of the use of data and artificial intelligence in society.

Over the last few years, I have been channelling my attempts to straddle the 'two cultures' into *Exponential View* – a newsletter and podcast that explores the impact of new technology on society. I founded it after my third start-up, PeerIndex, was acquired by a much larger technology firm. PeerIndex applied machine learning techniques (on which, more later) to large amounts of public data about what people do online. We grappled with many ethical dilemmas about what it was and wasn't appropriate to do with this data. After my company's acquisition, I had the mental space to explore such issues in my newsletter.

Exponential View has resonated with people. At the time of writing, it has a readership of nearly 200,000 subscribers around the globe, ranging from some of the world's most well-known founders through to investors, policymakers and academics in more than 100 countries. And it has allowed me to delve into the most thought-provoking questions raised by new technology. Through my podcast series of the same name, I've conducted more than 100 interviews with engineers, entrepreneurs, policymakers, historians, scientists and corporate executives. Over more than six years, I've read tens of thousands of books, newspapers and magazine pieces, blog posts and journal articles as part of my research. I recently estimated that I have read more than 20 million words in the last half-decade in my effort to understand what is going on. (Fortunately, this book is somewhat shorter.)

The conclusion all this research has led me to is deceptively simple. At heart, the argument of *Exponential* has two key strands. First, new technologies are being invented and scaled at an ever-faster pace, all while decreasing rapidly in price. If we were to plot the rise of these technologies on a graph, they would follow a curved, exponential line.

Second, our institutions – from our political norms, to our systems of economic organisation, to the ways we forge relationships – are changing more slowly. If we plotted the adaptation of these institutions on a graph, they would follow a straight, incremental line.

The result is what I call the 'exponential gap'. The chasm between new forms of technology – along with the fresh approaches to business, work, politics and civil society they bring about – and the corporations, employees, politics and wider social norms that get left behind.

Of course, this only raises more questions. What effects do exponential technologies have in different spheres – from work, to conflict, to politics? For how long can this exponential change continue – will it ever stop? And what can we all do, as policymakers, business leaders or citizens, to prevent the exponential gap eroding our societies?

The structure of this book tries to make my answers as clear as possible. In the first part, I will explain what exponential technologies are and why they have come about. I argue that our age is defined by the emergence of several new 'general purpose technologies', each improving at an exponential rate. It's a story that starts with computing – but also encompasses energy, biology and manufacturing. The breadth of this change means that we have entered a wholly new era of human society and economic organisation – what I call the 'Exponential Age'.

Next, I move on to the implications this has for human society more broadly – the emergence of the exponential gap. There are many reasons why human-built institutions are slow to adapt, from the psychological trouble we have conceptualising exponential change, through to the inherent difficulty of turning around a big organisation. All contribute to the widening gulf between technology and our social institutions.

But what effects does the exponential gap have in practice? And what can we do about it? Those questions are the focus of the rest of this book. I'll take you from the economy and work, through the geopolitics of trade and conflict, to the broader relationship between citizens and society.

First, we'll explore what exponential technologies do to businesses. During the Exponential Age, technology-driven companies tend to become bigger than was previously thought possible – and traditional companies get left behind. This leads to winner-takes-all markets, in which a few 'superstar' companies dominate – with their rivals spiralling into inconsequentiality. An exponential gap emerges – between our existing rules around market power, monopoly, competition and tax, and the newly enormous companies that dominate markets.

I'll also show how the prospects of employees are changing thanks to the emergence of these companies. The relationships between workers and employers are always in flux, but now they are shifting more rapidly than ever. The superstar companies favour new styles of work, mediated by gig platforms, which may be problematic for workers. Existing laws and employment practices struggle to cope with the changing norms surrounding labour.

Second, we'll explore the transformation of geopolitics – discussing how exponential technologies are rewiring trade, conflict and the global balance of power. Here, two great shifts are underway. The first is a return to the local. New innovations alter the way we access commodities, manufacture products and generate energy – increasingly, we will be able to produce all three within our own regions. At the same time, the increasing complexity of our economies will make cities more important than ever, creating tension between regional and national governments. If the story of the industrial age was one of globalisation, the story of the Exponential Age will be one of re-localisation. The second is the transformation of warfare. As the world gets re-localised, patterns of global conflict will shift. Nations and other actors will be able to make use of new adversarial tactics, from cyber threats to drones and disinformation. These will dramatically reduce the cost of initiating conflict, making it much more common. A gap will emerge between new, high-tech forms of attack and societies' ability to defend themselves.

Third, we'll examine how the Exponential Age is rewiring the relationship between citizen and society. State-sized companies are on the rise – and they are challenging our most basic assumptions about the role of private corporations. Markets are metastasising across ever-greater swathes of the public sphere and our private lives. Our national conversations are increasingly conducted on privately owned platforms; intimate details about our innermost selves are bought and sold online, thanks to the emergence of the data economy; and even the way we meet friends and form communities has been turned into a commodity. But because we remain wedded to an industrial-age conception of the role of markets, we don't yet have the toolkit to prevent these changes eroding our most cherished values.

In other words, an exponential gap is challenging many elements of our society. But that's something we can address. And so, at the end of the book, I'll explain the broad principles that we need to make sure we thrive in an age of exponential change – from making our institutions more resilient to rapid transformation, to reiterating the power of collective ownership and decision-making. The resulting book is, I hope, a holistic guide to how technology is changing our society – and what we should do about it.

As I wrote this book, the world changed dramatically. When I first started my research, there was no such thing as Covid-19, and lockdowns were the remit of zombie apocalypse movies. But as I was halfway through writing my first draft, countries around the world began shutting their borders and issuing stay-at-home orders to their populations – all to prevent a virus wreaking havoc on their health systems and economies.

On one level, the pandemic felt distinctly low-tech. Lockdowns have been used for millennia to prevent the spread of disease. Quarantines are nothing new: the word derives from the time of the Black Death, when sailors had to isolate for 40 days before coming on shore. That the global economy was brought low by a virus reminded us of how many ancient problems technology hasn't yet been able to solve.

But the pandemic also hammered home some of the key points of this book. The spread of the virus demonstrated that exponential growth is hard to control. It creeps up on you and then explodes – one moment everything seems fine, the next your health service is on the verge of being overwhelmed by a new disease. And humans struggle to conceptualise the speed of that shift, as shown by the lackadaisical responses of many governments to the spread of coronavirus, particularly in Europe and America.

At the same time, the pandemic revealed the full power of recent inventions. In most of the developed world, lockdowns were only possible due to widespread access to fast internet. Those of us locked at home spent much of the pandemic glued to our phones. And, most strikingly of all, within a year scientists had developed dozens of new vaccines – which, as we'll see, were made possible by new innovations like machine learning. In some ways, exponential technology proved its mettle with Covid-19.

Above all, the pandemic revealed that Exponential Age technologies – whether video calls or social media platforms – are now embedded into every part of our lives. And this will only become more pronounced. As the rate of change speeds up, the interaction between technology and other spheres of our lives – from demography to statecraft to economic policy – will become increasingly constant. Neat distinctions between the realm of technology and the realm of, say, politics will become unhelpful. Technology is remaking politics, and politics is shaping technology. Any constructive analysis of either requires an analysis of both. And for politics, one could substitute economics, or culture, or business strategy.

As a result of the constant feedback loop of technology, economics, politics and society, making stable predictions about the future is difficult. Even as I wrote this book, its subject matter was constantly shifting – no sooner would I finish a chapter than it would need to be updated to incorporate new developments. Such are the perils of writing in an age of exponential change.

But my hope is that this book remains a useful introduction to where new technology is taking us. We are living in an era when technology is getting better, faster and more varied at a greater speed than ever before. That process is undermining the stability of many of the norms and institutions that define our lives. And we don't, at the moment, have a road map that will help us get to the future we want.

This book, by itself, is unlikely to offer a perfect map. But it might help reveal the terrain, and point us in the right direction.

Azeem Azhar
London, April 2021

1

The Harbinger

Before I knew what Silicon Valley was, I had seen a computer. It was December 1979, and our next-door neighbour had brought home a build-it-yourself computer kit. I remember him assembling the device on his living room floor and plugging it into a black-and-white television set. After my neighbour meticulously punched in a series of commands, the screen transformed into a tapestry of blocky pixels.

I took the machine in with all the wonder of a seven-year-old. Until then, I had only seen computers depicted in TV shows and movies. Here was one I could touch. But it was more remarkable, I think now, that such a contraption had even got to a small suburb of Lusaka in Zambia in the 1970s. The global supply chain was primordial, and remote shopping all but non-existent – and yet the first signs of the digital revolution were already visible.

The build-it-yourself kit piqued my interest. Two years later, I got my own first computer: a Sinclair ZX81, picked up in the autumn of 1981, a year after moving to a small town in the hinterlands beyond London. The ZX81 still sits on my bookshelf at home. It has the footprint of a 7-inch record sleeve and is about as deep as your index and middle fingers. Compared to the other electronic items in early-1980s living rooms – the vacuum-tubed television or large cassette deck – the ZX81 was compact and light. Pick-up-with-your-thumb-and-forefinger light. The built-in keyboard, unforgiving and taut when pressed, wasn't something you could type quickly on. It only responded to stiff, punctuated jabs of the kind you might use to admonish a friend. But you could get a lot out of this little box. I remember programming simple calculations, drawing basic shapes and playing primitive games on it.

This device, advertised in daily newspapers across the UK, was a breakthrough. For £69, we got a fully functional computer. Its simple programming language was, in principle, capable of solving any computer problem, however complicated (although it might have taken a long time).[1] But the ZX81 wasn't around for long. Technology was developing quickly. Within a few years, my computer – with its blocky black-and-white graphics, clumsy keyboard and slow processing – was approaching obsolescence. Within six years, my family had upgraded to a more modern device, made by Britain's Acorn Computers. The Acorn BBC Master was an impressive beast, with a full-sized keyboard and a numeric keypad. Its row of orange special-function keys wouldn't have looked out of place on a prop in a 1980s space opera.

If the exterior looked different to the ZX81's, the interior had undergone a complete transformation. The BBC Master ran several times faster. It had 128 times as much memory. It could muster as many as 16 different colours, although it was limited to displaying eight at a time. Its tiny speaker could emit up to four distinct tones, just enough for simple renditions of music – I recall it beeping its way through Bach's Toccata and Fugue in D Minor. The BBC Master's relative sophistication allowed for powerful applications, including spreadsheets (which I never used) and games (which I did).

Another six years later, in the early 1990s, I upgraded again. By then, the computer industry had been through a period of brutal consolidation. Devices like the TRS-80, Amiga 500, Atari ST, Osborne 1 and Sharp MZ-80 had vied for success in the market. Some small companies had short-lived success but found themselves losing out to a handful of ascendant new tech firms.

It was Microsoft and Intel that emerged from the evolutionary death match of the 1980s as the fittest of their respective species: the operating system and the central processing unit. They spent the next couple of decades in a symbiotic relationship, with Intel delivering more computational power and Microsoft using that power to deliver better software. Each generation of software taxed the computers a little more, forcing Intel to improve its subsequent processor. 'What Andy giveth, Bill taketh away' went the industry joke (Andy Grove was Intel's CEO; Bill Gates, Microsoft's founder).

At the age of 19 I was oblivious to these industry dynamics. All I knew was that computers were getting faster and better, and I wanted to get hold of one. Students tended to buy so-called PC clones – cheap, half-branded boxes which copied the eponymous IBM Personal Computer. These were computers based on various components that adhered to the PC standard, meaning they were equipped with Microsoft's latest operating system – the software that allowed users (and programmers) to control the hardware.

My clone, an ugly cuboid, sported the latest Intel processor, an 80486. This processor could crunch through 11 million instructions per second, probably 4–5 times more than my previous computer. A button on the case marked 'Turbo' could force the processor to run some 20 per cent faster. Like a car where the driver keeps their foot on the accelerator, however, the added speed came at the cost of frequent crashes.

This computer came with 4 megabytes of memory (or RAM), a 4,000-fold improvement on the ZX81. The graphics were jaw-dropping, though not state-of-the-art. I could throw 32,768 colours on the screen, using a not-quite cutting-edge graphics adaptor that I plugged into the machine. This rainbow palette was impressive but not lifelike – blues in particular displayed poorly. If my budget had stretched £50 more, I might have bought a graphics card that painted 16 million colours, so many that the human eye could barely discern between some of the hues.

The 10-year journey from the ZX81 to PC clone reflected a period of exponential technological change. The PC clone's processor was thousands of times more powerful than the ZX81's, and the computer of 1991 was millions of times more capable than that of 1981. That transformation was a result of swift progress in the nascent computing industry, which approximately translated to a doubling of the speed of computers every couple of years.

To understand this transformation, we need to examine how computers work. Writing in the nineteenth century, the English mathematician and philosopher George Boole set out to represent logic as a series of binaries. These binary digits – known as ‘bits’ – can be represented by anything, really. You could represent them mechanically by the positions of a lever, one up and one down. You could, theoretically, represent bits with M&Ms – some blues, some reds. (This is certainly tasty, but not practical.) Scientists eventually settled on 1 and 0 as the best binary to use.

In the earliest days of computing, getting a machine to execute Boolean logic was difficult and cumbersome. And so a computer – basically any device that could conduct operations using Boolean logic – required dozens of clumsy mechanical parts. But a key breakthrough came in 1938, when Claude Shannon, then a master’s student at the Massachusetts Institute of Technology, realised electronic circuits could be built to utilise Boolean logic – with on and off representing 1 and 0. It was a transformative discovery, which paved the way for computers built using electronic components. The first programmable, electronic, digital computer would famously be used by a team of Allied codebreakers, including Alan Turing, during World War Two.

Two years after the end of the war, scientists at Bell Labs developed the transistor – a type of semiconductor, a material that partly conducts electricity and partly doesn’t. You could build useful switches out of semiconductors. These in turn could be used to build ‘logic gates’ – devices that could do elementary logic calculations. Many of these logic gates could be stacked together to form a useful computing device.

This may sound technical, but the implications were simple: the new transistors were smaller and more reliable than the valves that were used in the earliest electronic components, and they paved the way for more sophisticated computers. In December 1947, when scientists built the first transistor, it was clunky and patched together with a number of large components, including a paper clip. But it worked. Over the years, transistors would become less ad hoc, and more consistently engineered.

From the 1940s onwards, the goal became to make transistors smaller. In 1960, Robert Noyce at Fairchild Semiconductor developed the world’s first ‘integrated circuit’, which combined several transistors into a single component. These transistors were tiny and could not be handled individually by man or machine. They were made through an elaborate process a little like chemical photography, called photolithography. Engineers would shine ultraviolet light through a film with a circuit design on it, much like a child’s stencil. This imprints a circuit onto a silicon wafer, and the process can be repeated several times on a single wafer – until you have several transistors on top of one another. Each wafer may contain several identical copies of circuits, laid out in a grid. Slice off one copy and you have a silicon ‘chip’.

One of the first people to understand the power of this technology was Gordon Moore, a researcher working for Noyce. Five years after his boss’s invention, Moore realised that the physical area of integrated circuits was reducing by about 50 per cent every year, without any decrease in the number of transistors. The films – or ‘masks’ – used in photolithography were getting more detailed; the transistors and connections smaller; the components themselves more intricate. And this reduced

image
not
available

And, as ever, lower prices mean higher usage. In 1999, we had sequenced one messy genome. By 2015, humanity was sequencing more than 200,000 genomes a year.[11] One research group estimates that by 2025 as many as 2 billion human genomes may have been sequenced.[12]

There are a number of factors driving down the cost of coding a genome, and growing computing power is part of the story. Genome sequences are enormous chains of letters. Coding one human genome would require about 100 gigabytes of storage (enough to store about 25 high-definition movies) – a level that is much easier to secure now than it was two decades ago. But Moore's Law is far from the only cause of the price drop. There have been developments in the way we produce the reagents and 'amplifiers' required to turn a DNA sample into something readable. Over the years, these chemicals have become progressively cheaper and cheaper. Meanwhile, advancements in electronics have allowed scientists to create cheaper sensors, and developments in robotics have allowed for more automation of the manual parts of the convoluted process.[13]

Genetic coding is but one aspect of the revolution in biotech. Another is synthetic biology, a field that melds several disciplines, including computer science, biology, electrical engineering and biophysics, to create novel biological components and systems. It too is on an exponential march – one that is churning out breakthroughs in agriculture, pharma, materials and healthcare. Today, we can sequence and manipulate microorganisms. We can turn them into little natural factories to produce the chemicals and materials we need – something that would have been unthinkable even a decade ago. The impact of this will be transformative. According to some estimates, 60 per cent of the physical 'inputs' in the global economy could be produced biologically by 2040.[14] Harnessing nature in this way will let us produce completely new materials – biopolymers that won't hurt our oceans; electronic components that are lighter or consume less power.

And then there is the last of our four areas of exponential change: manufacturing. How we make things is in the process of transforming fundamentally, perhaps for the first time in millions of years.

We have been manipulating physical materials – matter – in pretty much the same way since long before there were *Homo sapiens*. The oldest-known flints were carved in Olduvai, in modern-day Tanzania, about 1.7 million years ago. Industrial-era manufacturing processes have much in common with those of our distant forebears. We too generally use a subtractive process – start with a block of stuff, and chisel away what we don't want. This is what hominids did with flints. It is what pharaonic stonemasons did to the stone blocks of the pyramids. And it is what Michelangelo did when he chiselled a block of marble to create *David*.

Today, we can do this at a grander scale and with greater precision, but the process is essentially the same. Even as the computer age heralded precise computerised machining, this was still a subtractive process: the early human's hammering of flint on stone was replaced by a diamond cutter controlled by a computer. Of course, there are other methods of making stuff, such as using casts to mould metals or plastic. These have the advantage over chiselling in that they don't waste any material. But a big disadvantage: casts and moulds only create copies of a single design. Want a new product and you need a new mould.

Additive manufacturing, or 3D printing – I'll use the terms interchangeably – is an exponential technology that delivers the individual detailing of subtractive manufacturing, without the waste. Typically, objects are crafted through computeraided design. The process involves creating a new object from scratch: by putting together layers upon layers of melted material, using a laser or a device a little like an inkjet printer. The material can range from glass to plastic to chocolate. It marks a fundamental break with many millennia of subtractive manufacturing, and thousands of years of casting and moulding.

Since the first 3D printers were developed by Charles Hull in the mid-1980s, additive manufacturing has improved dramatically. The process has become faster, more precise and more versatile – today, 3D printers can work with materials including plastics, steels, ceramics and even human proteins. In 1999, the Wake Forest Institute for Regenerative Medicine grew the first 3D printed organ for transplant surgery. And in Dubai in 2019, my friend Noah Raford spearheaded the then largest 3D-printed object: a single-storey 2,500-square-foot building.[15] Printed out of concrete in 17 days, Noah used it as his office for several months. The project used 75 per cent less concrete than a typical design and was built with unheard-of precision.

Additive manufacturing is still a tiny business. You'll find it in prestige products and in highly specialist sectors of the economy – lightweight parts for fighter jets, or medical implants. But the underlying technologies are on an exponential course. Researchers estimate that most additive manufacturing methods are developing at a pace of between 16.7 per cent and 37.6 per cent every year, with the average rate falling in the high thirties.[16] Over the next 10 years we will see performance improve 14 times – and, of course, see prices drop concomitantly. Terry Wohlers, an analyst of the additive manufacturing sector, tells me that the 3D printing market grew 11 times in the decade to 2019 – a rate of 27 per cent per annum.[17]

Why does the transformation of these four domains matter? After all, new technologies come along all the time. Researchers develop new ways of solving problems; engineers improve the methods we already have; occasionally we stumble upon a breakthrough. Even if the rate of change is increasing, you might think, the fundamental process is nothing new.

But the technologies in these four domains – computing, biology, energy and manufacturing – are special. To understand why, we must recognise a fundamental truth about innovation: not all technologies are equal.

Most technologies have fairly few uses – think of stirrups, or light bulbs. This doesn't mean their impact is small. The humble stirrup, really only of use for the horse rider, is credited with helping Genghis Khan sweep across Asia to create the world's largest land empire. The light bulb broke us free from the shackles of darkness. Society could function, at work or at home, after sunset. Narrow technologies can therefore have a broad impact. However, their uses remain relatively circumscribed.

Some innovations have a much broader utility, though. The wheel might provide power in a waterwheel, or serve as part of a pulley, or be used in a vehicle. Farmers, firemen and financiers might all have cause to call on a wheel. And a wheel could be used in every part of their trade. These wide-ranging inventions are known as 'general purpose technologies'. They may displace other technologies and create the opportunity for a wide variety of complementary products – products and services that can only exist because of this one invention.

Throughout history, general purpose technologies (GPTs) have transformed society beyond recognition. Electricity drastically altered the way factories work, and revolutionised our domestic lives. The printing press, which played a key role in the European Reformation and the scientific revolution, was much more than a set of pressure plates and cast metal type. GPTs upturn our economies, and our societies too – spawning changes far beyond the sectors in which they began.[18] As the economists Richard Lipsey, Kenneth Carlaw and Clifford Bekar put it, GPTs 'change almost everything in a society ... by creating an agenda for the creation of new products, new processes, and new organizational forms.'[19]

Part of the reason GPTs are so transformative is the way they have effects beyond any one sector. Consider one of the key GPTs from the start of the twentieth century: the car. To reach their potential, cars needed suitable roads – physical infrastructure that spanned nations. But cars also needed fuel and spare parts, and drivers needed sustenance – creating the demand for fuel stations and roadside cafés. Cars forced changes to the urban environment and so cities started to change, with precedence going to motorised vehicles. Over time, suburbs developed, and with them came the gradual reshaping of consumer practices: reasonably priced hotels for vacationers, and big-box retail stores. New rules slowly emerged, including a slew of safety regulations for drivers. In short, the GPT changed everything.

And this hints at why the exponential revolution in our four key sectors is so important. We are witnessing the emergence of a transformative new wave of GPTs. Not a single GPT, as in the time of the printing press. Nor even three GPTs, like the early twentieth century's offering of the telephone, the car and electricity. In the Exponential Age, we're experiencing multiple breakthrough technologies in the four broad domains of computing, energy, biology and manufacturing.

At this early stage, it is hard to pin down precisely what the GPTs in these areas will be. All we know is that the emerging technologies in each sphere can be applied in a massive array of ways. As we have seen, growing computing power has a seemingly endless range of uses. Gene engineering might be used to tamper with microorganisms, produce new screens for smartphones, or help us design precise medicines. And 3D printers let us create everything from precision car parts to new bodily organs.

The rapid evolution of these technologies does not presage instantaneous change. The revolutionary effect of general purpose technologies may take time to materialise. Consider electricity. In the words of the leading economic historian James Bessen: 'The first electrical generating stations opened in 1881, but electrification had little effect on economic productivity until the 1920s.'[20] GPTs take a while to have meaningful effects – as new infrastructure is built, ways of working change, and companies train their employees in novel techniques.

GPTs are integrated into the economy in a series of steps, best described by the economist Carlota Perez.[21] First comes the installation phase, when the basic infrastructure underpinning a GPT is developed. The roll-out of a GPT is an arduous process: to build an electricity network, you need to build generating systems and power lines and grids. In this phase, skills are limited, and know-how is scarce. It takes time to develop the knowledge required to apply a GPT at scale. These early years may be a low-productivity, discovery-led period in the life of a GPT. Existing, mature technologies are more efficient and pervasive than the new inventions – in some cases, using a novel technology may be more trouble than it is worth.

But the real revolution comes at the next stage: the deployment phase. After a laborious installation phase, our economies have rolled out enough of a new