The Exponential Age will transform economics forever

It's hard for us to fathom exponential change – but our inability to do so could tear apart businesses, economies and the fabric of society

Azeem Azhar 06.09.2021 06:00 AM

Nadav Kander

IN 2020, AMAZON turned 26 years old. Over the previous quarter of a century, the company had transformed shopping. With retail revenues in excess of $213bn (£150bn), it was larger than Germany's Schwarz Group, America's Costco and every British retailer. Only Walmart, with more than half-a-trillion dollars of sales, was bigger. But Amazon was by far and away the world's largest online retailer. Its online business was about eight times larger than Walmart's. Amazon was more than just an online shop, however. Its huge businesses in areas such as cloud computing, logistics, media and hardware added a further $172bn in sales.

At the heart of its success is a staggering $36 billion research and development budget in 2019, used to develop everything from robots to smart home assistants. This sum leaves other companies – and many governments – behind. It is not far off the UK government's annual budget for research and development. Tesco, the largest retailer in Britain – with annual sales in excess of £50 billion – had a research lab whose budget was in "six figures" in 2016.

Perhaps more remarkable was the rate at which Amazon had grown this budget. Ten years earlier, Amazon's research budget had been $1.2 billion. Over the course of the next decade, the firm increased its annual R&D budget by about 44 per cent every year. As the 2010s went on, Amazon doubled down on its investments in research. In the words of Werner Vogels, the firm's Chief Technology Officer, if they stopped innovating they "would be out of business in 10-15 years".

In the process, Amazon created a chasm between the old world and the new. The approach of traditional business was to rely on models that had succeeded yesterday. They were based on a strategy that tomorrow might be a little different now, but not markedly so.

This kind of linear thinking, rooted in the assumption that change would take decades, not months, may have worked in the past – but not anymore. Amazon understood the nature of the Exponential Age. The pace of change was accelerating. The companies that could harness the technologies of the new era would take off. And those that couldn't keep up would be undone at remarkable speed. This divergence between the old and the new is one example of what I call the exponential gap.

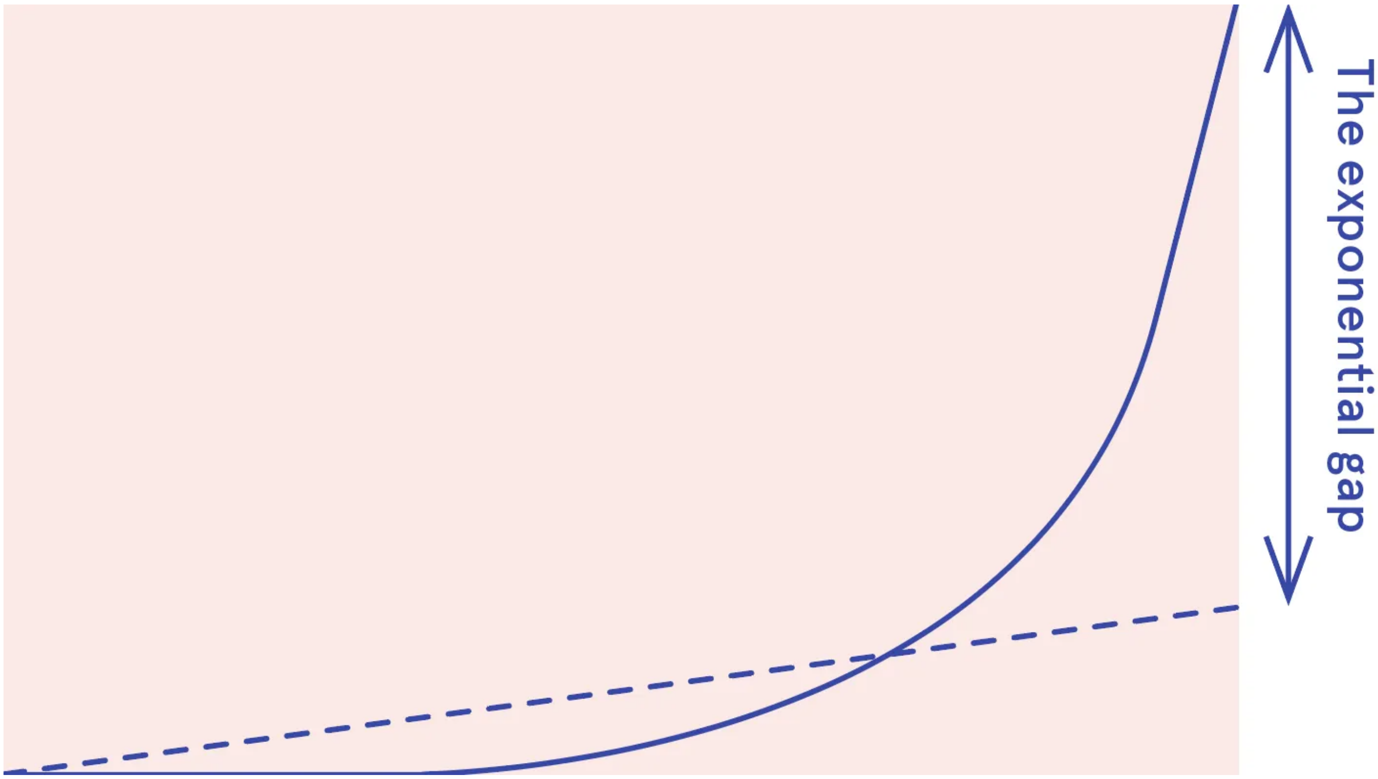

Linear institutions, exponential technologies and the exponential gap

The emergence of this gap is a consequence of exponential technology. Until the early 2010s, most companies assumed the cost of their inputs would remain pretty similar from year-to-year, perhaps with a nudge for inflation. The raw materials might fluctuate based on commodity markets; but their planning processes, institutionalised in management orthodoxy, could manage such volatility. But in the Exponential Age, one primary input for a company is its ability to process information. One of the main costs to process that data is computation. And the cost of computation didn't rise each year, it declined rapidly. The underlying dynamics of how companies operate had shifted.

Moore's Law amounts to a halving of the underlying cost of computation every couple of years. It means that every ten years, the cost of the processing that can be done by a computer will decline by a factor of 100. But the implications of this process stretch far beyond our personal laptop use. In general, if an organisation needs to do something that uses computation, and that task is too expensive today, it probably won't be too expensive in a couple of years. For companies, this realisation has deep significance. Organisations that understood this deflation, and planned for it, became well-positioned to take advantage of the Exponential Age.

If Amazon's early recognition of this trend helped transform it into one of the most valuable companies in history, it was not alone. Many of the new digital giants, from Uber to Alibaba, Spotify to TikTok, took a similar path. Companies that didn't adapt to exponential technology shifts, like much of the newspaper publishing industry, didn't stand a chance.

We can visualise the gap by looking at an exponential curve. Technological development roughly follows this shape. It starts off looking a bit humdrum. In those early days, exponential change is distinctly boring, and most people and organisations ignore it. At this point in the curve, the industry producing an exponential technology looks exciting to those in it, but like a backwater to everyone else. But at some point, the line of exponential change crosses that of linear change. And soon it reaches an inflection point. That shift in gear, which is both rapid and subtle, is hard to fathom.

Because, for all the visibility of exponential change, most of the institutions that make up our society follow a linear trajectory. Codified laws and unspoken social norms; legacy companies and NGOs; political systems and intergovernmental bodies – all have only ever known how to adapt incrementally. Stability is an important force within institutions. In fact, it's built into them. The gap between our institutions' capacity to change and our new technologies' accelerating speed is the defining consequence of our shift to the Exponential Age. On the one side, you have the new behaviours, relationships and structures that are enabled by exponentially improving technologies, and the products and services built from them. On the other, you have the norms that have evolved or been designed to suit the needs of earlier configurations of technology.

The gap leads to extreme tension. In the Exponential Age this divergence is ongoing – and it is everywhere.

Consider the economy. When an exponential age company is able to grow to an unprecedented scale and establish huge market power, it may fundamentally undermine the dynamism of a market. Yet industrial age rules of monopoly may not recognise this behaviour as damaging. This is the exponential gap.

Or take the nature of work. When new technologies allow firms and workers to offer and bid on short-term tasks through gig-working platforms, it creates a vibrant market for small jobs, but potentially at the cost of more secure, dependable employment. When workers compete for work on task-sharing platforms, by bidding via mobile apps, what is their employment status? What rights do they have? Does this process empower them or dehumanize them? Nobody is quite sure: our approach to work was developed in the nineteenth and twentieth century. What can it tell us about semi-automated gig work? This is the exponential gap.

Or look at the relationship between markets and citizens. As companies develop new services using breakthrough technologies, ever more aspects of our lives will become mediated by private companies. What we once considered to be private to us will increasingly be bought and sold by an Exponential Age company. This creates a dissonance: the systems we have in place to safeguard our privacy are suddenly inadequate; we struggle to come up with a new, more apt set of regulations. This is the exponential gap.

Transistors per microprocessor (billions). Moore's Law means transistors boom while computing costs fall

THE MOST BASIC cause of the exponential gap is simple: we are bad at maths.

Let's consider for a moment what it actually feels like to live in the Exponential Age – the answer, as many readers will know, is bewildering. Someone like me, born in the early 1970s, has experienced wave after wave of innovations. From the landline to the mobile, dialup internet to mobile internet, vinyl records to CDs to MP3s to streaming. I've owned the Bee Gees' *Saturday Night Fever* on at least five different formats: vinyl record, cassette tape, CD, MP3 download and, now, streaming access.

Human cognitive machinery does not naturally process such rapid change. The calculations bewilder us. Take the case of an atypical London rainstorm. Wembley Stadium is England's national soccer venue. It is about 8km north-west of my home, and I see it out of my window when I go to visit my in-laws. Its steel arch, spanning 315 metres and reaching 133 metres at its peak, soars above the silver-grey roof. It is an enormous edifice, seating some 90,000 people at capacity. Imagine sitting at the highest row of level three, the furthest above the pitch you can be – some 40m or so above the ground.

Rain starts to fall, but you are sheltered by the partial roof above you. Yet this is no ordinary rain. This is exponential rain. The raindrops are going to gradually increase in frequency, doubling with each passing minute. One drop, then a minute later two drops, then a minute later four drops. By the fourth minute, eight drops. If it takes 30 minutes to get out of your seats and out of the stadium, how soon should you get moving to avoid being drenched?

To be safe, you should start moving by no later than minute 17 – to give yourself 30 minutes to be clear of the stadium. By the 47th minute, the exponential rain will be falling at a rate of 141 trillion drops per minute. Assuming a raindrop is about four cubic millimetres, by the 47th minute the deluge would be 600 million litres of water. Of course, the rain in the 48th minute will be twice as large, so you are likely to get soaked in the car park. And if you make it to the car, the deluge in the fiftieth minute will comprise five billion litres of water. It would weigh five million tonnes. Frankly, if exponential rain is forecast, you're best off staying at home.

These exponential processes are counterintuitive. Psychologists who study how people save for the future have identified the "exponential growth bias" – which makes us underestimate the future size of something growing at a compounded rate. Studies in this area show how people are consistently befuddled by the compound growth of our savings, loans and pension plans. If you started investing in your pension a little late, you, like many of us, may have had a persistent bout of exponential growth bias.

One study tested how well Swedish adults could understand compounding growth processes. Researchers asked an apparently rather muted question: how much 100 Swedish kroner, if left in a bank account to earn seven per cent interest a year for 30 years, would grow to. Even that simple growth rate baffled respondents. The median answer was 410. The correct answer, 761 kroner, was almost double that. More than 60 per cent of the respondents underestimated the answer.And that was people's underestimate of an annual compounding rate of seven per cent. Imagine our predictions of an exponential technology that improves at ten per cent or more per annum. Three decades of compounded seven per cent change results in an increase of a factor of seven. A growth rate of around 40 per cent – roughly what Moore's Law describes – would see a 32,000-fold increase in that time. One peer-reviewed summary from 1975 pithily summarised the issue: underestimation of exponential growth was a "general effect which [was] not reduced by daily experience with growing processes."

This blindspot has a close relative – the “anchoring bias”. The Nobel Prize-winning economists Daniel Kahneman and Amos Tversky have explored how people make decisions amidst uncertainty. They find that when presented with a numerical challenge, people tend to fix upon some readily available number and adjust their responses around it. It’s a trick salespeople use: by starting at a particular price, they anchor our expectations about what the real value of something might be. But it fails when it encounters exponential growth. As the growth curve takes off, people's expectations remain anchored around small figures from early in the process.

Azeem Azhar

Nadav Kander

WHAT HAPPENS WHEN we try to make predictions about exponential change in the real world – rather than just in psychologists' experiments? I've seen the trouble this poses first-hand. Over the course of my career, I’ve watched well-informed people pooh-pooh mobile phones, the internet, social networks, online shopping and electric vehicles as niche playthings destined for eternal obscurity. Over two decades, I’ve observed executives in established industries regularly, perhaps deliberately, look at the spread of a new product or service and dismiss it. Often it was because the absolute numbers were small, in spite of signs of hockey-stick growth. Like spectators at Wembley Stadium during a period of exponential rain, they didn’t leave their seats until it was too late.

For example, in the early 1980s, companies started to operate the first cellular phone services. At the time, handsets were clunky, calls filled with static, data services non-existent and coverage patchy. Yet it was becoming clear that mobile telephony had clear, practical benefits. The giant American phone company AT&T asked the world’s top management consultancy firm, McKinsey, to estimate the future market size for this product. McKinsey put together a 20-year forecast: the US mobile phone market would approach 900,000 subscribers by 2000.

Not quite. The first phone, the cement brick-sized Motorola DynaTAC, cost $3,995 in 1984. The core components were getting cheaper every year, and phones followed suit: becoming better, smaller and cheaper. But in the year 2000, you could find a new cellphone for a couple of hundred dollars. And the capabilities of the networks were growing, too. In 1991, mobile networks were just starting to introduce data services – until then phones had been used almost exclusively for voice calls. In those days, if you bought a device to connect your computer to the phone network, you could use it to send data: at a rate of 9,600 bits (or about 1,000 words) per second. Had digital cameras been widespread at the time (they weren't), sending a single photo would have taken several minutes. By 2020, common 4G phone networks could deliver 30,000,000 bits per second or more to handsets. Mobile tariffs collapsed in line with the growing speed of the networks. Between 2005 and 2014, the average cost of delivering one megabyte of data, the equivalent of about 150,000 words, dropped from $8 to a few cents. In short: McKinsey had miscalled it. In the year 2000, more than 100m Americans owned a mobile phone. The most storied management consultants in the world had been wrong by a factor of 100. Predicting the future is hard – predicting it against an exponential curve is harder still.

This is not a problem limited to the private sector. The International Energy Agency is an intergovernmental organisation founded in 1974, in the wake of the global oil crisis of the previous year. The IEA's annual World Energy Outlook has, for years, predicted the amount of electricity generated by solar power. In one of its forecasts, made in 2009, the IEA predicted five gigawatts of global solar power by 2015. They were wrong. The actual number in 2009 – yes, the year in which they made that prediction – was eight gigawatts. In 2010, they upgraded that 2015 forecast to eight gigawatts. In 2011, they upped it again to 11 gigawatts. In 2012, they predicted 24 gigawatts. By 2014, they were predicting 35 gigawatts of solar capacity by the next year. The real capacity for 2015? 56 gigawatts. This global group of experts systematically misread the market for six years straight, right up until the year beforehand. But it didn't stop there. After six years of hopeless forecasts, they continued the trend for several more. In 2018, the IEA estimated global solar capacity was 90 gigawatts. And they predicted a rough stand-still for the next year, an estimate of 90 gigawatts for 2019. In reality, 2019's output exceeded 105 gigawatts. In that year, the annual growth they forecast was off by 100 per cent – or infinity, depending on how you do that maths. It was a decade of looking at an exponential technology, dropping in price and increasing in scale, and systematically getting it wrong.

The problem is not just that we underestimate exponential growth. Experts who are mindful of the power of exponentiality can also be prone to overestimate its power. In his 1999 book, *The Age of Spiritual Machines*, Ray Kurzweil predicted that by 2019 a $1,000 computer would be "approximately equal to the computational ability of the human brain." This proved optimistic. When trying to square rapid, exponential growth with an inordinately complex issue, a slight error in your basic assumptions can throw your whole prediction off. And with a neural network as complex as the human brain, it's near impossible to get these assumptions right. Our best current guess is that the human brain has about 100 billion neurons. Each neuron is connected, on average, to a thousand others, leading scientists to estimate there are 100 trillion connections in the human brain. If these estimates prove correct, and if we have properly understood the functions of the neurons, a machine that mimics the complexity of the brain could conceivably be built within a couple of decades. But those are big ifs. When our scientific understanding of a subject is still developing, predictions are sometimes little better than guesswork.

And these problems of underestimation and overestimation are confounded by a third difficulty – which we might call "mis-estimation". Exponentiality often has unexpected, arguably unpredictable effects. Take chewing gum. In the decade from 2007, American chewing gum sales fell 15 per cent – just as 220 million American adults bought their first smartphones. This was no coincidence. When people got into a shop's queue, they would once have spent the time browsing the goodies for sale at the counter – and gum was the obvious choice. Suddenly, they were spending that time playing with their phones. So gum sales plummeted. Nobody saw that one coming. Predicting the impact of the iPhone on grocery store gum sales would have needed a modern-day Nostradamus.

At A, the exponential growth is barely noticeable. At B, it is obvious

IF THE PRIMARY cause of the exponential gap is our failure to predict the cadence of exponential change, the secondary cause is our consequent failure to adapt to it. As the speed of change increases, our society remoulds itself at a much slower pace. Our institutions have an in-built tendency towards incrementalism.

The way fast technology runs ahead of our slow institutions is nothing new. This is arguably one of the key, inevitable consequences of innovation. In the 19th century, breakthroughs in industrial machinery catapulted the British economy into a position of global dominance. But there was a hitch. There was a 50-year period where British GDP expanded rapidly, but workers' wages remained the same. Those with capital to invest in new machinery did well initially, because it was technology that was driving the growth. It took decades for workers' wages to catch up.

The problem was not just wages. The industrial revolution eventually meant more wealth, a longer life span and a better quality of life for all. But for most labourers, the first effect of industrialisation was an often unwelcome change in working conditions. Starting in the late 18th century, technology moved millions of people from fields, farms and workshops into factories. In the 1760s, before the industrial revolution really took off, the average British worker toiled for 41.5 hours a week. By 1830, this had risen to 53.6 hours – an extra hour-and-a-half each day. By the 1870s, when the Victorian economy had largely completed its transition from agriculture to industry, the typical worker was nudging 57 hours a week.

One way to make sense of the social problems brought by industrialisation is as a gap – between the speed of technological and social change, and of institutional and political adaptation. The state's failure to regulate working practices reflected the preoccupations of a pre-modern, agrarian and aristocratic elite; Britain had a modern economy, but a distinctly pre-modern political order. As the former British Prime Minister Tony Blair told me, “There was a time lag between the change and the policymakers catching up with it”.

Like our Victorian forebears, today's institutions face the conundrum of keeping up with rapidly changing technologies. But this time, the gap will grow bigger and more quickly. In the Exponential Age, radical change takes place not over decades, but over years – sometimes months.

It's worth pausing to reflect on what these 'institutions' actually are. The word conjures up a sense of solidity: an imposing police building, a large church, or the towering UN headquarters in New York. Yet institutions are not buildings. They are more than that. They are the systems that govern our everyday life, our individual actions in public and in private, and how we relate to each other.

In sociological terms, institutions refer to all the lasting norms that define how we live. For our purposes, I consider an institution to be any kind of arrangement between groups of actors in society that helps them relate to each other. Some institutions might be obviously “institutional” in nature: a business is an arrangement between employees, bosses and owners; a state between its citizens and the machinery of government. Others, like the notion of the rule of law, or the body of international agreements and domestic legislation that makes up intellectual property law, are not groups of people, but are institutional nonetheless. And not all institutions need to be so formal. There are the habits and practices that guide our behaviour. Such unwritten rules can have as high, perhaps higher, levels of adherence than written ones.

All these institutions have something in common. They are largely not cut out to develop at an exponential pace, and in the face of rapid societal change. In the most extreme cases, they're not cut out to adapt at all.

Take one of the most institutional institutions in history: the Catholic Church. Nearly 2,000 years old, claiming divine origin, the Church is one of the longest standing organisations in existence. In 1633, it got into a dispute with the astronomer Galileo Galilei and his conclusions about the structure of the Solar System. Yet it was not until 1979, 346 years after the astronomer was condemned to house arrest until his death – and 22 years after the Sputnik satellite orbited the planet – that Pope John Paul II ordered a Papal commission into Galileo’s conviction. Some 13 years later the commission and the Pope overturned the findings of the 17th-century inquisition that denounced Galileo’s subversive ideas.

This is an extreme example. But if few institutions are as slow to adapt as the Catholic Church, almost none are particularly fast-moving. This is the inverse of the difficulty of predicting exponential change: even when our predictions are correct, our institutional responses can still be lackadaisical. Take the Kodak Corporation. In 1975, a Kodak engineer called Steve Sasson put together a device, about the size of a toaster, that could save images electronically. A 23-second process transferred the images to a tape where they could be viewed on a TV screen. At the time it was astonishing. Personal computers barely existed as a category. Kodak sold 90 per cent of the photographic film in the US and 85 per cent of its cameras. A couple of years later, Sasson was awarded US Patent 4131919A for an “electronic still camera”. He figured at the time, by extrapolating from Moore’s Law, that it would take 15-20 years for digital cameras to start to compete with film. His estimate was bang on. But Kodak, with a two-decade head start, did not clutch the opportunity. “When you’re talking to a bunch of corporate guys about 18 to 20 years in the future, when none of those guys will still be in the company, they don’t get too excited about it”, recalls Sasson.

Kodak did go on to develop a range of digital cameras; in fact, it was one of the first to market. The company even recognised the power of the internet, buying Ofoto, a photosharing site, in 2001, nine years before Instagram was founded. But institutional knowledge, the established consensus about what their business was about, held them back from attaining anything near Instagram’s reach. Kodak’s execs viewed Ofoto as an opportunity to sell more physical prints (their old business), not as a chance to connect people via photos of their shared experiences. Indeed, the rationale to hamstring Ofoto by linking the site to Kodak's traditional film-and-prints business was the same wrong-headed decision the firm took more than a decade earlier with the digital camera. Yet the market for film cameras had peaked in the late 90s – the world had moved on, propelled by technology. Kodak had not: the institutional memory was strong. Eventually, the firm hit the buffers, struggling and eventually folding in 2012. Kodak sold Ofoto as part of its bankruptcy process.

There are exceptions to this institutional slowness, of course. On occasion, institutions can lend themselves to very rapid change. Wars and revolutions help. It took less than a year to create the International Monetary Fund after it was proposed at the Bretton Woods Conference in July 1944. The visceral shock of World War II and the need to find a solid base for international co-operation provided the impetus to establish many other institutions, like the United Nations and the General Agreement on Tariffs and Trade. These are, in the language of institutional theory, moments of "punctuated equilibrium". But on one level, this caveat merely reveals the sheer scale of the disaster needed to shock institutions into rapid change. In the absence of a catastrophe, institutions tend to adapt more like the Catholic Church than the United Nations.

Moore's law (black) vs growth in computation for AI (blue) over 5 years

PUT TOGETHER, THESE two forces – the inherent difficulty of making predictions in the exponential age, and the inherent slowness of institutional change – and you have the makings of the exponential gap. As technology takes off, our businesses, governments and social norms remain almost static. Our society cannot keep up.

In the early years of the 21st century, this exponential gap was relatively trivial. The demise of Kodak is not a society-threatening problem. A company's failure to adapt is bad news for its shareholders. The company goes bust, leaving customers with defunct products and staff with fond (or not so fond) memories. But this is not the end of the world.

However, as the Exponential Age takes off, the gap will pose an ever more existential problem. At the turn of the 2020s, exponential technology has become systemically important. Every service we access, whether in the richest country or the poorest, is likely to be mediated by a smartphone. Every interaction with a company or our government will be handled by a machine learning algorithm. Our education and healthcare will be delivered through AI-enabled technologies. Our manufactured products, be they household conveniences or our houses, will be produced by 3D printers. Exponential technologies will increasingly be the medium through which we interact with each other, the state and the economy.

For the people and companies who understand this shift, the exponential gap creates a huge opportunity. Those who harness the power of exponentiality will do much better than those who don't. This isn’t simply about personal wealth. Our rules and norms are shaped by the technologies of the time – those who design essential technologies get a chance to shape how we all live. And these people are in the minority. We are witnessing the emergence of a two-tier society – between those who have harnessed the power of new technology, and those who haven't.

What is to be done?

On the one hand, we can close the gap by slowing the rise of the top line. But this is no mean feat. For one thing, this acceleration is already baked into the structure of our economies. The process through which technologies improve and accelerate is not centrally controlled. It emerges from needs of individual firms, and is met by a coalition of players across the economy. The virologist benefits from a faster genome sequence, and so seeks out better electrochemistry, faster processors and quicker storage for genomic data. The householder wants more efficient solar cells, the farmer more precision methods to fertilise her crops. The Exponential Age is a near-inevitable consequence of human ambition.

Even if we could slow the pace of exponential change, it's not clear this would be desirable. Many of the most urgent issues of our time can only be solved with exponential technology. Tackling climate change, for example, requires more exponential technology, not less. In order to decarbonise our economies, we will need to rapidly shift to renewable sources of energy, develop alternatives to animal proteins for food, and scale building materials that have a zero carbon footprint. What's more, figuring out how to deliver good quality healthcare, education, sanitation and power to the poorest billions of the planet is another problem that technological innovation can address. The expensive (and resource intensive) way that the developed world achieved those outcomes is not practical for poor countries at a time of environmental crisis.

So putting a brake on the development of technologies is hard to justify. On the other hand, we can close the gap by making the lower line rise faster. That means equipping our social institutions – from governments, to companies, to cultural norms – to adapt at pace. It would allow us to harness the power of exponentiality, and the rules and norms that can shape it, for the needs of our society.

This is an urgent need. In the Exponential Age the institutions that govern our economies will cease to be fit for purpose. New technologies will clash with our existing expectations, rules and systems. We need radical thinking to prevent the exponential gap eroding the fabric of our society.

Azeem Azhar's book, Exponential, *from which this extract was taken will be published in September (Penguin Random House)*

More great stories from WIRED

- 🌎 Sign-up to WIRED's climate briefing: Get Chasing Zero
- Hidden Arctic caves hold secrets about our past and future
- Fed up of your job? You could join the Great Resignation
- Zoom dysmorphia is following us into the real world

- On YouTube, you’re never far from a dying kitten
- Too much Google? Try these Google search alternatives
- A radical plan to treat Covid’s mental health fallout
- The 100 hottest startups in Europe in 2021
- 🔊 Subscribe to the WIRED Podcast. New episodes every Friday