




ARTIFICIAL INTELLIGENCE

How AI is Reshaping the World of Business



Julian Birkinshaw
Vice Dean and Professor,
London Business School

CIST

Foreword

Recently, I received a provocative question from Fast Company: What concerns me most about AI and leadership?

"My biggest concern is me! As leaders, we often rely too heavily on our past successes, our abilities, and even our intuition. I am actively unlearning and embracing the unknown, striving for a not-too-distant different tomorrow. With unwavering optimism and curiosity, I ensure these feelings outweigh any lingering fears. I feel positive about the journey while always remaining self-alert."

In Professor Julian Birkinshaw's paper, he delves into AI's profound implications for business, work, and society — a journey of unlearning and embracing the unknown.

Birkinshaw fosters a nuanced understanding of the long-term consequences of this technology, with thoughtful provocations, and practical insights on how businesses can strategically leverage AI.

Join me in co-thinking the immediate and enduring impacts of AI on the business world so that we can shape a future where businesses thrive in the face of change.

To the journey ahead,



Bob Wollheim


Partner, EVP - CI&T

Introduction

We live in exciting times, and the arrival of Generative AI as a consumer product opens up a wealth of challenges and opportunities. We shouldn't underestimate the long-term implications of these tools on our lives. But we also need to avoid scaremongering.

For most businesses, Generative AI will be a sustaining technology not a disruptive one, meaning that it will help you get stronger and more productive as long as you harness it thoughtfully.


My intention with this short paper is to lay out some of the likely consequences of this technology and to offer some practical thoughts about how you might use it.





Technologists and futurists have been discussing the impact of Artificial Intelligence (AI) on the business world for decades, but the launch of Chat GPT in November 2022 has brought these conversations into the mainstream. We now have instant access to AI; we don't need any advanced IT skills, we don't need any understanding of how the technology works; we simply use it and marvel at what it comes up with. The obvious comparison, in terms of historical precedent, is the World Wide Web, created in 1993 by Tim Berners-Lee. This invention turned the Internet – an obscure system linking up researchers and university scientists around the world – into something consumers around the world could access. Thanks to Chat GPT, AI is now a consumer product in its own right, creating enormous opportunities but also significant challenges.

There are some obvious consequences when you put AI in the hands of millions of consumers around the world. First, awareness and usage increase exponentially, opening up ways of applying the technology that its inventors could never have dreamed of. Second, commercial interests take over, with big-tech and start-up companies seeking out ways to exploit the technology for financial advantage.




Third, the limitations of the technology are exposed, sparking a useful debate about risks and biases. Regulation and oversight – some useful, some burdensome – quickly follow.

All these things are just starting to play out. The invention of the World Wide Web in 1993 kicked off the Internet revolution, the emergence of Amazon, Google and Facebook and others; and the reinvention of Microsoft and Apple. As the consumerisation of AI takes off, we will see some new players emerge (like Open AI, the company behind ChatGPT), though I expect the big tech players to retain their dominance in many respects.


There will also be as-yet unforeseeable opportunities. Consider the analogy to the Internet, where our openness to sharing information about ourselves and our intentions online created what Shoshana Zuboff calls a "behavioural surplus," which was then exploited by Google and Facebook through paid ads¹. What new behaviours will we develop, as we increasingly use Chat GPT to help us with our everyday lives, that entrepreneurs might harness for their own benefit? What new "killer apps" will emerge to disrupt existing sectors? Great riches await those who come up with the best answers to these questions.

One thing is for sure – the arrival of AI as a consumer product isn't the end of the story, it's just the inflection point that heralds even more change, as the underlying technology continues to improve, and as new possibilities come into focus.



For example, in July 2023 Mustafa Suleyman, one of the founders of Google DeepMind called for a new "Turing Test."

You might recall the original Turing Test, named after Alan Turing, an AI pioneer from the 1950s. It was based on a simple language conversation between two invisible partners. If the observer couldn't figure out which was the computer and which was the human, then the computer passed the test.



Suleyman said that in terms of text conversation (and indeed many other skills such as voice and image recognition), AI has already passed the Turing Test, so we need to raise the bar. It is not enough for AI to be able to process and generate information – it needs to show what it can do. He says:

"Go make \$1 million on a retail web platform in a few months with just a \$100,000 investment."²

Clearly this is a big ask – it would require the AI to design a new product, work with manufacturers, come up with marketing campaigns and so forth. My personal view, as I discuss below, is that some of these tasks may be forever beyond the realms of possibility. But the point is, if one of the leaders in AI is already talking in these terms, we should take the challenge seriously.

If AI is moving from the world of thought to the world of action, there are important consequences for all of us.



In this article, I explore some of the immediate and longer-term consequences of AI for the world of business. I certainly don't have all the answers, but based on my 25 years as a researcher and consultant in the fields of management, strategy and innovation, I can provide a sense of direction and I can put forward some useful provocations.

¹Zuboff, Shoshana. *The age of surveillance capitalism: The fight for a human future at the new frontier of power: Barack Obama's books of 2019*. Profile books, 2019.

²Suleyman, M. 2023. *My new turing test would see if AI can make 1 million*. MIT Technology Review. July 14th.



Generative AI

What does it do?

Most of us have no idea how AI works, but of course, we don't know how the iPhone or the Internet works either. However, it is important to know roughly what happens when we use Generative AI products like Chat GPT, to make sure we understand their limitations as well as their benefits. So a couple of points here.



First, Chat GPT does not think or understand in the way a human does. It uses an algorithm to generate responses based on an enormous body of data it has been trained on. It uses probabilities (what is the next most likely word in this sentence?) to create seemingly fluent text and may make factual errors and incorrect statements since it doesn't understand what it creates.

Second, Chat GPT is what techies call narrow AI – it is very good at one thing: excels at generating text responses, but it cannot do anything else (for example, driving a car). We are still a very long way from General AI, which can perform multiple functions operating at the level of a human. I don't expect to see evidence of General AI in my lifetime. But I could be wrong! The only thing we know for sure is that AI will continue to improve, and in sometimes unpredictable ways.



Business Productivity

How can we become more efficient?

AI offers the potential for huge gains in productivity, at both the individual and business level.

Consider the individual level first. Vast numbers of people are now using Generative AI in their daily working lives – a recent McKinsey study reported that 40% of respondents were using these tools regularly³.

³ The state of AI in 2023: a break out year. McKinsey and Co.

And we have all got our own examples of how our productivity has improved. For example, earlier this year I needed to write a 1000-word introduction to cloud computing. With a few careful prompts and a couple of iterations, about 10 minutes of work, Chat GPT created a draft text. I edited this text, adding in a few ideas of my own, and the job was done-and-dusted in less than an hour. Last year it would have taken four times as long. The same productivity boost is available in many professions. The same McKinsey study suggested marketing and sales, product development and service operations are the places Chat GPT is used most, for example in writing first drafts of documents, identifying trends and customer needs, and using chatbots (to talk to customers).

At the business level, we would expect the productivity benefits of AI to be even greater. But the story here is a bit more complex.

Let's take a step back, and consider the digital revolution as a whole, rather than focusing narrowly on Generative AI, as productivity improvements take time to transpire. Over the last thirty years, we have seen dramatic improvements in the collection, sharing and application of data, both within and across firms. This has made business more efficient – just think how easy it is to pay for your purchases, send money to a friend, or monitor your personal banking. It has enabled firms to become larger and flatter, as people on the front lines have much richer access to information about what they are doing and why. This data can





also be analyzed far more effectively than before – AI helps to optimize workflows, reduce errors, and focus effort where it is most needed.

As mentioned earlier, there is also a shift underway from knowing to doing. AI-enabled digital assistants arrange meetings without human intervention. Robo advisors don't just recommend stocks, they execute trades. So-called Smart Contracts are designed to execute transactions once pre-specified criteria have been met. There have even been experiments in creating Decentralized Autonomous Organizations (DAOs) that fulfil the function of traditional organizations but without human decision⁴.

In sum, digital technology is enabling the creation of structures and systems that get work done in a hyper-efficient way, which should have enormous benefits in terms of lower production costs, fewer errors, and less time spent (by humans) on repetitive tasks. However, there are a couple of important caveats and concerns.

First, we have the risk of Neo-Taylorism. You might remember FW Taylor was the father of Scientific Management, a technique developed in the 1920s to optimize the work of manual laborers by telling them exactly how to do their jobs – and essentially dehumanizing them in the process. With today's data sharing technology, Amazon drivers and call-centre workers can be monitored in real-time by their supervisors in ways that Taylor couldn't even have imagined.

⁴ Buterin, V. 2022. DAOs are not corporations. www.vitalik.ca. 20th September.

This may in the short term be good for productivity, but it is bad for morale and for our responsibility to employees to create meaningful work for them. I will return to this point later.

Second, despite all the gains in efficiency that we see for ourselves, it is a curious fact that the aggregate statistics on productivity don't show any improvements. In the UK, for example, productivity growth has stalled completely since 2008. US Bureau of Labor Statistics says TFP was 0.4% per year from 2005 to 2018. The UK was even worse, 0.1% per year over the same period.

Many explanations have been offered for this "productivity puzzle". Some observers note that it takes a long time for the benefits to be realized. Others say our measures don't capture the benefits of IT investment properly. But a big part of the story, I believe, is that we haven't made the necessary changes in management techniques or job design to capture all these potential benefits.





So what's the advice here?

How should you use AI technology to improve your business productivity?

It goes without saying that you need to invest in these technologies, but there are a couple of further points to make.

First, take a pragmatic fast-follower approach to new digital technologies.

It is tempting to seek some sort of first-mover advantage with new technology, and in a few isolated cases that might work. But the risks of investing in the wrong technology, or getting in before it's ready for use, are huge. Most new technologies are talked up long before they are commercially viable – think about the Metaverse, Blockchain, Web3, and Quantum Computing, all of which have failed to live up to their hype. I am not saying ignore these new technologies – you need someone to keep an eye on them. But you are better off focusing your energy and investment on the ones that definitively work, for example, predictive analytics, Robotic Process Automation, and (since November 2022) Generative AI.

Second, invest in workplace redesign.

There is academic research showing how the productivity benefits of Information Technology accrue only to those firms that also rethink how they do their work. Or stated differently, if you try to optimize an outdated process using technology, you end up frustrating your employees rather than enabling them.

Take the example of agile working – self-organized and multi-functional teams who work iteratively with their users to create outputs. It seems clear that agile teams will have a different view of how to use new AI tools, compared to waterfall-based employees, and you need their input to figure out what this new way of working might look like.

A case in point is the world of professional services, for example, audit, law and consulting. For decades, their business model has relied on huge numbers of trainees working on tiresome entry level jobs, an up-or-out system with professional exams, and pyramid-shaped teams that allow senior people to earn a handsome return. Artificial Intelligence is now challenging this entire model. Earlier waves of technology killed off manufacturing jobs and data-processing jobs, now it's the turn of low-level professional jobs. The knock-on effects in terms of the size and staffing of teams, the skills required, and indeed the profitability of these firms, will be sizable. The firms that adapt their management structures most effectively will be the winners.

The background image features several abstract, three-dimensional geometric shapes, primarily cubes and rectangular prisms, arranged on dark, circular pedestals of varying heights. The shapes are rendered in a dark, metallic purple color with some internal structures visible, suggesting a complex, layered design. The lighting is dramatic, coming from the side, creating strong highlights and deep shadows that emphasize the three-dimensional nature of the objects. The overall aesthetic is modern and architectural.


Business level choices **What is our strategy?**

Everything I have said so far is about using AI (and digital technology more generally) within an existing business model or strategy. It answers the question, how can we serve our existing customers more efficiently and effectively?

But strategy is about making choices, where to play and how to play. It's about beating the competition by coming up with something new that customers want to pay for. It's important to get this right, so we need to ask: how will AI affect business strategy?

One obvious risk is that AI-based technologies will lead to greater convergence among competing firms. They are becoming more efficient. More of their activities are being outsourced and put on the cloud.

Interfaces between firms are becoming standardized. Online aggregators and retailers are encouraging price comparison. The net result of these trends is that executives are making more constrained decisions than in the past, as they are operating within an established ecosystem. There is therefore pressure on them to take decisions that push them closer to their competitors. It's a low-risk, low-reward way of acting.



My point is that these 'convergent forces' have always existed in business but they are being amplified by the emergence of common IT infrastructure and by AI-enabled decision-making processes that push everyone in broadly the same direction.

So what's the role of the leader, the business person who is trying to do something innovative and distinctive? It boils down to doing the things AI cannot – taking a more 'human' approach to strategy that plays on such qualities as creativity, intuition, and empathy. What does this look like in practice? Here are three specific pieces of advice.

The **3 AI essential advices** for **C-level** running businesses




Take a long-term perspective



Focus on purpose



Nurture unreasonable behaviour





Take a long-term perspective

A rational, tech-driven calculus would favour putting investment in those areas where the payoff is most direct. But the most successful strategies are the ones that make calculated bets on the future. Every famous innovation story, from IKEA to the iPhone to Tesla, was a huge bet on something that wasn't obvious at the time, and required a lot of upfront investment.



Successful firms, in other words, take resources away from their short-term best use, in order to give themselves the chance to create even more value over the long term. This "one step back, two steps forward" logic manifests itself in many ways – risky R&D projects, pursuing sustainability goals, paying above-market wages to improve loyalty, and so on. We actually take it for granted that firms will do many of these things, but my argument is that they involve judgments that AI is ill-equipped to help us with. AI can devise seemingly-cunning strategies that look prescient (consider the famous case of AlphaGo) but only when the rules of the game are pre-determined and stable.

It is not easy to provide specific advice on how to do this, but the starting point is an innovative or disruptive frame of mind, or what Amazon's Jeff Bezos calls a "willingness to be misunderstood".

Focus on purpose

There is a second dimension to long-term thinking, and that is its impact on individual and team motivation. Purpose is a "moral or spiritual call to action" that leads people to put in discretionary effort – to work long hours,



and to bring their passion and creativity to the workplace.

This notion that a firm has a social quality, a purpose or identity, is well established. But it still arouses suspicion among those who believe that people are motivated largely through extrinsic rewards. My view is that you just need to look at charities, open source software movements, and many other not-for-profit organisations to realise that many people work harder when money is not involved. And it is the capacity of a leader to articulate a sense of purpose, in a way that creates emotional resonance with followers, that is uniquely human.

Successful firms, in other words, institutionalise a sense of identity and purpose that attracts employees and customers. For example, even though firms are experimenting with blockchain technology to build a system that cannot be hacked or misused by a few opportunists, its limited uptake at the current time suggests people still prefer to put their faith in other people.

Nurture unreasonable behaviour

There are many famous cases of mavericks who succeeded by challenging the rules, such as Steve Jobs, Elon Musk, and Richard Branson. With apologies to George Bernard Shaw, I think of these people as unreasonable – they seek to adapt the world to their view, rather than learn to fit in.



And if we want to see firms move beyond what is already known and proven – to create new market opportunities – more of these types of people would be useful.

Unreasonableness is antithetical to the world of AI. Computers work either through sophisticated algorithms or by inference from prior data, and in both cases, the capacity to make an entirely out-of-the-box leap doesn't exist.

As noted earlier, in the world of investment management, robo advisors are not just making trades, they are also providing investment advice to investors, and at a fraction of the cost of human financial advisors. But as the Financial Times said a few years back, "When it comes to investing, human stupidity beats AI." In other words, if you want to beat the market, you need to be a contrarian – you need to make investments that go against the perceived wisdom at the time, and you need to accept the risk that your judgment or your timing might be wrong. These qualities are –at the moment–distinctively human.

So one of the distinctive qualities of firms is that they nurture this type of unreasonable behaviour. Of course, many firms do their

⁵ Johnson, M. 2017. When it comes to investing, human stupidity beats AI. Financial Times. 10th April.



best to drive out variance, by using tight control systems and punishing failure. My argument is that as AI becomes more influential, through the automation of basic activities and simple contracts, it becomes even more important for firms to push in the other direction – to nurture unorthodox thinking, encourage experimentation, and tolerate failure.

Societal Impact

The future of work


In the final part of this article, I will take a quick look beyond the business world. How will AI affect jobs, careers, education, and society more widely?

We know many clerical and professional jobs are at risk through AI. However, it's important to remember that most jobs are made up of a multitude of different tasks, some of which will become obsolete while others won't. I already talked about the need for workplace redesign in business, but the challenge is much wider than that because we will need to rethink whole categories of jobs across the public and private sectors, including the training provided and the qualifications people gain to do them.

It is also well-known that new jobs emerge in the wake of other jobs being killed off,

the famous process of "creative destruction" sketched out originally by Joseph Schumpeter. An article in the Wall Street Journal showed the dramatic reduction in number of book-keepers, audit clerks and accountants from 1988 to today, following the mass adoption of spreadsheet technology (VisiCalc, Lotus 123 and Excel).

This reduction was matched by the dramatic growth in analyst and financial manager roles, as the application of spreadsheets to business challenges skyrocketed. There was also a significant growth in the number of accountants and auditors, to provide oversight of this boom in analytical work.



Which jobs are most at risk now? Clearly, those involving simple analytical or text-based work are under threat – transcription, translation, copy-editing, searches of legal documents, and basic writing tasks, to name just a few. Many call center and customer service roles will also be lost to chatbots, though we all know (from experience) that a human touch will still be vital for some elements of customer service. The number of repetitive manual jobs, for example, cleaners and check-out assistants, will continue to fall as automation becomes smarter.

Equally, there are many jobs unaffected by AI. High-end knowledge workers, from CEOs to lawyers to engineers, are paid to exercise judgment and to get things done through influence and persuasion – innately human qualities. Creative types in ad agencies, R&D and entertainment have nothing to worry about. There are also many poorly paid jobs, for example nursing and care work, where emotional intelligence and empathy are central to the work, and seem destined to remain in the hands of human workers.

Many books have been written about our innately human qualities, the things we are good at that AI cannot match: creativity, intuition, relationship-building, judgment, and empathy. Whatever job you have, it's good to emphasize and develop these qualities. These were always important capabilities, but with the AI revolution underway they are more important than ever.

So what are the implications of all this for government and society?

First, there is a need for new regulation.

Like any new technology, Generative AI creates huge worries for regulators, because they don't have a framework for defining the necessary guardrails around its use. Elon Musk and other AI experts called for a pause in developing these models in early 2023. The Italian data regulator briefly banned Chat GPT, before changing its mind. US and EU regulatory agencies (at the time of writing) are now looking into various issues, for example, whether Chat GPT infringed data protection or IP rules in building its model, and how to ensure its output doesn't stray into illegal and dangerous territory.

We can expect some significant new regulations over the next couple of years, and indeed this has been welcomed by many players, including Sam Altman the OpenAI founder. But it will be messy – we know from previous generations of new technologies that agile tech firms are good at finessing and sidestepping the efforts of regulators who try to rein in their activities.

Second, our education systems need upgrading and rethinking.

As touched on earlier, students can now produce the lowest-common-denominator output with almost no knowledge of what they are doing, so we need to come up with new ways of educating people on frameworks and concepts and how to apply them. Of course, we have been here before – the pocket calculator made slide rules obsolete and mental arithmetic less important, so the way math was taught and assessed adapted accordingly. Students will need to learn how to harness AI effectively, which shouldn't be that difficult. They will also need to learn how to be more creative, intuitive and empathetic, which is a much bigger challenge.

Third, creative destruction (new jobs emerging after old jobs are killed off) never happens without some pain.

The people who lose their jobs aren't necessarily able to build the skills to do the new jobs, some of the new jobs will be in the wrong places and at lower salary levels, and some companies will go bust. Again, we have been here many times before (for example, the collapse of the coal-mining industry), and the government has a vital role to play in helping individuals and communities through these difficult transitions.

Finally, there is a risk of further inequality in society

With those in less well-paid and less secure jobs the likely losers as AI becomes more pervasive across businesses and communities.

Best-selling author Yuval Noah Harari has warned of a "great decoupling" between intelligence and consciousness – he says we have reached a point, for the first time in human history, where intelligence (problem-solving ability) and consciousness (subjective experience of reality) aren't in lock-step. And the risks that those in the lower strata of society, who lack education and/or access to technology, get left further behind. This would exacerbate inequality problems further – and inequality is already the source of many of society's ills. No easy remedies exist here, but clearly, governments of all persuasions know they have to find ways of balancing the gains from technology and redistributing wealth to keep society together.

⁶ Harari, Y.N. 2017. Homo Deus: A brief history of tomorrow. Penguin Random House.



Julian Birkinshaw

Vice Dean and Professor,
London Business School

Professor of Strategy and Entrepreneurship at the London Business School. Fellow of the British Academy (FBA), the US Academy of Management (AOM), the Academy of Social Scientists (FACSS) and the Academy of International Business.

Professor Birkinshaw has a PhD from Western University, Canada. He is a Fellow of the British Academy, Strategic Management Society and the (American) Academy of Management. He has honorary doctorates from Copenhagen Business School and Stockholm School of Economics.

Professor Birkinshaw is an internationally-renowned academic who has made important contributions to the research on innovation, digital transformation and agility of large multinational firms. He believes strongly in the practical application of academic knowledge to business issues. He is committed to helping students and alumni of LBS become agents of positive change in the business world.



cist

ciandt.com