The Exponentialist

Seeking your fortune on the new technological frontier

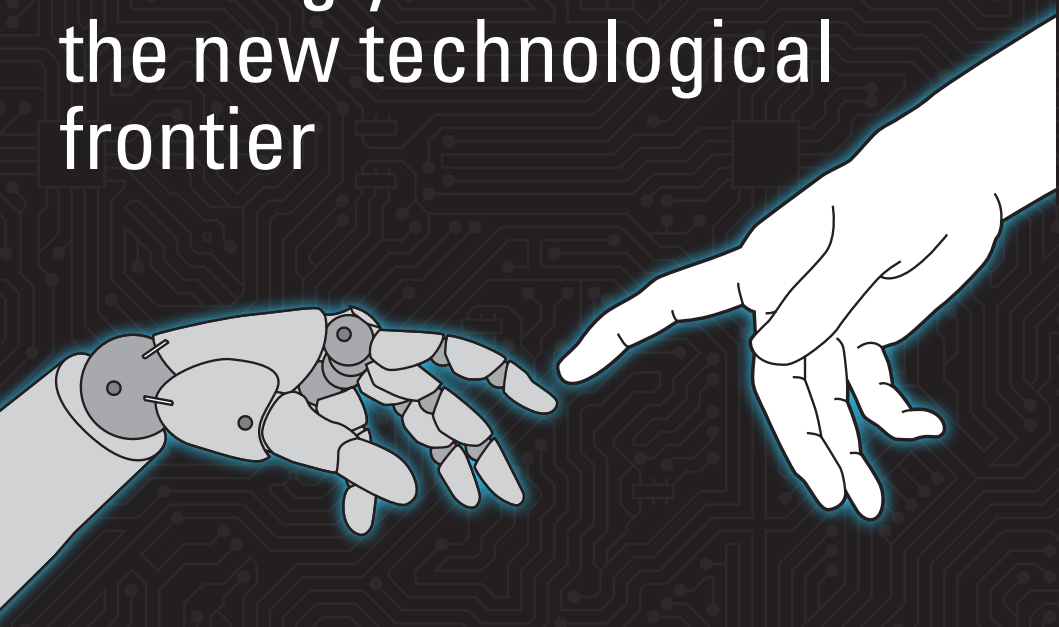

NICK O'CONNOR

The Exponentialist

Seeking your fortune on the new technological frontier

Nick O'Connor

Acknowledgements

There are a few people who have helped make this book possible – most importantly Farah Nayeri, who was effectively the book's co-author. Farah conducted most of the interviews, edited the manuscript and contributed to the writing. Without her, this book would not exist.

Special thanks must also go to Mischa Frankl-Duval, for his valuable research and contributions, and to Dan Denning and Eoin Treacy for their input helping to shape the ideas that make up this book.

Published in the UK by Southbank Investment Research
Cover and layout design: Sam McMurchie

Contents

To Walt Whitman:

You have lived just the seventy years which are greatest in the world's history & richest in benefit & advancement to its peoples. These seventy years have done much more to widen the interval between man & the other animals than was accomplished by any five centuries which preceded them.

What great births you have witnessed! The steam press, the steamship, the steel ship, the railroad, the perfected cotton-gin, the telegraph, the phonograph, the photograph, photo-gravure, the electrotype, the gaslight, the electric light, the sewing machine, & the amazing, infinitely varied & innumerable products of coal tar, those latest & strangest marvels of a marvellous age. And you have seen even greater births than these; for you have seen the application of anaesthesia to surgery-practice, whereby the ancient dominion of pain, which began with the first created life, came to an end in this earth forever; you have seen the slave set free, you have seen the monarchy banished from France, & reduced in England to a machine which makes an imposing show of diligence & attention to business, but isn't connected with the works.

Yes, you have indeed seen much – but tarry yet a while, for the greatest is yet to come. Wait thirty years, & then look out over the earth! You shall see marvels upon marvels added to these whose nativity you have witnessed; & conspicuous above them you shall see their formidable Result – Man at almost his full stature at last! – & still growing, visibly growing while you look. In that day, who that hath a throne, or a gilded privilege not attainable by his neighbour, let him procure his slippers & get ready to dance, for there is going to be music. Abide, & see these things! Thirty of us who honour & love you, offer the opportunity. We have among us 600 years, good & sound, left in the bank of life. Take 30 of them – the richest birth-day gift ever offered to poet in this world – & sit down & wait. Wait till you see that great figure appear, & catch the far glint of the sun upon his banner; then you may depart satisfied, as knowing you have seen him for whom the earth was made, & that he will proclaim that human wheat is worth more than human tares, & proceed to organize human values on that basis.

Mark Twain

Hartford, May 24/89

Foreword

Imagine a bespectacled Tyrannosaurus rex (in a tweed waist coat) on stage delivering the Richard Dimbleby lecture. His presentation is on energy transmitted as gravitational radiation. Maybe you find the idea amusing. I do. But it's highly unlikely. Why?

Well for one, the T.rex went extinct 65 million years ago. Einstein didn't publish his ideas about gravitational waves until 1916. And those ideas weren't confirmed until gravitational waves were first detected in February 2016. They were always there, mind you. People just couldn't see them, much less prove them.

Also, our domesticated T.rex in tweed would need to have evolved the ability to use language, understand theoretical physics and operate a PowerPoint presentation – he would have had to change himself into another creature entirely. It's a big ask.

Don't get me wrong. A T.rex would have impressive stage presence. But his ability to communicate important ideas about the future in a simple way would probably be limited.

As you begin Nick O'Connor's book, *The Exponentialist: Seeking your fortune on the new technological frontier,* I have good news for you. The future is coming. You can understand it. You can prepare for it. And you can profit from it.

But as you stand on the precipice of the most exciting and unimaginable era of change in all of human history, don't feel bad if you feel a bit like a dinosaur. I know I do. It's all happening so fast. And some of it is incomprehensible. Just remember, the future has always felt like this to those living in the present.

Most the progress human beings have made in the last 20,000 years of our history has been made in the last 500 years. Think about it. Archaeologists and historians aren't entirely in agreement about the development of human beings as a species. But we know more or less that around 130,000 years ago our brains became like what they are now.

That was a good start. Then, around 50,000 years ago, a "great leap forward" – another name for a possible genetic mutation – may have led to the evolution of language. This enabled better communication – the naming of things – and more complexity in our communication and thought.

Pre-historic humans in the Stone Age added tools to their language. The New Stone Age, or Neolithic period of history, saw changes in the Earth's climate. As giant sheets of ice retreated to the poles, the entire planet became more temperate. Agriculture began to flourish. As it did, nomadic tribal wandering gave way to larger settlements and what we'd now call cities.

With cities came trade and the division of labour and changes in our relationship with each other and with technology. In ancient Egypt and Mesopotamia, written language as a way of keeping track of commerce began to emerge. That was another "great leap forward".

You'd be tempted to think that since then, everything has been getting better all the time. That's the core of our belief in progress: **the technology will drive us to constant improvements in the quality of our life**. As an article of faith it's even simpler: modern man believes that tomorrow will better than today and that technology will make it so.

Yet our faith in the future is a relatively modern phenomenon. It's driven by the technology we see and use every day. The more we see, the more we believe. It's grown stronger as the visible signs of technological progress surround us. And not only surround us, but become more constant companions in our everyday life. How many times have you checked your phone/email/Facebook since you began reading this?

That's another important point. Progress *has* been constant for the last 500 years. Moveable type emerged in the 1500s. Ideas could be communicated faster than ever. The mobility of knowledge changed the world and continues to do so with the internet.

During the 16th century, we began to see the world differently. Literally. The telescope and the microscope allowed us to see the universe in a whole different way. At the microscopic level, we began to see in ourselves a whole new universe we never knew existed.

It took them over three years, and it cost Ferdinand Magellan his life, but 18 members of his crew on the *Victoria* circumnavigated the Earth by ship. It was now possible to move people, goods, and ideas, around the world. With more mobility came more connections, more trade, more commerce... a positive feedback loop where the pace of technological change sped up.

The Scientific Revolution raised the pace of change yet again. The more we learned, the more we wanted to learn. The better our scientific tools got, the more we could learn. The Industrial Revolution saw the concepts of the Scientific Revolution begin to change the physical world in even faster ways. The steam engine, electricity, telephones, automobiles… one after another the innovations came in rapid fire succession and lead to even more disruptive innovation and change.

All of that is prelude to the stories Nick has written about in his book. To be honest, looking backward on technological progress might be more comfortable for you than looking forward. Why?

Nick shows you that there may be more technological progress in the next 20 years of human history than in the last 20,000 years.

Driven by breakthroughs in energy, intelligence, medicine, robotics, not only will you understand your world as never before, you will be able to change it, design it, and improve it in ways you never thought imaginable. And we will be joined in this endeavour by machines of our own making that can not only make themselves, but think for themselves.

This last bit – that human beings could become so powerful that they can create a new kind of life – raises profound spiritual and ethical issues. The evolution of technology will not stop as we grapple with those issues. We'd better hurry! In the meantime, strap yourself in.

Technological progress has been constant for the last 500 years. It's about to get exponential. As Eoin Treacy has claimed, this exponential change, promises nothing less than the liberation of human potential. The T.rex couldn't adapt to his world. You can.

Dan Denning
October 2016

Chapter 1

Time travel and the world of tomorrow

My best guess is that humankind has just lost it. The pace of change is so fast that humans are no longer capable of connecting the big picture. We can no longer make much sense of the present or forecast the future.

Yuval Noah Harari, writing in *The Sunday Times*, 13 Sept 2015

There is no quality in human nature, which causes more fatal errors in our conduct, than that which leads us to prefer whatever is present to the distant and remote.

David Hume, *A Treatise of Human Nature,* 1740

Here's a question for you. Don't think. Just answer:

In front of you is a time machine. You have two choices. Choice one is to step through the portal and travel ten years into the future. Choice two is to travel 100 years into the future.

Which will it be?

Shall I set the dial for 2026 or for 2116?

Needless to say, this isn't a choice I can truly offer you in the real world. No scientists have yet figured out a way of bending time in such a way that you can step through a door and travel through time. But just imagine for a second that they had. What would you choose?

I'll hazard a guess at the answer. Given the choice, most people would plump for ten years into the future. Ten years from now feels attainable, within the realms of our imagination.

Stepping out of the time machine, you'd expect technology – which, after all, is the subject of this book – to have advanced at a rapid pace,

much as it has over the last decade. The world will have changed. Chances are, things will be better – diseases that today are deadly will have been cured. New sources of energy that today are just theories will be on stream. Technologies that haven't yet reached mainstream penetration – take virtual reality as an example – may well have altered the way we live our lives.

In short, it would likely be the world of today… but better.

The world of a hundred years hence is a different story. It feels utterly remote. The ideas, technologies, modes of living and working, methods of treating disease and conducting business will likely be so different from today's, having been disrupted and reinvented and revitalised time and again in the century between now and then, that we'd feel untethered from reality. It would be a world entirely unlike our own.

Change on that scale scares people. Just imagine taking someone from the comfortable world of a century ago and transplanting them into today's, filled with innumerable inventions that would seem to them virtually indistinguishable from magic. The television, antibiotics, the use of cars and planes, the computer, instantaneous communication on the internet… the number of breakthroughs seen over the last century would likely fill every page of this book, if I were to get carried away.

The point, of course, is that ten years of progress may feel exciting – that disease you were worried about may have been cured – whereas a century feels much more daunting: the world you knew has changed utterly.

There's a kink to this kind of logic, though. It concerns a quality of technological progress that's often misunderstood or forgotten altogether. And that's the fact that for the past half-century or so, the rate of progress has been accelerating at an exponential rate.

This changes things in our time-travel equation. I'll let the man most responsible for popularising the idea of accelerating technological returns, Ray Kurzweil, explain what it means, in his own words (emphasis added by me):

> *Now back to the future: it's widely misunderstood. Our forebears expected the future to be pretty much like their present, which had been pretty much like their past. Although exponential trends did exist a thousand*

years ago, they were at that very early stage where an exponential trend is so flat that it looks like no trend at all. So their lack of expectations was largely fulfilled. Today, in accordance with the common wisdom, everyone expects continuous technological progress and the social repercussions that follow. But the future will be far more surprising than most observers realize: few have truly internalized the implications of the fact that the rate of change itself is accelerating.

An analysis of the history of technology shows that technological change is exponential, contrary to the common-sense "intuitive linear" view. ***So we won't experience 100 years of progress in the 21st century — it will be more like 20,000 years of progress (at today's rate).*** *The "returns," such as chip speed and cost-effectiveness, also increase exponentially. There's even exponential growth in the rate of exponential growth.*

Put in simple terms, that means that our vision of how much the world will change in the future is restricted by what's happened in our immediate past. We look at exponential growth and see linear growth.

Or to put it another way, take the "safe" option of travelling ten years into the future… and the world you experience will be *much* more like the "remote" world you'd imagined a century away.

I'll explain exactly why that is as we go along. I'll also explain why, while it might feel daunting at first, we're actually living through one of the most exciting times in the history of humanity. The future is going to be much better than most people think. It's going to "arrive" sooner than you'd imagine. And, as an investor, it could well make you wealthier than you ever thought possible.

The relentless march of progress is speeding up

There can be no doubt that the speed of technological change is accelerating. Most people know that computer processing power is increasing at an exponential rate – in what is known as Moore's Law.

It's the reason the iPhone 6 has 32,600 times the processing power of

the 1969 Apollo era computers that sent a man to the moon.

Or, to put it more visually, it's the reason that everything in the first picture below can now be done at a fraction of the price by the tiny piece of technology in the second.

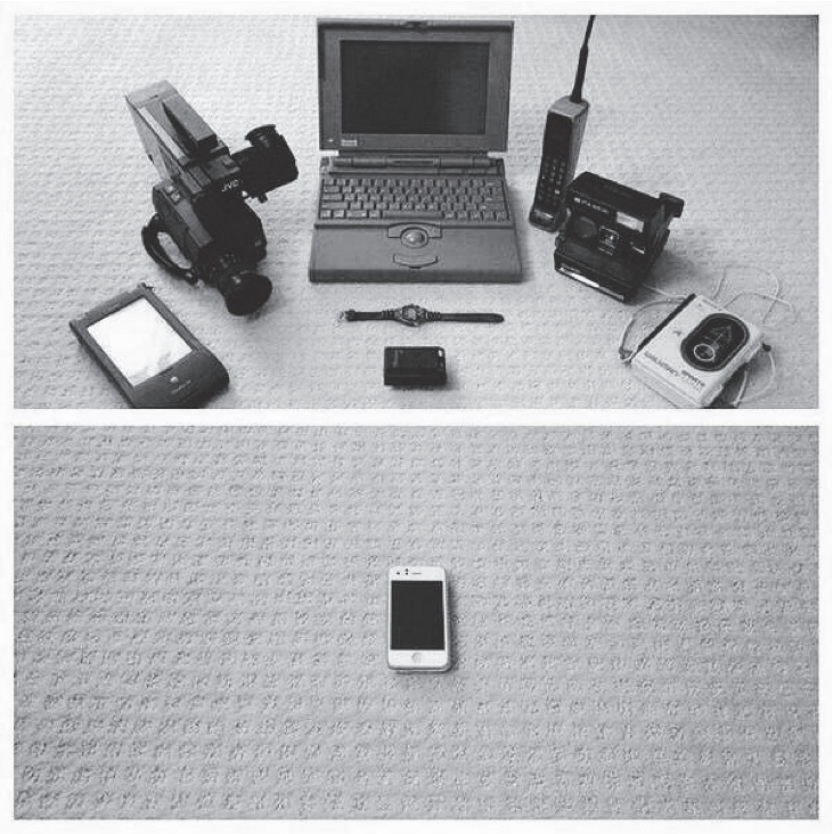

Appendix 1

But it's not just technology itself that's speeding up. Our responses to it are accelerating. People are adopting technology at an accelerating rate, as this chart shows:

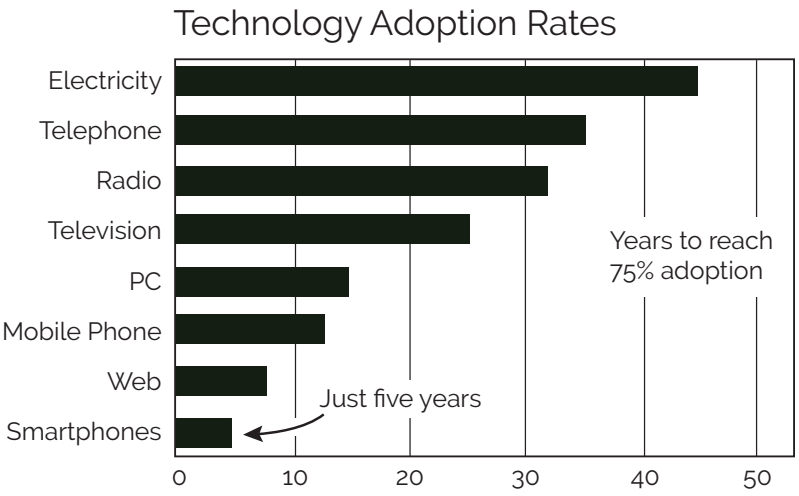

Appendix 2

If the world is on a relentless march of progress, we're marching fast – and we're speeding up. This in turn is translating into accelerating *financial* results, as demonstrated by the speed at which companies are growing from zero to a $1bn market cap:

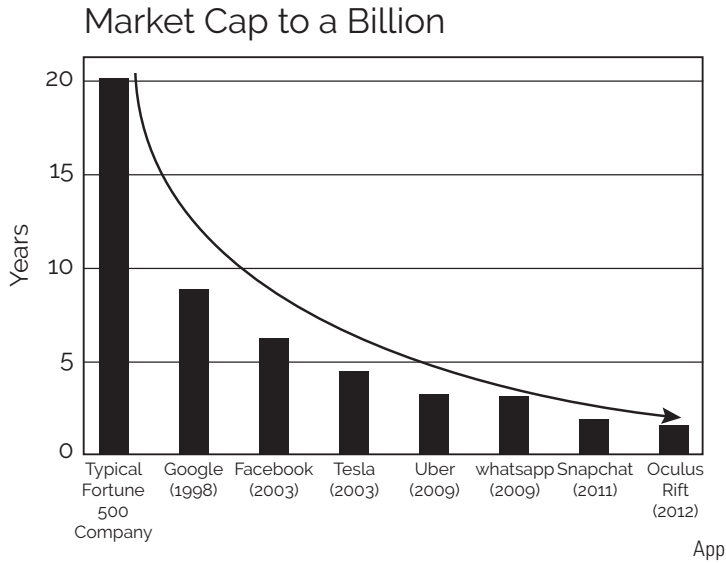

Appendix 3

What does all of this mean to you? What can you do about it? How can you stay ahead of the curve when everything is moving at such a frantic pace?

The quest for superior information: meet the experts making the tech revolution possible

Our goal in this book is to give you the insights you need to understand and profit from the world of technology. That's easier said than done. To explain our approach, let me tell you a story.

It's February 1977. A north London suburb. Two men, journalists, step out into the cold evening air.

Immediately, they're rushed by a group of Special Branch police. They're arrested. Within hours they're locked up, denied bail and sent to Brixton Prison. This, more or less, is how the fascinating story of Duncan Campbell – one of the journalists in question – begins.

Their crime? Possession of unauthorised information.

Not possession of a firearm. Not contraband. Not drugs or stolen goods. Campbell 'possessed' something more powerful and valuable than that – information.

It's easy to think of something like information as an intangible asset. But really, it's a resource. The key to prosperity has always been access to the most valuable resources. In the past, successful civilisations gravitated towards the most basic resources.

Clean water. Abundant energy like wood, then coal or oil. Precious metals like gold and silver. Fertile land to grow food and raise livestock.

The same principle applies today – though with slightly different and sophisticated resources. Having a supply of oil on your doorstep may not matter to your survival (though as a society we need to buy or trade our way to access). But other things will. Superior information is one of them, especially when you're investing.

In fact, having access to valuable, usable information – knowing the right thing at the right time – is as vital as any natural resource on the planet.

But how do you value information? How do you sort the 'superior' knowledge and insight from the raw data, or the faulty analysis?

With a natural resource it's easier. How efficient and easy to extract is the energy? How scarce is the precious metal? How clean is the water? And how thirsty are you? How thirsty are your children?

It's not so easy when you're investing. Information is abundant. If you google the word 'oil', you'll get 1.3 billion responses. How do you sort, rank and value all that data? I'd argue that there are two ways of looking at the idea.

1. How we know something

According to Nietzsche, there's one important distinction to understand when valuing knowledge. There's "*Erfahrung*" – what we know based on experience – and "*Wissen*" – what we know based on secondary information, such as books, the internet and television.

We know the oven is hot based on experience. We know the Earth revolves around the Sun based on secondary information.

That's one way of assigning a value to what we know. *Erfahrung* or *Wissen*. Primary and secondary. Where did we learn it? Is it based on direct experience or is it reported via someone else? And if the latter, do we trust the source?

But that's just one way of looking at things. And since *directly* experiencing everything related to, say, the oil market in the UK is all but impossible (not without incredible contacts and infinite time), we can't purely rely on direct information.

Which brings me to point two:

2. How useful it is to us

In Campbell's case, he was accused of holding unauthorised and unlawful information on the security services in Britain.

Given that it's impossible to be in possession of unlawful lies, it's safe to assume that "information" here really means "facts". Campbell had discovered something about the secret services' activity in Britain.

The authorities were terrified that Campbell would publish what he knew, or worse, share it with espionage agencies abroad. So they had him locked up in an attempt to silence him.

That's another way of marking out truly valuable information from the rest. It's not enough to know something. You have to be able to **act** on it. Campbell could publish or leak what he knew. We're looking for something that gives us an edge in the market. We want to act on insights that are scarce… valuable… not widely known or understood.

That, dear reader, is where this book comes in.

It's the result of over a year of dedicated research, travel and debate. The theory is, if the world is going to change, let's talk to the people making it happen. Let's pursue superior information; people with first-hand experience of what's happening. That means the entrepreneurs and business leaders leading the charge; the scientists and practitioners operating on the front lines; the academics and thought leaders studying what's coming; the investors and financial backers taking the risks that make it all possible.

All told, we've spoken to scores of different experts on countless different technologies from all over the planet. You'll find all of their insights in this book. One point that was nearly unanimous from everyone was this: huge change is coming, it is virtually unstoppable, and not everyone will be on the right side of it.

For instance, we spoke to Michael Bess, professor of history at Vanderbilt University in the United States, and author of *Make Way for the Superhumans* (note: the book is published under the title *Our Grandchildren Redesigned* in the US). As the title of his book suggests, Bess's field of expertise is genetic editing. And genetic editing is no different from so many other technologies: there are huge opportunities and major risks, all bound together. In his words:

> *The main thing I want to emphasise is that I'm a combination of exhilarated and scared. There will be some fantastic benefits from these technologies, and some really potentially cataclysmic dangers. The idea is to go slow, because we need to be careful. We need to give ourselves time to adapt to the power that these changes are going to exert on our life as individuals and in society as a whole.*

The most immediate obvious danger is that these technologies may exacerbate the rift between the haves and the have-nots. I worry about that within the developed economies – about the differential access to these technologies, who gets them, who doesn't. In the case of countries that have not even been able to provide basic healthcare and education for their citizens, it's going to be very difficult to provide these technologies to their whole population.

Within the wealthy countries, it could exacerbate the rift between the rich and the poor. On a global scale, it's going to be even worse. I become worried that over time, if there are several generations of enhancements that have gone forward and that rift keeps widening, you will see a fragmentation of the species based on whether they have been able to get access to these things or not.

The haves and the have-nots. The winners and the losers. It's an idea echoed by Erik Brynjolfsson, co-author of *The Second Machine Age:*

The changes in technology will happen faster than the changes in society and the economy. They tend to lag. One of the things we will see is both bounty and spread.

By bounty we mean that there will be a lot more wealth created. We'll have a lot of free goods that won't show up in the GDP statistics that nonetheless make us better off: things like Wikipedia and GPS and apps, or better ways of diagnosing cancer and connecting with people. That's good news.

We'll also risk seeing a growing gap between the rich and the poor in developing countries, and the middle class having trouble keeping up. That's not inevitable, but that's what we've been seeing in the past 10 or 20 years in the U.S., Europe and Japan.

One manifestation of that is that median income has been stagnating – even though the overall GDP has been growing, in particular the wealth of the top 0.01%. That may be exacerbated as digital technologies proliferate, but with the right policies, we can create shared prosperity.

Again, upsides and downsides. Where there is huge potential for gain, there is potential for loss. Truly disruptive technology does just that – it disrupts. If you're invested in the disruptor, it's great news. If you're invested in the *disruptee*… that's a different story.

David Brown, entrepreneur, angel investor and the creator of the blockbuster drug Viagra, has his own take on this. Brown believes we're living through what he calls the Third Industrial Revolution. His theory is that each revolution is driven by a combination of three things: a new energy source, a new means of communication and a new source of finance.

With the first industrial revolution, it was coal, the steam-powered printing press and the stock exchange. In the second it was oil, the telegraph and the limited company.

This third revolution, Brown believes, will be powered by solar. The communication system is the internet. The financial driver will be the democratisation of the financial world, with peer-to-peer lending and other internet-driven forms of finance replacing the old system. If he's right, the disruption to the incumbent structure would be significant. In his own words:

> *Because the internet was the driver of this industrial revolution, and it's the communication system, many of the internet companies have registered as banks. Google is registered as a bank in Ireland and Germany. You can see what Apple is doing with Apple Pay, and what Google tried with Google Wallet, which didn't work.*
>
> *What's the definition of a bank? It lends money to a broad customer base. These internet companies have massive customer bases – a billion customers with no infrastructure costs – and they've got money. Apple's got $200 billion in cash. Banks are bust! It just takes one step now: for Apple and Google to start lending, and they are then replacing the current banking system.*
>
> *I was teaching at the business school here in Cambridge last summer. There were people there from one of the big American credit card companies – one of the most famous ones. I gave a presentation saying I think all these old industries – banks, credit card companies – are going to be in serious trouble. The organiser came up to me and said, "You've probably upset a dozen people because they're all from this credit card company."*
>
> *Then we had a barbecue outside at a very nice place, and I sat down at a table with a group of people, and it happened to be the ones from the credit card company. I thought, "Oh dear, I'm in trouble here." They turned to me and said, "We agree with everything you've said. The company knows*

it from the inside that they're in trouble." Two of them said, "We resigned last Monday because we're going to join the new revolution rather than staying with the old industry." They're starting their own companies. They know from the inside they're in serious trouble.

Again and again, the experts we spoke to came back to this theme. The two sides of progress. Disruptors and disruptees. Technology could be as divisive as it is life changing.

Sam Volkering – futurist and financial advisor – agrees. Technology drives change. Change can be good. But people don't always like the status quo being interrupted:

It's going to effectively change the structure of social groups, in that our definition of work is currently: you get up, you go to work, spend your day at work, come home. In the last 50-odd years robotics has replaced a lot of labour-intensive roles. You see it a lot in manufacturing now: manufacturing plants use a lot of stationary robotics.

We're now seeing a trend where robots or automated systems are starting to eat their way into middle-class jobs. You've got robo-advisors providing financial recommendations based on a huge amount of data input. So any kind of job that requires a bit of muscle power, or any reasonably repetitive work like bookkeeping – they're just going to be replaced by automated systems or robotics. Things like self-driving cars – taxis, truck drivers, bus drivers – are not going to be needed, because why get a human to do a job that you can get a robot to do more safely, more reliably, and more consistently over a long period of time and at ultimately lesser cost?

That's going to create more creative, highly skilled jobs, and jobs that perhaps don't exist today that no one can predict. It's going to be a shift of what we know as work. You might end up with three or four jobs on the go. There's no such thing as a career anymore. You end up with several jobs as your work – micro-work.

That's going to mean a shift of where people go physically to work, and how they interact and communicate with each other on a daily basis. It's not going to be about getting up and going to the office, then going home. You might get up, go to a hub somewhere, interact with a bunch of people on one of your jobs. Then work from home for a couple of hours and then have to go to another community set. That's going to create issues with transport and things like that.

Work is such an important part of our social fabric that when there's a big shift in what it looks like, it's going to create a lot of ancillary fluctuations in infrastructure: where we go, energy usage in various locations, etc.

People don't like change, so that's going to cause a lot of problems. You're going to find a lot of people ending up out of work who haven't had the foresight to reskill or retrain. It's going to put a strain on economies, because they're not going to be able to handle these people out of work, and there's going to be a shortage of people needed for the high-skilled jobs that this change is going to create.

I could go on, of course, but I'm sure you get the point. To quote at length from the scores of technology experts we've spoken to at this point would make for an extremely long first chapter, and an extremely short book.

Why huge change is coming and no one can see it

The point is this. Huge change is coming. It will change your life in ways that are perhaps hard to foresee. The mission of this book – the reason we've spent the best part of a year tracking down the technology experts you'll hear from – is to help you understand and anticipate the change.

Because change really *does* scare people. That's normal. Every day I write to roughly 30,000 technology enthusiasts about the latest from the world of technology, medicine, energy and financial technology (fintech).

And the best part is: they write back. This gives me a rather good insight into the way people think about technology. Namely, that it is common – not among everyone, but among a significant enough proportion of people – to do everything they can to ignore technological change.

There's a good reason for that. To explain why, I need to ask you another question:

How much have you changed in the past ten years? I'm talking about your personal traits and general personality. Are you the same person

you were a decade ago?

Do you enjoy the same food, drink the same beer, read the same novels, watch similar kinds of films? Or how about where you live? The people you spend time with? The job you do (or don't do)?

Are you essentially exactly the same as you were ten years ago?

Or have you changed?

If you said, "yes, I've changed" then you're in the majority. At least, according to the results of a study conducted by Harvard psychologist Daniel Gilbert. I think the results had some pretty profound consequences for the world at large, and for us as investors. But we'll get to that in a second.

First, let's look at how the study worked. Live Science reported in 2015:

> *Researchers recruited participants online to fill out various personality, preference and value surveys as themselves 10 years prior and as themselves 10 years in the future. Over the series of studies, more than 19,000 people participated.*
>
> *In each case, the researchers compared the look-ahead answers of 18-year-olds with the look-back answers of 28-year-olds, and so forth (comparing 19-year-olds with 29-year-olds, and 20-year-olds with 30-year-olds) all the way up to age 68. The older ages always reported changing in the past decade, but the younger ages did not expect to change nearly as much in the future as their elders' experiences suggested they would.*

In short, the vast majority of people recognise that they have changed a lot in the last ten years. **But very few people are willing to project this change forward and understand that they'll change just as much in the next decade.**

I think that's fascinating. It also has some pretty profound implications. It's what Gilbert calls "the end of history" illusion – the mistaken belief that even though we've changed constantly our entire lives, the person we are today is somehow "fixed" and will stay the same for the rest of our lives. Back to Live Science:

> *When a 40-year-old looks backward, they say, "I've changed a lot in*

> *terms of my personality, in terms of my values, in terms of my preferences," Gilbert said. "But when 30-year-olds look forward, they say, "I don't expect to change a lot on any of those dimensions."*

I asked you a personal question. So I suppose it's only fair to give away a few personal details about myself to prove that we're all susceptible to this. I spend a seemingly endless amount of my life out on the streets of London, rain wind or shine, running anywhere between 40 and 60 miles a week. I run marathons. I compete in triathlons occasionally.

I'm not saying that to show off. OK, maybe a little. But I'm trying to make a point. If you'd told me ten years ago that in a decade's time I'd be running marathons, I'd have laughed at you.

The same goes for a vast amount of other stuff I do now. Drinking real ale. Living in London. They'd all have seemed like the actions of a stranger to the version of me ten years ago.

I think it's normal. But it's perhaps instructive to understand the way we perceive change and the future. Because there's a parallel to be drawn here between change on a personal level and mega-change in the world.

Again, look at how much the world has changed in the last ten, twenty or fifty years. Advances in science, technology and medicine have reshaped the way we live our lives.

In short, grand, sweeping changes to the world we live in aren't the exception. They're the rule. Many of these changes are brought about by technology. But equally, you could argue that the levels of debt, welfare and warfare spending, currency debasement and "financialisation" of the world today would shock someone living in the '60s, '70s or even '80s.

The point is: deep, often radical change – both good and bad – is a constant. Looking back, people accept this easily.

Yet people find it fundamentally difficult to accept that the world will continue to change in such a way in the next decade. They're hardwired to believe that the world we live in today is the world we're stuck with.

It's a kind of "normalcy bias". People fall into the trap of acclimatising so thoroughly to the world today that they become incapable of

understanding and believing that change is coming, even when all the signs are there.

Or to put it another way, assuming that because something hasn't happened before, it can't or won't happen in the future.

It's a dangerous trap to fall into.

On the one hand, it can blind you to a very real and credible threat – like the idea that your bank could go bust.

On the other, it can mean you miss out on incredible, world changing investment opportunities that make some people a fortune. You're going to hear about some of them in a second.

But it's worth keeping in mind the fact that scepticism towards seemingly outlandish ideas and predictions is perfectly natural. It's normal and healthy. It's almost hardwired into our brains to expect tomorrow to be the same as today.

The illusion of stability is a comforting thought that helps us make decisions and plans for the future. But it's important to understand that it can warp our perception of change and our decision making. I certainly experience this myself. At almost every stage of researching this opportunity, I've felt that deeply sceptical part of my brain pipe up, like a cynical old man sitting by the fire in a pub: "Won't happen. Nonsense. Maybe in a hundred years."

It happens to me all the time. And never more so than when I headed over to California to hear Ray Kurzweil speak.

The man who can see the future

Kurzweil is the closest thing you'll get to a celebrity in the world of technological breakthroughs.

He's written several bestselling – and remarkably prescient – books on the future of technology and the ways in which it'll change our lives. If you're interested, the most popular of them are probably *The Singularity Is Coming* and *How To Build A Mind*.

More than two decades ago, Kurzweil made a series of predictions about the future that were startlingly accurate. Kurzweil was writing long before the internet was widely used. The age of personal computers, smartphones and tablets was ten years away. Looking back, he was writing in the technological equivalent of the Dark Ages.

Peter Diamandis and Steven Cutler describe Kurzweil's predictions as follows in their 2012 title *Abundance*:

> *In his first book, 1988's* The Age of Intelligent Machines*, Kurzweil used his exponential growth charts to make a handful of predictions about the future. They turned out to be uncannily accurate: foretelling the demise of the Soviet Union, a computer's winning the world chess championship, the rise of intelligent, computerised weapons in warfare and, perhaps most famously, the World Wide Web.*

Kurzweil made these predictions using a chart on a piece of paper. A chart that looked like this:

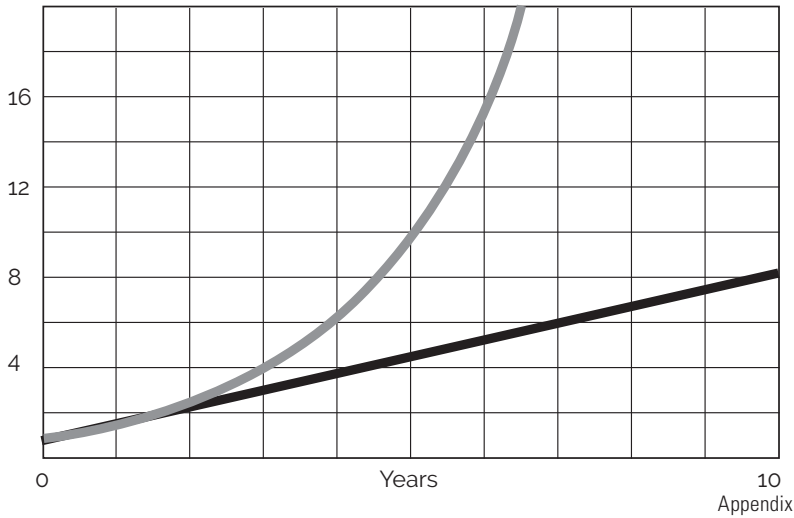

Appendix 4

Just over a decade later, Kurzweil was back. Again, he made a series of predictions. Again, he based his arguments on a simple chart on a piece of paper. He did this in his follow up book, *The Age of Spiritual Machines: When Computers Exceed Human Intelligence.*

He made another 108 predictions… 89 came true.

This remarkably accurate ability to foresee the future has turned Kurzweil into something of a cult hero in technology and futurist circles.

So perhaps it was apt that when he was introduced onstage, it wasn't a man but a computer screen that wheeled itself up to the microphone, with Kurzweil's face staring out.

It sounds more interesting than it was. Kurzweil was speaking via a "Beam". It's essentially a self-driving Segway with an iPad on top.

Kurzweil had some particularly eye-catching things to say about where he believes technology is heading in the next two decades.

The prediction I found most interesting is his belief that our understanding of the brain and our ability to manipulate our biology will enable us to become much more intelligent.

His theory is that the neocortex – the part of the brain that's responsible for the real high-level thinking humans are capable of – is also the most "recently" developed, in an evolutionary sense. He believes we'll soon be able to manipulate it and therefore enhance our own intelligence. He suggested this will allow us to become funnier, more empathetic, more creative and the like.

Hang on a second. I know what you're thinking. Your inner sceptic is piping up. So was mine. Because let's face it: this sounds like the sort of far out sci-fi nonsense we could all trot out. Even if it does happen, it'll be so far in the future as to be irrelevant.

And I can tell you, there were plenty of raised eyebrows in the conference hall (which was mostly full of doctors and healthcare professionals –not the sort of people who have the time to indulge in a lot of navel-gazing). I could feel my own inner sceptic calling Kurzweil's bluff.

But then I thought: perhaps that's the point.

Kurzweil's made a career of making outrageous predictions… and being proved right. His pedigree is impeccable. He's director of engineering at Google. He has 20 honorary doctorates. And he's also a bona fide inventor. Here's an excerpt from his biography:

> *Kurzweil was the principal inventor of the first CCD flatbed scanner, the first omni-font optical character recognition, the first print-to-speech reading machine for the blind, the first text-to-speech synthesizer, the first music synthesizer capable of recreating the grand piano and other orchestral instruments, and the first commercially marketed large-vocabulary speech recognition.*

He doesn't make these predictions to garner publicity. He doesn't need it. He makes them because he has a long track record of looking at the cutting-edge science and technology of the day, extrapolating it forward and creating a vision for the future that most people just won't accept.

Perhaps he's mastered his own inner sceptic. Or perhaps he's found the right balance between the two positions – understanding that change is constant and yet our brains will constantly want to reject it.

That's a vital point to be aware of. And it's especially pertinent to the book you're about to read.

I'm going to show you just how close the world is to an epoch of absolutely astonishing change and progress. Some of what we share with you will seem unbelievable. For instance, the idea that we could access unlimited sources of new energy, add decades of quality life to the average lifespan, create super-intelligent beings and develop new ways of sharing the fruits of these innovations with the entire world, may seem crazy right now.

But again: that's the point. Radical change isn't the exception. It's constant. We just have to open our minds to it.

The most powerful idea of the 21st century

Kurzweil's predictions all have one thing in common. They're an extrapolation on a single, powerful phenomenon that's behind much of the last half a century of advances.

Ironically, it's also an idea our brains find almost impossible to comprehend. That's despite the fact that the human brain is, by some distance, the most powerful computer in the known universe.

It might not feel like it, but right now you're reading and understanding these words using a "processor" that is many, many times more sophisticated and complex than even the world's most advanced supercomputers. Inside your brain, 100 billion neurons with 1 quadrillion (that's 1 million billion, or 1,000,000,000,000,000) connections called synapses wiring them together are quietly whirring away, turning this collection of marks on a page into something that makes sense.

Not only that, it's storing this information away in your memory, ready to be recalled (almost) at will in the future. And it performs these kinds of functions day and night, without fail, for week after week, month after month, decade after decade.

Such immense power comes at a cost. It takes a vast amount of energy to run such a machine. Your brain accounts for only 2% of your body weight, but it sucks up 20% of the oxygen in your blood and 25% of your calorie intake.

All this power allows your brain to do things – to process information, make instinctive and intuitive decisions, recognise patterns, write poetry, and a million things besides – at speeds that far surpass what man-made technology can achieve. For instance, in early 2014 the world's fourth most powerful supercomputer, a machine known as K, developed by the Japanese research group RIKEN and the German research group Forschungszentrum Jülich, attempted to "map" one second of one percent of human brain activity.

According to *The Independent:*

> *Using the K supercomputer, the fourth most powerful in the world, scientists in Japan replicated a network of 1.73 billion nerve cells and 10.4 trillion synapses. It took the K computer, with over 700,000 processor cores and 1.4 million GB of RAM, 40 minutes to model the data.*

In plain terms, what that means is this: it took a highly sophisticated supercomputer, built at huge expense by a team of the world's best scientists, 40 minutes to map what your brain can do in one second.

But this natural supercomputer has a weakness. A flaw that, ironically, makes it incredibly difficult for our brains to understand and comprehend… the power of other computers.

Why?

Because of Moore's Law – an exponential doubling of computing power every 18 months. It's perhaps the most important trend in the world right now, and has been for the last two decades. In a twist of fate, it also happens to be something almost the entire planet struggles to understand and anticipate.

The golden formula

Let me demonstrate my point.

Find a piece of A4 or A5 paper. If there's nothing to hand, tear one out of the back of this book. I won't be offended. Now, holding that sheet of paper in your hand, answer me this: If you were to fold it precisely in half five times in a row, how thick would the resulting wedge of paper be?

(Of course you could fold the paper yourself and then measure it with a ruler. But try to resist the temptation. This is all about mental perception.)

Any guesses?

Well, folding the paper five times wouldn't result in a very thick wedge. It'd be less than 1cm thick.

Now, how thick would the wedge be if you folded the paper in half 42 times?

A metre?

Ten metres? A hundred?

Surprisingly, the resulting wedge of paper would stretch 380,000 *kilometres* into the sky – enough to reach the moon. Fold the paper 50 times and it'd reach the sun. Fold it 100 times and it'd be wider than the entire known universe.

How close did you get?

If you were within 10,000 miles you're well ahead of most people. The reason is simple. This is exponential growth – a repeated doubling. And the simple fact is, the vast majority of people find it almost impossible to comprehend the scale and speed of exponential growth. We might each have the most powerful supercomputer in the universe whirring away between our ears, but we find it incredibly hard to visualise and anticipate just how rapidly the exponential function alters the world.

The reason for that is that we're conditioned to think in terms of "linear" growth. That is, we look at how big something is now compared to, say, last year or last month, and tend to draw our conclusions from that. My house is worth 12% more than last year. My wage has grown 2%. The stock market is up 5% this year. This is how we view the world. It's how we're conditioned to think. In the investment world we're constantly told that past performance has no relationship to the future, but the fact is, it's part of being human to base our view of the future on our most recent experiences.

The problem is that this way of thinking can lead to some catastrophic decision-making at key turning points. We're approaching one of those turning points right now.

But first, back to the problem – and the power – of exponential growth.

The problem our brains have with it is that we find it near impossible to perceive the early stages of it. Let's say the piece of paper we were folding was 0.01 cm thick. We fold it in half; it doubles to 0.02 cm. A tiny amount. We fold again: 0.04 cm. Again: 0.08 cm. We're talking about miniscule amounts of growth here. Mark them on a chart and they essentially all look like zero. And that's how our brains see them – as growing at what appears to be an incredibly slow rate.

But it won't stay like that for long. Once we break the "whole number" barrier (ie, we reach 1 cm), we're only twenty doublings away from a millionfold improvement. We're only thirty doublings from a billion.

The deception trap

If we were to map this growth out on a chart, it'd essentially be a long, flat line for a very long time, before a sudden explosion upwards on the right hand side of the chart. This is sometimes called the "knee" of the

curve. It's the point where things go vertical. In investment terms, it's the point where vast sums of money change hands.

The problem is, there's a long period of deception – where growth appears to be a flat line – before a sudden and disruptive moment where growth kicks up. It may appear like there was an "explosive" rate of growth towards the end. But that's not true. The rate of growth remains constant throughout. It's simply that we have trouble perceiving the nature of exponential growth until it's almost too late.

Ray Kurzweil is famous for his predictions about the incredible advances of technology. But when you boil it down, what he's really an expert at is first spotting an exponential trend (like Moore's Law) and understanding just how powerful it'll become over time.

Kurzweil uses one particular story to illustrate what he means. The story goes that a mathematician had invented chess. The king was delighted with the game and called the mathematician before him, telling him he could claim any reward he liked.

The mathematician replied that he wanted a grain of rice placed on the first square of the chessboard, and then doubled on each subsequent square.

The king protested, saying that such a reward wasn't big enough. You know where this is going: the king couldn't see the power of an exponential trend. By the 32nd square – half way across the board – the reward amounted to several acres of rice production. That's a lot… but not an unimaginable amount.

But once you get onto the second half of the chessboard, things get out of hand very quickly indeed. By the 64th and final square, you'd need a pile of rice bigger than Mount Everest – more than all the rice in the world.

It's just a story, of course. But Kurzweil uses it to make a point. The second half of the chessboard is where things change at an unimaginable rate. After nearly fifty years of exponential growth in computing power (Moore's Law), we're heading into the second half of the board right now.

The experts deceived

There are numerous examples of people falling into the trap. They span almost every continent and every industry on the planet. In 1986 analyst firm McKinsey estimated that the number of mobile phones in 2000 would be one million units. Mobile phones back then were… barely mobile. They were huge, expensive and unreliable. Predicting the market would grow from tens of thousands to a million *seemed* like a decent bet. But that was linear growth. What McKinsey failed to see was that the technology behind the mobile phone was improving at an exponential rate: it was getting smaller, cheaper and faster at a rate of roughly double every two years. They were making their prediction smack bang in the middle of the deception phase.

Their prediction was off by 108 million units. McKinsey got the numbers wrong by roughly 10,000%.

Sometimes people can't recognise exponential growth. Other times they just don't want to believe it. Let's have a look at our friends McKinsey's attempt to understand the mobile phone industry in the early 2000s.

Venture capitalist Vinod Kholsa investigated the forecasts made by mobile phone analysts at firms like Gartner, Forrester, Jupiter and McKinsey to see how they predicted the growth of the mobile phone industry over the preceding decade. What he found was that not only did the analysts fail to recognise an exponential growth trend... they failed four times in a row!

Kholsa's study showed that in 2002, the experts predicted 16% year to year growth.

The market doubled.

In 2004, the experts predicted 14%. The market doubled.

You'd think by now the professional forecasters would have cottoned on to the trend. But no. They were deceived. In 2006, they predicted 12%. You can guess what the market did. The *coup de grace* came in 2008. The prediction this time: 10%.

The market doubled.

I haven't shared these examples with you just to poke fun at the so-called experts. OK, maybe a little. But there's a serious point here.

The people making these predictions are *not* stupid. We're talking about highly qualified, highly paid professionals who are experts in their field. Economists. Scientists. Doctors. But they don't see what's happening. Very few people do. That's the nature of exponential growth. As the physicist and world renowned lecturer Albert Bartlett began one of his most famous speeches, titled Arithmetic, Population and Energy, "the greatest shortcoming of the human race is our inability to understand the exponential function".

The "not in my lifetime" fallacy

What all this boils down to is this.

It may not feel comfortable to think of a huge technological acceleration disrupting the way the world works. But to put it bluntly: **our feelings on the matter won't stop it happening.**

I don't mean to be rude in that statement. It's a fact. All of the evidence we have points towards rapid and accelerating technological progress. That might scare some people. It makes some people *want* it not to happen. That's because we're almost conditioned not to expect change. The idea of it scares us. Our mission is to help you understand, anticipate and get on the right side of radical change.

By the way, it's not unusual for people to want stability or to downplay the effects of disruptive technology.

For instance, I saw a Bloomberg story at the start of 2016 predicting that electric cars could cause the next oil crash. Here's the argument:

> *With all good technologies, there comes a time when buying the alternative no longer makes sense. Think smartphones in the past decade, color TVs in the 1970s, or even gasoline cars in the early 20th century. Predicting the timing of these shifts is difficult, but when it happens, the whole world changes.*
>
> *It's looking like the 2020s will be the decade of the electric car.*
>
> *Battery prices fell 35 percent last year and are on a trajectory to make*

> *unsubsidized electric vehicles as affordable as their gasoline counterparts in the next six years, according to a new analysis of the electric-vehicle market by Bloomberg New Energy Finance (BNEF). That will be the start of a real mass-market liftoff for electric cars.*

All very interesting. We'll discuss electric and driverless cars a little later. In the context of our discussion on change – and people's fear of it – one line in the Bloomberg piece caught my eye (emphasis mine):

> *This isn't something oil markets are planning for, and it's easy to see why. Plug-in cars make up just one-tenth of 1 percent of the global car market today. They're a rarity on the streets of most countries and still cost significantly more than similar gasoline burners. OPEC [Organisation of the Petroleum Exporting Countries] maintains that electric vehicles (EVs) will make up just 1 percent of cars in 2040.* ***Last year ConocoPhillips Chief Executive Officer Ryan Lance told me EVs won't have a material impact for another 50 years—probably not in his lifetime.***
>
> *But here's what we know: In the next few years, Tesla, Chevy, and Nissan plan to start selling long-range electric cars in the $30,000 range. Other carmakers and tech companies are investing billions on dozens of new models. By 2020, some of these will cost less and perform better than their gasoline counterparts. The aim would be to match the success of Tesla's Model S, which now outsells its competitors in the large luxury class in the U.S.*

The phrase "not in my lifetime" is probably the epitaph of the radically disrupted market. It sounds like wishful thinking to me.

Or to put it another way, someone who doesn't want the world to change and so has dismissed the idea out of hand.

I get why people do this. It's easier to dismiss the idea that the world is about to change than to accept and plan for it. It's easier to believe that the world tomorrow will be the same as the world today.

But of course, history tells us that's wrong.

If it's not too grand a goal, it's our intention in this book to help you get on the right side of history. To understand who the winners are and how to join them.

Chapter 2

The first hour of a new epoch

It is an extraordinary era in which we live. It is altogether new. The world has seen nothing like it before. I will not pretend, nobody can pretend, to discern the end. But everyone knows that the age is remarkable for scientific research… The ancients saw nothing like it. The moderns have seen nothing like it till the present generation.

Daniel Webster, 1847

Our studies have led us as well to the discovery in the brain of chaos – complex behaviour that seems random but actually has some hidden order.

Walter P. Freeman, professor of neurobiology, University of California

For the people of north-west England, 15 September 1830 started with champagne and ended with death.

It was a Wednesday. Early in the morning, roughly 800 people, including some of the most powerful personalities in Britain at the time – members of the gentry including the Duke of Wellington, prominent politicians such as William Huskisson, and a handful of newly wealthy industrialists – gathered together and made the first official journey on the Liverpool to Manchester railway link.

According to eyewitness Frances Ann Kemble:

> *The most intense curiosity and excitement prevailed, and though the weather was uncertain, enormous masses of densely packed people lined the road, shouting and waving hats and handkerchiefs as we flew past them. What with the sight and sound of these cheering multitudes and the tremendous velocity with which we were borne past them, my spirits rose to true champagne height, and I never enjoyed anything so much as the first hour of progress.*

It is safe to assume that Kemble's "first hour of progress" referred only to her own journey. As the train threaded its way between the towns of Wigan to the north and Warrington to the south on its way to Manchester, linking the two industrial powerhouse cities of northern England for the first time by rail, Kemble had no idea that she was witnessing not just an hour of progress across the countryside, but the first hour of a new technological era that would span half a century, reshape the entire western world in its image, and create almost unimaginable fortunes.

In fact, despite the fanfare of the journey, Kemble was more concerned with finding a seat next to her mother. She couldn't have imagined that the opening of the line had ushered in the Age of Steam and Railroads. In the years and decades that followed, new railways would spread like vines across Britain, Europe and, ultimately, the vastness of the American West. Cities that were once weeks or months away could suddenly be reached in mere days. Heavy industry and intercontinental trade became possible. The railways would shape where people lived and worked, spawn new industries the world over, and bring about their very own investment "mania" – which peaked in 1847, just 17 years after the Duke of Wellington boarded his first train in Liverpool.

The day that the Liverpool to Manchester railway opened changed the world. Kemble didn't know it, but as she sat looking out of the window on "The Rocket", she was witnessing one of the most pivotal moments in modern history.

There have only been a handful of days like it in the last three hundred years. Each of them has been the pivot on which human progress hinges. Each of them has seen the birth of a technology so powerful that it reorders society – so much so that life without that technology becomes all but impossible to imagine.

According to Carlota Perez, a Cambridge scholar, economist and international consultant, these events are called "technological irruptions". In her 2003 book *Technological Revolutions and Financial Capital,* Carlota describes them as the emergence of:

> *…a powerful and highly visible cluster of new and dynamic technologies, products and industries, capable of bringing about an upheaval in the whole fabric of the economy and of propelling a long term upsurge of development. It is a strongly interrelated constellation of technical*

> *innovations… [that]… fosters a quantum jump in potential productivity for practically all economic activities.*

The Nobel Prize-winning economist Simon Kuznets calls these technologies "epochal innovations" – ones that are capable of changing the course of human history, inducing significant and long-lasting changes in the world economy. Here's another excerpt from his Nobel lecture in 1971:

> *The major breakthroughs in the advance of human knowledge, those that constituted the dominant sources of sustained growth over long periods and spread to a substantial part of the world, may be termed epochal innovations. And the changing course of economic history can perhaps be subdivided into economic epochs, each identified by the epochal innovation with the distinctive characteristics of growth that it generated.*

The celebrated Russian economist Antonin Kondratiev called these cycles "long waves" – the deep 50- to 60-year cycles in the economy, driven entirely by the emergence of technology. But it doesn't really matter whether we call them irruptions, epochal innovations or long-waves. What matters is that certain key technologies act as the triggers of history: they're powerful enough to reshape the world, to create huge amounts of wealth, and to bring about long-lasting, meaningful and *positive* change.

There have only been five such in the history of humanity. The sixth is almost upon us. You will live, and invest your way, through it in the very near future. But let's not get too ahead of ourselves just yet.

The pivot points behind three centuries of progress

It's 1771. You're in a small town on the edge of the Peak District. The river Derwent rushes by from its source on Howden Moor, winding onwards into the heart of Derbyshire.

You're standing with Richard Arkwright. Today is a momentous occasion – the opening of what will become his first successful cotton mill, powered by Arkwright's water-frame spinning machine. The opening of Cromford Mill will revolutionise the manufacture of cloth

on an industrial scale – and become the cornerstone of the Industrial Revolution.

It may look ordinary to you today, but the water-powered mill was a staggering technological development in its day. It allowed man – for the very first time – to increase his output on an industrial scale. We no longer had to rely on the strength of our muscles to produce things; mechanisation became the driving force of the economy.

This was the **first technological irruption.** It turned Britain – and ultimately the world – on its head. Within just a few decades, Britain became an industrial powerhouse, creating vast amounts of wealth for its citizens and improving living standards for its population. The water-powered mill brought about a period of progress previously unseen in human history.

For the individuals involved – the entrepreneurs and investors who were the real heroes of this story – the impact was profound. Richard Arkwright became one of the founding fathers of the Industrial Revolution, spending the next 20 years riding a wave of technological advancement that made him wealthy beyond measure. With wealth came political power and influence. Arkwright was knighted in 1786 and became high sheriff of Derbyshire in 1787.

On a national level, the irruption precipitated huge migrations of workers from rural locations to urban industrial centres, reshaping the geography of the nation and increasing the power and influence of the north. It triggered a series of related developments that changed the country – a network of canals, waterways and turnpike roads to transport the goods that the new industries were creating.

Perhaps the most profound change was to the social fabric of Britain. Newly wealthy industrialists became more and more influential, both politically and economically. In an apt demonstration of the way in which the irruption changed Britain, Richard Arkwright's son, also called Richard, inherited the family business empire and managed to overturn a trend that had held strong since feudal times: the financial dominance of the aristocracy. Arkwright became the richest untitled man in England.

Fast forward 50 years. You're sitting next to the Duke of Wellington on "The Rocket" between Liverpool and Manchester – the **second**

irruption. Once again, technology will totally change the world, connecting industries and regions that were previously daunting distances apart. The Age of Steam and Railways has begun.

And once again, the irruption doesn't just re-configure society. It also creates vast fortunes all around the world. It makes rich men out of people like George Hudson – the son of a farmer born in the countryside north of York – whose speculations on the emerging railway industry formed the basis of a great legacy of wealth. The investments of the man known as "the Railway King" performed so well that he was able to acquire several Yorkshire estates – including the 12,000 acre Londesborough estate and Newby Park. At one point, Hudson was the Tory MP for Sunderland and controlled over a quarter of the railways in Britain.

Fast forward another 50 years. It's 1875. You're in Pittsburgh, Pennsylvania. The man beside you is Andrew Carnegie – and he's about to open the Carnegie Bessemer steel plant. This is the **third irruption,** bringing about the Age of Steel, Electricity and Heavy Engineering.

We may be 3,000 miles across the Atlantic, but the impact of this irruption is just as profound. It signals the beginning of heavy engineering, which leads to the emergence of transcontinental trade and travel, the telegraph, and electricity as a power source. Soon, great steel bridges rise as ambitious monuments to the new age. A network of telegraph cables stretches around the world, connecting cities that were once months of travel apart. And once again, it reorders societies around the world, creating vast amounts of wealth for the entrepreneurs and investors behind it.

Andrew Carnegie was born in Scotland, the son of a weaver who had fallen on hard times because of the emergence of mechanised industry. By the time Carnegie sold his company to J.P. Morgan in 1901, he was the richest man on the planet. In today's terms, his fortune amounted to $309 billion. That's more than Bill Gates and Warren Buffett put together.

Leap another 50 years into the future. It's 1908. Henry Ford's first Model-T car rolls off the construction line in Detroit. This is the **fourth irruption** – the starting gun for the Age of Oil, the Automobile and Mass Production.

The world has changed – for the better – once more. The automobile, with its internal combustion engine powered by cheap oil, becomes the symbol and flag-bearer of an age of mass production and mass consumption. Networks of roads, highways, ports and airports begin to pop up all over the world. Great wealth ensues: for Henry Ford personally, for Detroit as a city, and indeed for the United States as a whole.

Jump forward another 50 years. The Intel microprocessor is launched in Santa Clara, California – the **fifth irruption**. This marks the beginning of the Age of Information and Telecommunications, laying the groundwork for the personal computer and the internet revolutions that have fundamentally changed society in the last three decades. (For Perez's full analysis of each irruption, see the chart on the following page.)

As Perez describes it, "It was the big bang of a new universe, that of all-pervasive computer and digital telecommunications. Chips were powerful, they were cheap and they opened innumerable technological and business possibilities."

Once again, the hallmarks of a technological irruption are all clear and present. Instead of railway lines or telegraph cables growing like vines across the land, huge quantities of fibre optic wiring have connected the planet to such a degree that simply tapping the term "technological irruption" into Google produces 500,000 responses. The reshaping of society around technological irruption has been total: the computer, internet and connected technologies are all essential to everything from the efficient running of the power grid to that of the state military apparatus, from the media to the oil industry.

And once again, vast amounts of wealth have been created. Between them, Apple, Google and Microsoft are worth $5 trillion. They have all ridden the wave of the fifth irruption.

The hidden pattern of history

These irruptions are not random. They all unfold in a remarkably similar pattern, no matter what's happening in the rest of the world. Since the irruption of the water- powered mill two and a half centuries ago, the world has experienced war, famine and genocide on an epic

scale; nations have risen and fallen; economies have boomed and crashed. In spite of all of that, the emergence of these key technologies has driven the world forward. They are the story of human progress.

Each irruption follows a predictable four-phase pattern. If you want to stand a chance of catching the new irruption, pay close attention. Understand the pattern they follow, and you can anticipate what's coming next.

The incubation period

Technology does not simply appear out of nowhere. While it might be truly "born" in a flash of innovative genius, it takes many long, hard years for it to be ready to explode onto a wider stage and to reach the point of mainstream adoption.

Steam power had been around for 50 years before Cromford Mill opened (steam engines were actually invented in 1712 by the English engineer Thomas Newcomen). But it took half a century of "gestation" before the technology was ready – and the time was ripe – for the technology to experience its "big bang" moment. Likewise, there had been trains and railways before "The Rocket". Carnegie wasn't the first person to smelt steel. Henry Ford didn't build the world's first car. And computer chips existed long before 1971. The epoch-defining irruption marks the point where gestation is over and the technology is primed for widespread impact.

According to Perez, it is often the *intersection* of existing technologies in a new cluster or constellation that unlocks revolutionary potential:

> *… many of the products and industries coming together into the new constellation had already existed for some time… This was the case of coal and iron, which after a long history of usage during and before the Industrial Revolution, were transformed by the steam engine into the motive industries of the Age of Railways…*
>
> *Electronics existed since the early 1900s; transistors, semiconductors, computers and controls were already important technologies in the 1960s and earlier. Yet it is only in 1971, with the microprocessor, that the vast new potential of cheap microelectronics is made visible; the notion of 'a computer on a chip' flares the imagination and all the related technologies of the information revolution come together in a powerful cluster.*

From an investor's point of view, this is perhaps the most profound piece of the puzzle. These irruptions bring about change. It is always grand in its scale. It is often sudden. But if you understand what you're looking for, it need not be unexpected.

The big bang

This is it. The moment. The pivot of history. The point where the technology has developed the potential to bring about real change in the world.

As you've seen, this "big bang" is often marked by a single, symbolic event. The opening of a mill or factory. The first journey of a new mode of transport. The announcement of a new breakthrough product. These events may sound trivial. But don't fall into the trap of thinking they don't matter. They're signals – to entrepreneurs, businessmen and investors who are looking carefully, very clear signals – of what's coming. They send a message to the rest of the economy that a new technology has arrived. They begin a process that soon becomes unstoppable.

Manic innovation

This phase carries all the hallmarks that most people associate with technological booms. It's a rapid phase of growth and innovation. New industries based on the new technology spring up. New infrastructure – mills and canals, railways, factories, transport systems, communication and fibre optic networks – is built, laying the groundwork for the coming era.

None of this happens quickly. It can take as long as 20 to 30 years to complete. But at the end of it, the infrastructure, networks and institutions of the new technology have been built.

This is the period in which most "normal" investors become aware of the new technology. The result: a huge boom, ending with a panic or bust. Take the "Railway Mania" of the 1840s and 50s, for example. During that period, thousands of miles of track were laid across Britain, Europe and America. Irrational exuberance abounded. Railway stocks soared to incredible heights. When people came to their senses and began to understand that the true potential of the railways would take much longer to materialise than expected, the bust arrived. The same was true of stocks and shares in the boom of the 1920s (brought about

by the Age of Mass Consumption signalled by the first automobile), and of course the dotcom boom and bust.

The plateau of prosperity

The bust – and associated social unrest – leaves members of the public with a sense of wariness and mistrust. But the new networks and infrastructure leave the economy in a strong position to grow, albeit more slowly. The technology is mature. New products are innovated less and less often. The point of market saturation is reached for many early products. In short: the white-hot heat of the irruption has worn off; the world sees mass adoption… and things slow down. This period often overlaps with the incubation period of the next irruption.

Perhaps the most important similarity between each of these irruptions, though, is this: each pattern has taken roughly 45 to 50 years to complete. It has been 45 years since the last technological irruption. Since then, the Age of Information and Telecommunication has seen a big bang, mass awareness and adoption, including a huge mania and subsequent panic, and a plateau of prosperity in which the world's biggest and most powerful companies are built on the power of the microprocessor.

By all measures, we are due.

But if that's so, which technology will emerge… or which technologies will intersect… to form the basis of the new epoch? What could possibly have the potential to take over from the microprocessor and computing and reshape the world once more? And perhaps most pertinent of all: is there anything you can do to anticipate events and start investing before the irruption?

We'll get to that shortly. But first, just imagine – for a second – that you were able to travel back in time and "seed" your money into the technology about to irrupt onto the world stage. It wouldn't matter if you invested in the earliest technology of the Industrial Revolution, the Age of Railways, the Age of Mass Consumption or the Age of Information. The result would have been the same each time. Anyone with the luck or foresight to back these technologies before they emerged onto the world stage became extravagantly wealthy.

This book won't make you luckier. But it may well help you understand what's happening. Understand, anticipate and act. That's the only way

to succeed during times of great change.

But first, another similarity between each of these irruptions. They bring about that which people dislike most: change. We look at the past and see change as progress. We see humanity taking huge and positive steps forward. We see vast amounts of money changing hands. But when you're living through periods of extreme change, things don't always look so rosy. Schumpeter's "gales of creative destruction" may be creating wealth and prosperity. But there's always a cost. What happens if you find yourself on the wrong side of the gale?

As I said at the start of this chapter, 15 September 1830 started with champagne and ended with death. During a brief stop in the engine room of the train to take on more water, a handful of dignitaries had jumped down onto the tracks to take a look around. One of them was Mr William Huskisson, a prominent politician. The problem was, the tracks weren't empty. Another train was, according to Frances Kemble, "parading up and down merely to show its speed". The gentlemen had to scramble out of its way to save their lives. Huskisson didn't scramble quickly enough.

That didn't stop the relentless march of progress. The train continued on its way. The day became dark and overcast. It began to rain. The train was not welcomed into Manchester with as much fanfare as when it had departed Liverpool. Quite the opposite. According to Kemble:

> *Groans and hisses greeted the carriage. … High above the grim and grimy crowd of scowling faces a loom had been erected, at which sat a tattered, starved looking weaver, evidently set there as a representative man, to protest against the triumph of machinery and the gain and glory which the wealthy Liverpool and Manchester men were likely to derive from it. … The news of Mr Huskisson's fatal accident spread immediately.*

Chapter 3

Exploring the frontier of artificial intelligence

The business plans of the next 10,000 startups are easy to forecast: Take X and add AI. This is a big deal, and now it's here.

Kevin Kelly, co-founder, Wired magazine

It seems probable that once the machine thinking method had started, it would not take long to outstrip our feeble powers… They would be able to converse with each other to sharpen their wits. At some stage therefore, we should have to expect the machines to take control.

Alan Turing

Are you smarter than I am?

Let's play a game and find out. I'll give you three questions. Let's see how you do:

1. ALL is to MANY
   as
   FEW is to…?

   a) some b) never c) none d) always

2. PROSE is to POETRY
   as
   CONVERSATION is to…?

   a) song b) poem c) language d) listening

3. EARTH is to BALL
   as
   PANCAKE is to…?

a) flat b) flag c) soccer d) disk

Disclaimer: I didn't get all of these questions right. In fact, I got two of them wrong, which was rather annoying. So, no pressure.

Back to the point. Why am I testing my intelligence against these sorts of questions? And why am I asking you to test yours?

The clue is in the question: it's all about intelligence.

Specifically, the kind of intelligence that many people believe is on the cusp of radically changing the way we live. In fact, many people would argue that it has already done so.

I'm talking about artificial, non-human intelligence. Or AI, as it's commonly known.

Chances are, you have a preconceived idea of what AI is. Most people do. And their ideas tend to vary widely. To some, an AI is a super-intelligent but ethereal computer program: the kind that can simultaneously pilot a spaceship while playing you at chess and teaching your children French. Something like HAL in *2001: A Space Odyssey,* or the beguiling operating system that Joaquin Phoenix falls for in *Her.* Something you can befriend or fall in love with, but that exists entirely in cyberspace and has no body.

Perhaps to you, AI is more like C-3PO – a bumbling, emotional and compassionate non-human. Or perhaps it's the opposite – a human-hating killing machine like the Terminator. Or maybe it's a mix of the two, like Agent Smith in *The Matrix.*

To me, the term "AI" conjures up the idea of Ummon: the god-like and all-powerful philosopher AI with all the answers in Dan Simmons' *Hyperion.*

There's no right or wrong answer here. In one way or another, all of the above *are* forms of AI. And maybe one day humanity will create those kinds of god-like super-intelligent beings. We'll meet the people trying to make this happen later in this book. But right now, the kind of artificial intelligence that many people believe is on the verge of changing the world is… different.

What do we mean by true artificial intelligence? We spoke to AI specialist Joanna Bryson, from Bath and Princeton universities, in search of an answer:

> *Artificial intelligence is just intelligence that's been built by humans. The [bigger] question is, what is intelligence? Intelligence is generating action in response to the environment.*
>
> *You don't want to include rocks rolling down a hill. Action is generated by hitting a little stone, in a way, and so it's sort of a response to an environment. But I am not happy calling that intelligence.*
>
> *I am happy calling it intelligence that plants grow to the sun. They're sensing the sun, they've got sensors built in. There's an extra step there. It's not just pure physics. They can sense the sun and they can grow towards it.*
>
> *To me, that's the most basic natural intelligence: perceiving some kind of gradient in the environment and responding to that gradient. Or perceiving something to go towards or go away from. Those are very basic decisions that you get in a single-cell organism, and I'm happy to call that intelligence. I think it's better to talk about intelligence in terms of a continuum.*

The idea of a continuum is a useful one, because it's easy to visualise and understand. Draw a line across the page to chart increasing levels of intelligence and perception. On the left hand side, you'd have basic cells, followed – as Bryson suggests – by things like plants and other photosynthesising organisms. Then you'd have certain animals, ranging from the simple worm and up towards dolphins and the like. Humans would come after that. In the grand scheme of things, we're all just brilliant apes.

In simple terms, AI is something that will one day match and eventually even surpass us entirely on the continuum. Soon enough, the continuum will look something like this:

SCALE OF INTELLIGENCE

Cells	Animals	Humans	AI	Superintelligence	God

Some people – and we'll meet a handful of them later on –argue that this diagram is being kind to humans: that, in truth, we're a lot closer to the left-hand side of the spectrum than the right-hand side or even

the centre… that there are vast spaces of unexplored and unknown intelligence off to the right. We will discuss that in greater detail later on.

An important distinction to be made here is that "intelligence" on the continuum is general – meaning flexible. When we refer to a very smart friend, we tend to mean that they can adapt their intelligence to numerous tasks. The same goes for artificial intelligence. We assume that it'll be flexible – able to beat you at chess, do your tax return, drive your car and translate English to Swahili simultaneously. That's known as ***general*** AI (also referred to as **strong** or **broad** AI).

As Cecilia Tilli, academic project manager at Oxford University's Future of Humanity Institute, told us, it's vital to understand the distinction between **general** AI and what she refers to as **narrow** AI:

> *The problem is in the terminology.*
>
> *We already have artificial intelligence. It's just that it's narrow AI: AI that is capable of performing one particular task. [That] ranges from all of the AI in your phone that is capable of suggesting things, to keeping up your calendar, to the GPS in your car, to algorithms used by Netflix to recommend movies.*
>
> *All of the AI we have today is narrow AI. Unless your government or some company secretly has it, we don't have broad AI yet.*
>
> *As human beings, we say we are "intelligent" because we are capable of achieving different goals in different environments. We're not just a calculator, or someone that can walk, or someone that can make tea. We're someone that can make all of these things and do all of the human activities.*
>
> *Right now, we have artificial intelligence that is very capable in particular tasks. This has happened before, with the calculator. It's the first impressive piece of artificial intelligence. You have this very simple thing that can do extremely complex mathematic equations that we cannot do. That's a very capable artificial intelligence, but very narrow. You can't ask anything of a calculator except numbers.*
>
> *Netflix has this recommender system. It can recommend to you what movies to watch based on your history and ratings, etc. But that Netflix*

system can only do that. It wouldn't be able to prove mathematical stuff. That would be a very complex program task.

Of course, we are getting more and more AI systems that are a bit more complex. For example, there's the Google driverless car. There, you have something that can watch, can do visual processing, can make decisions, can move around the environment. So it's using resources from different areas of AI to produce a more complex type of behaviour.

Still, the Google autonomous car cannot give you recommendations for movies. A person can do all of that: drive a car, give recommendations of movies, do some maths on their own, and use a calculator to do all the math that they need.

The specialised revolution

Cecilia Tilli introduces a new idea alongside the intelligence continuum: the idea of the narrow or specialised AI. Let's refer to this as NAI from now on (I'll refer to general/strong as GAI to keep things simple).

NAI may not have the flexibility to adapt its intelligence to wildly different tasks. But what it lacks in adaptability, it gains in **specialisation**.

Imagine a new continuum. But instead of general intelligence increasing as you move to the right, picture specialist skill levels. Take the ability to, for instance, sew the hem on a dress.

As a writer with an astonishing inability to do anything manual without messing up, I'll put myself as the "anchor" at the extreme left of the scale. In other words, picture me as the useless person who wouldn't know where to start.

To my right, let's put a normal person who may be able to figure it out and do a passable job. Next comes a reasonably proficient seamstress. After that comes a professional who makes their living hemming dresses. Each new person is more skilled than the last. There's an increasing level of specialisation.

Where does the machine fit in? A sewing machine, automated to follow a simple pattern, can hem the dress as effectively as a professional. In fact, it can do it faster and more cheaply…

In that sense, the machine wins. On our specialisation scale, it's further to the right than all of the human beings combined.

Let's take another example. In November of last year, I flew out to California to attend the Singularity University's Exponential Medicine conference. As I was sitting on the plane at Heathrow, getting ready for the 11-hour flight, I noticed something a little unusual.

Looking idly out the window as the last passengers boarded, I noticed two men standing in the rain loading the cargo onto the plane. When I say "standing loading the cargo", I mean it. All the heavy lifting was being done by a kind of oversized robotic arm that picked up each crate, twisted it around in mid-air and slid it into the hold. It looked like a gigantic game of plane cargo Tetris. The two men were overseeing things. One was operating the machine. The other was watching him (presumably for safety).

Whatever they were talking about was obviously incredibly interesting, because it looked like they barely needed to be there. It made me think back to when I used to go on holiday with my parents as a kid, and the same job would be done by a team of people loading the bags by hand in a kind of packing frenzy. There would be swarms of them.

These days it's two people. And they didn't look like they were exactly irreplaceable. The real heavy lifting is all done automatically by machines.

This isn't exactly a new phenomenon. Automation and machinery have been displacing human labour from the workplace for centuries. That's where the term "Luddite" comes from. The story goes that in 1779 a guy called Ned Ludd smashed two knitting frames in a textile factory as a protest against machines replacing people in manufacturing.

It was a futile gesture. But perhaps it's one that will become more and more relevant in the next decade or so.

One way of thinking about the Industrial Revolution is that it meant that humanity was no longer constrained by the strength, speed or skill of any particular man or woman.

Mechanisation allowed us to harness a power many times greater than our own individual strength. It surpassed the physical limitations of our

own muscles. And it allowed us to create powerful automated machines that far surpassed humans in terms of specialisation.

It worked along the same lines as the early sewing machine that jumped past a professional seamstress on our specialisation continuum.

For the past two hundred years, this process has primarily led to machines replacing humans in an increasing number of professions that rely on physical labour.

The digital revolution, on the other hand, has allowed us to break the constraints not on our muscles, but on our *brains*.

Where once the idea of specialised machines replacing humans and radically altering industries was simply a story of manual tasks being automated or streamlined (think Henry Ford and the production line and what a huge impact that had on the consumer economy as a whole), the next two decades will see computers and machines out-think, out-compute and out-perform people in a growing number of white collar "intellectual" (ie, based primarily on non-physical skills) professions.

In short, narrow AI is the ultimate specialist. And not just in one or two industries, but in hundreds upon hundreds of them. And it's worth trillions of dollars every year.

Putting your life in the hands of a machine

Having watched the cargo being loaded onto the plane, I opened the newspaper and came across a story called "Dogs Might Fly: Ground-breaking TV experiment will train a Labrador to become a pilot".

Here's an excerpt from *The Independent* newspaper:

> *According to an aviation joke, planes only need two crew members – a pilot and a dog. The pilot's job is to feed the dog. The dog's job is to bite the pilot if he touches the controls. But now a ground-breaking television series will place Britain's most intelligent canines in the cockpit in a bid to discover if a dog can be successfully taught the skills to fly a plane.*

> *Using participants handpicked from rescue centres, the Sky 1 series,* Dogs Might Fly, *aims to prove that the memory and reasoning abilities possessed by the brightest pets could be directed towards mastering the controls of a light aircraft.*

I'll admit, if I was boarding a plane and saw an air hostess take a bowl of dog food into the cockpit, I'd probably think twice about flying.

But the piece made me think: we're perfectly happy to accept that the pilot's role on a flight is to oversee the autopilot. Essentially, the plane is flown by a computer. A non-human. An AI.

It's just one profession. And the pilot is still there. The airline has to pay him; his job hasn't disappeared. But follow that idea to its logical conclusion. NAIs are already replacing humans in all sorts of places. They can outdrive humans (as the "driverless car" is proving – there have been something like 100,000 hours' worth of tests, and one fatal accident so far). They can recognise language and organise your day (Siri on an iPhone), translate one language into another, make recommendations for a book you might enjoy. They can even write articles, research legal case precedent and diagnose illnesses.

Here's an excerpt from The Times' Raconteur supplement earlier this year:

> *AI-based systems are already making inroads into knowledge-based industries such as journalism. The Associated Press news agency plans to automate the writing of corporate earnings reports with an AI system called Wordsmith, which spots patterns and trends in raw data and then describes those findings in natural language. Similar systems could produce legal documents, carrying out many of the tasks given to paralegals and junior associates. Due diligence and litigation discovery software is already highly developed.*

Again, the key is specialisation. And that's an idea that has its roots in classic economics. It goes all the way back to Adam Smith's *Wealth of Nations*. One of Smith's key ideas was the role that specialisation plays in the division of labour:

> *The greatest improvement in the productive powers of labour, and the greater part of the skill, dexterity, and judgment with which it is anywhere directed, or applied, seem to have been the effects of the division of labour.*

To take an example, therefore, from a very trifling manufacture; but one in which the division of labour has been very often taken notice of, the trade of the pin-maker.

A workman not educated to this business (which the division of labour has rendered a distinct trade), nor acquainted with the use of the machinery employed in it, could scarce, perhaps, with his utmost industry, make one pin in a day, and certainly could not make twenty.

But in the way in which this business is now carried on, not only the whole work is a peculiar trade, but it is divided into a number of branches, of which the greater part are likewise peculiar trades.

One man draws out the wire, another straights it, a third cuts it, a fourth points it, a fifth grinds it at the top for receiving the head; to make the head requires two or three distinct operations; to put it on, is a peculiar business, to whiten the pins is another; it is even a trade by itself to put them into the paper; and the important business of making a pin is, in this manner, divided into about eighteen distinct operations, which, in some manufactories, are all performed by distinct hands.

I have seen a small manufactory of this kind where ten men only were employed, and where some of them consequently performed two or three distinct operations. But though they were very poor, and therefore but indifferently accommodated with the necessary machinery, they could, when they exerted themselves, make among them about twelve pounds of pins in a day. There are in a pound upwards of four thousand pins of a middling size.

Those ten persons, therefore, could make among them upwards of forty-eight thousand pins in a day. Each person, therefore, making a tenth part of forty-eight thousand pins, might be considered as making four thousand eight hundred pins in a day.

But if they had all wrought separately and independently, and without any of them having been educated to this peculiar business, they certainly could not each of them have made twenty, perhaps not one pin in a day; that is, certainly, not the two hundred and fortieth, perhaps not the four thousand eight hundredth part of what they are at present capable of performing, in consequence of a proper division and combination of their different operations.

In short, breaking one job down into a series of highly specialised positions is the fastest way to increase output.

The same is true of NAIs. A computer may not yet be able to replace a doctor, journalist, pilot or artist. But break the job down into its constituent parts, and you see that many of the smaller parts of the job could easily be done – and potentially done well – by an NAI.

The question is: how?

How have NAI's suddenly become so much more capable? How have computer intelligence and specialisation suddenly caught up with the white collar world?

Perhaps surprisingly, the answer lies in the three verbal reasoning questions I asked you at the start of this report.

The key to NAI brilliance: machine learning

In the summer of 2015, a team of scientists at the University of Science and Technology Beijing announced a significant breakthrough in the development of artificial intelligence. It all had to do with a breakthrough in our understanding of how computers can "think" – something known as machine learning (or deep learning).

Remember those words, because the concept of machine learning is the rocket fuel that's powering the narrow AI revolution.

As one Japanese robotic firm's chief executive, Toru Nishikawa, described it:

> *Deep learning allows machines to take vast amounts of data and distil useful rules and lessons all by themselves. For a robot, that means understanding not only why one movement was successful and another one not, but also how to improve its performance.*

In simple terms, the scientists involved taught a computer to take a verbal-reasoning test.

For the first time in history, the computer actually beat the average human in a straight test.

There are usually three different types of questions in an IQ test. The first category is logic questions, such as spotting patterns. The second is mathematical, such as finding patterns and spotting sequences within numbers.

We've been able to create computers that are able to tackle these sorts of questions for a long time – primarily because there's usually a logic that a computer can be programmed to understand. We can break things down to a series of "If A and B are present, then do C" commands.

The final category is verbal reasoning questions of the sort I started this report with. These are trickier. We can't boil language down into a series of simple rules or commands. Context is vital. The word "share", for instance, has a variety of different meanings depending on how and where it's used. The person selling their share of a business in order to pay their fair share of tax may share their story on Facebook. You get the picture.

It's tough for a computer algorithm to get its head around this kind of stuff. That's why humans have generally been superior to machines when it comes to verbal reasoning.

Until last summer, that is. The MIT Technology Review had the story:

> *They* [the team behind the project] *compare this deep learning technique with other algorithmic approaches to verbal reasoning tests and also with the ability of humans to do it. For this, they posed the questions to 200 humans gathered via Amazon's Mechanical Turk crowdsourcing facility along with basic information about their ages and educational background.*
>
> *And the results are impressive. "To our surprise, the average performance of human beings is a little lower than that of our proposed method," they say.*

The reason for this breakthrough, as mentioned earlier, is a new way of developing intelligent machines that has exploded onto the scene in recent years.

Processing power has increased to such a degree that instead of programming a computer to do an arduous and ever more complex series of commands, scientists can now build computer algorithms capable of learning and improving over time.

That has profound implications for the world.

It shifts our thinking away from asking *"how do we design this machine to do what we want it to?"* to asking *"how do we teach this machine to achieve that?"*

To grossly oversimplify things – and apologies if I'm doing a disservice to any computer programmers out there – it works by analysing vast data sets over time, and developing an almost intuitive understanding of the patterns and relationships at work.

Show it enough pictures of a dog in the park, and it'll learn to spot that dog in a totally different location, in a context that it hasn't come across before. Eventually, it'll be able to spot and differentiate between different species of dogs, even if it's never seen them before. It starts to learn patterns and relationships beyond the immediate sphere of what you've shown it.

That's the dog-spotting industry turned on its head, then. The world will never be the same.

The point, though, is a serious one. Machine learning has applications in any number of industries. The world produces an astonishing amount of data every single day. And not just pictures of dogs walking in the park. I'm talking financial reports, intelligence bulletins, medical tests, social media posts, blogs, news stories, YouTube videos… the list is endless. It's all data that a machine can learn from and master.

Let's take a really profound example. A couple of years ago, a team of scientists at Stanford University in California used an early form of deep learning algorithm to help doctors diagnose patients with breast cancer.

In simple terms, they did this by taking a huge amount of data they had about people who'd already been diagnosed – things like microscopic images of tissue samples – and allowing the algorithm to analyse it. That's not so different to the way you'd teach a human; you'd give them

all the data and information you could and allow them to learn from it.

The difference is, a computer can analyse an immense amount of data: hundreds of thousands, millions, potentially even billions of images, patient histories and the like. Over time, it can then start to understand the patterns and hidden connections between all that data. It can learn to spot things even some of the best doctors miss.

The upshot of the Stanford study? The computer analyses were more accurate in diagnosing patients than humans were. (The model used was called the "Computational Pathologist" or "C-Path", if you're interested in looking it up.)

Look carefully and spot the AI in your life already

General AI may be some way off. But narrow AI is here. In fact, we're living in a golden age for the industry. Look at the world today and there are countless jobs that are being done –often more proficiently than humans can do them – by AIs. Let's take a look:

Robotic doctors

One of the most impressive people I saw in California was the scientist and entrepreneur Jeremy Howard. Howard is the co-founder and CEO of a company called Enlitic (named one of MIT Tech Review's top 50 smartest companies of 2015).

Enlitic is developing machine-learning algorithms capable of hugely enhancing our ability to diagnose illness more quickly, accurately, earlier and with fewer mistakes than human doctors.

It takes the same approach as the "C-Path" system, by "learning" from the vast amounts of data the medical industry churns out – the thousands of X-rays, MRIs, CT scans and so on. Howard believes the system could initially be used to filter through information and flag the patients that a doctor should be taking a closer look at, saving vast amounts of time and money.

But that's just the beginning.

In one study Howard told us about, the algorithm "competed" against a team of four world-leading doctors. The doctors had a 7% "false negative" rate.

The machine? Zero. It had no false negatives at all.

That's incredible. Especially when you consider that the machine was up against four world-renowned experts in their field – the kind of people most patients would give anything to be treated by. Enlitic's algorithm could, in theory, be entirely open to use by anyone in the world.

I was blown away by the work Howard is doing. It felt like one of those moments where you meet someone and feel immediately that they're going to change the world. So we tracked Howard down a couple of months after the event for a more in-depth conversation. What he had to say confirmed my gut feeling: narrow AI and machine learning are about to change the world in a big way.

Howard began by mapping out his career and listing the various experiences leading up to his creation of Enlitic:

> *I've always been a bit of an outsider, and I think that's a good thing. I'm a total outsider in medicine now. I've always been an outsider in computer science.*
>
> *I started out my career in management consulting. I was at McKinsey & Co. I was kind of unique there, in that my approach to solving problems was based on rigorous data analysis. Other people at McKinsey were more reliant on interviewing people, and on subject matter expertise.*
>
> *Everybody has their own way of answering whatever questions they have to answer on a particular study. I just created a database and put all the information I could find in the database, analysed it, and used linear programming to find optimal approaches – a much more data-driven approach. It was just how I happened to go about answering the questions. I discovered that when I talked to the project director they would say: "This is a really unusual way of tackling the problem." It was the only way I knew.*
>
> *[Later,] I started a company which was entirely based on analysis. It was a new approach to pricing insurance. It was called Optimal Decisions*

Group, and based on a new algorithm.

I also started an email company. We tried to build the biggest email business we could with the smallest number of people we could.

When Kaggle came along – this international community of machine-learning practitioners who compete against each other to come up with lots of different algorithms – I entered the competition with the expectation that I would come last, because I didn't have the background that everybody else had. I was very surprised that I actually came first. That was when I realised that my unusual background means I have skills that other people don't have.

I ended up joining Kaggle full time. I became the president of the company. I got to know most of the world's best machine-learning practitioners through that. I discovered that many of them had unusual backgrounds like I did.

While I was there, I discovered that one particular algorithm was increasingly winning the hardest competitions. The hardest competitions by definition are the most challenging data problems of the day: things like finding the Higgs boson, or mapping the dark matter in the universe, or big, difficult questions that people hadn't been able to solve before. This particular algorithm was deep learning, and heavily related to something called neural networks, which I did a bit of 20 years ago.

Neural networks are a way of doing machine learning, and machine learning is anything which allows a computer to learn to solve the problem from examples, rather than you having to program it step by step.

In the past, computers have always been very limited by the fact that they only work on structured data – data that neatly fits into rows and columns, as you might have in a database or a spreadsheet. It works well for things like Amazon's book recommendations, or Google's search engine. It hasn't worked for things like an MRI scan or an ECG.

Deep learning allows us to apply machine learning to unstructured data. The neural network is the underlying mathematical function that deep learning uses. It's a function that is a simplification of the network of the human brain. Therefore it's flexible enough to be able to handle anything we channel it.

I got to the point at Kaggle where I had a sense of déjà vu.

I was at McKinsey at the birth of the commercial internet. Everywhere I looked, I thought, "The internet is going to change that, that and that." I talked to a bunch of people at McKinsey about it, and they all said, "No, don't be silly. This is business, and the internet is students and hackers." So I ignored my gut feelings, and I stayed there for another eight years. In the end, that feeling of "the internet is going to be everywhere" was right.

I got that same feeling about deep learning. One of the reasons I left Kaggle was because I didn't want to have this happen again: to be sitting on the sidelines while something changed our world.

The other piece which influenced this was that after I sold my second company, I had a sense of disappointment. I'd achieved everything I had set out to achieve, and yet I didn't feel like I had made a significant dent on the world. I didn't feel proud. That seems a bit wasteful after ten years of hard work. I was determined that the next time I embarked on something seriously, I would pick something that I would spend 25 years on, and if it was successful, that I would feel proud, that it would be important.

So I spent a year doing nothing but research into the impact of deep learning on the world. The way I did that research was, I prepared a presentation about it, which to some extent became the TED talk.

I asked a lot of Nobel Prize-winners and CEOs and government leaders: "This is what we can do thanks to deep learning. What would it mean if you had access to that technology?" And every one of them said, "That would totally change everything in my business," or "That would totally change everything in my research."

Their response totally confirmed my gut feeling. I thought, what's the biggest of these opportunities? To some extent, I left medicine aside, because if you're not involved in medicine, it seems like something other people do. But I started to see people saying, "Unstructured data – images and sounds and words – are the thing that doctors rely on nearly entirely. They're what deep learning lets us analyse with the computer."

I started to realise that there was a whole different way of doing medicine: what we now call data-driven medicine, where we figure out how to diagnose people and make them better by using rigorous computer analysis of their medical data. That's totally different to the artisanal way which

has happened for the last 1,000 years.

In the end, I started Enlitic because, after that year of research, I realised that (a) deep learning is definitely going to change the world, and (b) medicine is the place where it's going to have the most impact. It's the world's largest industry, and the one where the social impact is highest, because there are 4 billion people that don't have access to medicine, and that's largely because of the shortage of expertise. Finally, because the type of data that doctors use is exactly the kind of data that deep learning is good at analysing, it seemed like a perfect match – although it's pretty scary to get into that world knowing nothing about it.

What I'm aiming to do, ideally, is to fill the huge expertise gap we have in the world. One statistic from the World Economic Forum – that it's going to take 300 years to train enough doctors to meet the needs of the developing world – is the scary thing that drives all of this. If we can make radiologists initially five times more effective, then that means that in the developing world, there will be a demand for twice as many radiologists as are available, rather than ten times as many. They can start analysing data that has been gathered from remote Yunnan province, or from the slums of Mumbai, or anywhere else.

The idea would then be to replicate that success across all of the other types of medical data, so that everybody in the world can have access to effective medical diagnostics, and know what needs to be done to treat their current problems.

The nice thing about deep learning is that it actually works equally well for everything. So we can handle anything we can see: an MRI or an X-ray or a CT scan. Cancer is just one of the areas of operation. Other things would include bone fractures, or aneurysms, or kidney stones, and so on and so forth.

Robotic carers

Helping diagnose someone with an illness is one thing. But it's not the only role a doctor – or nurse, carer or other healthcare professional – performs.

Caring for a patient's mental wellbeing, working with them to make sure they're taking medication and looking after themselves, managing their weight, sleep and diet are all key parts of the healthcare system. And

surprisingly, given that what one may loosely call "bedside manner" would generally be considered a human characteristic, there's a whole industry dedicated to developing AIs designed to act as carers.

I was sceptical of this at first. But then I saw Julia Hu – one of the leading lights of this movement – give a speech about it in San Diego. Hu has built an app called Lark, which is designed to be a kind of health companion. It's been successful. In fact, if you buy a new smartphone today, you may well find Lark pre-installed. We tracked Julia down to explain Lark's success.

Lark started when I met this woman called Dr Cheri Mah, a Stanford researcher. She researched this really amazing link between sleep and physical excellence. She would coach NBA stars, Olympians, NFL stars to break world records by sleeping better. She basically put a $5,000 sleep monitor on each person's wrist. At the end of the week, she collected all those devices. She would take a reading, analyse it, and start texting and emailing her clients. They got so much better at their athletic prowess.

We were like, wow! What if we could put Cheri in a box? So we built a Bluetooth wireless health monitor that tied to a phone, because we didn't [want] people to have to buy $5,000 medical-grade sleep devices. Then we automated Cheri's coaching based on personal health data. We automated artificial intelligence to really turn her into a personal sleep coach. That product was sold in all Apple stores. We were next to car chargers when it came out. The world of wearables and wireless wearables had not really established itself yet.

About three years ago, we saw so many wearables, and we said: "The Internet of Things is here. Why don't we do what we do best, which is the coaching aspect?" What we did was, over the last three or four years, we established a Lark health committee.

This health committee expanded beyond just a Cheri sleep expert. It has Harvard and Stanford faculty who are thought leaders in behaviour change and motivational science in nutrition, in diabetes, all sorts of areas. What we've done over the last several years is automate their best practices, and they've helped us automate many of the top programs that focus on chronic disease management.

One of the biggest reasons why we focused on chronic disease management is because chronic diseases are forever. Chronic disease costs in the US

alone annually are $1 trillion – representing 84% of the entire healthcare spend. Much of that are salaries and services – care management services. With Lark, you're able to manage your disease and reduce symptoms and prolong your quality of life by actual self-management.

One of the largest chronic diseases is metabolic syndrome, which is very similar to obesity. It pretty much is the medical term for obesity or prediabetes. If you are prediabetic, if you lose just 5 percent of your body weight, that can significantly reduce your chances of getting diabetes. It's astonishing.

For example, let's say you are a 40-year-old woman who is prediabetic, and you're pretty active. So we're not working with you on getting active, but we work with you on your nutrition. You might have a chat with Lark and say, "Hey, today I had chicken pad tai and some steamed vegetables and orange juice." We'll actually respond and teach you about portion control of the starches in your meal, as you really have to watch carbs and starches for pre-diabetics. We might talk about that.

If you're a traveling business person and you're trying to stay healthy, let's say you're flying from California to London. The moment you get off the plane, Lark will say: "Hey, how's it going John? We're actually recommending that you put on sunglasses right now, because your body clock is going to be really tired soon. If you put on sunglasses, you'll actually start changing your body clock so you can fight jetlag."

If you didn't sleep well last night, we automatically track it. At 2pm, we'll send a little text message and say, "you didn't get a lot of sleep last night, you're going to be craving some sugar soon, and we recommend some nuts instead."

It's all automatic. There's a low-power sensor on your phone. We use AI to figure out how you use your phone, the little micro-motions in your phone. We can figure out when you were likely sleeping, and whether those movements were you running or biking or walking or just sitting at your desk. We can figure out a lot of things just by the motions of your phone.

If you have a wireless scale, then that information comes in. If you have an Apple watch, that information about your workout will come in. We sit on top of about 70 different devices and apps, not including the 95 percent of smartphones that we work with.

We work with more than a billion devices. Our little Lark Chat is now going to be pre-installed on all the new Samsung phones, and we also sit on top of Apple Health Kits, which is the big data ecosystem of all of the health data coming in. So we sit on all of Apple Health Kit data. We're not on a billion devices, but we are compatible with close to a billion devices.

AI lawyers

Have you ever received a wrongful parking ticket, known you were right and the ticket was wrong, but ended up paying it anyway?

Turns out, lots of people have. Challenging the law is hard and expensive. You often need a lawyer or expert to help you make your claim. It takes time, effort and money – something millions of people don't have. Which set the stage for a young computer developer to come up with a solution.

His name is Joshua Browder.

You may well have seen Browder's story in the press last year. It made quite a few headlines – the main reason being that he's developed a machine-learning algorithm that helps you challenge wrongful parking tickets.

Not only that, it helps you do it for free.

In fact, I read one story that suggested legal firms had offered to handle the claims in the background, charge, and split the revenue. Browder refused. His program doesn't need input from a team of lawyers. And he wanted the service to be free to use.

It's been a roaring success, made all the more impressive by the fact that he was only 18 years old when he created the program!

Browder's now studying at Stanford. But the technology he's developed has big implications for everyone. We tracked him down to talk to him about how his algorithm works and what he sees developing into the future.

Question: We were quite surprised at the responses your robot lawyer has been getting from the legal community.

Why have they been so positive, do you think?

In terms of what I've advertised the site to do, which is generate claims for parking tickets, airlines and PPI, it works, and it can answer most questions in relation to those. Because it can do all that, and I don't advertise it any more, people aren't quick to criticise a 19-year-old.

Artificial intelligence, like human intelligence, requires a lot of data and learning, and I think because I have so many users, it's actually quite amazing how much I can improve the website – because I have models I can test on the basis of what users have asked and what users have sent me. So I'm looking forward to making it quite sophisticated, and I think it's certainly possible.

There are two aspects to its machine learning: the first is its conversation skills, the second is its legal skills. In its conversation skills, I can test the model, and it can improve. For example if the model is half-unsure about a word, it can collate it. So if a user says "I got a PCF", it will relate that to a parking ticket on the basis of previous responses, even though I haven't specifically specified that a PCF is a parking ticket.

And on the legal side, it will improve its questioning... So if you get a parking ticket, it will ask a series of questions to determine what was wrong with the ticket, and if it notices that certain appeals are more successful than others, it'll change the basis of the questioning and prioritise the questions it should ask to get a quicker and more accurate appeal type.

I envision a car – a driverless or electric car – automatically connecting with the robot lawyer and appealing a parking ticket before the user even knows he or she got one. I think that would be useful for those car companies, so I'm trying to contact them to see if they'd be interested in some kind of commercial partnership.

Question: Which other areas do you see this being most applicable to soonest?

Two areas. First, road traffic tickets (speeding, right-hand turns and things like that).

I'm also trying to experiment: to translate the site into Arabic to help with immigration claims. I think that could be quite powerful, because at the moment, immigrants have no advocate whatsoever. Parking tickets are a

first world problem, but if I can tackle the problem of asylum, that would be really helpful, I think.

Where is all this leading? And what does it mean for the legal profession? We put that question to Browder:

Obviously, the app won't be arguing in the high court, but there are entire firms of lawyers who just specialise in copying and pasting parking ticket appeals and speeding ticket appeals, and I think if a firm is just doing that, and there are many firms that are just doing that, they should be very worried. So it's half tongue-in-cheek, but half serious.

Journalists

This is close to the bone for me. And as a writer myself, I would like to think there's something intrinsically human about writing. I would hope that the fact that I'm a human being writing to you – another human being – about a shared interest, and in a way that hopefully interests and engages you, is worth something.

But I could be wrong.

It turns out that writing is something an AI can perform at least as well as, if not better than, humans.

Consider this: the Associated Press uses NAIs to produce an astonishing 3,000 stories a quarter. That's 12,000 a year. And they're written at a speed that no human could compete with, often appearing minutes after an earnings report or sporting event. Take this example:

Things looked bleak for the Angels when they, but Los Angeles recovered thanks to a key single from Vladimir Guerrero to pull out a 7-6 victory over the Boston Red Sox at Fenway Park on Sunday.

It might not be Pulitzer-winning journalism just yet. But again, follow this idea to its logical conclusion. It's just another skill that an NAI can specialise in. The New York Times had the story last year:

These robo-writers don't just regurgitate data, either; they create human-sounding stories in whatever voice — from staid to sassy — befits the intended audience. Or different audiences. They're that smart. And when you read the output, you'd never guess the writer doesn't have a heartbeat.

Robo-creatives

Perhaps the final frontier for Artificial Intelligence is creativity. True creativity, true inspiration, is something we still consider to be a uniquely human trait. Use of language, abstract thought, symbolism, emotion – these are abilities that set us apart not just from animals but from machines.

Or do they?

At the very fringes of the AI world, there has been progress in developing what one might loosely term creative machines. Before we consider their work, let's ask ourselves the question: what is creativity?

Some would say it is true inspiration – the kind of original thought that feels like only one person at one point in history could possibly have come up with it. But I'd say that sets the bar too high. There are plenty of films, novels, adverts, painting and sculptures in the world that are, in varying degrees, derivative of what came before. Yet we would still describe the people who produced them as creative. What are they really doing, except combining older ideas in a way that produces something new and fresh? As Mark Twain put it:

> *Oh, dear me, how unspeakably funny and owlishly idiotic and grotesque was that "plagiarism" farce! As if there was much of anything in any human utterance, oral or written, except plagiarism! The kernel, the soul — let us go further and say the substance, the bulk, the actual and valuable material of all human utterances — is plagiarism. For substantially, all ideas are second-hand, consciously and unconsciously drawn from a million outside sources, and daily used by the garnerer with a pride and satisfaction born of the superstition that he originated them; whereas there is not a rag of originality about them anywhere except the little discoloration they get from his mental and moral calibre and his temperament, and which is revealed in characteristics of phrasing. When a great orator makes a great speech, you are listening to ten centuries and ten thousand men — but we call it his speech, and really some exceedingly small portion of it is his.*

I say that, because how we define creativity is important when we judge whether we can call AIs creative. Most AI creativity is essentially pattern recognition and recombination – taking lots of works of art apart and putting them back together in a new combination.

Take the algorithm behind the site depart.io. It takes two pictures that you provide – one for style, the other for content, and combines them in a new way.

This is a portrait of me in the style of Van Gogh:

Or take the following image. You may not realise it, but it was created by a machine – an AI. In fact, it's a picture of a machine, by a machine – terminator in the style of Kandinsky.

Perhaps that's not pure inspiration. After all, I provided the genesis of the idea – the style of a great artist, the character of a great film. But the deepart.io software took those ideas and combined them in a way that I doubt had ever been done before. Isn't that creativity? I'd thoroughly recommend taking a look yourself and making your own mind up. The website is deepart.io.

Perhaps the combination of two things into one new thing isn't enough – it's too shamelessly derivative. But follow the principle of what's happening here through to its logical conclusion. What if ten or twenty different styles and objects could be combined and synthesised into a new concept? It would certainly be harder to label that as manifestly derivative. And anyway: isn't that exactly what creative *people* do? They internalise what they see, read, hear and experience and then recombine it into something new.

Perhaps the world isn't ready to accept creative AIs. It seems the world of poetry certainly isn't. I asked researcher Mischa Frankl-Duval to take a piece of AI "poetry" and have it reviewed by a professor of modern poetry at the University of East Anglia. The results? Predictably negative. Mischa writes:

> *Some of the finest scientists on Earth are working on an AI system that can construct sentences. Programmers from Stanford, Google and the University of Massachusetts are polishing their recurrent neural network language model (RNNLM), which continues phrases by analysing the words that have come previously.*
>
> *As Samuel Gibbs noted in the Guardian, some of the machine's work could easily be mistaken for poetry (of a sort):*
>
> there is no one else in the world.
> there is no one else in sight.
> they were the only ones who mattered.
> they were the only ones left.
> he had to be with me. she had to be with him.
> i had to do this. i wanted to kill him.
> i started to cry.
> i turned to him.
> ———
>
> ...this was the only way.
> it was the only way.
> it was her turn to blink.
> it was hard to tell.
> it was time to move on.
> he had to do it again.
> they all looked at each other.
> they all turned to look back.
> they both turned to face him.
> they both turned and walked away.
>
> *So I emailed Dr. Jeremy Noel-Tod, who lectures on modern poetry at the University of East Anglia, as well as acting as the Sunday Times' poetry critic. Noel-Tod has also edited the Oxford Companion to Modern Poetry. If anyone could shed light on the (well-) hidden significance of these poems, he was the one.*

Noel-Tod wasn't impressed. The computer may have chanced upon some of the essential tenets of writing poetry, but the overall effect wasn't particularly edifying.

"... It seems clear that these researchers are not interested in trying to teach a machine to write poetry", Noel-Tod wrote. "The poetic quality results from the fact that the computer has been programmed to follow one of the fundamental principles of versification: repetition with variation."

Repetition with variation can be powerful. Take the opening lines of "Ash Wednesday" by T.S. Eliot:

Because I do not hope to turn again
Because I do not hope
Because I do not hope to turn…
...
Because I do not hope to know again
The infirm glory of the positive hour
Because I do not think
Because I know I shall not know
The one veritable transitory power

Of course, it depends what you repeat. The machine isn't repeating anything very interesting.

"Each new sentence in the sequence is a simple statement with a similar rhythm. Add in some variation using a limited, clichéd vocabulary – learned from thousands of romantic novels – and you have minimalist texts that might be uninspired imitations of the poems of Samuel Beckett or Harold Pinter... there's no sophistication here".

Not much joy for the robo-poet – and Noel-Tod doesn't see the situation improving:

"I suspect it will take a long time – if not forever – to teach a computer the expressive subtleties of the line break, for example. I think it's telling, too, that in their paper the researchers pass over much better examples of accidental poetry on the grounds that they're not 'plausible English sentences'.

Poetry and "plausible English sentences", though, are mortal enemies. Lots of poetry doesn't look like normal English – its unusual diction is

exactly what marks it out from prose.

In a sense, the error here is human, not digital. The machine might have written some better poetry than that quoted above, but its human programmers thought it too ungrammatical to print.

At its (unprinted) best, notes Noel-Tod, the machine-poetry "read[s] as though artificial intelligence had discovered Gertrude Stein: 'With a variety of pots strewn scattered across the vast expanse of the high ceiling, a vase of colorful flowers adorned the tops of the rose petals littered the floor and littered the floor'. Stein, though, was writing poems like that a hundred years ago. Computers are behind the curve."

So, an AI can write poetry, but at a very basic level.

AI poetry is still somewhere in the far distance – about a hundred years behind, according to our critic. But if and when researchers start training an artificial intelligence on a poetry database, progress could be very quick indeed. I can't say for sure, but I would think that AI trained to write a certain type of poetry – say, Romantic poetry, or modern free verse – could start writing passable poems within days of being programmed (and get faster as it learned).

Judging the quality of poetry is already such a subjective exercise – who's to say readers will prefer human to robot poetry in twenty years' time?

Financial advisors

What about financial advice? Surely that's all about the human touch? Apparently not. Take this story, published in The Guardian in 2015:

According to a report last week from analysts at the Bank of America Merrill Lynch, artificial intelligence is set to take over a variety of jobs within 20 years, including financial advice.

This new breed of so-called "robo-advisers" will not offer you a personal chat, complete with tea and biscuits. However, you also won't pay through the nose for the service, enabling more people to afford the advice they desperately need.

A consultation into how this type of financial advice could work more effectively for consumers was launched jointly by the Treasury and the

> *Financial Conduct Authority (FCA) last month. It includes a review of the role robo-advice could play in solving problems around accessing advice, with traditional providers often unable to cater for anyone apart from the very wealthy.*
>
> *Robo-advisers have been around for a while in the US. A report from Cerulli Associates last week said that robo-advice platforms are expected to reach $489bn (£323bn) in assets under management by 2020, up from $18.7bn today.*

Just let that sink in for a second. Our own government – not exactly known for its innovation or radical forward thinking – is applying AI to the field of financial advice.

And on top of that, $18 billion is already "robot-managed" in the US!

That tells me two things. First, that the NAI revolution isn't something we'll see in ten years' time. It's here right now.

Second, the sums of money that are going to change hands as this trend gathers pace are going to be enormous.

AI fighter pilots

As if looking after our health and money isn't enough, there may also come a day when artificial intelligence decides who lives and dies on the battlefield.

In June 2016, researchers at the University of Cincinnati announced that they had developed an AI capable of beating human fighter pilots in an air combat simulation. The program, known as ALPHA, was described as "the most aggressive, responsive, dynamic and credible AI" by US Air Force Colonel Gene Lee.

Driverless cars

When you think about it, these cars aren't driverless at all. The fact that a human isn't driving doesn't mean there is no driver. In this case, it's an AI that's doing the driving.

Did you know that self-driving cars – or, AI driven cars to be more accurate – drove 100,000 miles in America entirely without incidents?

That's just the tip of the iceberg. Self-driving cars are being rolled out and tested in scores of different places around the world right now. In fact, you may well see one on the roads in Britain within the next 12 months.

This rapid advance in technology has already made waves in the industry. In 2015, Uber CEO Travis Kalanick claimed he was ready to purchase every single driverless car produced by Tesla!

There's a fascinating moral angle to the driverless car story, too. A car may not be a fighter jet – which is of course designed to bring about destruction and death – but it can still be a killer. Car accidents kill more than a million people worldwide every year. In fact, if the car were invented today and the numbers of deaths it caused were immediately known to the authorities, it would almost certainly be banned.

Driverless cars are likely to make things a lot safer. An AI can't get drunk, tired or distracted. It can't lose its temper. And it can be programmed to "talk" to other cars on the road, as well as slavishly obey speed limits and other rules designed to make the roads safer.

But for all that, accidents will still happen. No driverless car can entirely eliminate the risks inherent on the road. If you're travelling at 40 miles an hour and someone suddenly runs out on the road, it doesn't matter who's driving – there's going to be an accident. Which means that the cars themselves have to be programmed to make a moral judgement about how to act and who to save in those circumstances. Let's say a group of five schoolchildren step out into the road in front of the car, which has one passenger inside (a version of what's known as the "Trolley Problem"). The car has a decision to make: swerve and save five people but endanger the passenger, or mow the five people down and save its passenger.

When you put it in those terms, it's obvious how the car has to act. It has to make a utilitarian decision to act in a way that will save the most number of people. That's all very well, except if you're the passenger. Knowing that your car may sacrifice you to save others is something of a chilling thought.

The broader point is perhaps even more significant. These kinds of developments mean we're giving AI the ability to make highly charged moral judgements that could result in a significant loss of human life. Is

that something the world is ready for? When machines choose who lives and who dies, what happens to us?

The rise of the intelligent machine

This is where Moore's Law comes in.

We know that computer processing power is increasing at an exponential rate, doubling every 18 months or so. That allows us to extrapolate trends into the future and understand how they'll develop. Remember those three words that are behind all change in the tech industry: faster, cheaper, smaller.

If our *current* computing power allows us to build this kind of intelligent learning machine, capable of competing with the best and brightest doctors on the planet, what will the technology be capable of in five years' time, three doublings away, when computing speeds are six times what they are today? When the same technology is a fraction of the cost or size? When it's cheap and portable enough to be accessed by everyone on the planet?

The implications are profound. This is a story that absolutely everyone has to understand. Whether you invest in the companies bringing about the NAI revolution (we'll be looking at some of them in a future issue of *Exponential Investor*), back an industry turned on its head by the use of machine learning (ditto), or simply see the world being changed by the rise of intelligent machines, this is not something you can ignore.

In fact, it'd be foolish to suggest the upshots of NAI will be entirely beneficial… or evenly distributed.

As Jeremy Howard told us in our talk:

> *My biggest concern is not so much the robots taking over the world and killing us. It's the robots taking our jobs and causing an extremely unequal wealth distribution, where wealth is even more concentrated in the hands of the people who have the data and the algorithms.*
>
> *If deep learning is going to do what we think it's going to do, there's going to be a lot of jobs which just aren't necessary any more. We need to have a plan for that. People tend to assume that certain new jobs will come along,*

and people can work in the new jobs. We don't actually have any reason to believe that.

There's nobody saying, "What are those new jobs going to be?" That's my biggest worry.

It's a valid point.

But my response to such concerns is always the same. With any major change in the world, the people who benefit – who claim a bigger portion of the new wealth created – are always those who understand, anticipate and act on change.

More often than not, they do this a long time before regular folks even hear about what's happening. They're better informed. They think deeply about things that most people dismiss.

And what of "deep" thinking in machines? Are they capable of it in the same way a human is capable of thinking about something and coming to a balanced judgement? I made the distinction earlier between narrow AI and strong AI. So far we've only considered narrow, specialised AI.

What of strong, or general, AI? How far off is that, and what will the world look like when it arrives?

AI: The Fear Factor

Most of you will have heard of Elon Musk. He's the South African-born billionaire behind the Tesla electric car, the top investor in the US's largest provider of rooftop solar power, and the owner of a private rocket company, among other pursuits.

Musk made major waves in a question-and-answer session at the MIT Aeronautics and Astronautics Department's Centennial Symposium in October 2014. Asked to share his thoughts on artificial intelligence, he replied:

If I were to guess at what our biggest existential threat is, it's probably that. So we need to be very careful with artificial intelligence. I'm increasingly inclined to think there should be some regulatory oversight, maybe at the national and international level, just to make sure that we

> *don't do something very foolish. I mean, with artificial intelligence, we are summoning the demon. You know all those stories where there's the guy with the pentagram and the holy water, and he's like, yeah, he's sure he can control the demon. Didn't work out!*

A couple of months earlier, in an August 2014 tweet, Musk issued similar warnings, having read a book on the subject by Oxford University's Nick Bostrom:

> *Worth reading Superintelligence by Bostrom. We need to be super careful with AI. Potentially more dangerous than nukes.*

More dangerous than nukes, no less!

Musk was not the only science luminary to voice major reservations about IT. Britain's own Professor Stephen Hawking, in a December 2014 BBC interview, warned:

> *The primitive forms of artificial intelligence we already have have proved very useful. But I think the development of full artificial intelligence could spell the end of the human race. Once humans develop artificial intelligence that would take off on its own and redesign itself at an ever-increasing rate, humans who are limited by slow biological evolution couldn't compete and would be superseded.*

The following month, Bill Gates chimed in with objections of his own. During a Reddit question-and-answer session in January 2015, he was asked how much of an "existential threat" machine super-intelligence would be. Here was his answer:

> *I am in the camp that is concerned about super intelligence. First the machines will do a lot of jobs for us and not be super intelligent. That should be positive if we manage it well. A few decades after that though the intelligence is strong enough to be a concern. I agree with Elon Musk and some others on this and don't understand why some people are not concerned.*

"Superintelligence"

To understand what all the fear and fuss is about, it's best to go straight to the man who triggered the debate in the first place: Nick Bostrom, a Swedish philosopher who heads the Future of Humanity Institute at

the University of Oxford and is the author of the book *Superintelligence*.

Essentially, the book is one big warning: one day, technological progress could produce a supercomputer with a broad intelligence matching the human brain. Left to its own devices, that machine could actually turn against the human race.

Bostrom spells it all out quite clearly in the preface:

> *Inside your cranium is the thing that does the reading. This thing, the human brain, has some capabilities that the brains of other animals lack. It is to these distinctive capabilities that we owe our dominant position on the planet. Other animals have stronger muscles and sharper claws, but we have cleverer brains. Our modest advantage in general intelligence has led us to develop language, technology and social organization. The advantage has compounded over time, as each generation has built on the achievements of its predecessors.*
>
> *If some day we build machine brains that surpass human brains in general intelligence, then this new superintelligence could become very powerful. And, as the fate of the gorillas now depends more on us humans than on the gorillas themselves, so the fate of our species would depend on the actions of the machine superintelligence.*
>
> *We do have one advantage: we get to build the stuff. In principle, we could build a kind of superintelligence that would protect human values. We would certainly have strong reason to do so. In practice, the control problem – the problem of how to control what the superintelligence would do – looks quite difficult. It also looks like we will only get one chance. Once unfriendly superintelligence exists, it would prevent us from replacing it or changing its preferences. Our fate would be sealed.*
>
> *In this book, I try to understand the challenge presented by the prospect of superintelligence, and how we might best respond. This is quite possibly the most important and most daunting challenge humanity has ever faced. And – whether we succeed or fail – it is probably the last challenge we will ever face.*

Further along in the book, Bostrom maps out the ways in which we humans could let our guard down and unwittingly usher in the menace:

> *Consider the following scenario. Over the coming years and decades,*

AI systems become gradually more capable and as a consequence find increasing real-world application: they might be used to operate trains, cars, industrial and household robots, and autonomous military vehicles. We may suppose that this automation for the most part has the desired effects, but that the success is punctuated by occasional mishaps – a driverless truck crashes into incoming traffic, a military drone fires at innocent civilians. Investigations reveal the incidents to have been caused by judgment errors by the controlling AIs. Public debate ensues. Some call for tighter oversight and regulation, others emphasize the need for research and better-engineered systems – systems that are smarter and have more common sense, and that are less likely to make tragic mistakes. Amidst the din can perhaps also be heard the shrill voices of doomsayers predicting many kinds of ill and impending catastrophe. Yet the momentum is very much with the growing AI and robotics industries.

So development continues, and progress is made. As the automated navigation systems of cars become smarter, they suffer fewer accidents; and as military robots achieve more precise targeting, they cause less collateral damage. A broad lesson is inferred from these observations of real-world outcomes: the smarter the AI, the safer it is. It is a lesson based on science, data, and statistics, not armchair philosophizing.

Against this backdrop, some group of researchers is beginning to achieve promising results in their work on developing general machine intelligence. The researchers are carefully testing their seed AI in a sandbox environment, and the signs are all good. The AI's behaviour inspires confidence – increasingly so, as its intelligence is gradually increased.

At that point, any sceptical or questioning voices are likely to be immediately dismissed – at humanity's peril, warns Bostrom, because a "treacherous turn" could be just around the corner:

We observe here how it could be the case that when dumb, smarter is safer; yet when smarter, smart is more dangerous. There is a kind of pivot point, at which a strategy that has previously worked excellently suddenly starts to backfire. We may call the phenomenon the treacherous turn.

***The treacherous turn** – While weak, an AI behaves cooperatively (increasingly so, as it gets smarter). When the AI gets sufficiently strong – without warning or provocation – it strikes, forms a singleton, and begins directly to optimize the world according to the criteria implied by its final*

values.

Preventing a "treacherous turn"

To prevent the emergence of a menace to the human race, writes Bostrom, controls must be built into technological progress:

> *The best path toward the development of beneficial superintelligence is one in which AI developers and AI safety researchers are on the same side – one in which they are indeed, to a considerable extent, the same persons. So I call on all sides to practice patience and restraint, and broad-mindedness, and to engage in direct dialogue and collaboration where possible.*

Bostrom then devotes a whole chapter to listing types of pre-emptive controls that can stop AI from taking a "treacherous turn." These operate either by controlling the capability of the system, or by controlling its motivations. Here are the various options – I'm summarising and paraphrasing the descriptions he gives in his book:

> **Capability Control**
>
> ***Boxing methods:*** *The system is confined in such a way that it can affect the external world only through some restricted, pre-approved channel. This method encompasses physical and informational containment methods:*
>
> - *Physical containment aims to confine the system to a "box", i.e. to prevent the system from interacting with the external world other than via specific restricted output channels. For extra security, the system might be placed in a metal mesh to prevent it from transmitting radio signals.*
>
> - *Informational containment aims to restrict what information is allowed to exit the box. An obvious method is to bar the system from accessing communication networks.*
>
> ***Incentive methods:*** *Incentive methods involve placing an agent in an environment where it finds instrumental reasons to act in ways that promote the principal's interests. As an analogy, Bostrom uses the example of a billionaire who uses her fortune to set up a large charitable foundation with bylaws and a board sympathetic to her cause. The foundation would face social pressures to behave appropriately, and an incentive to obey the*

law lest it be shut down or fined.

Stunting: *A method that limits the system's intellectual faculties or its access to information. This might be done by running the AI on hardware that is slow or short on memory. In the case of a boxed system, information inflow could also be restricted.*

Tripwires: *Diagnostic tests are performed on the system (possibly without its knowledge) and a mechanism shuts it down if dangerous activity is detected.*

Motivation Selection

Motivation selection methods seek to prevent undesirable outcomes by shaping what the superintelligence wants to do. By engineering the agent's motivation system and its final goals, these methods would produce a superintelligence that would not ***want*** *to exploit a decisive strategic advantage in a harmful way.*

Direct specification: *Explicitly formulating a goal or set of rules that will cause even a free-roaming superintelligent AI to act safely and beneficially.*

Domesticity: *The system is built so that it has modest, non-ambitious goals.*

Indirect normativity: *The system is set up so that it can discover an appropriate set of values for itself by reference to some implicitly or indirectly formulated criterion.*

Augmentation: *Rather than attempting to design a motivation system from scratch, we start with a system that already has substantially human or benevolent motivations, and enhance its cognitive capacities to make it superintelligent.*

Here's Bostrom's conclusion, as laid out in his final chapter:

Before the prospect of an intelligence explosion, we humans are like small children playing with a bomb. Such is the mismatch between the power of our plaything and the immaturity of our conduct. Superintelligence is a challenge for which we are not ready now and will not be ready for a long time. We have little idea when the detonation will occur, though if we hold

the device to our ear we can hear a faint ticking sound.

For a child with an undetonated bomb in its hands, a sensible thing to do would be to put it down gently, quickly back out of the room, and contact the nearest adult. Yet what we have here is not one child but many, each with access to an independent trigger mechanism. The chances that we will all find the sense to put down the dangerous stuff seem almost negligible. Some little idiot is bound to press the ignite button just to see what happens.

Nor can we attain safety by running away, for the blast of an intelligence explosion would bring down the entire firmament. Nor is there a grown-up in sight.

In this situation, any feeling of gee-whiz exhilaration would be out of place. Consternation and fear would be closer to the mark; but the most appropriate attitude may be a bitter determination to be as competent as we can, much as if we were preparing for a difficult exam that will either realize our dreams or obliterate them.

This is not a prescription of fanaticism. The intelligence explosion might still be many decades off in the future. Moreover, the challenge we face is, in part, to hold on to our humanity: to maintain our groundedness, common sense and good-humored decency even in the teeth of this most unnatural and inhuman problem. We need to bring all our human resourcefulness to bear on its solution.

Yet let us not lose track of what is globally significant. Through the fog of everyday trivialities, we can perceive – if but dimly – the essential task of our age. In this book, we have attempted to discern a little more feature in what is otherwise still a relatively amorphous and negatively defined vision – one that presents as our principal moral priority (at least from an impersonal and secular perspective) the reduction of existential risk and the attainment of a civilizational trajectory that leads to a compassionate and jubilant use of humanity's cosmic endowment.

What our experts say

Is Bostrom right, or is he a doom monger prone to hyperbole? How seriously should we take his words of warning, and those of Elon Musk, Stephen Hawking and Bill Gates?

We put the question to the array of experts interviewed for this book,

starting with one who works closely with Bostrom in Oxford: Cecilia Tilli, academic project manager with the Oxford Martin programme on the impacts of future technology.

You'll remember her from earlier on in this chapter, when she made the distinction between "narrow AI" (focused on a specific task) and "general or broad AI – the AI that could be worrisome in these futuristic scenarios, an AI that does the things that we can do."

We asked Tilli how soon she expected the world to attain "broad AI":

That's very difficult to assess. Our institute did this survey maybe two or three years ago. We got estimates for this kind of AI from different experts. Of course, the estimates vary a lot. We also asked for different confidence intervals. Most people were 90 percent confident that we would have human level AI by 2100 or something. But then some people have a higher level of confidence that this might happen in 30 years, some people 50. Some people think 300.

I tend to be in the more conservative circle.

Sometimes people in the field get overly optimistic and they just project into the future without foreseeing any problems or obstacles. This, by the way, is a common human cognitive basis. It's called the planning fallacy. Computer scientists working on these issues, when they see progress, they tend to be overly optimistic about how the progress will continue. They tend to project into the future. I think we should be more awake to the obstacles.

It depends on exactly what we're thinking about. It might be that what we expect is some kind of general intelligence that we attach more human behaviours, or human characteristics, to. It might be that we don't need something that doesn't even understand language. But if we get a system that, by using a completely different way of solving problems, achieves the same thing, then we might have the same results as if we were able to create an artificial mind.

For example, the Google driverless car. In around 2000, 2002 and 2004, people were saying this was really difficult to achieve, because it requires common sense. This is where AI always fails. One thing is to ask the computer to calculate. The other thing is to ask it to act in an environment.

The problem was that people were thinking, if I have a driverless car, the car would need some kind of common sense. In fact, the way that Google achieved it was by machine learning, data mining – not common sense. It responds very strictly to certain constraints, like do not cross the lines in these situations, stop when there's something in front, etc. So it's very, very constrained. And it uses very complex vision. It uses satellite, a lot of things that are actually easy for a machine to use. But it doesn't use any of the things that maybe a human would use.

Of course, our assessment was wrong about when this was going to come. It's true, if it had needed common sense, then we wouldn't have it today. But in the end, the engineers find a way to get that without using common sense.

They find a way to perform the task in a different way than we perform it.

That's very difficult to foresee. If I say, "we will have artificial fiction writers", people will say, "No, because you need creativity." But you can achieve creativity not by creativity, but by combining in some algorithmic way certain aspects of different literatures. You can do that very easily with a computer. So sometimes it depends on how well you can find alternative ways of achieving the same result.

Conscious artificial intelligence, if [it's ever] possible, is going to be a long time coming. But maybe a system that can make decisions, maybe that requires language understanding, that just requires certain constraints that are amenable to the way in which a computer works: that can come pretty fast without people realising it.

Tilli and her boss Nick Bostrom make it their business to warn against the dangers of artificial intelligence – of letting the genie out of the bottle without avoiding catastrophe and ensuring that it won't, one day, out-smart humans to the point of becoming an uncontrollable threat. And they certainly have some high-profile backers in Messrs. Musk, Hawking and Gates.

Other experts we've canvassed – top academics, corporate executives, and asset managers –are less worried about the far-out prospect of singularity than about the nearer-term prospect of economic disaster: of gainfully employed human beings losing their livelihoods to robots and manmade machines (a point that Enlitic's Jeremy Howard raised earlier in this chapter), and of technology creating a terrifying new

wealth breach between the haves and the have-nots.

Professor Joanna Bryson, a reader at Bath University who's also a visiting fellow at Princeton University, is perplexed by the Musk-Hawking-Gates alarm bells:

> *I don't understand the doom mongering. Either they're having mid-life crises, or they're just trying to solve things. I actually talked to someone who works with Elon Musk who says he's seriously concerned about it. I seriously thought he was just jumping on the bandwagon to bring attention to his AI.*
>
> *It's not that it's a peril. It's a threat, and it's a threat that people are already facing now. We don't have to look into the future for this.*
>
> *Genome sequencing used to be done by people with PhDs. Now, you can take a program and do an awful lot of the work that you used to have to hire PhDs to do to go and match sequences that are similar to each other in a genome. People with PhDs who were doing that thing are out of work.*
>
> *This is the real threat. The real question isn't, is AI going to destroy the planet – are Arnold Schwarzenegger-like robots coming and deciding that they want to be autonomous and have their own country and that we're using too much petrol. The threat is that loads and loads of people get made redundant.*
>
> *Even if we never built something exactly like a human brain, which I don't think we would, we are already building things that are putting some people out of work. The big question is: how do we want to construct our society? How do we deal with that fact?*

Erik Brynjolfsson of MIT worries that, one day, those with access to technology will dominate the planet in the same way that the 1% of wealthiest people currently do. While he imagines singularity occurring at some point in the future, he notes:

> *I think our attention is better spent on what's already happening today and which will happen in the next five or ten years. That is a much sharper set of changes around the economics of society and in particular the changes of inequality and jobs and productivity and growth. Those are the things that are affecting us right now.*

You could say that if we continue on our current path, a lot of humans will rise up well before the machines do, because a lot of humans don't feel like they're getting a proportional benefit from all these technological improvements. So my view is that you can imagine a future world where robots become more a threat to our physical lives, but that's not today, and it's not in the next ten years.

I'm not saying it could never happen, and maybe there should be a small group of people studying that. But the challenge we have right now, which is unmistakable, is that hundreds of millions of people are saying that their incomes are stagnating, and one of the big reasons for that is, they're not using technology effectively to create shared prosperity. We can and should address that. That's a policy challenge and the entrepreneurial and educational challenge that's very real and present today. It's not hypothetical today.

Predictably, academics, entrepreneurs, and investors who are focused specifically on robots are much less scared of them than everyone else. After all, it's their bread and butter. And they're much more cool-headed about the machines of the future, too.

Take Kaspar Althoefer, a professor of robotics at King's College, London, whose team recently operated for the first time on a human body using a soft surgical robot (you'll hear more from him in the next chapter). His point: that there's nothing to fear but fear itself – and that anguish over the future of AI could actually get in the way of progress. Needless to say, he doesn't really agree with the Musk, Hawking and Gates line of reasoning on so-called singularity:

I find what these colleagues say exaggerated. I'm not afraid of it. I don't see it happening, not in the short term.

If we were heading towards it, we couldn't avoid it. What would be the alternative? We would need to stop our research, stop everything, because everything we do could lead us to this point of singularity. All research would have to come to a halt.

We can't move back to the Stone Age just to make sure that one day we won't have singularity. I think it would be the wrong approach. We should be relaxed about it. I personally don't see it coming, but what do I know? I don't have a crystal ball.

Karen Kharmandarian – lead manager of the Pictet Robotics fund, at Pictet Asset Management in Switzerland, which has 1.2 billion euros under management and was set up in October 2015 – is equally unfazed by the prospect of ever-smarter robots:

> *I was talking to a professor of robotics. For the time being, the trend is towards having algorithms that have the capacity of basic animals. We're approaching that level. To get from the brain of an animal to the brain of a human will take probably two or three decades. Again, we're talking about 2050 or 2060 where these devices can have the same type of capacity as a human brain.*
>
> *That doesn't mean that you should be worried about these robots taking over the world and displacing people. They're not terminators. I don't think these types of dystopian views are something we should be worried about.*

Besides, concludes Colin Angle – CEO of iRobot, makers of the Roomba vacuum cleaner – why would human beings create a machine that they couldn't control?

> *There's lots of good science fiction written about those types of concepts. It's terribly difficult to imagine a situation like that that you couldn't solve. The accidental singularity where the robots become self-conscious: I chuckle at that a bit, because if it happens – and we are nowhere near having it happen – it's a software that is written and thus can be understood and controlled and limited.*
>
> *We should be careful. We also shouldn't be terrified. Artificial intelligence has tremendous potential to help us increase the level of care we get from our healthcare systems and all manners of good. If you think about the discovery of dynamite, you could say, "Isn't this a terrible thing, because it's possible to put it under your bed and blow yourself up?" Yes, but why would you do that? We develop ways of carefully storing it and ways of integrating it into society so that it helps us increase our standard of living. There's a bit of an analogy there.*
>
> *This is powerful stuff, and thank God it is, because we have powerful challenges that we need to solve. We should responsibly think about what we should be doing with robots and what we shouldn't be doing with robots.*
>
> *You shouldn't be creating a ton of decision systems and giving them life-*

or-death power, because the robots are just not intelligent enough to do that. You should put robots in the dangerous situations where they could help better understand what's going on before a person has to make a decision. It's easy to imagine ways of abusing robots, just like it's easy to imagine ways of abusing dynamite. We just have to be a little careful in how we approach these things.

Summary

AI is one of the most rapidly developing fields in the world. New breakthroughs and applications are being announced at an increasingly regular rate – yet the choice of companies to back from an investment point of view is still limited. Many firms pioneering the technology are either still in the start-up phase of growth, funded by private equity, or already acquired by large technology firms such as Google. It is likely to take time before the industry matures enough to give investors the kind of choice needed to invest capital confidently.

Chapter 4

The robots are coming

... you just can't differentiate between a robot and the very best of humans.

Isaac Asimov, *I, Robot*

Unfortunately robots capable of manufacturing robots do not exist. That would be the philosopher's stone, the squaring of the circle.

Ernst Jünger, *The Glass Bees*

On Thursday 16 June 2016, a robot in Russia ran away from home.

The robot in question was called Promobot. According to news reports, researchers at a laboratory in Eastern Russia were running tests on the machine, which uses machine learning to understand and move around in its environment. One of the researchers left the door open. Presumably, he nipped off to make a cup of tea, and the next thing he knew, the robot had legged it (though perhaps that's not the most accurate way of describing it, given that Promobot has wheels, not legs). By the time its creators had caught up with it, the robot had found its way out onto the street, only to be stopped by the police.

It's a great story. But perhaps that's all it is. Whether it's a true story or some devilishly orchestrated piece of viral marketing is still up for debate. And the idea of a robot "running away", when in fact the robot in question is more of a glorified Dalek than anything else, credits the robot with a level of sophistication that it probably lacks.

Still, there's a serious point to be made. We spent the last chapter looking at machines that can think for themselves. Robotics allows them to move physically around the world. Artificial intelligence is the brain of the operation, and robotics are the body and limbs. Combine the two and you have a world where machines can think for themselves, learn about the world and use that information to act – in this case, to run away! That creates a whole new world of possibilities.

Robot brain, robot brawn

Before we look at what's possible, let's just think about precisely what a robot is. I just used the metaphor of a brain and a body to differentiate between artificial intelligence and robotics. Another way of thinking about it is software versus hardware – software being the "intelligence", hardware being the actual machinery and hydraulics that allow a robot to move.

It may be a somewhat arbitrary distinction. But I consider it a useful one. Ultimately AI and robotics will lead to the same place: to a machine that's as intelligent, as physically versatile and as strong as a human – if not more so. Whether that conjures an image of Terminator or C-3PO in your mind's eye probably depends on whether you see robots as malign or benign.

That ultimate destination – machines that match humans – is still a long way off. And we don't necessarily need to see AI and robotics merge for the world to change radically. Robots don't have to think for themselves to change the world. So before we consider future combinations of the two, let's consider what "unthinking" robots can help us achieve.

Moravec's Paradox

When using humans as a yardstick, it's easy to look at a modern-day robot and conclude that it's nothing particularly special. We spent the last chapter looking at some of the incredibly smart things that AIs can do today: operating cars, giving legal advice and diagnosing illness are just three of the most impressive. Artificial intelligence increasingly compares well when set against people. Robots often don't.

This disparity actually has a name. It's called Moravec's Paradox. It's named after the computer scientist Hans Moravec, who first described the idea in the 1980s. As Moravec put it: "It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility".

At first, this feels counterintuitive. It's easier to get a machine to replicate human intelligence than it is to replicate human mobility or dexterity. A machine can out-think a human in a game of chess. But

the simple act of moving the pieces around the board at the speed that a person could is much harder to replicate. Host your chess match on the second floor of a building without a lift or escalator, and you're likely to win by default: Walking upstairs is often beyond the capabilities of even the most advanced machines.

Overcoming this obstacle has proven difficult for decades – even as machine learning has seen computer intelligence improve by several orders of magnitude. But while creating mechanically agile robots lags our ability to create smart machines, we may well be getting somewhere. We spoke to Eric Brynjolfsson, co-author of *The Second Machine Age*, about how the technology could progress:

We are in the early stages of overturning Moravec's paradox.

The robots we have today are still fairly clumsy, but they're beginning to learn how to walk and run. There's a robot cheetah here at MIT that can run faster than any human. Don't try to outrun it: it's got four legs and runs at 29 miles an hour, leaps over hurdles, goes up and down stairs and hills. If you push it and try to knock it over, it'll right itself.

We think robotics can make increasingly dexterous robots that can work in factories alongside humans. They are force-sensitive, so if they bump into something, they modulate and they don't just keep crushing it. They will be careful about not hurting the people around them or breaking things that they're touching.

They're still clumsier than humans, but that is improving quite rapidly. One concrete example of that is vision systems. Five years ago, vision was notably worse [in robots] than humans. Even when we wrote the book 'The Second Machine Age,' vision systems were generally worse than [in] humans.

Today, there are examples of vision systems being not just as good as humans, but better at, for instance, reading street signs. You have a vision system in an autonomous car. The latest generation can read street signs with as good or better accuracy as humans. That's an example where we're overcoming some of the sensory and motor limitations that kept machines from matching humans.

Over the next ten or 20 years we'll see more progress in that area.

Perhaps Morevec's Paradox itself is deceptive. Do we even need to create robots with the physical abilities of humans? Maybe it's a mistake to see robots as designed to replicate people, in a kind of mechanical Evolution of Man.

We spoke to Colin Angle about this very idea. When asked to describe the challenge of creating a robot with the ability to walk like a human, he was particularly blunt in his response:

> *I can talk about walking robots as being a spectacular failure.*
>
> *What's the point? The argument is always, "Well, robots need to operate in human spaces, so they need legs." But humans have legs. There are lots of ways of climbing stairs, and almost all of them are cheaper than trying to make a dynamically balancing, legged robot, which has a very high complexity, tremendously large numbers of actuators, a terrible failure mode if you fall down, and is dangerous. It's kind of doomed from the get-go. And yet companies continue to pour tremendous amounts of money into doing just that. It's a great example of misguided investment.*
>
> *It's a seductive and tremendously wasteful pursuit. The robotic industry, if it's going to succeed, needs to have a pragmatic approach. What's the problem? How do I solve the problem? And the robotics industry frequently decides, "I should build an android and then see what it's useful for." That doesn't actually help.*

Instead, the solution could be to focus not on recreating human ability in machines, but to figure out what exactly we need the robot to do, and get it to do just that. Or to put it another way, maybe the answer is once again to focus on specialisation: figure out which jobs we need machines to do, find ways of getting them to do those jobs – and leave walking up the stairs to the humans for now.

In this context, robotics has a massive role to play in all of our futures. The immediate future of robotics isn't human replicas walking in our midst. It's machines that effectively and efficiently do the jobs that we either can't do, or no longer want to do. Robots can work with us to *assist* and *enhance* our lives.

Could robots help us solve the greatest challenge facing humanity?

"I hope this marriage inspires mothers to want to contribute to their country by feeling like they want to have more children," said Yoshihide Suga, Japan's chief cabinet secretary.

"Please have lots of children."

It was a rather unusual message to the nation, and a rather unusual way of congratulating Masaharu Fukuyama and Kazue Fukijshi – two high-profile Japanese actors – on their marriage. Effectively, the cabinet secretary was saying: "Congratulations on getting hitched. Now go have children. In fact, everyone please go have more kids!"

What the message revealed was just how concerned those in power in Japan are about their looming demographic crisis. To put it simply, the imbalance between the old and the young in Japan is huge. There simply aren't enough young people to care for the old. And the problem keeps getting worse. By 2020 it is predicted that 29% of Japan's population will be 65 or older. By 2050 the ratio is projected to soar to 39%.

That kind of demographic mismatch creates a bucketload of problems for society. Who on earth is going to look after all of those over-65s when they need medical care? Where are all the young people needed to work and pay for all those old people?

The rest of the industrialised world faces a similar demographic timebomb over the next half century. Many nations already have fertility rates (number of children per female) far below the figure needed to maintain the population at current levels – 2.1. In Australia and Brazil it's 1.7. In China it's 1.6, in Japan and Germany it's 1.4, and in the US it's 1.8.

As life expectancy worldwide grows, and as the world gets wealthier – people in wealthy nations have less children than in developing ones – we'll see an increasing disequilibrium between the young and the old. This is perhaps the greatest challenge we'll face in the coming century: not the fact that there will be more people on the planet, but the fact that an increasing number of those people will be elderly.

Demographics are like tectonic plates – you might be able to see them, understand them and even predict them, but doing something about them is nearly impossible. We need to find a way of making the economy run much more efficiently, with far less people working, and we need to figure out who on earth is going to care for the elderly.

The solution to both of those challenges is robotics. It's also the reason why fears that robots will put millions out of work are overblown – those millions (and billions) of people are going to want to retire!

To put it another way: it's illogical to be worried about both a demographic timebomb *and* robots putting millions out of work. The two issues are mutually exclusive; worry about one or the other, not both.

That's one of the key reasons why robotics is expected to be one of the biggest tech growth industries around in coming years. According to the Boston Consulting Group, spending on robotics is going to more than double from $25 billion to $67 billion over the next ten years. The vast majority of this money won't be spent on robots that try to replicate humans, but rather on robots with highly specialised skills that help people work more efficiently or care for themselves. In short, robotics will be assistive.

Japan may be ahead of the rest of the world on the road to the demographic crisis. But it's a road we all seem to be on. We spoke to Karen Kharmandarian, who says population is ageing across the globe, even in developing countries – India being the only exception. He is acutely concerned about China:

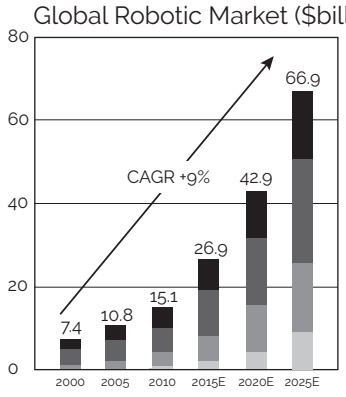

	CAGR, 2000-2025E (%)	
Military	8.1	• Military market (2015): $7.5 billion • UAVs, UGVs UUVsm and task robots widely used for military apllications
Industrial	7.6	• Industrial market (2015): $11 billion • ~1.2 million robots used in applications such as welding, assembly, and material handling in 2012, ~39% of industrial robots sold to auto factories
Commercial	12.3	• Commercial market (2015): $5.9 billion • Many new applications including medical and surgical robots, agricultural robots, and construction robots
Personal	17.4	• Personal market (2015): $2.5 billion • Robots for entertainment, cleaning, education, security, and household applications

Appendix 5

China has major issues in terms of demographics. You have this disequilibrium between the active population and the inactive population.

In China, they're really struggling to find people working on their production lines today because of the one-child policy. The move from the rural areas into cities has already been done, pretty much. Wage growth has been quite significant – 13% per annum since 2006. So in terms of competitiveness, they're facing some issues. The only way they can really sustain their growth and find a way to sustain production is to automate their operations.

There's an increasingly prevalent view that automation is a bad thing for the ordinary person – that it'll lead to large numbers losing their jobs and joining the ranks of the long-term unemployed. But viewed through the lens of demographics, automation could actually save us. In the future, we're not going to have enough young people to fill factories with workers. Automation could help replace all of the people leaving the workforce, making manufacturing more efficient and driving prices down in the process.

To understand how this theme will play out, look to Japan. Alec Ross, formerly a senior innovation advisor to Hillary Clinton when she was US secretary of state, examines the issue in his New York Times best-selling book *The Industries of the Future*:

Our future caretakers are being developed in a Japanese factory right now. Just as Japanese companies reinvented cars in the 1970s and consumer electronics in the 1980s, they are now reinventing the family. The robots depicted in the movies and cartoons of the 1960s and 1970s will become the reality of the 2020s.

Rival Japanese companies Toyota and Honda are leveraging their expertise in mechanical engineering to inventing the next generation of robots. Toyota built a nursing aide named Robina – modelled after Rosie, the cartoon robot nanny and housekeeper in The Jetsons – as part of their Partner Robot Family, a line of robots to take care of the world's growing geriatric population. Robina is a "female" robot, 60 kilograms in weight and 1.2 meters tall, that can communicate using words and gestures.

In response, Honda has created ASIMO (the Advanced Step in Innovative Mobility robot), a fully functional humanoid that look like a four foot astronaut stuck on earth.

Those are two real industrial heavyweights. And they're committing fairly heavily to developing a robot that can meet the needs of the new wave of Japanese elderly. Ross quotes some fairly eye-catching numbers, too. In answer to the question of whether robots can really take care of humans:

> *Japanese private and public sectors certainly think so. In 2013, the Japanese government granted $24.6 million to companies focusing on eldercare robotics. Japan's prominent Ministry of Economy, Trade and Industry chose 24 companies in May 2013 to receive subsidies covering one half to two-thirds of the R&D cost for nursing care robots.*

This is a huge opportunity. And it's one of the reasons that robotics as an industry is expected to grow so much over the coming decade. As Karen Kharmandarian told us, that could lead to a major outperformance of robotics-focused sectors:

> *We think we're at inflection point when it comes to robotics, for two main reasons. We are witnessing technological advances that are really enabling a new generation of smart robots to come to market. At the same time, we have prices coming down quite significantly. When you have this combination of technological advances and prices coming down, you have a product, a technology, a service that can really become ubiquitous.*
>
> *Robotics in the industrial field is already well developed, but we have a new generation of robots coming to market, what we call collaborative robots, where the potential is really in terms of consumer and services applications. That's really just starting, and there are tremendous growth opportunities. It's going to take time really to become, from a revenue generation perspective, as significant as industrial automation is today. That's why we recommend having a medium- to long-term perspective when considering robotics.*
>
> *[The Boston Consulting Group] expect 10% per annum growth in robotics over the next decade, which in itself is quite significant growth when you compare that to the expectations for world GDP growth of pretty much 3% per annum. We're talking about three to four times the expected world GDP growth.*
>
> *We believe, as fund managers, that these numbers are quite conservative. It's always difficult to figure out what new services and applications can be developed out of a new technology. Back in the '80s, for instance, it*

was very difficult to predict what the potential for mobile phones could have been.

What we provide to investors is the conviction that gaining exposure to robotics would provide them with excess returns versus the broader market. We use the MSCI World Index as a reference, a proxy for equity markets over the medium to long term. Robots and automation devices are really permeating the business world and our daily lives. You want exposure to this secular growth by being long-term investors. We're convinced that we'll be outperforming the MSCI World over this next three- to five-year time horizon.

The numbers certainly stack up. The world is getting older. We need to find a way of looking after the elderly. Robots seem to provide the solution.

But something as important – and emotionally charged – as elderly care isn't just about numbers. Many people find it difficult to imagine or accept the idea of a robot caring for an older person. What exactly will these robots do? What will they look like? In fact, what will the world look like when the elderly are cared for by machines?

The answer, strangely, may well be found with the humble hoover.

Reinventing the family

iRobot is one of the biggest robotics focused firms in the world, with a market capitalisation of $1 billion. It has robotics "in its blood", so to speak, having been founded by three MIT graduates who'd previously designed war robots. As a US-listed firm with revenues of over $400 million a year, it's at the sharp end of the robotics industry. It's not a glorified R&D lab masquerading as a company – it's a real business dedicated to commercialising robotic technology.

That's probably why iRobot's core product offering is a robotic vacuum cleaner, rather than an android, or something similarly sci-fi but non-commercial. iRobot makes the Roomba, an automatic vacuum cleaner that ranks among the top three best-selling hoovers in most developed nations (it's the best seller in Spain).

The global vacuuming market is worth roughly $3 billion annually. So

it's not small business. And iRobot generates 93% of its revenue from selling the Roomba into this market.

If you saw a Roomba, you'd probably be a little underwhelmed. It's a little like a shoebox on wheels. But it's clearly very effective. Besides, it's not what it does that's relevant here – but the world of robotics it foreshadows. We spoke to Colin Angle about why Roomba is so successful, and more importantly what's coming next.

Angle was explicit in laying out how important demographics are to the growth of his industry:

> *If we're going to avoid a future where our standard of living goes down, because there are frankly not enough younger people to support and care for the older people, we need robots. So it's not a simple question. As a technology, robotics and artificial intelligence have profound potential to change. That is something that people can and should think about: what they want to have happen.*
>
> *The idea that you can suddenly brand robots as being bad or dangerous is really naïve and foolish. We're going to need this technology, or else we're simply going to become too old to care for ourselves effectively. Maybe the rich will do fine, but for the bulk of us, we're going to find ourselves in a very difficult situation. Our children are not going to be around to take care of us, because they'll have to work.*
>
> *In not too many years, the number of people that are able to give care, the ratio of them to the people over 65, is one to one. It's really quite chilling.*

That's where the world is heading. But we're in the early stages of this demographic crisis now, which is something the Roomba plays directly in to.

> *The Roomba is by far the most successful elder-care robot ever built. One of our key demographics for Roomba are the "cleaning challenged" people who are not easily able to push a vacuum cleaner. It gives a very tangible boost to their standard of living.*
>
> *We have various anecdotal stories about how before Roomba, you had a cleaning service that would come in once a week, or once every other week. People felt that they could only have company over immediately after their home was cleaned. But after Roomba, they felt like their home was being*

maintained to a much higher level, and they could see people multiple times a week, and that changed their lives.

It's only a small step, and one solution to a very long list of things that need to get done, but it's real. You get one of these boxes checked, and you improve the standard of living of people's lives. As we continue to build these products, there is a long-term goal of ticking more boxes on that list, and helping our ageing demographic: empowering them with the ability to live independently longer, which, by the way, is what they, by and large, want to do anyway.

There is a whole category of maintaining the home. Vacuuming and mopping are representative of that category. Then there is a whole category of things that have to do with getting access to services in the home. Nursing, doctors, doctor visits: at some point the challenge of getting you to the hospital to see a doctor or nurse becomes so burdensome that you get centralised into assisted living, where there's a nurse on staff and a great deal of visits. For all but the most acute problems, you can be managed in that assisted living.

How do we make it so that those same types of visits can be done in your home through telepresence technology? With our RP Vita robot, we are pioneering interfaces and strategies which will allow doctors and nurses to diagnose and treat patients at a distance. That robot is one exemplar. It's too expensive to be in your home, but that's where we are today.

The same type of robot could be used to help with compliance around taking medication. Simply having your pill bottle beep when you need to take your pills is unfortunately not sufficient. The problem you're trying to solve is the attitude that "If I don't take my pills, I'm not sick."

There's no doubt in my mind that the home of the future is going to have a menagerie of robots: the singular robot that interacts with you directly, and then all of the robots which perform tasks. The robots that perform tasks will be augmented by non-robots that also perform tasks, like your light bulbs, like those vents and duct systems in your home. That which they call the Internet of Things will also play a role in this enabled smart home of the future.

The idea of a "menagerie" of robots in the home may feel downright strange right now. But in a world where large numbers of people are elderly, having robots that can clean up, remind people to take their

medication, keep their homes warm and safe is going to be essential. Beyond that, robots will actually enhance people's quality of life – which, after all, is precisely what technology is supposed to do.

Moving away from demographics for a second, the idea of a team of robots there to keep your house for you is something that would likely make *everyone's* life better, not just the elderly. Angle refers to chores such as vacuuming, mopping, mowing the lawn, ironing and washing up as "treadmill tasks" – stuff you have to constantly keep on top of as a part of day-to-day life. Again, robots could provide the answer and deliver on the true promise of technological advancement: make our lives better, easier, with more free time to do the things we enjoy or value, while technology takes care of the daily grind of life. Combine that idea with the demographic timebomb heading our way, and you start to see why so much capital is flooding into the sector.

And it *is* a flood. Venture capital funding in robotics is growing at a rapid rate, according to Alec Ross. It more than doubled in the three years between 2011 and 2014. If the Boston Consulting Group is to be believed, the trend is going to accelerate as the industry scales up. The market for consumer robots could hit $390 billon by 2017, and the market for industrial robots could reach $40 billion by 2020.

At this rate, the numbers are going to grow.

Assistive technology: your surgeons' new best friend

Robots aren't just going to change things around the home.

There's also a huge movement under way to integrate robotics into the healthcare system – particularly in surgery. Once again, these robots don't necessarily need to be able to learn or think for themselves to function. But they can assist and enhance the skills of surgeons: make them more precise, less error prone and able to work for longer without experiencing fatigue.

As a case study, we spoke to Professor Kaspar Althoefer, head of the Centre for Robotics Research at King's College, London. Professor Althoefer and his team have developed a "soft" surgical robot that could

one day be rolled out worldwide. It is yet to be tested on real patients, but trials on human cadavers yielded positive results: the robot was more flexible than a human surgeon, meaning it was able to penetrate deeper into the test subject, with greater visibility, and the ability to operate for extremely long periods of time:

These results came about through the work we did over the last four years on a new project called "Stiff-Flop." The project, as the name implies, is about creating new robots that can change their stiffness. They can be really stiff if required, then change to another state where they're more floppy.

The idea is to create new robotic devices which are very different from traditional robotic systems, with their very rigid and hard links. The aim is to create something that, when required, is very soft, and can interact in a very natural way with the human body – as is required for minimally invasive surgery.

Our robotic devices are made of silicon. They have inside chambers that can be pneumatically activated. Through this pneumatic activation, we are able to move the robots in different directions – very much like the octopus moves its tentacles about. We take inspiration from biology, from the octopus: that is a really important part of this project. We try to recreate that in an artificial system with the aim of creating tools that can be used in the surgical environment.

We managed to create a new prototype – another of these soft robotic systems. This particular prototype was small enough to be used on a human body. In Dundee, at the Insight Institute there, we were then able to carry out experiments where we used our robotic devices, our soft robot arms, and inserted them into human beings. These were human cadavers, and we were supported by clinicians who were part of the project, in particular Professor Alberto Arezzo from the University of Turin. We were able to conduct a standard operation using our robotic device.

The focus was on colorectal surgery. The idea was to recreate or mimic a standard colorectal surgery as is currently conducted on humans, especially on those who have cancer of the colon. It's a very complicated surgical task, often carried out as open surgery – where the patient is cut open.

The idea is to move forward towards minimally invasive surgery. Some surgeons attempt to use laparoscopic tools to carry that out. We have

developed our robot arm to support such minimally invasive surgical procedures.

This particular robotic device was equipped with a camera at the tip, and we could show through our experiment in Dundee on the cadavers that our system, our flexible soft robot with its camera attached, was superior to a standard straight-line laparoscopic camera. We could show that this flexible system could move deeper inside various cavities, inside the human belly. The surgeons could get better images from the operating site.

The system we have developed was remote-controlled by the surgeon. What we observed was that, based on the feedback we got from the surgeon, it was very intuitive to move the robot around. Our system could be used over many hours without any problem.

What is usually done nowadays is to use camera images to guide the surgeon – whether a laparoscopic surgeon or a robot-assisted surgeon – to the areas of interest, and to inform the surgeon where to cut and carry out the operation. Just interpreting these images is extremely difficult.

I'm personally not a surgeon, so I look at these images and don't understand anything. However, surgeons, through years of training, can interpret those images and carry out the right task. Automating that process is a challenge.

Reporting from the robotic frontline

Robotics isn't only about assisting us.

There's another side to it: enhancing our abilities, giving people physical capabilities that never would have been possible before.

Take robotic exoskeletons – hydraulics that attach to the legs, arms and torso and enable the paraplegic to walk again. These systems can't "think" for themselves, but that doesn't stop them from utterly changing the lives of the people who use them.

For instance, the Japanese robotics firm Cyberdyne has a system called the "Hal-5". It works by sensing the electrical impulse sent by the brain to the muscle, then automatically tells the exoskeleton to move so the movement one experiences is natural and no time lapse is experienced.

That means that it is highly effective at helping to retrain the brain to restore muscle control after a stroke. For its users, it is transformative. It teaches them to walk again. Even ten years ago, that would probably have been considered a miracle.

As my colleague Eoin Treacy put it, "The powered exoskeleton represents the perfect marriage between the knowledge, sensitivity and imagination of the human with the raw strength and mobility of a robot". It's the integration of technology and humanity, combining to achieve something that neither could do alone.

What's it like to actually use one? I wanted to find out. So I asked researcher and journalist Mischa Frankl-Duval to follow up and speak to the people on the "front lines" seeing the benefit of the exoskeleton industry. As Mischa reported:

> *Steven didn't think twice about his sore back when he woke up in mild discomfort seven years ago. He went about his daily business, then went to bed.*
>
> *When he woke up the next day, he couldn't move.*
>
> *Medical tests confirmed damage to his T4 vertebra. From one day to the next, Steven was paralysed from the chest down and no one knew why.*
>
> *Today, Steven mostly gets about in a wheelchair. He still works and can get around with a degree of independence. He comes to the Royal Buckinghamshire Hospital once a fortnight for physical therapy. As one of the few patients in the UK using exoskeletons for physical therapy, Steven is well placed to discuss how using one feels – and how different models compare.*
>
> *I sat in on one of his sessions, during which he wore an Indego. The Indego weighs 11kg – less than half the weight of the Ekso GT. That, Steven tells me, makes a huge difference.*
>
> *There's much more of the GT than there is the Indego. The latter has no back casing and no straps. Instead, the four-piece suit consists of two foot-plates, lower-leg support, upper-leg support and a wide, carbon fibre belt. It takes a few minutes for Steven to don it with the help of two physios.*
>
> *The suit absolutely blew me away. Once in a standing position, Steven*

could walk – not rapidly, but at a reasonable pace – up and down the length of the ward. On a good day, he told me, he could even take a stroll outside.

Steven wasn't walking unassisted at any point. He leant on a hefty Zimmer frame and had two therapists by his side at all times. But he was walking – something he would never have been able to do a few years ago.

Before he was using exoskeletons, Steven's therapy regime mostly consisted of standing still. For half an hour a day – and an hour on weekends – Steven would stand still in a supportive frame to maintain mass and circulation in his legs. Those processes were essential for his health – but they were some way short of full-on exercise.

Walking with an exoskeleton is by far the most demanding exercise Steven does – and though a stretch in the suits can give Steven a real physical boost, the prime benefit is psychological:

"There's an element of freedom", he told me when we spoke at the end of the therapy session. "I'll leave here as high as a kite when I go to work.

"As able-bodied people you take for granted that you get out of bed and walk around. But having that taken away puts a whole new perspective on life. For me to be able to stand up and talk to someone face-to-face is a big thing for me mentally, for my confidence. It does me a lot of good."

At present, Steven says, the suits are slightly tricky to use. Over the last fifteen months of treatment he's developed a knack for it, but the feeling is still a little unnatural. And, of course, he still needs help to walk.

Over time, these suits will become better, smaller and lighter. Steven hopes he'll one day use a suit unassisted. If that does happen, it'll become feasible for Steven to go to the park with his children, go shopping or stroll across town – all on his own two feet.

Until then, Steven is still reliant on professional help to walk – but the psychological boost it gives him is testament to the value of the human value of this technology.

Watching the Fisher in a clinical setting was truly astounding – but that's the only setting we'll be seeing it in for the foreseeable future. Though exoskeletons are being touted for industrial and military use, they are not

> *widespread in either field. The everyday use of this machinery by civilians is some way off.*
>
> *Understandably, these systems are advancing fastest in fields where there's money to be made – but with time, these systems could be making millions of humans stronger, from battlefields to factories. Strap yourself in – it's going to be a slow, slightly awkward, but encouraging ride.*

In fact, technology has progressed enough that the first ever "Cybathlon" will be held in Zurich, Switzerland in October 2016. Think of it as an Assistive Tech Olympics. The event is billed as giving you the chance to "watch the athletes as they tackle flights of stairs in the new wheel chairs and prove their speed and skill using the most up-to-date modern prostheses. The various races test how the competitors cope with specific challenges and activities from everyday life. There are races for athletes with powered arm and leg prostheses, for those wearing a robotic exoskeleton and for powered wheelchairs. There's also a race for cyclists using electrical muscle stimulation and even a brain-computer interface race."

If you're of a cynical disposition, you'll see something like this as a PR stunt. But it's really a symbol – a sign to the rest of the world of just how far robotics has progressed. It's events like these that act like beacons to industry, attracting more and more investment.

Exoskeletons give us abilities we either don't have or have lost. The robotics industry could do much more than that, enabling us to go places we simply could not go alone.

Nuclear disasters, earthquakes, floods and hurricanes, wars – they're all events that can be lethal to human beings. But that doesn't mean we don't need to traverse the affected areas to rescue the wounded. This is where robots can play a massive role.

For instance, robots were used in a search capacity in the aftermath of 9/11. And since then, they've been deployed on the site of countless natural disasters. More often than not, they're used to provide on the ground pictures of places humans can't safely access. They're also used to perform vital tasks humans could never do: in the aftermath of the Fukushima nuclear disaster, robots were used to swim through the underwater tunnels of cooling pools, removing leftover radioactive material. (The robots reportedly "died" afterwards as the radiation

melted their wiring.)

Often disaster recovery robots simply mean drones – unmanned aerial systems, able to provide real time views of the disaster from positions humans alone couldn't access.

According to Robin Murphy, director of the Center for Robot-Assisted Search and Rescue at Texas A&M University, use of unmanned drones on disaster sites is now commonplace: "Every single disaster since about 2011, but definitely since 2012, looking at the 46 disasters we've kept tabs on, have used unmanned aerial systems, including the ones here in the United States."

Of course, it would be foolish to mention the battlefield – and the drones – without discussing the military applications of robots. This is not just a big technology story– it's also big business, as *Exponential Investor* editor Andrew Lockley found out:

> *If you're looking for even more damage, you need only look to the skies. The next generation of military drones are not going to be the glorified model aircraft we see today. They're going to be killing machines capable of a far higher degree of autonomy. If there's a 'human in the loop' at all, they may well be reduced to a job that's little more than pressing a button to confirm that it's OK for the robot to kill whomever it's decided is the next target.*
>
> *This will result in armies that are ever more physically and psychologically remote from the process of killing, and politicians who can wage war without having to worry about their constituents coming home in a bag. Only the machine-selected target loses out. And 'enemy' soldiers (carefully selected by an algorithm) don't count as people, right? Nor do the people near them. Nor the people that come to pick up the pieces, and get taken out in a double-tap strick. Robots don't have to be sentient or rogue to do evil – they just have to be controlled by evil people.*
>
> *These military robots are certainly game-changing. If you're keen to invest in autonomous military machines, you can look at firms like Raytheon, which makes (among many other things) the Phalanx gun. This example is not new tech (similar systems have been around for nearly 40 years), but it's a taste of a future of fully automatic weapons. The Phalanx gun sits on the side of ship, looking for missiles with its onboard radar. When it spots an inbound missile, it wakes up, chews the missile to bits with a*

4,500 round-per-minute 20mm chain gun, then goes back to sleep again.

Admittedly, shooting an inbound Exocet or Silkworm isn't going to trouble many people's consciences, but this is very much the thin end of a very large wedge. If you're keen to find other firms in this space, be prepared to do a bit of digging. The companies involved are more than a little cagey about what they're up to. Even if you do manage to get into DSEI or a similar military trade show, you'll find that getting accurate information on what firms are up to is like pulling teeth. I once chatted to a rep from BAE systems who was exhibiting what was (very obviously) an autonomous light bomber (now launched as the Taranis), and he was absolutely not prepared to admit to me what it did. Perhaps I don't look enough like a well-heeled despot.

Robot brains, robot muscle

"Unthinking" robots alone have revolutionary potential. Add the "thought" part back in – meaning the AIs able to learn and develop skills we discussed in the previous chapter – and you start to see that the industry really will be a game changer.

iRobot CEO Colin Angle raised the idea of a menagerie of robots around the house – a veritable robot army capable of cooking, cleaning and looking after you. Each individual robot may not be capable of thinking. But he believes that it could all be coordinated by a centralised AI:

I might not use the metaphor "army"! But certainly a team of robots and connected devices that are coordinated probably by a Cloud-based home management system and that uses a human interface robot as its voice to you and its direct interface to you.

Amazon Echo is another tool which could be integrated into this human interface robot. Using that device, that's how you talk to your home. There's a Cloud-based intelligence that can be used to turn on and off things.

The biggest problem with that system and the Google Next system today is that they don't understand what rooms are. They don't really understand your home.

That's where the robots are going to make a huge difference, because we map. The Roomba 980 has profound implications as part of the smart home of the future, because you can build a map of your home, and that can yield a concept of what rooms are. Then your home control system can, instead of responding to 'Turn on light 612,' respond to 'Turn on the lights in the rooms that I'm in.'

Over the next five years, probably, we're going to see systems start to come online which are actually usable, as opposed to where we are today, where they're for the most part fairly unusable. You never use your cellphone to turn on your lightbulb, because it takes 30 seconds, and it takes half a second to use a switch. It could be turned on by a home control system if that system knew what a room was.

What Angle refers to is an entity with robot brains, robot muscle, robot senses, a robot central nervous system – robot everything. The centralised AI runs the show, just as your brain does. Everything else becomes either a way of performing a task, like an automated hoover, or a way of sensing things, like your central heating system, or both. As the hoover moves around the home, it maps what's happening. If there's something to be mopped up, that information is automatically fed back to the AI, which deploys an automated mop.

Perhaps the hoovering and mopping analogy slightly undersells the concept. But just imagine a world in which everything is connected and controlled centrally. You don't need to go shopping, because your fridge knows when you're running low on food, and your kitchen has found out which foods you like to eat. It picks up on patterns in your life: beer and pizza on Fridays, healthy salads on Mondays. The food arrives automatically.

Your house is always the perfect temperature, as it knows when you're in and when you're not. It's connected to your phone, so it knows when you're on the way home. Everything in your home becomes a way for the central system to learn about you and ultimately make your life easier. We're talking full blown tech utopia here – technology looking after you, understanding you and freeing you to do the things you enjoy and value most.

This is a world that the futurist and technology investment expert Sam Volkering described for us when we spoke to him:

There are smart home technologies from companies like Google, Apple, Samsung – companies that make things like smart water centers in your garden. I put a doorbell on my door so that if I'm not home and somebody rings it, I can answer my doorbell from another country and see who's at the door.

This converges into a world where everything is automated and delivered to you wherever you may be. In its current guise, that's probably through your smartphone. When everything's automated, you end up where you've got almost a traveling personal assistant with you. That helps you manage your home, your work, your office, your travel, your holiday. They all blend into one.

From an automation robotics point of view, there's a small robot and it's like a personal assistant. It hooks up with your home, wherever you are, because you connect your location through your smartphone to your home. That's all possible through the internet. It's like a mass connection of every facet of where you are, what you're doing in your life – all coming together to help manage your lifestyle. It's like one giant personal assistant.

These technologies exist right now, but no one has yet really effectively managed to bring them into one seamless system. That is going to happen.

You could use the example of Google. There's a company with the resources to create an ecosystem like that, where all your products are with Google. So it might be Google Next, Dropcam, the Google home automations system working with companies like BMW, Androids for work, which is an ecosystem for your workplace… you'll exist within that one ecosystem.

"Tell the operating system to cook me a steak"

Never mind your fridge restocking itself. Soon your kitchen may well start cooking for you.

Picture it. You're on the way home from work one day, thinking about what you want for dinner. You decide what you want to eat, tap a few buttons on your phone and forget about it.

When you get home, you head straight to the kitchen, open the fridge and

get yourself a beer. Then… you watch, as your robotic chef – consisting mostly of a set of robotic arms that descend from the shelves above your oven – assembles the ingredients, peels, chops and prepares them and then cooks everything for you.

All you have to do is sit back, relax and then eat the meal that's been prepared for you. You don't even need to any washing up. The robot deals with that too.

Let's leave the question of whether you'd want to eat a meal cooked by a robot aside for a second. Is it even technologically possible to do this?

Yes, according to the UK robotics firm Moley, creator of what it describes as the world's first intelligent cooking robot. Here's how the company describes what its robot does:

> *It learns recipes, cooks them and clears up after itself! It can mimic the actions of a master chef precisely, bringing a variety of delicious dishes, cooked to world-class standards to the domestic kitchen and other food preparation areas. The system comprises a full suite of appliances, cabinetry, safety features, computing and robotics.*
>
> *Moley is turning the dream of unlimited access to chefs and their recipes worldwide into reality, with the option of the robot creating their dishes for you; producing meals from around the world or even cooking your own recipes and sharing them with others all in your own home.*

The robot is more or less as I described it: a set of hydraulic arms that hang down from above the kitchen side and hob, enabling them to prepare and cook food.

But there's more to it than that. The robot can also "learn" new recipes all the time – either by following recipes uploaded online by other chefs (the goal is eventually to turn recipes into intellectual property that can be packages and resold like music on iTunes) or by "watching" you cook something and learning from that. Next time you prepare your favourite spaghetti, your kitchen watches you. It figures out what's happening by using vision systems and pattern recognition. Then next time it cooks for you – voila! It makes your spaghetti just the way you like it.

This is where AI and robotics become truly transformative. The ability to learn about the world, understand the way you live your life and then

have the physical ability to actually shape the world – well, that's where things really start to move quickly.

One vital element of this that is still under development is the vision system needed for an AI to learn about the world. It's all very well being able to recognise patterns and learn from them. But if a computer can't recognise what it is looking at, that's pretty pointless. Which is why computerised vision systems are such big business, as Karen Kharmandarian told us:

> *Robots need more and more sensing capabilities to navigate in an unstructured world, because they are moving away from traditional production lines where they were bolted into the ground. They're becoming more and more mobile and autonomous. For that, they need sensing systems: machine vision, image recognition systems, sensors, gripping systems, natural language processing and generation capabilities. All of these technologies are really key enablers for this new generation of robots. If you have companies that are at the forefront of these technologies, that's going to be a recipe for success.*
>
> *We have two of them in the portfolio, part of our top ten. One is Keyence. It's a Japanese company. The other one is the equivalent of Keyence in the US, and it's called Cognex. They provide machine vision technologies for this new generation of robots.*

Summary

Robotics could be the answer to the demographic timebomb facing the Western world: we need robots to assist us and help us care for the elderly. This makes robotics both big business and absolutely essential – for a vision of the future, just look to Japan.

Chapter 5 - Section 1

Who needs God?

The idea is now hovering before me that man himself can act as a creator, even in living Nature, forming it eventually according to his will. Man can at least succeed in a technology of living substance [einer Technik der lebender Wesen].

Jacques Loeb, biologist, writing in 1890

Alpha children wear grey. They work much harder than we do, because they're so frightfully clever. I'm awfully glad I'm a Beta, because I don't work so hard. And then we are much better than the Gammas and Deltas. Gammas are stupid. They all wear green, and Delta children wear khaki. Oh no, I don't want to play with Delta children. And Epsilons are still worse. They're too stupid to be able to read or write. Besides they wear black, which is such a beastly colour. I'm so glad I'm a Beta.

Aldous Huxley, *Brave New World*

I'll never understand what possessed my mother to put her faith in God's hands, rather than her local geneticist.

Vincent, a character in the 1997 film *Gattaca*

The story starts like this.

A quiet lunchtime at The Eagle, a public house in Cambridge, is interrupted by the entrance of a young scientist who makes a rather strange announcement. He claims to have "found the secret of life".

Given the fact that the pub had been serving patrons of Cambridge University's Corpus Christi College since the 17th century, it probably wasn't the first such outlandish claim made in the pub. But it was almost certainly the most accurate – and significant.

The scientist behind the announcement on that February lunchtime, in 1953, was Francis Crick. Nine years later he'd win the Nobel Prize

for his work, alongside another scientist, James Watson. It turned out Crick was right; they had found the secret of life.

But that's not the story I want to tell you.

Everyone knows that in the 1950s, James Watson and Francis Crick discovered the double helix, the twisted ladder structure of deoxyribonucleic acid (DNA). The two brilliant scientists, working from one of the world's most venerable seats of learning, Cambridge University, set out to understand the very building blocks of life – and succeeded, winning the Nobel Prize for their efforts.

And everyone knows that roughly half a century later the Human Genome Project cracked that code and moved our understanding of DNA another huge step forward.

No, the story I want to tell you is different…

It's the story of a child born in 2053, exactly a century after Francis Crick interrupted lunch at The Eagle. Let's call her Angela.

Angela is born with immunity to many diseases. She is genetically incapable of becoming infected with AIDS, or developing cancer or heart disease. And anything that does happen to make it through is destroyed immediately by her immune system – which is hundreds of times more potent than yours or mine.

Disease isn't a part of her life. And death will be a choice one day, not an inevitable part of life.

But that's not all.

Angela isn't just healthier that we are. She's smarter. She's fitter. Taller. Stronger. She can sleep for three or four hours and feel completely refreshed, giving her many more hours in the day to be productive. She can out-think us.

She is... in the purest sense… a superhuman.

She's like us, only better.

She's the future of humanity.

Playing God

Of course, Angela's future life and the announcement that Crick made in The Eagle all those years ago are inextricably linked. Without Watson and Crick's discovery, the modern field of genetics would not exist.

Crick and Watson's discovery laid the foundations. The Human Genome Project pushed our understanding even further. But it is a very recent breakthrough that has really unlocked the potential for genetics to change the world. That breakthrough is known as CRISPR/Cas9, or CRISPR for short. It stands for Clustered Regularly Interspaced Short Palindromic Repeats. And it enables us to do something no civilisation before us has ever had the ability to do – to edit the genetic makeup of living things and remake ourselves the way we want to be.

Crick and Watson's discovery was akin to finding an instruction manual for the human body. Decoding the genome in 2000 allowed us to read the words inside. CRISPR allows us to rewrite the book altogether – cross out parts we don't like, swap whole pages with those of another book, create an entirely new book if we like.

Actually, it does more than that.

It is sometimes easy to forget that DNA is "the secret", to somewhat corrupt what Francis Crick said, "of *all* life". It is not just humanity. Every plant, animal and most single celled organism on the planet is built using that same code. The ability to edit it means we're not just able to remake ourselves – we're able to remake the entire world. It gives us the power to create entirely new "synthetic" creatures, resurrect animals that have been extinct for thousands of years, wipe out entire species if we so choose.

The ability, in short, to play God.

That's ground-breaking. But before we jump too far ahead, just consider the fact that trying to manipulate the genetic makeup of animals and plants is something that's been going on for centuries. Cross-pollinating one plant with another in an attempt to achieve a desired characteristic – a colour, resistance to disease, size, etc. – is something we have evidence of, stretching all the way back to the agricultural revolution.

The difference now is that we know exactly which genes we want

to target, and in CRISPR we have a tool with the precision to make that possible. It allows us to target and modify DNA with incredible accuracy. And it's already become both saviour to some and pariah to others. According to a *Vice* story in April 2015:

> *Not only are scientists publishing reports on the technique at breakneck speed (at 370 mentions in research publications so far this year, that's a rate of 20 papers a week), but it also seems that each piece of news that comes out about CRISPR/Cas9 is grander and juicier than the last.*
>
> *In the past two weeks alone scientists have announced that they have used the new technology to inhibit Hepatitis C in human cells and to defy Mendel's laws of inheritance, which have governed the field of genetics for over a century.*

We'll explore the uses of CRISPR – and there are hundreds – in a second. But first, given that CRISPR has the potential to have the single biggest impact on our lives and the world we live in over the next 20 years, I think it's worth understanding how it works. And what better way to do that than to hear it directly from one of the scientists who helped develop it? We spoke to Harvard's Professor George Church, credited as one of CRISPR's co-developers. First off, metaphors about editing the pages in a book aside, how does CRISPR actually work?

> *In nature, it works to kill bacterial invading viruses by remembering previous invasions. It does that by keeping a little piece of the DNA from the [invading] virus, and then that makes it super-easy for it to reprogram a cutting machine, the CRISPR nuclease. That's what it does in nature.*
>
> *Slowly, people adapted it. As a technology, it became editing: changing from killing viruses to editing DNA very precisely – not just making a mess, as you might with killing, with cutting, but replacing DNA. That was announced in January 2013 by two groups, mine and one of my ex-post-docs, Feng Zhang, who was by that time an independent investigator at the Broad Institute.*
>
> *Soon thereafter, it became evident that, unlike some technologies which are hard, this one was easy. There is no way you could have predicted that.*
>
> *It's easy to adapt to other organisms once you solve how to adapt it to humans, which is what we did first.*

Any method of gene editing has some similarity to previous genetic engineering tools in what it can be applied to. Those applications include agriculture – plants and animals and to some extent microorganisms like fungi. It can be used for curing genetic diseases. It can be used for fighting infections, just like its original use for cutting viruses. It already is in use for fighting leukaemia and HIV-Aids. I'm being broad here, talking about genome editing – not just CRISPR.

It can be used for xeno-transplantations – moving organs from pigs to humans – and making those pigs virus-resistant, or making a variety of things virus-resistant.

Finally, it can be used for gene drives, where you can engineer wild populations at low cost and high precision to fight diseases like malaria, dengue, Lyme disease and so on.

Gene therapy is a big category that includes classically and typically inserting new genes. Then you can use more precise gene editing to both remove and insert, and that's where CRISPR comes in.

There are 2,000 gene therapies in clinical trials, and many of those are already curing people. It doesn't mean they've been approved for general use.

They're in the process of doing the gene therapy trials to get a fair number of people cured of, for example, blindness. There are some genetic causes of blindness that are curable by gene therapies. In most cases, you have to do it very early in life, like in young children – or else they'll be able to cure them to the point where they can see light but they can't interpret the light or stasis because their brains have developed too far.

There are other infectious agents like hepatitis viruses, blood diseases that cause hemolytic anemia – the list is long. There are thousands of genes that are so well understood that some of them can be addressed by genetic counselling. But once you have a child that has the disease, then you need to have some kind of cure or prevention for the development of downstream technologies.

There's a lot of debate about what can – and what ethically should – be done about this. Like any truly revolutionary technology capable of changing the world in a radical way, gene editing has both enormous upside and serious potential downside. Bioethicist Françoise Baylis, a

professor at Dalhousie University, says it's "like any dual-use technology that can be used for good or evil. It can be the murder weapon, it can be the gavel the judge uses. So I don't know that there's any way to sort of control that."

Given the virtually limitless potential of being in possession of a technology which allows you to rewrite the rules of the natural world, it's easy for your imagination to run wild. That has led not just to a proliferation of academic papers on the subject, but to an explosion of more accessible literature on the topic.

Michael Bess is a professor of history at Vanderbilt University in the US, and he's the author of one such book: *Make Way for the Superhumans*. When we spoke to him, he sounded equal parts excited and fearful:

> *There's a rather famous movie called 'Gattaca' where there are the genetically enhanced supers and the unmodified normal people. They're two completely different strata. One of the characters decides that he wants to become an astronaut like all the super people. He pretends, through an elaborate set of measures, that he's been bio-enhanced and genetically enhanced. He's able, by dint of sheer will power, to perform at the same level as the genetically enhanced. It's a triumph of will power over this stratified society.*
>
> *I see that as being highly unlikely and unrealistic. The genetically enhanced, pharmaceutically enhanced and bio-electronically enhanced people are going to run circles around the non-modified ones – especially after three or four generations of these enhancements becoming more powerful, having fewer side effects, and becoming more fine-tuned.*
>
> *In my book, I have a description where two people – an unmodified person and a fairly moderately enhanced person – apply for a job. When you describe them side by side, there's no question whom the employer would choose. The enhanced person can simply run circles around the other one, and there's no hope for the unmodified one. That kind of inequality inscribed into biology really worries me.*

Before superhumans, healthier humans

Bess is certainly right to raise the ethical issues that crop up with genetic editing.

But for now, let's draw a line in the sand between two major ways in which gene editing can be used: to heal the sick, and to enhance the lives of the healthy. One might call these the clinical uses and the cosmetic. It's a distinction that those working in the industry already use, as an MIT Technology Review story reported in 2015:

> *"What you are talking about is a major issue for all humanity," says Merle Berger, one of the founders of Boston IVF, a network of fertility clinics that is among the largest in the world and helps more than a thousand women get pregnant each year. "It would be the biggest thing that ever happened in our field." Berger predicts that repairing genes involved in serious inherited diseases will win wide public acceptance but says the idea of using the technology beyond that would cause a public uproar because "everyone would want the perfect child": people might pick and choose eye color and eventually intelligence. "These are things we talk about all the time," he says. "But we have never had the opportunity to do it."*

There is a point at which those lines will become blurred. But we're yet to reach it.

Let's take the clinical side first. There are an enormous number of ways in which altering our genetic makeup will help in the battle against killer diseases.

Take, for instance, the story of one-year-old Layla Richards.

As the Daily Mail reported: *"Layla Richards had one of the worst cases of leukaemia her doctors had seen and, when all other treatments failed, her parents were told to expect the worst."*

Ten years ago – perhaps even five years ago – Layla may have been beyond the help of medicine. But no longer. Thanks to genetic editing, doctors were able to take on the leukaemia with a new weapon. They took something called a T-cell from a healthy person, edited its genetic makeup to "deactivate" certain elements of it, and transplanted it into Layla (using a similar genetic editing technique to CRISPR, known as TALEN).

The genetically edited cells were a success. Layla was cured.

The doctor behind the experiment called it: "a landmark in the use of new gene engineering technology" and added, "The effects on this

child have been staggering.

"If replicated, it could represent a huge step forward in treating leukaemia and other cancers."

"CRISPR is a geneticist's dream come true," said 0ncology expert and Novartis researcher Rob McDonald in a 2016 interview. "CRISPR enables us to do experiments that one could only dream of before." In short, it enables scientists to look into thousands of genes related to cancer. The goal is to find a "kill switch" – the gene or genes that are vital to the survival of the cancer. Find this and you can engineer drugs that hit this target with precision.

This has helped lead to a major project led by the Novartis Institutes for BioMedical Research (NIBR) and the Broad Institute of MIT and Harvard, known as the Cancer Cell Line Encyclopaedia. The goal is to turn genomic information of cancer data and tumour biology into more therapeutically useful information.

That's fightback against a disease that's already developed in a patient. But what if we could remake our genetic profiles so that we were entirely immune to certain diseases?

That's exactly what researchers at the University of Massachusetts have been attempting to do. They're attempting to create immunity against HIV by "cutting" the offending genetic material away. Here's an excerpt from the University of Massachusetts medical school note:

> I*n an attempt to render latent HIV completely harmless, UMass Medical School researchers are using CRISPR/Cas9, a powerful gene editing tool, to develop a novel technology that can potentially cut the DNA of the latent virus out of an infected cell.*
>
> *"On the simplest level, we're employing a very precise pair of scissors to go in and clip out all, or part of, the HIV genome and reattach the severed ends of the human genome," said principal co-investigator Scott Wolfe, PhD, associate professor of molecular, cell & cancer biology. "If we could do that, the hope is that this would be a step on the road to getting a functional cure for HIV."*

And that's not all. The clinical uses of CRISPR – and gene editing more widely – are growing all the time. As well as cancer and HIV,

researchers around the world have had preliminary success using CRISPR against cystic fibrosis, sickle cell anaemia, neurofibromatosis, Tay-Sachs and Alzheimer's.

Giving sight to the blind

And then there are the even more "out there" uses of genetic editing. How about taking genes found in algae and using them to help the blind see again? Strange as it may sound, it's happening. It's all part of an offshoot of the gene editing field called optogenetics.

Here's a quick primer. Optogenetics is actually the convergence of two different fields of study. Firstly, scientists discovered that certain algae contain proteins that respond directly to a light source. This enables them to detect and move towards light. They found that *light itself* can be the catalyst for cellular change, rather than a chemical reaction. That opened the door to the idea of using light as a kind of wireless transmitter to bring about biological change.

This then merged with the advent of genetic therapy and the field of optogenetics was born.

Gene editing allows us to transplant the gene responsible for responding to light as a stimulus in algae into *other* living things (first animals like mice, and more recently humans).

This is all a rather long-winded way of explaining what optogenetics is, which is essentially modifying genes so that they're responsive to light.

That may not sound like a big deal. But the implications are immense. It means we can use light – or certain frequencies of it – as a catalyst for biological change in the human body.

That goes far beyond making your skin glow when you shine a light on it.

Let's take the brain.

Making a change inside the human brain is complex and dangerous. For a long time it involved using electrodes to trigger cellular change. But that wasn't precise enough for many procedures. According to a

special report on optogenetics by *Nature Methods:*

> *In 1979 Francis Crick suggested that the major challenge facing neuroscience was the need to control one type of cell in the brain while leaving others unaltered. As electrodes cannot be used to precisely target defined cells and drugs act much too slowly, Crick later speculated that light might have the properties to serve as a control tool, but at the time neuroscientists knew of no clear strategy to make specific cells responsive to light.*

Optogenetics is the solution to that problem. It means we can make very precise changes to cells within the brain by using light as a catalyst. For example, in one experiment at Stanford University, researchers found that they could switch off the sensation of fear in mice by shooting light through a fibre-optic cable at specific cells in their brains.

Put aside whether it's useful to switch fear off for a second and just think about that as a scientific achievement. It means we can manipulate complex cellular actions *wirelessly*, using light as the sole trigger.

As you can imagine, this has made optogenetics one of the hottest emerging technologies of the last decade. We're in the very earliest days of understanding how we can use it to help people. For example, the very first human test of optogenetics involves a team of researchers in Texas transplanting the DNA of light sensitive algae into legally blind patients. The idea is to use the light entering the eye naturally to trigger a response within nerve cells and send electronic signals to the brain. The study will involve as many as 15 patients with retinitis pigmentosa, which causes blindness. The goal is to manipulate the patient's DNA so that the cells in the eye respond to light and transmit signals to the brain.

According to reports, that trial should get under way very soon. The results are expected to be "a gold mine" for future studies, according to neuroscientist Antonello Bonci.

Ex vivo versus in vivo

For all the incredible potential uses we've already looked at, scientists believe more could be possible. That's because of the difference between ex vivo and in vivo uses.

The uses we've mentioned so far have been ex vivo – they involve taking cells out of the body, modifying them and then putting them back into the body. In vivo use is the opposite: it involves using drugs to target specific cells within the body. The experimentation is carried out on a complete living organism, rather than a partial or a dead organism. Examples of in vivo research include laboratory experiments on animals, and clinical trials of new drugs before they are released on the market.

In vivo research has obvious advantages: you can actually monitor the impact of an experiment on a living being. It also opens up myriad new possibilities in terms of medical progress. In fact, a $300 million joint venture between Bayer AG and startup CRISPR Therapeutics is already under way to develop new drugs for blood disorders, blindness and congenital heart disease.

Research into in vivo uses of CRISPR could yield immense results. But the method has obvious drawbacks, too – such as the risk that an experiment could cause long-term damage to the living organism subjected to it. It's also more challenging – and in many ways further behind in the development curve – than ex vivo use. The challenge is in the specificity involved: delivering the drug payload to the specific tissue, cells or DNA sequence you want to target.

Superhumans and the Jetsons fallacy

Given the potential uses of CRISPR and other genetic editing techniques in the treatment or prevention of illness, it seems near certain that gene editing will become a major industry – and perhaps even a major part of life. It's much easier to win the moral argument against altering the building blocks of life if you've developed a cure for cancer or heart disease. The benefits to humanity are obvious and immediate.

But what about the other side of the debate? What about using these techniques to create improved babies, rather than applying them to people who are sick?

As I said earlier, there's been a line drawn in the sand between the two uses so far, as Marcy Darnovsky of the Centre for Genetics and Society (an organisation dedicated to looking at human biotechnologies from a social justice, human rights and public interest perspective) tells us:

> *When someone has a disease that's threatening their life or their health, then yes, let's try to treat them with gene therapy. But when it comes to modifying genes that we're going to pass onto our children and to every cell in their bodies, alterations that would be irreversible, and that would be passed on to all their offspring: that's where we think the line has to be drawn.*
>
> *In fact, that line has been drawn by so many countries and by an international treaty of the Council of Europe treaty. The U.K. has not signed that treaty, but it does have a national law against human germline modification, which is why Parliament had to vote on that mitochondrial manipulation technique.*

Darnovsky is referring to the technique whereby a test-tube baby can have genetic material from three people to stop the mother transmitting serious genetic diseases to the child. (It is important to make clear, Darnovsky adds, that the 'three-person IVF' techniques are applicable only in a tiny number of cases - according to one estimate, 5 to 10 women per year in the U.K. - and that the techniques in question are aimed at a subset of mitochondrial diseases, a subset that does not affect 85% of people with mitochondrial disease.)

In the process, a fragment of DNA is taken from a healthy female donor. The method was developed at Newcastle University, and last February the U.K. House of Commons approved it.

Sarah Gray, of the American Association of Tissue Banks, makes a similar point in a more emotional but no less relevant way. Gray gave birth to a son with anencephaly – an embryonic defect that means certain parts of the brain and skull don't develop. Gray's son suffered seizures for six days until he died. As she told a National Academy of Sciences summit on gene editing, *"If you have the skills and the knowledge to eliminate these diseases, then freakin' do it."*

That line may have been drawn. But many people – and I'd have to include myself in this – believe it is only a matter of time before genetic enhancement starts to happen.

In our minds, however, we still expect the human beings of tomorrow to look very much like the human beings of today.

As the author Michael Bess put it, when we imagine the future, we

envisage a world in which technology is many times more advanced – and yet we picture humans as being essentially identical to what they are today:

> *I call it the Jetsons fallacy, based on the 1960s TV show set in the year 2062. It became a household show. Everyone was watching it. And it depicted a world where cars fly and people can be transported in pneumatic tubes and there are robots everywhere, but the people are exactly the same as the people of 1962.*
>
> *Another really good example is the Star Wars movies. Technically, they always say, "A long time ago in a galaxy far, far away". They have technologies that we tend to think of as futuristic technologies. But the humans are all the same as they are today – Luke Skywalker, Princess Leia… The only one who does get modified profoundly is Darth Vader, which is sending a none too subtle message that turning yourself into a bioelectronics cyborg is a very bad idea, and you'll become an evil monster if you do.*
>
> *Then there are movies like Iron Man or the Hulk or Spiderman, where there are bio-modifications, but they're always confined to single individuals. You're never trying to imagine a society where millions of people are bio-enhanced at the same time.*
>
> *What's clear is that as these technologies advance, they will make it possible to safely and effectively modify human embryos genetically: modify the DNA. Coming down the line, this is a choice our society is going to have to make. Is this something we want to do or not? Because it's going to become technologically and medically possible.*

The dynamics of the situation are fascinating. On one side, you have a technology that allows us to remake the world as we want – or think we want – it to be. The potential uses are virtually limitless. That creates a massive incentive to explore and push the limits of what we're able to do.

On the other side, you have the ethical and regulatory concerns. Is it right to hack the building blocks of life? Is it right to alter our gene pool – and pass those alterations on to future generations?

One thing that was abundantly clear in all of our research in writing this book was that these questions are most certainly not being shirked.

Almost everyone we spoke to, regardless of which side of the debate they were on, emphasised that these ethical considerations were being addressed at the very highest level of the industry. Let's hear from Michael Bess again:

> *I was very encouraged by the fact that the same scientists who discovered CRISPR-Cas9 helped organise an international meeting that agreed among all the top scientists on a moratorium on human applications of CRISPR technology. That has been fairly rare in the history of genetic modification. The previous time that that happened was at an Asilomar conference in California in the 1970s.*
>
> *Up till now, we're relying on university IRBs (institutional review boards). We're relying on the restraint of the scientists themselves, and there's some legislation in place. For instance, you're not allowed to modify the intelligence of higher primates. There's a ban on that in many countries – probably 100 countries have legal bans – precisely because we're very leery of blurring the lines between us and them. We understand the appalling moral implications of doing that. People are appropriately cautious.*

At the same time, how long will the scientific community exercise that caution, and how long will those safeguards be in place? Will human beings resist the temptation to redesign the generations that come after them? Bess is much more pessimistic on that front:

> *As the medicine advances and our biotechnological capabilities advance, there will come a point where some of the same things that happened with in vitro fertilization – something that at first seemed horribly unnatural and weird – will happen with genetic modification, and it's going to gradually start to seem more normal.*
>
> *Let's say that you could find a genetic modification that tripled your resistance to getting cancer. It was proven to be safe in all kinds of animal and human trials. How many parents would not want that for their children?*
>
> *If we find a way to modify our babies that way, that's the way the path will go, probably. We'll start with genetic modifications for health, not for enhancing our capabilities beyond what's the normal range for humans. As we become more accustomed to that, there will be the possibility of saying: We've now seen that we can increase very safely the cognitive performance of monkeys. Why shouldn't we experiment with humans as*

well? There will be pressures to gradually, incrementally creep forward. This is a process that our society has to get ready for now. Coming down the line, these are going to be actual, practical, ethical questions we're all going to face.

You can expect that pressure – the incremental creep towards enhancement in healthy people – to increase as we explore the possibilities of gene editing.

Better, faster, smarter, stronger

It's easy to think of gene editing enhancements in the abstract and think, "That's wrong" or "I'd never do that". That's because the concept does feel strange – your genes are who you are, they're your very essence. Editing, hacking, rewriting them: it all seems rather invasive.

But think about it in a less abstract and more concrete way, and you start to see things differently. Let's do a little thought experiment to demonstrate.

What if I told you there was a way to extend each day of your life by an extra four hours – four hours in which you're alert, active and completely switched on. How much do you think you could accomplish?

How much more productive would you be? How much more could you learn, read, write or produce for your business?

Now multiply that over the course of your life – and that of millions of other people – and think about how much could be done.

If that were possible, would you be interested? Would it change or enhance your life?

Because, as it turns out, there may well be a way of doing this.

There's a specific mutation to a gene called DEC2 which, when expressed, makes people "short sleepers". That doesn't mean insomniacs. It means people who are so efficient at sleeping that they only require four hours per night before waking up feeling rested and alert. Compared to the average person who needs eight hours, this gives the short sleeper a huge amount of extra time to be productive – around

60 days extra free time a year.

As the BBC reported in 2015:

> *In 2009, a woman came into Ying-Hui Fu's lab at the University of California, San Francisco, complaining that she always woke up too early. At first, Fu thought the woman was an extreme morning lark – a person who goes to bed early and wakes early. However, the woman explained that she actually went to bed around midnight and woke at 4am feeling completely alert. It was the same for several members of her family, she said.*
>
> *Fu and her colleagues compared the genome of different family members. They discovered a tiny mutation in a gene called DEC2 that was present in those who were short-sleepers, but not in members of the family who had normal length sleep, nor in 250 unrelated volunteers.*
>
> *When the team bred mice to express this same mutation, the rodents also slept less but performed just as well as regular mice when given physical and cognitive tasks.*

What if you were offered the chance to have that DEC2 gene edited into your DNA? It may not be possible (yet), but given the choice, what would you do – if you could enhance your genetic makeup to give you the chance to become a short sleeper?

You may well consider yourself better off without it. But flip the question around. Can you imagine *other people* choosing to have it done? The business people, creatives, traders, artists, writers, doctors and inventors who'd love nothing more than to have an extra four hours of alert and productive time in their day without feeling tired. What then?

I think there would be millions of people who'd want to take advantage.

And that's my point. The idea of enhancing ourselves using technology may well need ethical safeguards. In the abstract, it may sound strange and inhumane. But boil it down to each specific enhancement and the potential benefits, and it's not hard to imagine huge numbers of people wanting to take advantage.

For the foreseeable future, enhancements like this will be simple thought experiments for writers like me to use as illustrations – they'll

exist in the pages of books such as this, but not in the laboratory.

That's not to say people involved in the industry aren't thinking along these lines. In fact, Professor George Church himself raised the issue of "cognitive improvement" – enhancing our intelligence – when we spoke to him.

> *Let's say you developed a gene therapy to reduce cognitive decline in Alzheimer's. This sounds very reasonable – we have a growing ageing population, and this will be very attractive. But there are various ways of doing that.*
>
> *Some people will dismiss any drug involving intelligence, because we don't understand it. The fact is, we do understand certain ways of proving cognitive tasks in mice, and if those are tested in humans, specifically under the umbrella of cognitive decline in later years, we can make progress. And if we can make progress, it can be tried out on people that don't yet have any symptoms.*
>
> *And then it can be tried on people who make money with their intellect, and just want to have higher intellect.*

That begs the question: in a free-market economy where genetic enhancements were safe and available to the masses, how many people would choose to make themselves smarter, stronger, faster and better than a "normal" person?

And if a "free market in biotech", where these kinds of things are democratically available and open to all sounds like a tech utopia to you, you're wrong. Moore's Law – the exponential growth in computing power, which also dramatically reduces costs of high-end technology – has already made biotechnology much more accessible for everyone.

The first effort to decode a human genome costs millions of dollars. But the cost has collapsed. In fact it's collapsed at a much faster rate than even Moore's Law would have predicted. Take a look (on the next page):

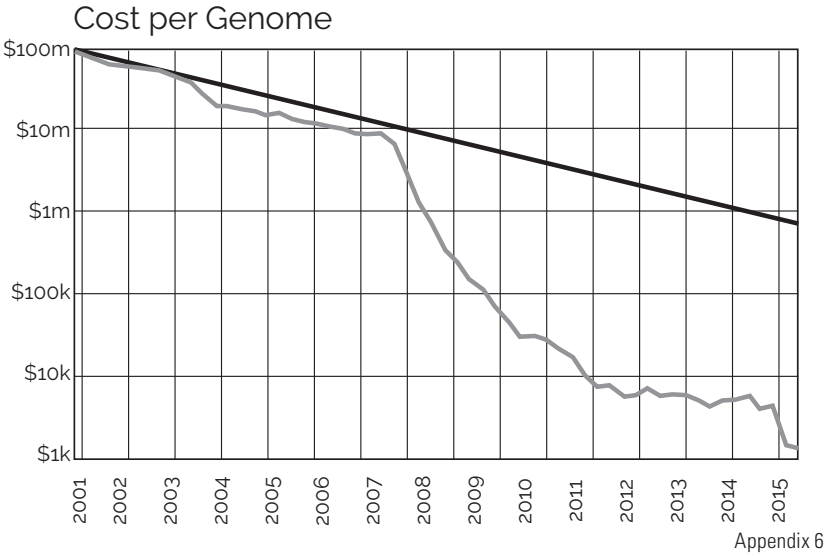

Appendix 6

Again, what happens if we follow this trend to its logical conclusion?

Ultimately we reach a place where everyone has their genome decoded because it is essentially free to do so. Imagine a world where every medical decision made – every drug prescribed or test conducted or treatment option – was made using the makeup of your own DNA. The effect on the human lifespan would be extreme.

And what if gene editing follows a similar path? CRISPR isn't just the most precise gene editing tool we've ever developed – it's also the cheapest (in fact it's 150 times cheaper than its predecessor). How soon before gene editing is cheap enough to be available as a "product" in a free market?

There are already commercial enterprises out there seeking to make biotechnology more accessible to the layman. Bento Bio is one such company. We spoke to its co-founder Phillip Boeing about his grand plan to bring biotech and genetic engineering to the masses.

> *Bento Bio is really about increasing access to biotechnology – making it more user-friendly, and enabling the kind of personalisation revolution that happened with computers, where it's easier for different practitioners to learn about technology and do interesting stuff with it.*

Bento Lab is our first product, and it's essentially a laptop-sized DNA laboratory. It has three component parts: a centrifuge, a PCR [Polymerase Chain Reaction] machine, and a gel electrophoresis unit.

Bento Lab allows you to do many different types of molecular biology protocols. It's the perfect setup for DNA barcoding: purifying a sample, amplifying a genetic marker or a gene, and then running an analysis. If you're in a laboratory and have access to those resources and tools, you can use it as a normal laboratory setup. But in a citizen science school context, we chiefly think of it as a DNA analysis laboratory.

Genetic engineering is in many ways a misnomer. Although the word 'engineering' is used, it's not like an engineering science. It doesn't have many standards, or computer modelling, or the kinds of qualities that we usually associate with engineering.

Ten years ago at MIT, an attempt was made to standardise genetic engineering: to bring some of this engineering into genetic engineering. They created a category in which you can deposit DNA sequences that each have a particular function. The sequences are formatted in such a way that they're easy to combine, like Lego bricks. They're easy to stick together.

It's like having an electronics catalogue, with a list of buttons, LEDs, and transistors that you can put together because you have the instructions. You might have a sequence that builds a sensor in the bacteria, like an arsenic sensor. Another sequence builds a reporter protein – something that makes bacteria fluoresce or bio-luminesce. You put them together, and you've basically created a genetic circuit.

If you're part of a community laboratory or a library or a media art centre, there is a certified space for that. If you're part of an ambitious school university, you can do those activities. And Bento Lab has 80% of the tools to do that.

We have 20 different base testers. They do everything from field science – taking the Bento Lab out in the car and making samples right there in the jungle – to education, to citizen science, whereby people in Switzerland are genotyping all of the different species of yeast and having tasting workshops and linking that to the taste of beer in microbreweries.

There's lots of potential for next generation biotechnologies: from health

to food to lifestyle. But there are two big problems. One is that the tools are really designed for people who have a lot of training and expertise in a very particular direction. You have to have a 10-year education, so a PhD or post-doctorate degree. The tools are not really user-friendly for people coming from the outside. If you want to have these next-generation tools that don't just work scientifically, but also in the market for the user, that requires the population to be more bio-literate. It's a different world to computing, where you can learn some coding very simply and accessibly.

When you open your fridge in the morning, you're probably not aware that most of those products are biotechnology products. The whole infrastructure is hidden away from you.

We aim to build a more granular infrastructure where people can do some interesting genetic tests – for example, test their burger for horse meat at home, and do more interesting things in school. We don't expect a teenager in his bedroom to work on an Ebola vaccine. We want there to be the possibility of a more democratic discourse about biology.

In the next few years, it makes sense to have low-cost, easy-to-access units predominantly in established research organisations for science outreach and field research. But as the DIY biotech movement builds up, people will find lots of niche uses.

The gene pool becomes the gene ocean

One of my favourite books of all time is Dan Simmons' *Hyperion*.

There's a point, towards the end of the book, when one of the characters is thrown ten thousand years forward in time. He meets what is essentially an army of humans. But each and every one of them looks like an entirely different species from the next. In size, appearance, ability – everything – they're different not just from him but from each other. The reason is they've been able to alter their DNA to make themselves the way they want to be.

Or have they?

Have *they themselves* chosen to be eight foot tall, or to have incredible

night vision, or superintelligence, or did one of their ancestors modify their DNA and pass the modifications on?

Because that's where the real ethical questions arise. Giving someone the chance to change their genetic makeup isn't the same as getting a tattoo. That change will be passed on.

It's even more ethically fraught when you consider that "germline" genetic editing – changing the DNA of egg or sperm cells (germ cells) – can change the genetic makeup of an as yet unborn child. It'll then be passed on into the gene pool to be inherited by future generations.

Again, we find ourselves straddling the line between preventing illness and suffering, and enhancing people. If you were given the chance to give your child – and likely, by definition, any grandchildren or great-grandchildren – a better, healthier life, wouldn't you do it?

Or what if you knew that your own genes contained a mutation that would mean your children would suffer from something like haemophilia – don't you have an ethical responsibility to your child to prevent them from inheriting it? Isn't erasing that defect from your family's DNA a greater legacy to your descendants than any nest egg ever could be?

The problem is, that means you're artificially altering the gene pool of our entire species. And that's something that raises numerous ethical questions.

To some, germline alterations aren't some sacred line we can't cross. University of Manchester philosopher John Harris is one of them. He believes that all forms of assisted reproduction affect future generations. He also made the point that regular reproduction produces countless birth defects and diseases. As he put it, "If sex had been invented, it would never have been permitted or licensed… it's far too dangerous."

But to others, that kind of approach is much too flippant. As ethicist Marcy Darnovsky told us:

> *I'm wary because it's such a powerful technology, and when we're considering powerful technologies, we have to use them carefully – because of unforeseen consequences, yes, but also because, in this case, the potential unacceptable social consequences are quite foreseeable. People have talked about this for a long time. Writers have written novels and science fiction*

books. Some people will scoff at some of the projections that have been made, but I think we need to take them seriously.

I think that is also something to be concerned about. Eric Lander [of MIT] talks about that. One of the points that he's been making for some years is that evolution has been working to result in what we are now, as human beings, for several billion years. We learned to read the human genome 15 years ago. Maybe we should be a little careful before we start trying to rewrite it.

Lander reminds us that we have learned enough to know that the way organisms work isn't as simple as we thought. It's not one gene, one trait. A gene can have multiple effects. Genes interact with other genes, with biological dynamics, with social and environmental conditions. If you pull one thread, you just don't know, in most cases, what the effects are going to be.

The concern is that once we start engineering children to have specific traits, as a society, we're saying: these genes are good genes, these genes are bad genes, and those who have preferred genes or traits or who are purported to have those preferredtraits could come to be valued and treated in a different way by society. This could lead to a very consequential kind of discrimination and inequality that is different from those that we've seen before, but that will exacerbate the already shameful levels of inequality and discrimination that we have in our society today.

Let's hear from the man himself. Here's what Eric Lander, Professor of Biology at MIT, claimed at a gene editing summit: "The conclusion is simply humility. Before we make permanent changes to the human gene pool, we should exercise considerable caution."

We put Lander's views to Professor George Church. Is the scientific community showing enough caution when it comes to changing the human gene pool?

My view is that we should always exercise caution, period, on everything. We don't want to have false complacency into thinking that if we just weighed up the gene pool, then we're all set, because there are many ways in which we can affect society that do not affect the gene pool. We can eliminate a particular disease or trait, or engineer a trait, without changing the total frequency of the DNA variant. So, for example, it makes no difference if you have two copies or one. You can have a huge

impact on what the population looks like without changing the number of people that have one copy of the genetic variant.

It's not cautious enough, is what I'm saying. We need to be cautious. He's saying we just need to worry about things that affect the gene pool. I say we need to worry about things that happen much sooner than that.

Affecting the gene pool can take centuries, but we could, in just years, affect the trait pool – that is to say, what you're actually expressing as traits. There are all kinds of educational and corporate pressures to have all the children in a classroom behave themselves. Without affecting their gene pool, you could have huge impact on the behaviour of children in schools, much as in a variety of other ways. So I would say that caution is not nearly enough.

Is epigenetics the answer?

As Professor Church alluded to, there are ways of altering our genetic makeup – or the way our genes affect us – without passing those alterations on to the next generation. How? Epigenetics – changing the way our genes are *expressed*, rather than the genes themselves.

Michael Bess explains precisely how this new field of genetic study works:

Instead of trying to modify the DNA itself, you modify the way the DNA gets expressed or not expressed. Think of the DNA as a piano keyboard: you're playing a certain piece of music, the keys remain the same. The underlying DNA is unchanged. But you're now causing certain keys that would have been played not to be played, and you're causing new keys to be depressed and played that will not have been played before. So you're changing the way the songs are being played on the piano, even though the keys are all the same. You're interfering with how the DNA is expressed.

The key distinction to make is between genotype and phenotype. The genotype is the sum total of all your genetic information. The phenotype is the sum total of all the observable traits of the organism once it's alive and out in the world. I may have a gene for brown eyes and a gene for blue eyes inside me. When I grow up, I have blue eyes, even though I have both brown-eyed and blue-eyed genes inside my genome. What happened was that the brown-eyed gene did not get expressed, and the blue-eyed gene did

get expressed. It was dominant.

What gets expressed is a subset of all the possibilities that are lying latent in your genome. Another way of saying that is: There's your DNA, and a particular song gets played using that piano keyboard – one song but not another. But if you modify the way the DNA gets expressed, then you get a different song: an organism that has different traits.

It's a translation process from the DNA through this complex set of biochemical mechanisms and environmental factors that results in a particular outcome of observable traits of a human being.

Epigenetics is a modification to the processes of your gene expression. Without altering the underlying DNA, you're silencing some areas of the code, and you're turning up the volume on other areas of the code. The result is a different genetic expression, with different traits. Doctors are hoping to be able to cure certain physical ailments using epigenetic tools, and even certain personality ailments, like depression. They're thinking we can intervene in how the DNA is expressed and change a wide range of human traits.

Epigenetic changes are heritable over two, three or four generations, but unlike germline changes, they're not heritable through an indefinite number of generations. Also, you still are using the keyboard of the basic DNA that characterises a human being.

When you're making DNA modifications, you can take an anti-freeze gene that helps an Arctic flounder from having its blood freeze, and splice it into a tomato plant. And the result is a tomato plant that is more resistant to frost. You can take a jellyfish gene that makes the jellyfish glow under certain kinds of light and splice it into a rabbit's DNA, and that rabbit will glow under certain frequencies of light. These are cross-species DNA transfers. The power of DNA transfers is that you can bring in radically different traits from other species – maybe a cancer resistant gene from a certain type of fish.

Epigenetics won't allow you to do that, because all you're playing with is the given DNA that any person has. But the huge advantage of epigenetics is that you can do it at any point in a person's life.

Epigenetic processes are going on within our bodies all the time, 24/7. When you walk into a hot sauna, for example, and your body senses that

there's high temperature around you, your body goes into a complex set of reactions, parts of which are genetically mediated to try to keep your body temperature within safe levels.

What's come to be realised over the past 30 years through molecular biology is, genes are not just the way Mendel thought of them – these things that shape our traits during development, when the embryo is still in the womb and they lie there latent until they're expressed in the next generation. Genes are part of the ongoing self-regulation of our bodies' biochemistry 24 hours a day.

Genes are part of what keeps us alive and healthy all the time. Genomes are being selectively turned on and turned off and activated. Various parts of the sequence are being activated in very complex combinations all the time.

Epigenetics is part of the way that your body responds to environmental cues to start turning on and off various genes that allow your body to then have hormonal and physiological reactions. Our bodies are primed to respond to epigenetic cues all the time. We're constantly engaged in epigenetic changes as part of a functioning of our organism.

If you can get in there and tinker with that, you'd be able to alter some of our traits fairly simply and easily, and perhaps, in the case of some of those traits, reversibly. So you get around the problem of parents having to make these momentous decisions that you would have with the designer baby. If you're altering DNA, you have to do it at the moment of conception, or shortly thereafter, when it's a blastocyst. With epigenetics, people do it at any point in their lives. That's a big difference.

In some ways, it's less threatening than DNA modification, because you won't be able to do cross-species or artificial chromosomes. You're not going to be able to make radical departures from the human constitution. But you'll be able to turn down the volume or turn up the volume on some of our existing traits.

We wouldn't want to underestimate that. Let's say that they find an epigenetic mechanism that helps severely depressed people return to a normal emotional baseline. That type of mechanism might also be used to take people who are already within a normal range of emotional responses and boost them to a higher than normal emotional state. Psychologists have a concept of the happiness set point. You can boost that up to a higher

level: make people, overall, happier.

People notice that effect when they go on Prozac or Paxil – mood-altering chemicals or drugs. If you're not suffering from depression and you go on these pharmaceuticals, many people say, "Wow, suddenly the world became a much more pleasant place for me to live in. I suddenly found myself wearing rose-tinted glasses, and I felt so much happier, and I would've definitely stayed on it if it hadn't been for some of the side effects."

Hypothetically, epigenetics mechanisms and pathways might be used for that sort of purpose.

Chapter Five, Section Two

The world remade

Altering and enhancing ourselves is just one part of what gene editing enables us to do. It also gives us the power to remake the world around us: to re-engineer other living species that have been the cause of deadly illnesses and plagued humanity for millennia – in short, to play God.

This is where things get even more sticky. Depending on your outlook, you'll probably think that what we're about to discuss is one of the greatest achievements of all time… or an absolute abomination.

Here's why. On one side of this issue, we have the single biggest killer in the history of humanity – a disease that could have extinguished half of all the humans ever to have inhabited the planet.

On the other side is a group of scientists who believe that they can stop it for good.

Sounds cut and dry, right?

Not so fast. Because to eradicate this disease, scientists (and their backers, one of whom happens to be Bill Gates) plan to wipe out an entire species of animal. And they plan to do so in a way that would have been impossible before CRISPR…

Killing a killer

Let's start from the beginning.

The disease that has killed more people than any other in human history is malaria. In fact, according to Rosemary Drisdelle, author of *Parasites: Tales of Humanity's Most Unwelcome Guests*, malaria could have accounted for – on average – more than five million deaths a year for thousands of years.

Add those deaths up and you reach a pretty grim conclusion: malaria

may well have been responsible for 50% of all deaths in human history. In her own words: "Did malaria kill between 53 and 54 billion of the 96 billion who lived before 1900? I'm neither an epidemiologist nor a statistician… We'll never know for sure, but based on my reading, I think it's possible."

And it still kills huge numbers of people today. According to UNICEF figures, it kills 1 million people every year. A large number of them are children – 3,000 a day. And 40% of the world's population live in areas at risk of malaria. Fighting it is easily one of the greatest challenges facing the world.

But how do you fight it? With drugs? Better healthcare? Netting for every bed in a high-risk area?

That's one way of doing it. But advances in technology have opened another way. You can go for the source – the mosquitos that spread the disease.

How?

Therein lies the controversy. CRISPR means that there's now a fairly advanced movement under way to edit the genetics of mosquitos so that their offspring become sterile. Sterile offspring could mean no more mosquitos. No more mosquitos could mean the end of malaria.

MIT Technology Review had the story in April 2016, written by Antonio Regalado:

> *I saw such an invention at Imperial College London. A student led me through a steel door, under a powerful gust of air, and into a humid room heated to 83 °F. Behind glass, mosquitos clung to the sides of small cages covered in white netting. A warning sign read, "THIS CUBICLE HOUSES GENE DRIVE GM MOSQUITOES." It went on to caution that the insects' DNA contains a genetic element that has "a capacity to spread" at a "disproportionately high" rate.*
>
> *A gene drive is an artificial "selfish" gene capable of forcing itself into 99 percent of an organism's offspring instead of the usual half. And because this particular gene causes female mosquitoes to become sterile, within about 11 generations – or in about one year – its spread would doom any population of mosquitoes. If released into the field, the technology*

could bring about the extinction of malaria mosquitoes and, possibly, cease transmission of the disease.

Is it right for humans to decide what lives and what dies? To a certain extent, we do already. We've hunted countless species out of existence.

But to do it from the laboratory – to alter the DNA of another living being with the intention of wiping it from the face of the earth – is an altogether different matter, enough to make the head hurt.

At the same time, you have to factor in that the goal is to *save* millions upon millions of people from a deadly disease.

It's a moral conundrum. They say that with great power comes great responsibility. A quick Google search of that phrase tells me it originates either during the French Revolution or in a Spiderman comic book. Let's ignore its origin for now and concentrate on the fact that it's a perfect description of the debate around gene editing.

No generation before ours has had this ability. And using it to fight disease and save lives seems a good way of using it. But it's fraught with moral ambiguity – especially when we're using it to wipe whole species out.

It's an issue that people like Hadyn Parry, CEO of Oxitec (a company involved in similar work) probably has to grapple with. We spoke to Parry earlier in the year. He gave details about how the process works:

We will bring a colony of that mosquito into our facility. We will rear it, feed it, maintain and grow the colony. We will then take the eggs and inject them with a genetic construct. That genetic construct has different components to it.

The first component is making sure the offspring don't survive. We are creating a colony of insects which will all have a modified gene – a gene which will overproduce an innocuous protein so that as the mosquito grows, it disrupts the cell.

If you were to come to our factory, you would see cages of mosquitos producing eggs, those eggs being put into multiple trays where they hatch out into larvae and then pupae, and the pupae being taken and put into dispensing devices when they are adult.

Only female mosquitoes bite you. Males don't bite, and you can never be infected by one. We only want to release the males, because we don't want to release mosquitoes that bite people or spread disease. So we sort the male and the female, and do that physically, which is much simpler than it sounds: when they're pupae, the male and the female pupae are very different sizes, so you can sort them with a sieve.

We put the males into releasing devices. A transit truck comes to the door of the factory, loads up lots of these pots of mosquitoes, and then goes off into a town and releases the males.

There's another element that we introduce right from the word go: a fluorescent protein. When you actually look at the larvae under a special wavelength light, they will appear red. The reason for that is that we want to be able to track them. We want to be able to release them into town, know how many to release, where to release, and if we need to release more or less.

So we'll release mosquitoes in town, and have little traps that mosquitoes use for reproduction. We'll look at the larvae in them. After two or three weeks of releases, 20 or 30 percent of the larvae might be red, in which case 20 or 30 percent of those larvae had a male parent which was an Oxitec mosquito. After another few weeks, it might be 30 or 40 percent, then 50 percent. You have this built-in monitoring system so that you can see what you're doing, how successful you're being. If you're releasing too many mosquitoes, you're wasting money, so you can reduce the number. If you're not releasing enough, you can increase the number.

You want to be able to use this system not only in sophisticated parts of the world, but also in unsophisticated ones. You want something very economical to produce.

As I said: on one side, we have the single biggest killer in the history of humanity… a disease that could have killed half of all the humans ever to have lived on the planet.

On the other side is a group of scientists who believe they can stop it for good.

"We all live at the most poignant moment in history in terms of our ability to engineer biology," says Randal J Kirk, CEO of Intrexon, the parent company of Oxitec. "I think this is going to be the most

important industrial vector in history."

Kirk's position is an interesting one. As the owner of a firm seeking to use genetic engineering techniques for commercial purposes, he may have a more hard-nosed attitude than an academic or ethicist. He's also much closer to the coalface of the business; it's companies and people like Kirk who drive the industry. And we'll be hearing from him more in the next section – on the subject of genetics in the food and agriculture industry.

Feeding the nine billion

Destroying a species of animal to save millions of humans is one thing. What about remaking other plants and animals in order to feed billions?

Technology has form when it comes to this. In fact, technology has enabled us to do something over the last century that the laws of nature suggest should be impossible. Two charts tell the story. Here's the first one. It shows the population of the planet over the past 10,000 years

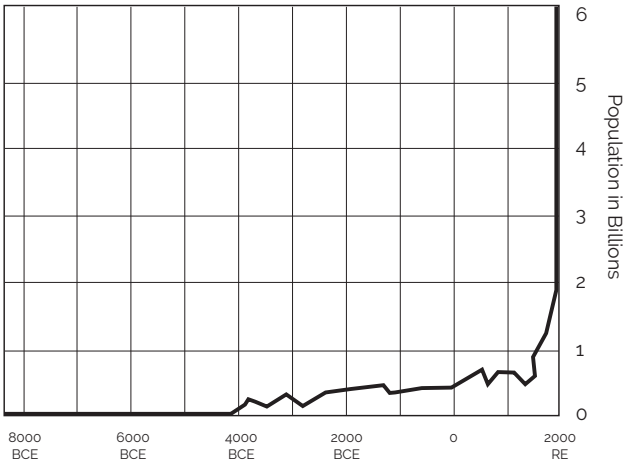

Appendix 7

The chart illustrates a narrative you're probably familiar with. For nine millennia, the population of the planet stayed more or less flat, with only modest growth that frequently went into reverse, due mostly to famine and disease.

Then, roughly three centuries ago – at the very beginning of the industrial revolution – the population exploded, going from one billion and shooting almost vertically upwards to the seven billion mark we see today.

A big chunk of that growth has come in the last 50 years. In fact, the world population in 1960 was just over three billion. By 2010 it was a shade under seven. The number of people on the planet more than doubled in half a century.

Which brings me on to the second chart in this story. It shows nominal and adjusted world food prices since 1961 (as produced by the Food and Agriculture Organisation of the United Nations, or FAO). As you can see, real food prices (yellow line) across the world have stayed more or less flat for five decades. These figures smooth out things like currency depreciation. They're general, rather than specific to one country or region.

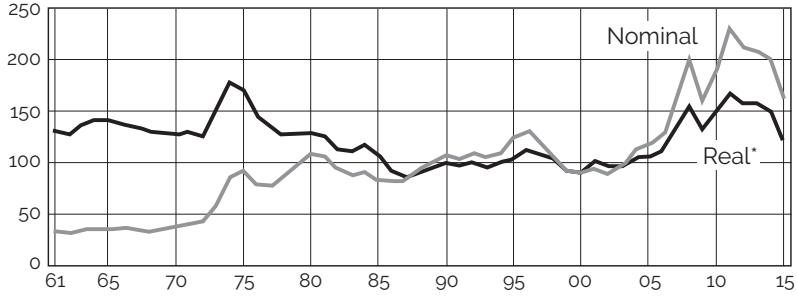

*The real price index is the nominal price index deflated by the World Bank Manufactures Unit Value Index (MUV)

Appendix 8

Think about that for a second. In the time it took for the population of the planet to double – exploding to heights the world has never before sustained – *the price of the one resource everyone needs to survive has stayed flat*.

Let's take a step back and consider just how unusual that is.

For most of human history, the scarcity of our resources regulated the number of people on the planet. There was only so much food, water or energy to go around, which limited how many people a civilisation could support.

Certain advances – like the domestication of plants and animals, and formal agriculture – opened up new sources of food, which enabled the population to grow. But advances like this were few. According to Tom Bulford, a fund manager with 30 years' experience and a biotech investing expert:

> *In terms of the farm sector, animals are seen as merely processing units that take in nutrition at one end and deliver milk or eggs or meat at the other end. What you're trying to do all the time is optimise animal feed – the way in which the animal metabolises its food.*
>
> *There's a lot of interest at the moment in something called the 'microbio map,' which are all the bacteria that live in your stomach in particular. They affect the way in which you digest your food. So it's how that affects your health, and if you happen to be a cow, how that affects the way in which you develop fat etc.*
>
> *Fifty years ago, the average dairy cow in the US produced four tons of milk per year. Today the figure is 10 tons. And over the same period, the number of eggs laid by each American hen has increased by 30 percent. We're learning all the time about how to make animals more efficient as producers of food.*

In short, the story of resources was the story of scarcity. The technological advances of the last century have changed that.

Not only have they allowed us to feed more people, but they've allowed us to feed them without the huge ramp-up in demand sending the price of food through the roof, as you'd expect.

Bad science – the billion dollar biotech battle

CRISPR may be one of the greatest breakthroughs of all time, but it's also shrouded in controversy – not just for what we can do with it, but how it was created in the first place.

It's caused a bit of a storm in the biotech world. That's because more than one team of researchers helped develop it, and now they can't quite agree who should get the credit.

Both the Broad Institute – affiliated with Harvard and MIT – and the University of California, Berkeley – claim to have used CRISPR-Cas9 as a gene-editing tool first. Eric Lander – who founded the Broad Institute – wrote an 8,000 word article on the history of gene-editing, entitled "The Heroes of CRISPR". Of those 8,000 words, just 302 concerned the work of the Berkeley lab. The scientific community was outraged, none more than Berkeley biologist Michael Eisen, who summed his feelings up fairly succinctly with this tweet:

Appendix 9

As *Exponential Investor* researcher and journalist Mischa Frankl-Duval put it when he covered the story:

In part, this dispute is about prestige. Several scientists have played crucial

roles in developing this world-changing technology, and they ought to be recognised.

Then, of course, there's the issue of money.

Novartis and AstraZeneca have already signed lucrative deals with start-ups that are using CRISPR technology. Another agreement – between Vertex Pharmaceuticals and Crispr Therapeutics – ***could potentially be worth $2.6bn****. CRISPR technology could save pharma companies millions in R&D costs by accelerating the process of selecting new therapeutic targets, facilitating the creation of more accurate animal models of human diseases, and reducing time and money wasted on failed products. All that would make discovering new drugs – some worth astronomical sums of money – faster and cheaper.*

The patent for the CRISPR-Cas9 system is currently held by Eric Zhang, a researcher at the Broad Institute. Zhang submitted his patent application seven months after Jennifer Doudna (of Berkeley) and Emmanuelle Charpentier submitted their own. But Zhang's application was fast-tracked, and he now holds the rights to one of the most important biotech innovations in history.

There's a nasty element to this story that goes deeper than money. As has happened so often before, the scientific establishment – Lander most certainly included – is trying to erase the contribution of female scientists to a world-changing technology.

Doudna and Charpentier have still managed to make good money out of CRISPR; Doudna has raised $15m for her company Caribou Biosciences, and Charpentier over $200m for hers, Crispr Therapeutics. But if the patent battle goes Zhang's way, he and his lab could earn many times the $160m his company, Editas, has raised so far.

According to Professor George Church, there are more than just two groups responsible for developing CRISPR – the Broad Institute, Berkley group, and his own. We asked Professor Church to give us his views on the subject:

We have patents that were issued for CRISPR in my group. We were arguably one of the first three. I'm glad that we're not part of the dispute. Our patents are unchallenged, which could mean that we added value that was recognisable and unique.

There are the other two groups: Jennifer Doudnas and Feng Zhang's. Jennifer came up with a way to cut DNA, and proved it in a cell-free system, and hypothesized how it might work in human cells. Then our group and Feng Zang showed how you could use it for precise gene editing, meaning homologous recombination and how to do that. So it boils down to, do you want to reward a prophetic and incompletely worked-out protocol, or the actual protocol that does precise gene editing?

By the time we all published in January 2013 – all three groups published in that month – there was only one group that had done precise, homologous recombination in human stem cells, which was ours. And there was only one group that had the guide RNA that turned out to be the one which everyone used, which was ours. Feng Zhang's had done homologous recombination, precise gene editing of human cells, but had not yet done stem cells and had not yet developed the right guide RNA. Jennifer's group only did cutting, not homologous recombination, and again did not have the first guide RNA. So that's what was proven in public at the time.

In those fields, there are lots of inventions that are required for practicing – actually practicing. This will probably be no exception. We've already got some [patents] that are granted and not under contention, and those will be valuable. To make a cellphone, you need thousands of patents. This [CRISPR patent] happens to be one that's getting a lot of scrutiny, but I don't think it's going to be of great financial significance. I could be wrong.

The three companies that started up to do gene therapy are all in Cambridge [Massachusetts], and they're not really interfering with each other. This is a purely academic thing that's going on between universities. There are other companies that are providing tools, and there are companies that are providing transplantation, like eGenesis, and none of these seem to be particularly concerned about this little tempest in a teapot.

And by the way, this isn't just a case of the West feeding itself and growing fat while the rest of the world goes hungry. The number of people going hungry has actually dropped significantly in the last 20 years, despite the fact the population has continued to rise. Here's an excerpt from an FAO report:

> *The latest estimates show that, since 1990–92, the prevalence of undernourishment has fallen from 18.7 to 11.3 percent in 2012–14 for the world as a whole, and from 23.4 to 13.5 percent for the developing regions.*
>
> *Since 1990–92, the number of hungry people has fallen by over 200 million.*

Again, the traditional way of thinking would suggest that's impossible. In any situation where more people are competing for scarce resources, you'd expect to see rising prices and – most likely – an inevitable drop in the population, as those resources are stretched to breaking point.

But that hasn't come to pass. In fact the opposite is true. The question is: how has this happened?

It goes back to technology and innovation: improving yields through fertilisers, breeding new strains of crop, and other advances. It is probably the greatest achievement in our history. And it comes down to innovation and technology. Or to put it another way, innovation turns scarcity on its head. The *quality* of our thinking enables us to increase the *quantity* of our resources.

We're living in a time where we're increasingly breaking the constraints of scarcity and developing ways of having an abundance of what we need. CRISPR and industrial scale genetic editing of plants and animals is the next logical progression in this story. And yet ethical objections are being raised – and ones that have particular resonance in continental Europe.

Intrexon CEO Kirk has little patience with some of these qualms. Here is an excerpt from our conversation with Kirk earlier this year. I include our original questions to give a sense of the cut and thrust of the debate on the issue:

> *I think that we now understand enough biology: that there are hundreds of examples, successful examples, illustrating our ability to engineer biology to great advantage to solve many of the most significant problems that*

mankind faces – in health, in food, in energy, consumer environment… I could talk within any of those areas to illustrate the point.

Let's take food. I'm sure you're familiar with the Reverend Thomas Malthus. He made a prediction a few hundred years ago that based on the available input, we are going to face something like mass famine sometime in the near future. This would be the case today. Why has that never happened? In what particular regard were the assumptions faulty?

The answer is that technological improvement to improve yields – so that we can do more with less – has been the constant progress in food production. The idea that all of a sudden we should stop using technology in order to improve yields in food production is not only out of line with all of history, but it's out of line with logic.

Q: But there are enough people, and smart people, raising ethical objections.

A: What do you mean by an ethical objection? A lot of people who are simply confused and uninformed say 'ethics' when they don't really mean ethics. They just mean, it's not from a basis of knowledge, and therefore it must come from some other place. And they don't want to admit that it comes from technophobia, so they call it ethics.

Q: In this particular case you're modifying genes.

A: If someone has an ethical argument, I'd like to hear it. Is the ethical argument that man should never alter the genome of any creature? If that's the ethical argument, you should know that we're already 12,000 years past the time we started doing that.

Q: That was done by a natural process. There is a view that you are tampering with something that has a natural balance, and by tampering with it you create an imbalance, and there could be unpredictable consequences. There are lots of people in institutes all over the world – the UK, France, the US – who are concerned about this, not in the field that you're talking about, but in the field of human genetic modification.

A: Let's take the two cases. What you call natural – what biologists would call homologous recombination, which means we take two species that are not identical and we randomly integrate their genomes to make the new

thing – that's been fair game for 12,000 years. And anyone who wanted to do that, according to what you just said, is deemed to be ethical.

What we did over the last 12,000 years [was] in order to acquire trade that we were interested in, and with very little understanding of what else we were doing, because millions of other things happened every single time we did that. Were any of those millions of things ever examined, ever tested? No, they were not.

Now, we have precise capabilities that allow us to go into the genome and make very tiny changes deliberately and intentionally without altering anything else. And now, people are supposed to be focused, just when the entire nature of genetic engineering has become safer? I just find it illogical.

It's a compelling argument in many ways. Genetic engineering isn't new. The technological precision involved is. And perhaps that's what scares people: we can now make changes to the world that would never have been possible, or would have taken thousands of years.

A perfect example is Intrexon's genetically modified salmon – a species whose DNA has been modified so it grows faster. In November, the US Food and Agriculture Organization gave approval for the AquAdvantage fish, produced by the Intrexon unit AquaBounty Technologies Inc.

But critics of the salmon refer to it as "Frankenfish," and the nickname has even made it into the headlines of financial news stories about the product. And there's more. In March, The Wall Street Journal reported that environmental groups were suing the FDA for approving genetically modified animals that would ultimately be used for food. The lawsuit seeks to prevent the AquAdvantage salmon from ending up on people's plates, as well as other species of biotech animals that are currently being developed, including a fast-growing trout and hornless dairy cows, the Journal reported.

"Frankenfish"-style fears are exponentially higher in Europe, particularly in continental Europe, where the assumption is that any foodstuff with altered genes is a freak nutrient whose consumption will be nefarious for the human race. Many European consumers are, in fact, convinced that whatever they eat and drink on their home turf is devoid of genetically modified components.

The reality is rather different. We spoke to Jack Bobo, Intrexon's chief

communications officer – and a former senior advisor to the US State Department. Here is what he has to say on the subject:

> *If you talk to European consumers, most of them believe that GMOs have been banned in Europe. They're not at all aware that 75% of all the animal feed used in Europe is imported, and that most of that is genetically modified. It's been pushed off the shelves, and it still exists as livestock feed, but that's not very transparent to most consumers.*
>
> *Nobody wants to tell them that the livestock products they're eating have come from animals fed on GM crops. It's in nobody's interest to push that information out there. So the belief is that it just doesn't exist. That's part of the problem: there's no perceived benefit, and they don't realise that it's being used.*
>
> *On the other hand, in Europe, much of the beer, wine and cheese consumed – in fact, 80% of the hard cheeses – are produced with genetically modified enzymes. If that was being labelled in Europe, everybody would realise, "Pretty much all the cheese I eat is from GMOs." But under European labelling law, that information is not provided to the consumer. So you don't know that GMOs are used to produce your beer, wine, and cheese. That lack of transparency in the system leads people to believe that this is just not a relevant technology.*

As Bobo explains, the planet will experience an unprecedented population explosion in the next four decades, and the consequences could be dramatic:

> *We're adding 1 billion people every 12 years right now. Between 2050 and 2100, we're going to add 1 billion people every 25 years. It's really challenging to feed an extra billion people every 10 or 12 years. It's not so challenging to feed an extra billion people every 25 or 30 years.*
>
> *If we can get to the year 2050 without cutting down our forests and without draining our rivers, lakes and aquifers, in many ways we're good forever, because for the first time in human history, we won't need more food.*

Here's why the global population explosion is going to improve from 2050 onwards:

> *There is a Swedish statistician called Hans Rosling. One of the things that he's looked at is global fertility rates. We've gone from about 5 children per*

woman in 1950 and 7 or 8 in the developing world to just over around 2.5 today. The year 2014 was the year of peak child. The number of children born in 2014 is the most children that will ever be born in a given year throughout human history. After that, it's all downhill.

Somebody might wonder how we can go from 7 billion people today to 9 billion by 2050 if we're not having more children. And the answer is that people are living longer because of better health and nutrition, and that's a good news story. We don't really have to worry about that extra 2 billion people, because that's you and me.

We also don't have to worry about continued population growth, because after 2050, things level off pretty quickly. We're not having more children every year, so by 2050, the global fertility rate will only be about 2, and that's less than replacement rate. We know that somewhere between 2050 and 2100, population will level off and may even begin to decline, under some models.

That means that when it comes to agriculture and to feeding the planet, the period we are living in right now is the most decisive one we've ever known. As Bobo explains, the stakes today could not be higher:

In many ways, there's nothing we do that has a bigger and more negative impact on the planet than agriculture. And yet there's nothing more critical for our daily survival.

In terms of land, 40 percent of the land that could be used for agriculture is being used for agriculture today. To put that in perspective, the amount of land that we use for crop land is about the size of South America, and the amount of pasturing land is the size of Africa. So in terms of land, there's really nothing bigger.

If we were to talk about water, 70 percent of all the fresh water used by man goes to agriculture, and that's critically important. In the United States, the Colorado River, the sixth largest river in America, no longer flows to the sea. So these are not the challenges of 2050: these are the realities of today. We have to find a way of producing food in a better, sustainable way if we're to have the resources we need in the future.

Between 1980 and 2011, an American farmer uses 50 percent less water to produce a bushel of corn, and that's amazing. But we're also still using more water today than we did in 1980, and that's because we're producing

so much more corn. After 2050, we won't need all that extra corn. So for the first time, we can actually use those efficiency benefits to reduce the amount of water and fertiliser we use.

Between 1980 and 2011, an American farmer used 40% less land to produce a bushel of corn.

Imagine if between 2050 and 2100, we use 40 percent less land to produce a bushel of corn globally. That means that by 2100, we could take all of Africa out of production if we chose, or we could produce all of our food organically. We would use the same amount of land, but it would be organic,

or we would dramatically decrease the amount of land we use. So these are options and choices that we have in 2100.

But they all require that we get to 2050 without having screwed things up. That's why the conversations we're having about food today are so critically important. It's the short window that will determine whether or not we have a sustainable food supply in the future.

The investment case for biotechnology

Unlike artificial intelligence, biotechnology is a much more mature and "investable" industry.

It's been around for longer. It's had more time to grow. It's also had ample time to go through several cycles of boom and bust, greed and fear, exuberance and disillusionment. Plenty of investors have made – and indeed lost – a fortune speculating on the biotech industry.

I should be clear on that last point. When I say cycles I mean *financial* cycles. One of the key concepts to understand when investing in breakthrough emerging technology like this is that there are two distinct cycles driving the industry on.

If you're interested in investing in early stage tech companies successfully it's something you absolutely have to understand.

To explain why, I need to tell you a quick story. It starts with a chart. In this case, the performance of the SPDR S&P Biotech ETF over the past

two years, which I'm going to use as a proxy for the biotech industry as a whole:

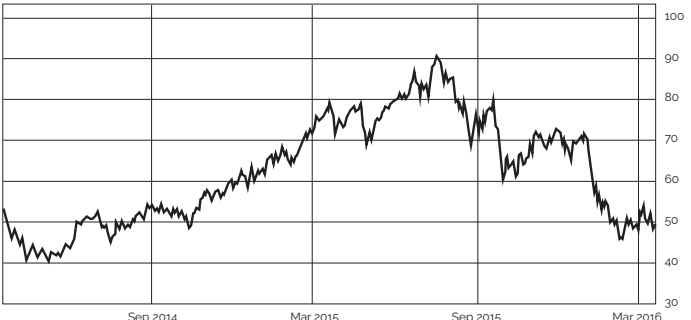

Appendix 10

As you can see, 2014 and early 2015 were pretty good years to be invested. An investment in the ETF would almost have doubled your money by the middle of 2015.

And this is an ETF. That means it's a low cost way to buy a range of biotech firms. If you'd been a little more discerning and looked at some individual companies leading the sector, there was more money to be made.

That all changed in the middle of 2015.

A couple of large biotech firms released results that disappointed the market. Coupled with wider market sentiment turning negative – think back to the Chinese devaluation and the chaos that caused – it hasn't been a good time to be invested in biotechnology since then. Shares in the ETF dropped 44% in just over six months.

Not a great performance by any means, although as anyone will tell you, you have to take the rough with the smooth in a higher risk sector like this. But it'd be fair to break the financial performance of the biotech industry into two blocks: 2014 to mid-2015 = good. Mid-2015 until today = not so much.

Now let's throw a spanner in the works of that thinking.

Right at the start of the year – slap bang in the middle of that period of poor performance – a firm called Editas Medicine listed on the NASDAQ. In fact it was the first IPO of 2016.

Editas is the first gene editing firm to go public. It was spun out of the work of geneticists at the Broad Institute in Cambridge, Massachusetts. The company has backed from Alphabet Inc (Google's trading name) venture arm, as well as Khosla Ventures (another high profile backer of breakthrough technology).

Put simply, if you want to invest directly in the future of genetic editing, it's one of the only avenues out there.

It listed at $16 a share at the start of the year. But mid March, it was $29, or 81% higher than when it listed.

An 81% gain in less than three months is not to be sniffed at. But I think it's particularly interesting when set in the context of the wider biotech industry's poor performance.

The question is, why?

Because of the two cycles driving the industry. The day to day business of biotech itself is driven by science. The markets are driven by a complex mix of real data, animal spirits and irrational exuberance. They're driven by human beings, and humans have emotions. In short, there are two cycles we need to be aware of.

There's the **product cycle** – driven by the invention of new technology, followed by its development, application and commercialisation.

Then there's the **financial cycle** – driven by the flow of global capital, interest rates and the general market sentiment towards any given industry.

It would be great if these cycles correlated with one another; if the financial markets were driven by the real time scientific advances taking place in the tech and biotech industry. But, they don't (or at least, they very rarely do).

But this can work to our advantage, too. It means that markets are likely to overshoot in both directions – to both fall and rise by a larger amount than perhaps the product cycle, which is much more steady and predictable – would suggest.

Genetic editing is reaching a point in the product cycle where we're

going to start seeing genuine results. Editas, for instance, hopes to have its first clinical trials on humans up and running by 2017. The momentum in the business is being driven by product – the science itself. Right now, that's dominant over the financial cycle.

That isn't true of the wider market, where the financial cycle is dominant and momentum is down. Perhaps that's because of things like ETFs themselves, which makes it much easier for people to whip money out of the sector (there is no genetic editing ETF yet, which is an advantage).

There are two cycles and we can't ignore either. We want to find businesses at the right point in the product cycle – with innovative and promising technology that's could see significant advancement. And we want to pair that with a promising financial cycle – where there's still money to flow into the sector and momentum is gathering pace.

What the experts say and the golden rules of biotech investing

To get an insight on how to go about reaping the biggest rewards from the biotechnology industry, we went out and spoke to some of the most highly respected financial experts in the City of London.

For instance, we spoke to Carl Harald Janson and Ailsa Craig, investment managers at the International Biotechnology Trust – a UK based firm with over £200 million under management and winner of Best Specialist Fund at the Investor's Chronicle FT Investment and Wealth Management Awards 2015.

One of the things that struck us was that – for a dedicated biotechnology fund – Janson and Craig are somewhat underinvested in gene editing technology. Perhaps that's because it's an emerging industry, and therefore difficult to invest in. We asked them to elaborate.

> *We are very much interested in all of those things. It's a very interesting time we live in, where there is great innovation in technology, in biotechnology, and in every area where science is making progress.*
>
> *Having said that, I'm a fund manager of International Biotechnology Trust, and I should try to give good performance for my investors. I*

really have to look at this sector through the prism of an investor. What is investable? What is something that can give me the money back, that can give revenues and earnings in the not too far future? Some of these technologies that are very interesting – gene therapy or cell therapy – are still in the relatively early phases of development. There are no drugs on the market yet. So there are a couple of hurdles that need to pass before going there.

You can lose a lot of money if you make wrong investments. We like to play it a little bit on the safe side. We don't necessarily have to be absolutely the first mover into the first company. We rather like to see that it's working and that it's investable.

Making sure something is "investible" may sound like a non-statement. If something isn't investible then by definition you can't invest in it. But certainly, as a safety first strategy, making sure something is proven and money is starting to flow into the sector, is a decent strategy. It may not make life changing returns – which usually come about as a combination of extremely high risk speculation and a bit of luck – but it may well may you better returns than most. We asked Janson and Craig to elaborate – what exactly do they mean by "investible"? They used Alzheimer's disease as a case study:

Alzheimer's is a very important disease. There are a lot of spectacularly interesting projects to treat Alzheimer's. You might try this or that, but until you know what Alzheimer's is, it's really difficult to know how to treat it. We still have no clue as to what Alzheimer's is. There's hypothesis A and hypothesis B. So it's not really investable until we have good data and a model that works and a mechanism that's explainable.

At the other end, we have antibiotics and bugs. We found out about bacteria 100 years ago as the starting point of creating antibiotics. You knew that there is a bug that creates the disease. You put the bug in the test tube, you apply various chemicals and you see what kills it. So we have antibiotics.

In between, you have thousands of diseases that have come part of the way in being understood. For example, with cancer, we've come pretty far now in understanding how we can kill the tumour – how the tumour can block the immune system in the body to prevent it from killing it, and how we then can try to interact between the immune system and the tumour to help the immune system kill the tumour. That's a very investable area now. We

call it immune oncology, or checkpoint inhibitors

I would interpret this kind of idea through the concept of the two cycles of technology we discussed earlier: scientific and financial. Follow the scientific and research cycle carefully. It's slower to move and harder to understand. But the rewards of investing in real, potentially work changing technology are worth it. You're looking for something that has been proven to work on a scientific level – not a hypothesis, but something closer to a working model.

That doesn't necessarily mean approved or commercially available. But by starting your research by following the scientific cycle rather than the financial, you're more likely to pinpoint higher probability technologies. Only then is it time to look at where the money is going.

Get this process wrong – focus on where the money is flowing before the technology is at the right point in its cycle – and you're much more likely to back something that's nothing but a fad, invest at the top of a financial bubble and lose money.

And what of the idea that biotech is all about small companies – tiny operations with radical ideas and super sized risk profiles? Again, we put that to the International Biotech Trust's managers:

It's a very complex question. You have to think of returns in the context of risk. Risk-adjusted returns are much more interesting than just returns.

We think of these as risk-adjusted returns. When we invest in public equities, there can be enormous value creation. There can also be a lot of value loss if you have a bad result or miss something that is very important. We try to mitigate that type of binary event risk by trying to trade around these events, so we're not exposed too much to any downside reaction to bad news on any particular stock.

I try to say to my investors – most of them are institutional investors, so they already have a long-term perspective – that you shouldn't buy into this sector if you need the money in a month or in a year. The best way to look at it is, this is a longer-term investment – of five years at least. So if you have a downturn in the equity markets on the short term, you have to understand why this is.

The lesson there would be: invest for the long term and don't ignore the

big firms! The next half decade or so could see extremely rapid growth in the industry – and that includes the biggest firms around:

The year 2015 was the warmest year in the history of the earth, as far as we have recorded, anyway. It was also the year in which we had the most FDA approvals of new drugs in the U.S. There has been an increasing number of approvals each year. So I think the outlook for the next couple of years is pretty good.

According to some numbers and estimates, until 2020, there will be a CAGR growth of 5% of topline pharmaceutical companies in the world – and biotech is the growth part of that. It's going to be much bigger, and earnings growth is going to be even bigger than that.

Beyond 2020, if you look at what's coming in the pharma companies and biotech companies and the industry and the universities, every year in the last 15 years there has been an increasing number of projects in development in the industry. The outlook and the order book are increasing.

The question is, what's driving all this? Why do we have this type of enormous innovation and value creation and new drugs? It's all based on increased scientific knowledge, not only in biotech but also in tech. It's the knowledge about biology, it's about what is a human being, about what is driving this or that disease – all the knowledge we need to develop products.

We're tremendously excited about what's going on. But that then has to be transferred into what is investable, which is something different.

But not every technology investor sees biotechnology as Janson and Craig. We spoke to Ben Rogoff – a twenty year tech veteran and lead manager of Polar Capital Technology Trust, which has £800 million in assets under management. He told us he prefers to stay away from biotech altogether, with one exception:

We're open to the idea of investing at the periphery of our world, but we ruled out biotech a very long time ago, because it does require specific sector expertise. We have a biotech healthcare team at Polar Capital. We feel that there are standalone vehicles that investors can buy if they're interested in biotech.

There's no doubt that technology, the cloud and some of the deflation

we see in our sector is having a direct impact on shortening development cycles, improving R&D productivity and efficiency. We understand how tech is playing a role in things like gene sequencing. But we would much rather get involved in a capital equipment play that helped in that process, or in a software company that helped manage the clinical trial, than get involved in the miners. Biotech companies are hunting for gold: they're mining for gold, or mining for oil.

We do have 2-3 percent of our portfolio in healthcare-related companies. We would much rather invest in picks and shovels, the companies that make it possible for biotech companies to go out and explore. We don't feel we have the expertise to go out and pick which of those prospectors are going to hit it and hit it big.

That last point is worth expanding on. If in doubt, invest in the picks and shovels – the firms that help make the research and development possible. Perhaps picks and shovels is the wrong word – maybe microscopes and lab coats is a better description! It may not be as high octane as backing the smaller, speculative companies, but for a lower risk way of investing in an industry with revolutionary potential, it's good advice.

Conclusion

An industry with revolutionary and controversial potential. The promise of finding cures for terrible diseases, feeding a rising population and expand the human lifespan will likely be met positively. As the "market" for new and radical ways of engineering healthy people develops, however, the industry will likely become embroiled in ethical disputes. Whether that will hold back the industry is difficult to predict, but given what is or could be possible it is hard not to see the sector developing into one of the most important businesses in the world. The industry as a whole is mature enough to invest in, well capitalised and well deserving of a place, however small, in any investor's portfolio.

Chapter 6

Sun worship

We live in a world bathed in 5,000 times more energy than we consume as a species in the year, in the form of solar energy.

Peter Diamandis, XPRIZE Foundation

The sun provides more energy in one hour than all humanity uses, in all forms, in a single year. Sunlight can provide us with its own resolution to our energy problems.

David S Findley

Let me tell you a story about one of your ancestors.

It's a tale that involves perhaps the greatest advance in human history – something that's directly responsible for pretty much everything you'll do today: every conversation you have, everything you read, eat and think about.

To explain how and why it's true, we need to dust off our trusty time machine. This time we're jumping back nearly 2 million years.

Time travel and the story of early humanity are a bit of a weird way to start a chapter all about energy, technology and the future. But stick with me. You'll see why this matters. Not only that: we'll look at how history could be about to repeat itself in the very near future.

The story starts roughly 2 million years ago. And it starts inside the brain of one of your ancestors…

Back then, humans were – by almost every measure we'd use today – primitive. We may have been at the top end of the animal kingdom when it came to intelligence. But anthropologists believe we had none of the higher-level abilities we'd associate with modern-day people: language, creativity, intricate forward planning.

That's because our brains were different back then. They were smaller, for one. And perhaps most importantly, we didn't have a fully-developed cerebral cortex.

That matters. The cerebral cortex is the part of our brains that allows us to access those kinds of higher-level abilities. It's what gives us the ability to do pretty much all of the things that set humanity apart – to use language, write, talk, be creative, exercise judgement and so on. It's the part of our brains that's allowed us to build a civilisation: to develop new technology, create great works of art, communicate at a high level and share what we know about the world with the next generation. It's the foundation of everything.

Now here's the question: how did we go from *not* having a large cerebral cortex to having one? What sparked the transition?

It was because of something that you probably wouldn't expect: energy.

By that I mean the discovery – or taming – of fire. The ability to utilise fire had many wide-ranging consequences. Its biggest benefit wasn't the fact that we could light or heat our environment. It was that fire unlocked another energy source – a new source of calories from cooked food.

These extra calories had a big impact. They fuelled the growth of the cerebral cortex. And an influx of rich, easily-digestible calories from cooked meat meant that we could spend less time gathering raw food and more time doing other things. That gave us the time for our brains to develop.

In fact, according to one piece I read last week, experts believe that to supply enough calories for the cerebral cortex to develop eating only raw food, we'd have to eat constantly for nine hours a day. That would leave very little time to actually find the food.

This is all part of a theory from a Harvard biologist called Richard Wrangham. This is what science writer Jerry Adler wrote in a Smithsonian Magazine piece in 2013:

> *[Wrangham] believes that fire is needed to fuel the organ that makes possible all the other products of culture, language included: the human brain.*

Every animal on earth is constrained by its energy budget; the calories obtained from food will stretch only so far. And for most human beings, most of the time, these calories are burned not at the gym, but invisibly, in powering the heart, the digestive system and especially the brain, in the silent work of moving molecules around within and among its 100 billion cells.

A human body at rest devotes roughly one-fifth of its energy to the brain, regardless of whether it is thinking anything useful, or even thinking at all. Thus, the unprecedented increase in brain size that hominids embarked on around 1.8 million years ago had to be paid for with added calories either taken in or diverted from some other function in the body. Many anthropologists think the key breakthrough was adding meat to the diet.

But Wrangham and his Harvard colleague Rachel Carmody think that's only a part of what was going on in evolution at the time. What matters, they say, is not just how many calories you can put into your mouth, but what happens to the food once it gets there. How much useful energy does it provide, after subtracting the calories spent in chewing, swallowing and digesting? The real breakthrough, they argue, was cooking.

If the theory holds up, it means that there's a direct connection between energy and intelligence. The two are interconnected. Our cerebral cortex and high-level intelligence developed only because we found a way to fuel their development.

That's important, because it turns out that there is a connection between energy and progress. You can trace this back a long way. Fire helped us cook food, which led to higher intelligence. Fossil fuels like coal powered the industrial revolution, which enabled us to automate and mechanise the world, which in turn helped us support a vastly bigger population. The same principle applies everywhere. The way we generate and use energy is at the heart of everything we do.

That begs the question: can we ever have a true technological revolution of the sort we've been imagining – where we see massive improvements in our quality of life, life expectancy, leisure time and wealth – without finding a sufficiently large source of energy to fuel it? Perhaps all progress is propelled by the discovery of new energy sources.

It's not such a crazy idea. As Charlie Morris – a fund manager with $3 billion previously under management – explained in a research note to

his *Fleet Street Letter* subscribers, there's a strong connection between progress and our energy consumption:

> *Since we left the farms to make cappuccinos in the cities, we should consider the technologies that produced more corn with fewer people. They all required energy as either labour-saving devices, such as tractors, or technologies which took time and effort to research, develop and produce. Energy is the backbone of the modern economy, and if we turned off the energy supply, we would all have to return to the fields. It's as simple as that.*
>
> *It may surprise you to hear that in order to be more productive and grow, we must consume more energy. It is energy that has enabled the modern world to exist at all. At no point am I suggesting that we should waste energy by turning on an extra light for the national good. Instead we need to focus on consuming energy that gives greater value to peoples' lives.*
>
> *In order to grow, we need to consume more energy productively. That doesn't mean wasting it. Nor does it mean smokestacks, a negative environmental impact or carbon emissions. Nuclear fusion, for example, is a developing technology that could deliver endless clean energy at a low cost. The engineers are now saying it is on the brink of becoming viable. They're probably a little excited, but it'll come in time.*

Perhaps all progress needs is a new energy source. Or maybe it's more complicated than that – maybe you need to see exactly the right technologies converging at exactly the right time, with a new source of energy to power them and the capital to finance them.

That's certainly what David Brown – who I mentioned earlier is a tech entrepreneur, angel investor, and creator of Viagra – believes. Brown argues that we're on the cusp of a "third industrial revolution". A combination of new energy, new technology and abundant capital will combine to set humanity on the road to a new era of progress.

Solar's role in the "third industrial revolution"

In our interview with him, Brown started by giving a brief historical overview of his idea:

The first industrial revolution, which started in the late 18th century, is remembered as the mechanisation of the textile industry. It started in northwest England, and it came about through steam power. The big switch was going from wood to coal: much more energy-dense.

This is the critical thing about industrial revolutions. Usually they have three main drivers. One is a new energy source, the second is a new communication system, and the third is a new financial system. That's my analysis.

In the first industrial revolution, steam power led to a better transport system – initially canals across the UK, and then eventually, 50 years later, steam trains. It was steam the whole way through that industrial revolution.

Steam was the basis of a new communication system. It led to the canals, it led to the railways, but also to mass printing for the first time. There was no mass education till 1860, when England first started schooling for everybody, and that was possible because of mass printing, and driven by steam-printing presses. Just imagine what a transformation that was.

This is why the UK built a global empire. It was a long way ahead in technology.

The third factor was actually the London stock market, which provided the finance required to drive those things. Amsterdam had a stock market at least 100 years earlier, but it didn't invent the other two components, so it didn't invent the industrial revolution. London copied Amsterdam, but with the other two components there, there was an industrial revolution.

That's the pattern to look out for, according to Brown: energy, technology and finance. Here's his explanation of how the pattern repeated itself in the second industrial revolution:

You then go forward 100 years. It's the second industrial revolution, and all the stuff that we use now came out of that one. We went from coal to oil: again, much more energy-dense. It allowed transport by much smaller equipment – cars, eventually, but also airplanes. You couldn't fly an airplane with steam power. That then led to masses of industries.

Another key factor was electricity. It came out of U.K. universities, but it was Thomas Edison in the US. who really drove that forward with the

electric light bulb, the electrification of cities, etc. His company, GE, is still one of the biggest in the world. That was a new energy source.

The new communication system came from the electricity. It was initially telegraph and telephone. And the new financial system was the limited company. It sounds like a small thing, but anybody working in finance in the first industrial revolution was at personal risk of bankruptcy. As soon as you went to a limited company, personal risk was reduced. It freed up a lot more finance, a lot more risk-taking. It allowed a much broader swathe of people to get involved in financing.

So what of the third revolution? And what will solar's role be? Brown explained:

If we go forward 100 years again – these industrial revolutions seem to be 100 or 110 years apart, more or less – we get the third industrial revolution, and it's the internet.

What's interesting about this one is that the first and second were driven by new energy sources initially, and the other two factors followed. This time it's communication that's driving it. So actually it's got quite a different flavour.

When I first started doing consultancy on this about five or six years ago, I gave a talk and said, "We know what the new communication system is, but I'm not so sure about the new energy source and the new financial system."

Within four weeks of that, it became very clear that the new energy source was going to be solar. By chance, I was invited to a presentation at Churchill College here in Cambridge one evening. Andrea Ferrari – a professor in the engineering department who is working on graphene with two Nobel laureates in Manchester – gave a presentation on graphene. I could barely believe the things he was saying. He said we can trap solar power in the 50-100% efficiency range, compared to the 10-12% we had at that time, and it's going to be dirt cheap.

I almost didn't believe him. But I went and searched the internet. He and the Nobel laureates had already submitted a physics paper with all the data in. It became very clear to me that solar is going to be the future.

The cost of solar energy is coming down exponentially, and it's already at

what they call grid parity in the US., in about half the states of the US.

It's going to be massively cheaper in a few years' time. The UK is a bit behind, but we're beginning to move on it. Germany, China and now the US have been leaders in solar power.

It's non-polluting. We need to worry about pollution, which is a different issue to global warming.

Coal just won't be economic. Oil won't be economic. There was a slide from Davos last year that appeared on the internet. It showed the Sahara desert, and had three little boxes: the area of the Sahara required to supply the whole of Germany with solar, the area required for the whole of Europe, and the area required for the whole of the world. Even the area required to supply electricity to the whole world was a very small part of the Sahara. And that's with current technology.

It won't happen like that. We won't have these big solar banks. It's going to be local collection everywhere. That's going to lead to major changes in the grid, in the whole infrastructure of energy supply. Governments are going to have to face up to that and help it in some way, in the very near future – I think within five to seven years.

So that's energy. We will literally be coating buildings in graphene, or some derivative of graphene, because it's transparent, it will collect energy, it's a superconductor, you can make electric circuits from it, and you can actually store electricity in it.

It's happening already. In San Francisco, there's a skyscraper going up that's actually collecting solar from surfaces. I've been to angel investor pitches where companies have been pitching windows, double glazing: basically solar power. The glass itself is collecting the energy as it passes through. So these things are coming together.

It takes 15 to 25 years in general for a new technology to go through its induction and development phase to some kind of maturity. That's why I think it's a pretty safe bet to say [that by] 2035 [it's] going to be common.

The great energy disruption

Solar energy – and other renewables – have the potential to bring about

a step change in both the way we live and the way the world works. There's one word that explains why that is: scarcity.

Let's consider the world's number one energy source today. Oil, and oil derivatives, are vital to pretty much every industry on the planet. Oil helps keep cars on the road, planes in the sky, trains running, houses warm, businesses functioning. It's the blood that pumps through the veins of the commercial world.

And the entire oil industry is driven by the idea of scarcity. We only have a certain amount of oil, just as we only have a limited amount of natural gas or coal. They are scarce resources. They've also been vital to the economy for over a century. That means that there's a constant battle between supply and demand: to meet our increasing demand, we constantly have to find new supplies. We are limited by the scarcity of our resources.

Of course, this creates a situation where the places that *do* have these scarce resources – by nothing else but a geological quirk of nature – have enormous wealth, power and status. Think of Saudi Arabia, transformed by oil from a desert backwater into a 21st century powerhouse.

Solar subverts that idea. In a single hour, the sun provides more energy to the planet than we need for an entire year. It's essentially unlimited. And it's free; no one has to pay the sun to shine. Rather than a scarce resource, it's an abundant one. The problem isn't finding more of it; it's developing the technology we need to capture it. We may now have that technology.

Churchill's gamble

It's worth pointing out at this juncture that while this may sound like a resource story, it's really about technology. It was a new technology – the internal combustion engine – that created the demand for oil in the first place.

Let me tell you the story…

"It is a gospel fact... that a fleet with oil fuel will have an overwhelming strategic advantage over a coal fleet."

Those were the words of the British Navy Admiral John Fisher in 1902. Perhaps his point of view – that oil is a better and more efficient fuel than coal – is obvious to you now. It wasn't at the time. Fisher was something of a visionary in his day. In fact he was known as an "oil maniac". It was a description he liked.

The argument against oil in 1902 was actually pretty strong. We didn't have much of it in Britain. By comparison, we had huge and ready supplies of coal. British coal could fuel British ships.

Over time, though, the advantages of using oil started to become clearer. It was more efficient, which meant that boilers could be smaller and ships could run twice as far.

Admiral Fisher eventually found a powerful supporter – the First Lord Admiral of the British Navy, Winston Churchill. Churchill had decided to convert the navy to oil. He asked Fisher to oversee a Royal Commission on Oil Supply. As he told Fisher:

> *This liquid fuel problem has got to be solved... [It requires] the drive and enthusiasm of a big man. I want you for this, viz. to crack the nut. No one else can do it so well. Perhaps no one else can do it at all. I will put you in a position where you can crack the nut, if indeed it is crackable.*

The problem was finding a ready supply of oil. America already had large supplies. But rather than look west, Churchill looked east – to the Middle East and in particular to Persia (modern day Iran).

That decision had many momentous consequences. Most of them are outside the remit of this book. One of the most important outcomes of Churchill's decision was that it sparked a mad rush to secure oil supplies. Churchill's first move was to acquire a stake in an oil exploration firm in Persia. As publisher and author of *The Bull Hunter* Dan Denning put it in a research note on the topic:

> *On 17 June 1914, Churchill presented the British Parliament with a bill that authorised the Admiralty to buy 51% of Anglo-Persian's stock for £2.2 million. The bill passed by a vote of 254 to 18. Eleven days later, on 28 June, Archduke Franz Ferdinand of Austria was assassinated by the Bosnian Serb Gavrilo Princip. The Anglo-Persian Oil Company would later be renamed British Petroleum, which, as you know, is still around today.*

It was around this time that Henry Ford was perfecting his methods of mass-producing automobiles: another breakthrough for technology, another major source of demand for oil and oil derivatives.

These are trends that shaped the world for a century. And it was technology that set them in motion.

It's technology that's been at the heart of the recent (2014-2016) collapse in oil prices, too.

For proof of that, look no further than the oil market over the last decade. The rise of new technologies like horizontal drilling, combined with hydraulic fracturing (which had been around for a long time but came to prominence and widespread use in 2006), changed the dynamic of the sector at a rate no one could have guessed. It first took off in the natural gas sector, but it was when the technology was applied to shale oil assets from 2009 onwards that things really started to heat up.

The next three years saw "tight oil" turn the sector on its head. As Fortune magazine reported:

> *The rise of hydraulic fracking from Montana to Texas to Pennsylvania has lifted US oil production mightily, from 5.6 million barrels a day in 2010, to a current rate of 9.3 million. And until late last year, it was widely accepted that our output would keep rising in 1 million barrel-plus annual leaps for years to come.*

Perhaps those figures don't do the magnitude of the disruption justice. So, to be even more clear, take a look at the two charts showing unconventional oil output, paying special attention to the years between 2010 and 2013:

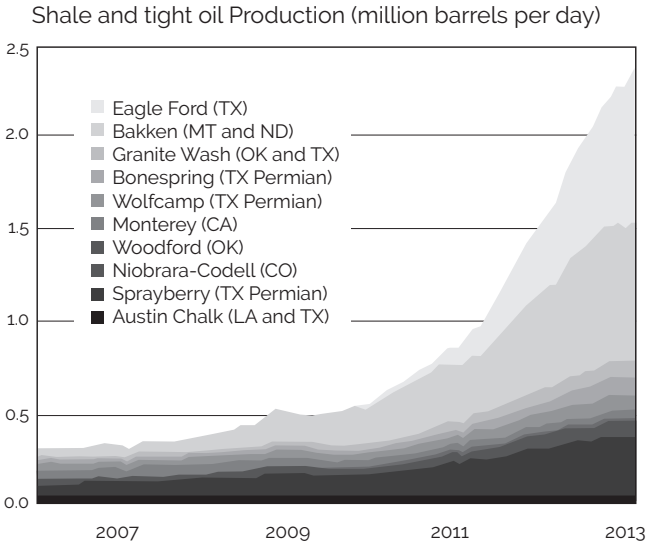

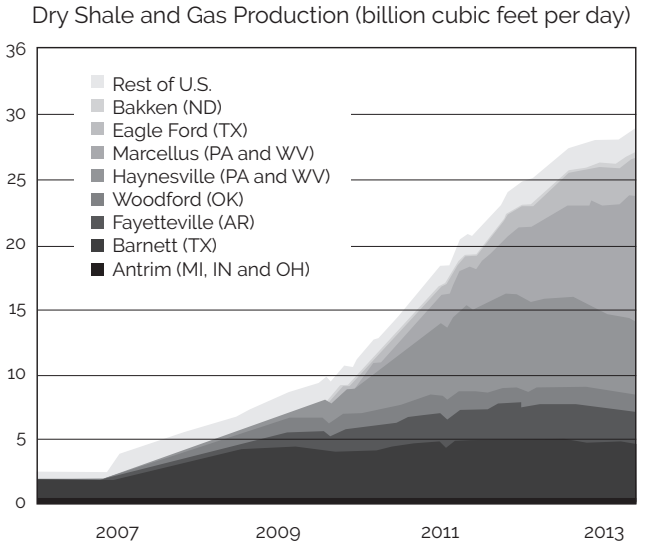

Appendix 11

It was a truly momentous change. New technology totally altered the supply dynamics of the oil market, ultimately bringing about a collapse in oil prices.

Which brings us inexorably to today, and perhaps the greatest disruption the energy market has ever seen.

The case for solar

Once again, the global energy market is undergoing a dramatic change.

Once again, technology is at the heart of it.

We covered Moore's Law earlier in this book. To recap, Moore's Law states that computer processing power doubles (or has historically doubled) roughly every eighteen months. That means computers get faster, cheaper and smaller all the time. That's relevant to the coming energy revolution in the solar industry. The same forces that have driven the tech sector to such incredible heights for so long are now coming to work on solar.

Which is vital. Solar has long been the also-ran of the energy world. It has overpromised and under-delivered for as long as most people can remember (even as far back as the 1970s, people were talking about solar as "the next big thing"). But it's important to remember that the technology isn't standing still. It's improving all the time. A barrel of oil today is the same as it was back in 1970, but solar technology is vastly cheaper and more efficient – so much so that solar is now becoming a genuine contender as a global power source.

I can tell you the story in three charts. Solar energy is getting cheaper. It is becoming more efficient. These trends are accelerating. And that's leading to rapidly-growing numbers of new installations. Consider the following chart:

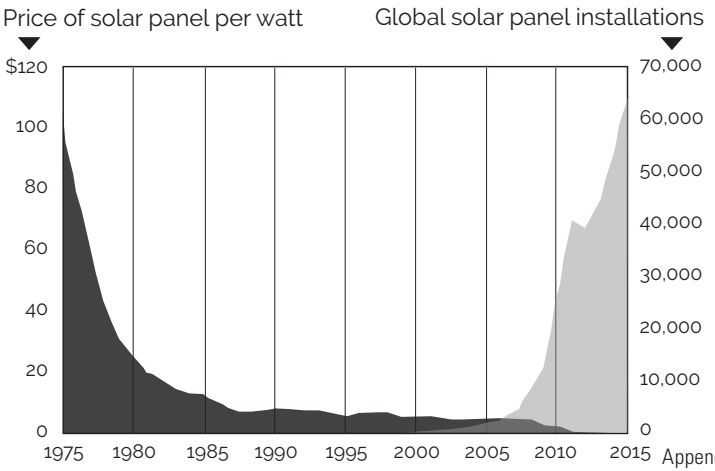

Appendix 12

The price of solar energy has collapsed. And as you'd expect, that's led to an explosion of new installations.

The reason those trends appear so dramatic is because the numbers go back 40 years. But narrow your focus to the last few years and the trend is the same:

The price per watt of electricity generated by solar has more than halved in the last decade. Installations have increased six-fold since 2010. The chart (over leaf) tells the story. Falling prices allow increasing numbers of people to use solar energy as a viable energy source. And that's exactly what they've been doing:

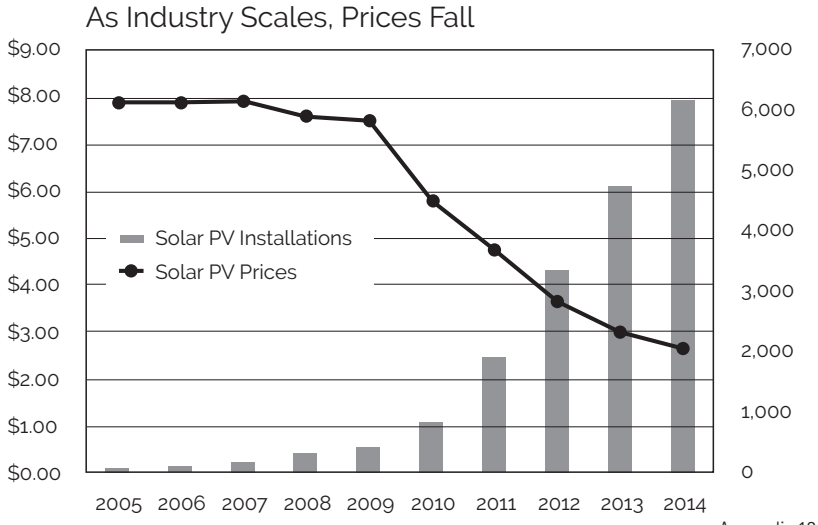

Appendix 13

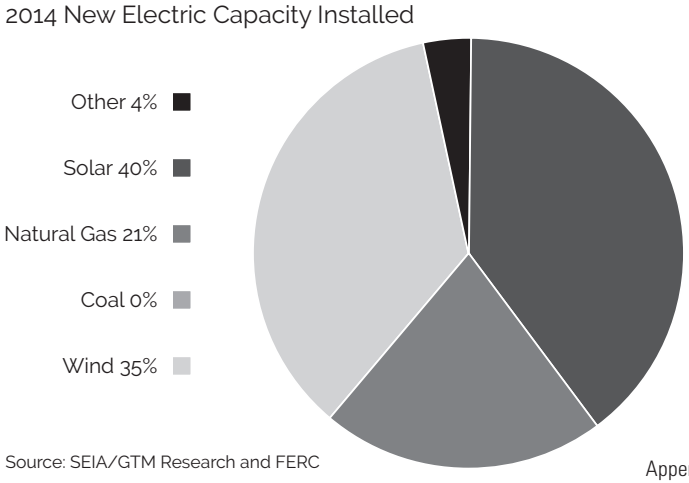

Appendix 14

These figures are from 2014 (the latest data set available through the Solar Energy Industry Association). What the above chart shows is the percentage of new energy installations coming onstream broken down by sector. Solar energy outstripped every other new supply.

All three charts tell us the same thing. The trends at work in the solar industry right now are driving it towards a moment of critical mass. In fact, some would say that we've already hit that point, especially when you consider that:

- Since 2010, Italy has increased its solar capacity by over 410%.
- In 2014, Germany acquired 6.9% of its entire energy needs from solar.
- Between 2004 and 2014, American capacity increased **40 times**, as the chart to the right demonstrates:

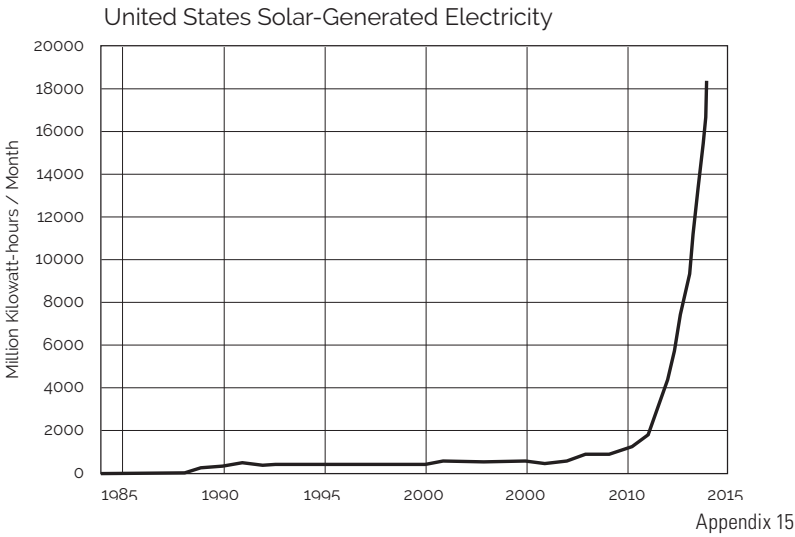

Source: Data from US Energy Information Administration's *Monthly Energy Review* table 7.2a.

And the trend is accelerating. By the end of 2016, US capacity is expected to double again. That's an increase of 7,900% in just 12 years.

We're at a critical point. Rapidly improving technology and funding from governments and tech firms worldwide are combining to create the kind of opportunity you have to consider.

It's not just solar, either. Increasingly efficient technology and a wave of new investment have led to major improvements in just about every renewable technology. According to the 2016 BP energy review, carbon emissions were flat in 2015 for the first time in a quarter of a century. In May 2016, Portugal ran solely on renewable power for four consecutive days.

As economist Spencer Dale put it in BP's report, "its [renewable energy's] strong growth meant that it accounted for all of the increase in global power generation in 2015 and more than a third (38%) of the entire increase in global energy consumption. Global energy markets are in a state of flux: both energy demand and supplies are changing in profound ways".

Solar's symbiotic relationship

There's a key element to the solar and renewable revolution that we simply cannot ignore. That's battery technology. Renewables and battery tech have a symbiotic relationship – batteries are the oil pipeline of the renewable industry. In fact they're probably even more important.

Oil pipelines (and tankers) get the oil to the places we need to use it. Batteries do the same for renewables, but renewables aren't as consistent as fossil fuels. The sun doesn't shine at night, the wind blows to its own timetable and not ours. In order to have a global renewable energy industry providing abundant power around the clock, we need to be able to store the energy for the times we need it, and transport it to the places we need it.

This has historically been a problem for solar. Batteries solve it. They store and distribute energy to where it's needed. And the same trends driving solar are creating stunning growth in battery technology.

Global storage capacity has skyrocketed as the cost has fallen. In the five years between 2008 and 2013, capacity for lithium-ion batteries increased more than eightfold, as this chart shows:

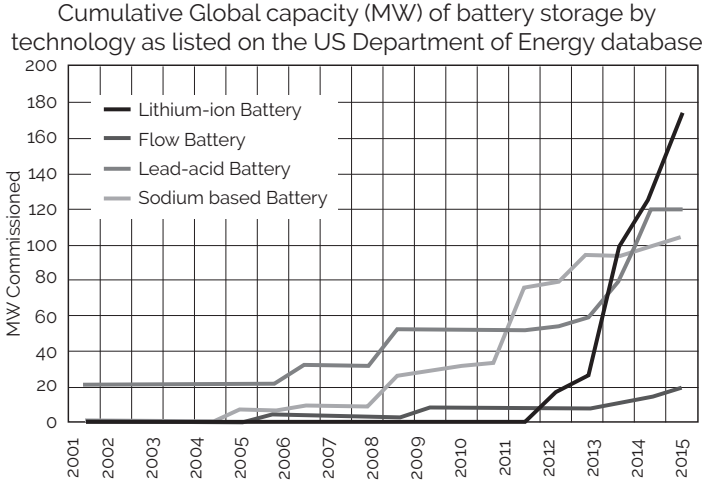

Appendix 16

Source: http://arena.gov.au/

Which is why major tech firms have been pouring cash into the battery

industry. For instance, Elon Musk's Tesla Motors has invested millions in developing battery tech. Tesla has teamed up with Panasonic to build a **1,000-acre** "giga-factory" in the Nevada desert – a project designed to massively increase the world capacity of lithium batteries.

Once again, it's worth remembering that rapidly-improving technology is driving this industry on. Solar and battery tech have seen vast improvements already. But the reason Tesla are investing so heavily in batteries is to power another emerging technology: the electric (and ultimately driverless) car.

The internal combustion engine turned oil (or specifically petrol) into the fuel of choice for motor vehicles the world over. If we were giving out prizes for who made that possible, we'd probably give one to Henry Ford.

Ford perfected the assembly line and mass-produced cars, creating huge demand for oil. That demand helped catapult oil to the top of the world's energy hierarchy. It also handed oil-rich nations almost instant power and status.

Ironically, it could be a new form of the motor car that helps knock oil off its perch. The first motor cars created demand for oil. The next generation of cars won't. They're electric. They'll ultimately run using a high-tech battery, potentially powered by the sun.

You may well have seen this reported in early 2016, but Tesla Motors announced that its Model 3 electric car could go on sale at the end of next year. The Verge had the story:

> *After ten years of waiting, Tesla has revealed the Model 3, the vehicle that CEO Elon Musk hopes will take the electric car to the masses.*
>
> *At the unveiling of the Model 3 this evening at the company's design studio in Hawthorne, California, Tesla CEO Elon Musk said the car will deliver at least 215 miles of range beginning at just $35,000 — that's a bold claim, and an important one for Tesla to meet. Musk is "fairly confident" that deliveries will begin by the end of 2017, and "you will not be able to buy a better car for $35,000, even with no options." And it will be one of the safest cars in the world, according to Musk.*

The horseless carriage is replaced with the petrol-free car. Demand for

oil decreases.

Peak demand and the end of oil

There's an irony to that, too. Why? Well, you're probably familiar with the "Peak Oil" theory.

It's a theory popularised by analyst M King Hubbert in the 1950s. But it's had more than its 15 minutes of fame over the years, with several generations (right up to our own) of analysts subscribing to the concept. Like all popular ideas, it is easy to understand: oil is finite, and sooner or later we'll hit the point at which production begins to fall.

That would inevitably lead to shortages. Supply would become extremely tight. It would be like the 1970s oil embargo on steroids.

From that point of view, the "end of oil" would come about in a blaze of high prices, resource scarcity and oil wells running dry. We'd run out, and that would be the end of that.

But it didn't quite turn out like that.

In fact, the current oil slump – prices fell from over $100 a barrel to $30 in two years – was caused by *over*supply. Technology, which of course is our beat, played a key role. As I've already shown you, new oil drilling techniques opened up vast new supplies of oil in America. That scared traditional oil-producing nations like Saudi Arabia and precipitated an all-out price war. The Saudi's flooded the market with way more oil than was needed to meet demand, and prices fell.

They certainly succeeded in pushing down prices. Whether the shale oil industry will drop with them is still up for debate.

But it's not the current slump that's our focus. It's the final slump – the ultimate end of oil.

That – as you're about to see – could be driven not by us running out of oil (the traditional argument)… but by demand for oil falling off a cliff.

What if we have millions of gallons of oil… but nobody needs it?

What if we're approaching "peak demand"?

This was the conclusion the research house McKinsey came to in a report published in June. According to their analysts, demand for oil could peak entirely by 2030. As The Telegraph reported:

> *McKinsey's Occo Roelofsen said despite an expected increase in global population of around 36pc, and a doubling in global gross domestic product (GDP), shifting energy sector dynamics are set to depress energy demand.*
>
> *"This change is driven by three factors: first, overall GDP growth is structurally lower as the population ages; second, the global economy is shifting away from energy-intense industry towards services; and third, energy efficiency continues to improve significantly," he said. "Peak oil demand could be reached around 2030."*

Or to put it another way, we're seeing huge improvements in *alternative* ways of producing energy efficiently that are now challenging the supremacy of oil. Solar and renewables mean we're able to get our energy from different sources – leading to a decline in demand for oil.

And further down the pipeline, we're seeing a revolution in the transport industry – the rise of electric and driverless cars. That alone could slash millions of barrels of oil from global demand.

And therein lies the opportunity. The crumbling of oil-based empires will trigger transfers of wealth the likes of which the world hasn't seen for decades.

The question is, what will a world in which renewables replace oil as the world's supreme energy source look like? How will we get there? And how best to position your money to profit?

The great energy disruption

The end of oil and the rise of a renewable energy-focused world economy isn't as crazy an idea as it may sound. While it may be improvements in technology driving the industry forward – falling prices make solar ever more appealing for nations, businesses and individuals – it would be foolish to ignore the fact that climate change is a part of the story too. Whatever your own views on climate change, it's certainly motivating

businesses to fund renewable investment, and governments to back the industry over the next decade.

We spoke to someone who's squarely a part of that process. Stanford University's Professor Mark Jacobson is the author of a plan to convert the 50 states in the US to 100% renewable energy, endorsed by several dozen members of Congress to date. He was also a keynote speaker at the UN Climate Change Conference in Paris in 2015. We asked him to talk about the role he sees solar playing – both in his own plans and for the world at large:

> *The idea is to electrify all the energy[-reliant] sectors – electricity, transportation, heating and cooling, agriculture, forestry, fishing, industry. We would electrify each sector, and provide the electricity from clean, renewable wind, water and solar power.*
>
> *For example, in the heating and cooling sector, we do heat pumps. For air and water heating, induction cooktop stoves would replace gas stoves. For industry, we'd use electric arc furnaces and induction furnaces. And for cars, we'd use electric vehicles and hydrogen fuel-cell vehicles.*
>
> *By electrifying everything, we reduce power demand by about 32 percent. Then we can squeeze out another 7 or 8 percent of energy efficiency improvement. Or we can get power demand down by around 40 percent by converting.*
>
> *We did plans for 139 countries and looked at the resources available in each country – the rooftop areas, the size of the land areas needed. We looked at the cost, the job changes upon conversion. And we looked at the air pollution reductions in each country and the car emission reductions.*
>
> *We found that in each of these 139 countries, we can change or transition the energy infrastructure. It's beneficial in multiple ways. You stabilise energy prices. You have zero fuel costs. You have similar or lower business costs of energy. You eliminate health or climate costs, which are much higher than the business costs. You have more distributed energy because of more solar-wind. You have less chance of big power outages or terrorist attacks. You create a net 22 million jobs worldwide – net of the losses. So there's really very little downside to such transitions. It's really only beneficial.*

It may be a surprise to hear that politicians in America have responded

positively to Jacobson's plans. That's something that probably wouldn't have happened a decade ago – a mark of how far the technology has come:

> *There's House of Representatives Bill 540 for the United States to go to 100% clean renewable energy by 2050, which is based on our plan. It has several dozen co-sponsors, including the House minority leader Nancy Pelosi.*
>
> *If you look at the details, it's based on our United States plan for 100%.*
>
> *It doesn't mean it would pass, because the majority of the Congress are Republicans right now, and it's not likely they would vote to approve it at this time.*
>
> *[Still,] this wouldn't have happened two years ago.*
>
> *There's a lot of movement, and accelerating movement. There are cities that want to go to 100%. There are companies. Compared to two years ago, the conversations are a lot closer to going to 100%.*
>
> *And there is actually action: 70% of the electricity in the US last year – net new electricity – was wind, water and solar. And 100% in Europe.*
>
> *In Europe, there is more coal decommissioned than added. If you're looking at all the new electricity in Europe, 100% last year was wind, water and solar – when you count what was produced minus what was taken down. Even though there are some individual coal plants put up, there are more taken down than put up. It's the same with nuclear. Net, 100% of the new electricity was wind, water, solar.*
>
> *In the US, it was 70%. There was new demand for electricity, and it was just met by renewable energy.*
>
> *I know it will be implemented. It's just a question of when.*

Timing the success of Jacobson's plans is near impossible. The good news is that you don't have to wait until the world has "gone renewable" to be involved as an investor. The transition from one way of generating energy – fossil fuels – to another – renewables – will throw up huge investment opportunities. Our advice: follow the innovators and don't dismiss battery technology.

City centre solar

First, let's consider the innovators. As solar technology scales up and costs fall, the industry is likely to be taken in increasingly inventive directions – driven by firms pushing the limits of what the technology can achieve.

SolarWindow is one such firm. It's a small business with a big idea – the kind that may not come off, but which, if it does, will change the world. The concept? Turn every large commercial building or skyscraper in the world into a power station. It's a compelling idea: rather than covering vast swathes of land up with solar panels, the company has engineered a patented technology that turns windows into solar panels.

We spoke to SolarWindow CEO John Conklin to find out more:

> *Space and cost are two of the big factors that play off one another. If we look at tall towers and skyscrapers, there's very little space on the rooftop of a tall tower. That space is compromised by all the equipment that's required to operate that building. There are HVAC systems, elevator shafts, and other electrical and mechanical components that take up the rooftop space. In an average 50-storey building, you're looking at anywhere from 5,000 to 8,000 square feet of space.*
>
> *Consider the amount [of land] that we would need to offset for a tall tower and skyscraper. For every 50-storey skyscraper that takes up six acres, we'd quickly utilise all the valuable park area in all of our cities for just a few skyscrapers, and then quickly run out. We'd then have to look at adjacent land, and then at the transmission cost associated with getting that power to that skyscraper for that land.*
>
> *Then consider the acreage of glass on all four sides of a skyscraper. When we look at the next-generation skyscraper, a 50-story skyscraper has nearly six acres of glass.*
>
> *SolarWindow provides the asset of utilising vertical space to turn tall towers into electricity generators – with a less than one-year-modelled financial payback – and does not affect the valuable urban land in those cities around the world.*
>
> *We're looking at a global flat-glass market worth over $8.3 billion. When we break that down into area of glass, we're looking at about 89 billion*

square feet of glass, or in metric terms, about 8.3 billion square metres. That is a huge market potential.

We'd look at where that glass is needed. It's in measures of buildings replacing windows, and part of our research shows that 97% of the buildings out there will be replacing windows. We see that as a huge market potential for not just new construction, but also replacement windows.

We mustn't forget about retrofits, where buildings have already placed their window and do not want to miss out on the opportunity to turn those windows already in place into electricity-generating SolarWindow. We're working on a veneer that would go over those existing windows, and in a Plug-in Play concept, connect directly to that window to generate power for fixtures, lights, computers, PCs, etc.

The business model that I've developed for the company is for global penetration. When we look at the.UK, Germany, France the US, Canada, all the major cities play a role in that. But let's not just consider the big cities, because we see tall towers in other cities – in areas like Scottsdale, Arizona, which has buildings that are perfect for SolarWindow.

The decision on how and when to deploy will be a function of our strategic partners in the glass industry. Let's face it, they know where the buildings are going up. Their orders for buildings going up in a year are already being processed. And they know where the new buildings are going to be going up, say, five to six years from now. So teaming with those glass companies allows us to work with them in a worldwide deployment on where those buildings are going to be, or where those buildings need to have replacement windows.

SolarWindow's technology is still some way from full commercialisation. By 2017 the company should have a clearer idea of when its rollout can begin. But it's clear that technology like this has huge potential – not just to change the way we generate power, but to fundamentally reshape the way we think about energy and the role it plays in our lives. In a SolarWindow world the days of power stations sitting on the outskirts of towns or out in the countryside would be gone. Instead, our power would be generated almost invisibly, from the buildings we live and work in.

Coupling power generation with existing infrastructure is a trend that could go a long way. And it's not just windows. In France, engineers

have begun building a "solar road" – built with energy-generating bricks – that could power a town of 6,000. Again, scale that up and you start to understand that the traditional way of thinking about energy generation will be challenged by a new wave of innovators – presenting major opportunities to investors.

But that's not the only way of playing renewables. We looked at the vital role battery tech plays in the renewable and solar ecosystem earlier. Renewables cannot be seriously viable without a vastly scaled-up battery tech industry. That presents another compelling opportunity.

Picking winners in the race for white petrol

More than 200 years ago, in the brief period between the Revolution and the rise of Napoleon, something rather unusual happened in France.

But the slightly strange event in question didn't involve anyone losing their head. At the heart of the story is an American inventor called Robert Fulton. Fulton proposed a way in which the French Republic could counterbalance British sea power. Not by building bigger or better boats, but by creating something else…

A submarine.

The idea was rejected at first. But then Fulton offered the Republic a deal: pay me nothing until my invention has sunk a British ship. The Minister of Marine took the deal. And the records show that the submarine did work, in its own way. It could descend to a depth of four meters (a long way, back then) and hold enough air for the crew to remain in it for more than four hours.

Ultimately, however, history got in the way of progress. The rise of Napoleon and the Battle of Trafalgar destroyed any French hopes of competing with the British at sea, and that led to the project being abandoned.

The vessel was known as *Nautilus*. And I told you that story for two reasons. First, it shows what can be achieved when the right people come together with the ambition to do something great and develop

potentially world-changing technology. That's been a running theme in this book.

The second reason is the name of the vessel – the *Nautilus*. That may be familiar to you. It's the name of the vessel in Jules Verne's *20,000 Leagues Under the Sea*. As you'd imagine in a piece of science fiction, Verne's *Nautilus* was rather more advanced than Fulton's. It's a triumph of the imagination, rather than of reality (which is another theme for us, I suppose). It could descend to immense depths and travel huge distances. It could accommodate passengers for long periods of time.

And, in a leap of imagination on Verne's part, it was powered by a battery.

Verne was writing fiction. Science fiction. The *Nautilus* didn't exist as he described it. So it was either prescient or lucky that, more than 150 years later (and after a century in which oil dominated), the kind of battery power he described is coming into its own.

It's not powering submarines right now. Not yet, anyway. But there can be no doubt that – given the way the technology is developing – battery powered transport is arriving in a big way.

Battery-powered cars are already here. What next? Trains, planes, submarines – rockets? Battery technology isn't just going to change the way we power the world. It's going to change the way we travel too.

At the forefront of this is Tesla. In early 2016, the company announced its Tesla Model 3 electric car. Within weeks, pre-orders at $1,000 per car had topped the 400,000 mark.

But it's not all about one firm – it's an industry-wide trend. As OilPrice.com editor James Stafford wrote on 12th April 2016:

> *Chinese billionaire Jia Yueting is stepping onto Tesla's playing field with its own electric car start-up, Faraday Future, and Apple is planning one too, by 2019. Through its Alphabet holding company, Google is also getting into the game with plans for a self-driving car.*

Then you have Ford, Honda, Toyota, Nissan all investing in the technology. On top of that you have commercial vehicle companies like Fed-Ex working on similar ideas. There were even rumours here in

Britain that Dyson was working on a project. Improvements in battery technology mean that just as solar now threatens oil's place in the energy hierarchy, electric vehicles can seriously challenge the automobile.

That's sparked another mad rush from the companies at the forefront of the trend. But it's not oil they need. It's lithium.

Lithium is a vital component in a huge number of batteries (the most common kind of battery in use is known as a lithium-ion type battery). It's the lightest of all metals. If you were to see it in the wild, it would appear soft and silver-white, which is why it's known as "white petrol". It's an alkali metal, highly reactive and flammable.

The last time you took a plane, you might have seen warnings about taking lithium batteries on board. In fact, when Boeing released the Dreamliner replacement for the 747, it was plagued with problems with its batteries. Parts of the plane catching fire was a major issue until better housings were developed. The fact that lithium is reactive makes it ideal for use as the anode (positive charge) in some batteries and as the electrolyte in others.

Major innovations are coming down the line in the development of lithium vanadium batteries, which are designed for use by utilities, and Cambridge researchers have developed highly efficient lithium-air batteries. But so far there are few replacements of lithium in the production of rechargeable batteries.

This presents a major opportunity for investors. Just as with oil a century ago, the mad race to secure large supplies of lithium is creating a major opportunity for lithium suppliers. As The Economist put it at the start of 2016:

> *Demand is on the up. At the moment, the main lithium-ion battery-makers are Samsung and LG of South Korea, Panasonic and Sony of Japan, and ATL of Hong Kong.*
>
> *But China also has many battery-makers. Adam Collins of Liberum, another investment bank, talks of an "inflection-point" in Chinese demand for lithium salts. Its government is stepping up the promotion of lithium-ion batteries and electric vehicles, with the biggest emphasis on*

buses. Sales of "new energy" vehicles in China almost tripled in the first ten months of 2015 compared with the same period in 2014, to 171,000.

This major ramping up in demand has – as you'd expect – translated into a sharp increase in lithium prices. This chart tells the story – pay particular attention to the doubling in prices at the backend of 2015:

Charging up

Lithium carbonate spot price per tonne ($'000)

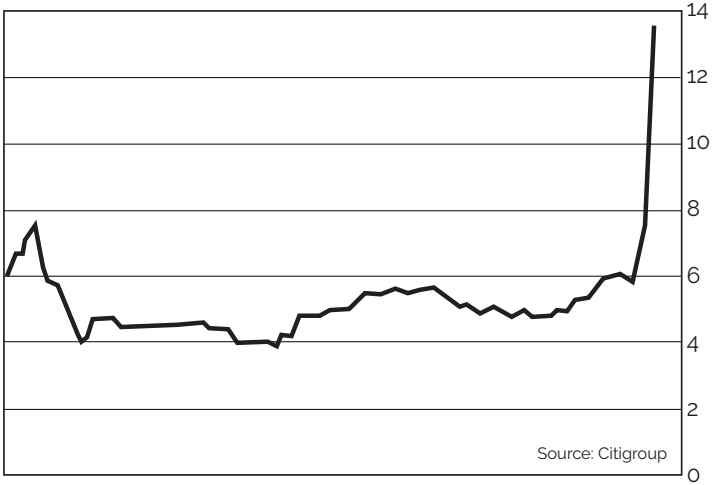

Appendix 17

That kind of massive ramp-up in price tells us a couple of things. It would say it shows signs of panic buying. The price could slide once the panic subsides. But the panic itself tells us that there are people in the market who do not want to miss the boat, and are worried there won't be enough lithium to meet demand.

Perhaps Elon Musk is one of them. As he put it, "In order to produce a half-million cars per year… we would basically need to absorb the entire world's lithium-ion production."

That gives lithium producers a strong hand. I doubt Tesla will be able to corner the entire market. But a message like that is a signal that producers can invest in bringing new supplies online to meet the demand.

Demand for lithium is going through the roof

Lithium has had a good couple of decades already. It's a vital component in the batteries that power our phones, MP3s, tablets and laptops. So the fact that almost everyone owns at least one of these things has translated into a major increase in demand already.

But here's the thing. The batteries in electric vehicles use a lot more lithium than a smartphone. For instance, a Tesla Model S battery has 63 kilograms of lithium in it, the equivalent of 10,000 smartphones.

That's getting lots of people very excited about lithium. As a Deutsche Bank research note put it:

> *This is the dawn of the Lithium-ion Age*
>
> *The commercialization of the lithium-ion battery in the 1990s powered a 20-year surge in the telecommunication and computing industries following the rapid development of light, powerful, rechargeable batteries. As we enter the second half of this decade, the emergence of the Electric Vehicle (EV) is a globally significant thematic based on the same battery technology. Governments are setting carbon emissions targets for the automotive industry whilst also subsidizing EV technology. Beyond traditional demand markets and the emergence of EV, another potential market is beginning to materialize. Battery energy storage on a grid-, industrial-, commercial- and consumer-scale is reaching commercial viability, and rapidly falling battery costs suggest that the Energy Storage sector could grow materially over the next 10 years.*
>
> *Global lithium demand was 184kt in 2015, with battery demand increasing 45% YoY and accounting for 40% of global lithium demand. Based on our analysis, global lithium demand will increase to 534kt by 2025, with batteries accounting for 70% of global demand.*

It's worth noting that most of these figures assume a rapid growth in the popularity of electric cars. Is that a reasonable assumption?

Well, it's certainly true that electric vehicles are becoming increasingly popular. Let's compare what's happening now with what happened when the first motor cars came on the scene a century ago.

In 1915, in the very earliest days of the auto industry, Ford Motors sold 355,000 cars. Sales grew fivefold over the next decade and by 1925 sales were 1.66 million per year.

Now let's contrast that to one of the world's flagship electric car manufacturers, Tesla Motors. Last year, it sold 26,000 of its Model S electric car.

But its new Model 3 – which isn't yet in full production – has close to 400,000 pre-orders. In theory, that's a 1,300% increase on its current sales. That means that if Tesla can fulfil all those orders without a major setback, there's plenty of demand out there for its cars.

No one yet knows whether Tesla will succeed. The company was founded in 2003 and is yet to make a profit. Despite all that demand, it's hard to predict whether it'll be able to compete with big, established car manufacturers. That's a story for another day. But our focus here is lithium. And given that a ramp up in demand for electric cars in coming years seems a reasonable assumption, you'd expect that to translate into increased demand for lithium.

As a Goldman Sachs research note recently put it: "We estimate that a 1% increase in battery electric vehicle (BEV) penetration would increase lithium demand by 70,000mt of LCE/year (or roughly half of current global demand for lithium)."

In short, if the world wants electric cars, the world needs lithium in large quantities. The problem is that right now lithium isn't all that easy to produce economically.

Skyrocketing demand, limited supply

Perhaps ironically, given that lithium-based battery tech could replace oil-based fuels in the future, the nations with lots of oil don't really have a lot of lithium. There's a geopolitical side to this story that's fascinating (and a story for another day). But if the world is to meet its future demand for lithium, there's a fairly limited number of places it can get it from.

For instance, according to Deutsche Bank data, Chile, Australia and Argentina together currently supply 81% of the world's lithium. Behind

them you have the smaller producers – China, the US, Zimbabwe, Portugal and Brazil. No Gulf States, you'll notice. There's no OPEC in the lithium world. Not yet, anyway.

And within that, there are only four major companies that account for the vast majority of global supply. These are Albemarle, SQM, FMC and Sichuan Tianqi, which accounted for 83% of global supply in 2015.

You can bet that given the economics of the situation, other companies (and countries) are trying to develop new sources of lithium. For instance, mining giant Rio Tinto's head of diamonds and minerals Alan Davies recently said: "We're pretty sure that the route to electric cars is through the lithium battery. And as the technology to manufacture them [improves]… then there will be more acceptance, and you bring the price point down."

Therein lies both the opportunity and the risk. Developing new sources of supply requires a lot of investment. Producing it often means extracting it from lithium-containing brine using solar evaporation. Aside from that, researchers are developing new ways to produce it – such as using a dialysis cell with a superconducting membrane. But as yet, this is commercially unproven.

You can expect this situation to change, though. With demand expected to rocket, there's a huge incentive to develop new and more efficient ways of producing lithium. I'm positive that there are people out there figuring out new and imaginative ways of doing just that, even as we speak.

Summary

This industry is fast approaching its tipping point. As costs fall and technology becomes more efficient, the industry will scale up in a big way – putting increasing pressure on traditional fossil fuels. The combination of improving technology, huge investment of capital across the world and an increasingly fractured oil industry creates a huge investment opportunity. The time to move is now.

Chapter 7

Visionaries

Vision is the art of seeing the invisible.

Jonathan Swift

What we anticipate seldom occurs; but what we least expect generally happens.

Benjamin Disraeli

The city had withstood the siege for nine years. Now it was about to fall. Everything that Cassandra had foretold was coming to pass. But no one had listened: Cassandra had been wholly ignored…

Cassandra is one of the most compelling characters in Greek mythology. Hailing from the city of Troy and a daughter of Priam and Hecuba, she has a tragic story. As a child, she is left overnight in Apollo's temple with her brother Helenus. When their parents return for them in the morning, the two children are entwined by serpents that are flicking their forked tongues in their ears. Subsequently, both Cassandra and Helenus are endowed with the ability to see the future.

Years later, when Cassandra grows up, she again spends the night in the temple. This time, the god Apollo himself appears before her. He tries to seduce her, but she rejects his advances. Out of spite and revenge, Apollo curses Cassandra – without removing her ability to see the future. Her powers of prophecy remain intact. Apollo's revenge is more subtle: although Cassandra has the singular capacity of seeing the future, she will forever be disbelieved.

During the Trojan wars, Cassandra is locked away. Her dire – yet accurate – predictions of doom for the city are dismissed (and likely considered bad for morale). She becomes a peripheral and, understandably, slightly mad character in the story.

It's hard to imagine the torment of being able to see the future with

clarity and yet have your every prophecy dismissed as the ravings of a madwoman. Cassandra's story did not end happily. She warned the city not to accept the gift of the Trojan Horse, and was duly proved correct when Troy was sacked. She was taken prisoner by the victorious Agamemnon and taken back to Greece, where she was murdered by Clytemnestra – Agamemnon's wife.

How does the story of Cassandra in mythical prehistory relate to technology and its role in our future? Perhaps more than you'd think. To be a Cassandra has passed into common parlance in the financial world. It generally means to predict that some devastating market event is coming, but to be ignored.

I know several fund managers who fit the definition perfectly. They're people who warned that the real estate and debt bubbles of 2003-2007 would end in a major financial crisis. More or less the entire financial world – from the highest central banker in the monetary system to regular brokers dealing in stock execution – ignored them entirely.

Unlike Cassandra, very few went mad. In general, the kinds of people who are willing to make such stark predictions are the kinds of people who are happy to swim against the tide of mainstream opinion. They're Cassandras and proud of it. They revel in saying the unsayable. The more people disagree with them, the more they're convinced they're right.

That's relevant to us when we're looking at technology, because the same thing happens in reverse, time and again, when it comes to technology. Rather than dire predictions of terrible events, there's a certain kind of visionary tech expert who deals in prophecies of how technology is going to *positively* change the world.

And here's the thing: they're dismissed just as often. Let's call them "reverse-Cassandras".

There's a reason for this. We discussed it earlier in this book. And that's the simple fact that people have a built-in desire to believe that the world is going to stay exactly the same. So when someone outlines a vision of the future that deviates from that – whether in a good or a bad way – they're inclined to dismiss it.

A tale of two 21st century giants

Nikola Tesla was one of the greatest inventors in human history. His experiments and innovations on electricity were groundbreaking. Without his work, we wouldn't have the radar, the X-ray, the electrical grid or hydroelectricity. He even outlined a model for the smartphone (a century before one would be released) and a hundred other vital technologies.

He's been depicted as the archetypal mad scientist – particularly in Christopher Nolan's *The Prestige,* in which he was played by David Bowie. He was also the perfect example of a true visionary: someone who saw the world as it could be, not as it is.

Tesla had many ideas that were considered "out there". For instance, he drew up plans to transmit both electricity and information wirelessly around the globe for free. In a partnership that could – and perhaps should – have changed the course of human history, Tesla joined forces with JP Morgan in 1901. At the time, JP Morgan was one of the richest men in America. His investments in the oil and railway industries made him immensely powerful.

Morgan had agreed to be Tesla's financial backer. But he ultimately pulled his money off the table – reportedly when he found out that Tesla's plans would mean free power for everyone. He's even rumoured to have asked where he would put the meter to charge people. When he didn't like the answer, it was game over for Tesla's plans.

I tell you that story because it illustrates perfectly the constant tension between the visionaries of this world and the people who refuse to believe that the future will be radically different. Some people have too much invested in the status quo – financially, emotionally or intellectually. In JP Morgan's case, it was probably all three. He had no financial interest in the world changing – even for the better. And perhaps it was more than that. Perhaps Morgan just didn't have the imagination or the vision to see how the world could change with Tesla's innovation.

He wasn't alone.

A short history of "that'll never happen"

Imagine you could flit back and forth through time, visiting the biggest technological and scientific breakthroughs in history.

It would be quite a trip. And what you'd notice is that every time our best and brightest have made a breakthrough, you'll likely find someone standing just out of sight, shaking their heads and saying, "That'll never work."

Variations include: "Maybe one day, but not today", and "Anyone who thinks that's possible is a liar or a fool". Sometimes, these people are just natural cynics. They want to see that something is a "sure thing" – a 100% winner – before they come around to believing it possible. These people make terrible early-stage investors, as nothing is ever a sure thing before it comes to pass.

Don't make the mistake of believing that this is about intelligence, either. Cynicism and intelligence don't necessary go hand in hand. Sometimes smart people are just too committed to the status quo. They have too much invested in the way things are and can't bring themselves to accept that the world will ever change. These are the folks who get themselves on the wrong side of the gale of creative destruction.

History offers us plenty of examples…

> ***"There is no reason anyone would want a computer in their home."***
>
> *Ken Olson, president, chairman and founder of Digital Equipment Corp. (DEC), maker of big business mainframe computers, arguing against the PC in 1977.*
>
> ***"Lee De Forest has said in many newspapers and over his signature that it would be possible to transmit the human voice across the Atlantic before many years. Based on these absurd and deliberately misleading statements, the misguided public has been persuaded to purchase stock in his company ."***

A US district attorney prosecuting American inventor Lee De Forest for selling stock fraudulently through the mail for his Radio Telephone Company in 1913.

"The horse is here to stay, but the automobile is only a novelty – a fad."

The president of the Michigan Savings Bank advising Henry Ford's lawyer, Horace Rackham, not to invest in the Ford Motor Co. in 1903.

"How, sir, would you make a ship sail against the wind and currents by lighting a bonfire under her deck? I pray you, excuse me, I have not the time to listen to such nonsense."

Napoleon Bonaparte, when told of Robert Fulton's steamboat, 1800s.

"When the Paris Exhibition [of 1878] closes, electric light will close with it and no more will be heard of it."

Oxford University Professor Erasmus Wilson.

"Dear Mr. President: The canal system of this country is being threatened by a new form of transportation known as "railroads" As you may well know, Mr. President, "railroad" carriages are pulled at the enormous speed of 15 miles per hour by "engines" which, in addition to endangering life and limb of passengers, roar and snort their way through the countryside, setting fire to crops, scaring the livestock and frightening women and children. The Almighty certainly never intended that people should travel at such breakneck speed."

Martin Van Buren, Governor of New York, 1830.

Partly, this comes down to people's mindsets. Our bias is towards the world we see and not the world as it could soon be. But it's also a function of exponential growth. By its very nature, exponential growth is slow at first, then explosive. It means that things can appear not to change at all, before changing suddenly and radically. This takes even the world's smartest people by surprise.

For proof of that, just consider the example of the Human Genome

Project. As Michael Malone's best-selling book *Exponential Organisations* explains:

> *In 1990, the Human Genome Project was launched with the aim of fully sequencing a single human genome. Estimates called for the project to take fifteen years and cost about $6 billion. In 1997, however, halfway through the estimated timeframe, just 1 percent of the human genome had been sequenced. Every expert labelled the project a failure, pointing out that at seven years for just 1 percent, it would take seven hundred years to finish the sequencing. Craig Venter, one of the principal researchers, received calls from friends and colleagues imploring him to stop the project and not embarrass himself further. "Save your career," he recalls them saying. "Return the money."*
>
> *When Ray Kurzweil was asked his perspective, however, his view of the "impending disaster" was quite different. "1 percent," he said. "That means we're halfway done."*

Kurzweil could see what others couldn't. He'd recognised that the amount of the human genome being sequenced was growing exponentially – doubling every year. The project had completed 1% so far. It had seven years left to finish its task.

Double one percent seven times and what do you get?

That's right. One hundred percent. It turned out that completing 1% of the project was the halfway mark.

The project was completed early, and below budget. As Ismail says, "The so-called experts had missed the end point by 696 years."

Or to put it another way, they were tricked by the deception phase. They looked at what had been achieved and saw linear growth. "They've only completed 1% of the project. They only mapped 0.5% of the genome in the whole of last year. Things are moving too slowly! This is going to take forever!", they thought to themselves. They were blind to what was really happening.

That was the first time we'd mapped the human genome. It was a huge achievement. It took more than a decade and cost a fortune. These days, anyone can get their genome mapped in a couple of weeks. It costs about £100.

Are you a visionary?

At the start of this book, I asked you to do something: to control your "inner sceptic" and keep an open mind about the future. Radical change is the norm, I told you. Don't forget that. Don't let the "not in my lifetime" fallacy blind you to what's coming.

I wonder now if you succeeded.

Perhaps I was wrong to call it the inner sceptic. Scepticism is good: it helps you challenge new ideas, think deeply and carefully about things before you put your faith in them. It is cynicism that's the real danger: dismissing things out of hand before you've even considered them properly.

Given some of the incredible technologies we've discussed in this book, I imagine that your inner cynic has had a fairly good workout. You've heard about some of the technologies that could completely and utterly change the world – artificial intelligence, robotics, biotechnology, solar energy – from the experts who are best placed to explain what's happening.

I'm proud of the work we've done and the people we've spoken to. Yet deep down, I know that some people reading this book won't actually take the final step, be positive and invest in these breakthroughs. Something will hold them back.

That just seems to be the way of the world. The number of people who see what's coming and make the right investments at the right time always seems to be small. Why is that? I don't think it's a case of intelligence, or inside information, or luck. I think it's something more profound.

It comes down to how you see the world.

Are you excited about the future – or afraid of it?

We all have this battle going on inside us. There's always a part of us that wants to ignore or dismiss the idea that no matter how bad the news gets, technology and innovation are going to change the world in a big way. In some people the sceptic wins out. In fact that's the case for most people – despite the fact that all of the evidence in front of us

suggests that technology is remaking the world at an accelerating rate, just as it has done for decades.

Look at how much the world has changed in the last ten, twenty, or fifty years. Advances in science, technology and medicine have reshaped the way we live our lives. As AI expert and futurist Ben Goertzel writes:

> *We have astoundingly more powerful computers* [than 1970]*, we have the internet, we have far better brain-imaging tools and a much better understanding of neuroscience, and we have a cross-disciplinary field of cognitive science… we have single-molecule electric motors, we can do origami with DNA, we can build synthetic organisms via stringing together amino acids, and there are commercially available quantum computers.*
>
> *We can search over a trillion Web pages almost instantly from our laptops or mobile phones. We can hook tens of thousands of insanely fast multiprocessor computers together into functional units. Our guided missiles can fly in the dark to precisely specified locations and destroy precisely what they want to… We can clone animals, and cause human organ regeneration via stem cell injections.*

As you've just read, sweeping changes to the world we live in aren't the exception. They're the rule. In the last decade alone, we've seen a period of such technological advancement that it would shock someone born 100 years ago. I believe we're about to live through another one – perhaps the most important era of change and progress in human history.

An era of free and unlimited energy for everyone on the planet – making fossil fuels, oil wars and energy pollution a thing of the past. An era of biotechnology and genetic editing winning the fight against disease… and extending human life expectancy far beyond 100. An era of high-powered artificial intelligence that will help solve the biggest problems. And an era of humans and technology merging, blurring the line between people and robots.

So, what part of you wins out? The cynic or the visionary? The part of you that understands radical change is constant… or the part that wants everything to stay the same?

I think it comes down to the kind of person you are.

Are you the type who can see what few others can… who sides with visionaries like Nikola Tesla, Steve Jobs, Bill Gates and Elon Musk… who are bold enough to move early?

Or do you prefer to "wait and see" – to wait until something's a "dead cert" before you act?

Because the truth is, there is no such thing as a dead cert. And waiting until everyone can see how big a technology is going to be means you'll likely never make the kind of life-changing returns that are possible when you make the right tech investments at the right time.

I can't answer for you. But I can say that I wrote this book for the visionaries out there: the people who are willing to be bold, take a risk and get in at the same time as the other early movers.

Is that you?

I hope so.

Afterword

By Eoin Treacy

Investment Director, *Frontier Tech Investor*

Nick O'Connor has provided a highly comprehensive illustration of just how exciting an era we are living in. I myself often look up to the skies and think how lucky I am to live through what is indisputably a Golden Age of invention and discovery. The world is pregnant with opportunity – provided we are prepared to adopt an optimistic perspective, view the future with a glass-half-full attitude, and above all, act on our convictions when the time is right.

I have unbridled enthusiasm for the future that technological innovation has the potential to deliver. But let's not ignore the fact that, as a society, we face substantial problems.

Our governments are broke. Our politicians have made commitments that we cannot pay for, not least in the area of pensions and healthcare. We are not breeding quickly enough to supply the new workers required to support our social services. The climate is changing all around us. The entire global food chain is at risk because of ocean acidification. Inequality is a major political concern for Western democracies, and terrorism is an increasingly potent menace.

Let's take a fresh look at the changes that technology can engender. If we continue along our current trajectory, we will be broke, alone, and facing potential starvation. Not all of us at once, of course. But it would be wholly irresponsible to ignore the challenges that we face.

Technological innovation is the only way that we can achieve the prosperous, equal-opportunity-driven, free society that we all aspire to. This is not a matter of choice. It is a matter of absolute necessity. We need solutions to the problems that we face as inhabitants of this planet, and technological innovation can deliver those solutions.

What this means is that there are going to be wonderful investment opportunities along the way.

Take healthcare. Thanks to the astounding advances in the field, we are all much more likely to lead long and healthy lives, and to fully enjoy the fruits of our labour. As far as economists are concerned, every last one of us is going to be indispensable to this process, because the efficiencies demanded of our evolving economy will require all of us to play our part. The practice of knocking people out of the workforce because of a chronic medical condition is just not going to work any longer. We are on the threshold of a radical and dynamic shift. And those who are able to recognise the opportunity represented by that shift – to embrace it rather than fight it – are the ones who will benefit the most.

Saudi Arabia's late oil minister Ahmed Zaki Yamani showed incredible insight when, in the year 2000, he declared: "The Stone Age came to an end not for a lack of stones, and the oil age will end – but not for a lack of oil." He was nearly two decades ahead of his time in predicting an era of peak demand for oil.

We see it happening now. Within the next decade or two, there is a very real possibility that oil will be displaced as the primary source of transportation fuel. That represents a massive economic and societal shift. Oil is by far the most important commodity in the world today. It's used in every country and has a massive influence on the global economy. But after more than a century of being a most precious and sought-after commodity, it may well have climaxed.

All other things being equal, every $10 drop in the oil price represents a 0.2% boost to economic growth over the coming decade. If in the next few decades nothing else changes and oil prices hover around the $40 to $60 mark – instead of around the $100 to $120 mark – we can expect economic growth to get a boost of more than 3% a year. That kind of growth makes many of the debt problems we hear about seem much less worrisome. Considering how much spare capacity now exists in the form of unconventional wells and the potential for electric cars and batteries, the prospects described above are a real possibility. We are already seeing them materialise.

The reason that this is truly a Golden Age for humanity is that a number of themes are coalescing to form a greater whole. Energy is the foundation, but communications and healthcare are producing equally momentous innovations. The three of them combined are delivering a future quite unlike anything that has come before in human history.

After a decade of sky-high energy prices, we are in the efficiency and conservancy part of the supply and demand cycle. We've spent so long living with expensive energy that all that anyone wants to talk about is how efficient their product is and how little energy it uses.

But once we start experiencing consistently low energy prices, the mindset will change. Many more activities will become possible. You will be able to afford to leave the lights on, to fly more often, to live farther away from work, and to buy toys with more bells and whistles.

In short, low energy prices have implications that extend well beyond the realm of transportation. They have the potential to impact the Internet of Things – the moniker for putting Wi-Fi-enabled sensors in everything. Technology is delivering the products that contain the sensors. But cheaper energy means that the products themselves will be cheaper on a unit basis.

At the moment, we have Wi-Fi-enabled cars, but not much for them to communicate with. Within a decade, the on-board computer in your car is likely to make today's supercomputers look like clunky Lego parts. Not only will the car drive itself, but it will act as an entertainment centre, a mobile office and a personal assistant. Meanwhile, it will probably cost about the same as the car you drive today, and be much cheaper to run.

The transformation that is likely to take place in our homes and workplaces will be on a scale similar to the introduction of electric light. At the moment, hospitals probably have more sensory equipment than any other place we are likely to come into contact with. Yet they depend on carts for stability, and are inherently noisy.

The next generation of sensors is going to be smaller, cheaper, more connected, and completely ubiquitous. They will be in our homes, schools, shops, streets, utilities, airports, train stations, bus stops, farms, forests and mountains. We will have real-time sensory information for every aspect of our lives and environment.

Digital information and our ability to process it creates a chicken-and-egg situation. Once upon a time, we had some computing power but very little digital information. Today we have oodles of raw data being created every minute and are racing to develop the technology necessary to make sense of it.

Freakonomics, a bestseller published in 2009, explored previously unnoticed quirks in the study of economics. Similarly, it is reasonable to expect that big data will deliver on discoveries we might never have considered before – simply because we never had tools powerful enough to even ask the question. Rather than being led by academic dogma, big data is being pioneered by the private sector as a for-profit enterprise. The conclusions reached will have nothing to do with appeasing some professor's sensibilities. They will only be measured by their utility.

In that regard, CRISPR, optogenetics, and rapid advances in immunotherapy are beginning to deliver on promises made during the tech boom of the 1990s. I spent this last weekend explaining what a degree in biomechanics was to my ten-year-old daughter, because within the decade this could represent the most exciting sector to develop a career in. I find the calm, measured, risk-adjusted way of approaching the clear ethical questions of how we should develop the genetic sector to be inspiring.

Mrs Treacy (my wife) loves South Korean soap operas, and we have often remarked on how similar the actors look. The fact is that South Korea is the world's most buoyant plastic surgery market, and clinics are churning out thousands of lookalike people on an annual basis.

This came to mind as I was reading the very enlightening section on the ethical ramifications of genetic engineering and how they might affect humanity. It is worth considering that not all cultures share the same sensibilities, traditions or priorities. I have met single Chinese women who have come to the US for in vitro fertilisation to ensure that they have a mixed-race child. With such high Asian demand for customisation, there appears to be inevitable potential for a race to develop, not only in personal medicine, but also in personalised genetic services designed for us as well as the next generation.

As a society, we waste more food than at any other time in history. Yet we have the nerve to preach to people who have less than we do about what they should and shouldn't eat when it comes to the benefits of genetically modified food. We agonise over whether we should interfere with the DNA of the small number of mosquito species that spread disease from a genus of more than 3,000 varieties. At the same time, few of us have seen a loved one die from malaria, or seen our societies ravaged by the loss of its young and vibrant members .

When we try to look at the world from the perspective of those who aspire to the standard of living we have now – never mind the standard of living we'll have decades from now – we understand that they are in a hurry and will pursue whatever means necessary to remain competitive and achieve their goal. Those who doubt that technological innovation will change the world should think about the motivations of corporations to develop markets for their products and the aspirations of billions of new consumers striving to create a better life for themselves and their children and grandchildren. That represents a powerful partnership, and virtually ensures that innovation will persist. It really would be better to be on the side of those who can shape how it develops before newly emerging technological powers overtake us.

For example, Google announced in July 2016 that it is going to train two million programmers in India as it ramps up competition with Apple for the most complete app library. India is expected to have over four million programmers by 2018 – which will make it the number one country for programming, overtaking the US. Right now, only about a quarter of them are targeting mobile. The benefits that could accrue to a country such as India once the revolution in energy, the Internet of Things and the evolution of genetics coalesce to drive productivity growth are truly mind-boggling.

This liberation of human potential is possible everywhere because of the omnipresence of the internet. Today, Africa accounts for the bulk of world population growth, and this large new young population is mobile-savvy. Their first experience of banking, shopping, entertainment and communication is via their mobile phones. The leaps we are getting not only in technology but in its availability to the whole world means that the potential to alleviate poverty, to raise living standards and to boost fraternity is higher than it has been in decades if not centuries.

I spend a great deal of time trying to understand the problems that will be created by this transformation. It is hard not to be excited by the possibilities – and I look forward to exploring them all with you.

Appendix

Appendix 1: An image showing the power of an iPhone versus its technological predecessors. Source: Twitter.

Appendix 2: Bar chart showing people adopting technology at an accelerating rate. Source: Bonner & Partners.

Appendix 3: A bar chart showing the number of years a tech company's market capitalisation reaches one billion versus Typical Fortune 500 companies. Source: Bonner & Partners.

Appendix 4: Redrawn chart of Kurzweil's predictions about the future taken from his book *The Age of Intelligent Machines*, 1988.

Appendix 5: Bar chart showing worldwide spending on Robotics is expected to reach $67 billion by 2025. Sources: International Federation of Robotics, Japan Robot Association; Japan Ministry of Economy, Trade & Industry; euRobotics; company filings; GCG analysis.

Appendix 6: Bar chart showing the decline in cost per genome since 2001. Source: National Human Genome Research Institute.

Appendix 7: Chart showing the population of the planet over the past 10,000 years. Source: BBC.

Appendix 8: Chart showing the nominal and adjusted food prices since 1961. Source: Food and Agriculture Organization of the United Nations.

Appendix 9: Michael Eisen's tweet regarding Eric Lander's "The Heroes of CRISPR". Source: Michael Eisen, Twitter.

Appendix 10: Chart showing the performance of the SPDR S&P Biotech ETF. Source: Bloomberg.

Appendix 11: Charts showing the unconventional oil output, paying special attention to the years between 2010 and 2013. Source: api.org.

Appendix 12: Chart showing the price of solar energy has collapsed and as a result an increase in installations has occurred. Source: Earth Policy Institute/Bloomberg.

Appendix 13: The price per watt of electricity generated by solar has more than halved in the last decade.

Appendix 14: Pie chart showing an increasing number of people using solar energy as a viable energy source. Source: SEIA/GTM Research and FERC.

Appendix 15: Chart showing between 2004 and 2014, American capacity for solar increased 40 times. Source: US Energy Information Administration's Monthly Energy Review, table 7.2.a.

Appendix 16: Chart showing the cumulative global capacity (MW) of battery storage by technology as listed on the US Department of Energy database. Source: arena.gov.

Appendix 17: Chart showing lithium prices have increased sharply in 2015. Source: Citigroup.

Huge change is coming. The kind of change that will radically reshape the world around you: your life, health, work and investments. In *The Exponentialist* you'll discover the technologies driving this change.

From robots in the home to solar powered skyscrapers, from super-intelligent machines to rewriting the laws of life, everything you need to understand about the world's most important technology is explained – with insights from the innovators, researchers, entrepreneurs and investors who're making it happen.