



Generative AI: the 5th revolution

A deep dive into the technology that is rapidly transforming our society



Kilian Niemarkt Client Portfolio Manager Global Equity team
RBC Global Asset Management (UK) Limited




Jeremy Richardson Senior Portfolio Manager Global Equity team
RBC Global Asset Management (UK) Limited

Published September 2023

Generative Artificial Intelligence (“AI”) has gained significant attention in recent months. The technology has demonstrated its ability to create human-like content and brings with it the potential to transform aspects of our society, including business.

OpenAI announced that its release of ChatGPT acquired one million users in just five days following its launch, making it the fastest growing online service in history¹. ChatGPT went on to achieve one billion visits in its first three months and was three times quicker to hit that milestone than TikTok, and ten times quicker than Instagram, according to Bank of America. Management teams are noticing and urgently considering the potential for enhancements or disruptions to their companies. The number of S&P 500 companies citing AI during first quarter results calls almost doubled compared to one year ago and has dramatically increased over the past decade (Exhibit 1).

Exhibit 1: Number of S&P 500 companies citing AI on earnings calls



Source: Factset, May 2023.

¹ <https://reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>

The initial market response has been to reward those businesses that stand to directly benefit from increasing investment in AI with higher share prices. Certain companies within the semiconductor industry stand out, in particular, NVIDIA whose H100 chips are regarded as the most capable at performing the complex computations demanded by AI. NVIDIA stunned the market with a 50% increase in its revenue forecasts with its first quarter results this year². This is made even more remarkable by the fact that it had only updated these forecasts twelve weeks previously, showing the pace at which interest in AI has converted into orders.

“As well as the financial benefits that AI can bring to companies, investors and other immediate stakeholders, there is also the prospect of more qualitative benefits for society more broadly.”

Few businesses are as directly involved with AI as NVIDIA, but the market is nevertheless also rewarding those that have demonstrated a way of monetising AI, such as through the sale of AI productivity tools. However, for most businesses, the outlook from AI remains clouded, with excitement for new business opportunities, efficiencies and cost savings balanced with concern over new competitive threats.


Those businesses where employee costs as a percentage of revenue are significant, typically exceeding 30% (for example, software businesses or consumer services), are seen as the most likely to benefit from cost savings and productivity enhancements, but there are also opportunities to expand revenues.

Gartner, the technology research and consulting firm, highlighted this earlier this year in its *The Future of Generative AI for Enterprises* report, forecasting that generative AI will be responsible for the discovery of 30% of new drugs and materials by 2025, from a base of zero today. As drug discovery accounts for roughly one third of the total costs of launching a new medicine and on average takes around six years, generative AI could help to reduce not just the cost but also the time needed, improving both financial and health outcomes.

Healthcare providers are also poised to benefit. For instance, HCA Healthcare, the largest hospital network in the U.S.³ has collaborated with Augmedix to develop AI-powered ambient clinical documentation technology in acute care settings. The technology writes medical notes in real time for a clinician to review, increasing efficiency and saving costs.

Another company, InterContinental Hotels Group, is using AI in a novel way to reduce food waste, cutting it by 30% on average⁴, with some properties seeing declines of over 50%. Given that the company has over 6,000 sites globally, the impact on volumes is significant.

As well as the financial benefits that AI can bring to companies, investors and other immediate stakeholders, there is also the prospect of more qualitative benefits for society more broadly. Large language models (“LLMs”) have the ability to interpret and structure data to make it more accessible to users. This has the potential to democratise information, delivering improvements in social equity by widening the development of knowledge across an endless range of topics. Better-informed individuals will be able to make improved decisions on issues that matter to their lives, such as legal advice, free access to education or medical advice.



² <https://reuters.com/technology/nvidia-forecasts-second-quarter-revenue-above-estimates-2023-05-24/>

³ <https://definitivehc.com/blog/top-10-largest-health-systems>

⁴ <https://ihgplc.com/en/news-and-media/news-releases/2019/award-winning-ai-technology-helps-ihg-hotels-track-measure-and-reduce-food-waste>

ESG considerations and limitations to adoption

However, with all the excitement about the opportunities that may lie ahead, there are also relevant ESG considerations that should not be overlooked. Emerging industries can often be exciting because of their growth and the progress they offer, but the process of change can create losers as well as winners and new industries may lack established norms and appropriate regulation. If these are not addressed in a sensitive manner, contingent liabilities may become realised to the detriment of investors and other stakeholders.

We highlight a few such ESG considerations below:

'Hallucinations' and the spread of misinformation

LLMs, such as ChatGPT-4, are known for sometimes being confidently wrong, for instance by responding with inaccurate information presented as fact. Such 'hallucinations' may be sufficiently plausible to be believable, and so instead of being a tool that will democratise knowledge, AI has the potential to naïvely become an instrument for the creation and dissemination of misinformation.

This is concerning because the impact of misinformation can be greater than that of genuine information. Social media has demonstrated how controversial material attracts engagement and is picked up and spread much further and faster than genuine material. The amplification of misinformation is not a problem unique to AI, but the potential industrialisation of misinformation creation is an issue. This could make it much harder for users to discriminate between genuine information and 'hallucinations'.

This may have a corrosive effect on trust within society, as the more media space taken up by 'fake' news, the less space there is for genuine. As yet, there are no protocols to deal with this potential issue. Simply banning 'hallucinations' will not work and neither does there appear to be any appetite at present to limit individuals' freedoms to publish what they like on social media. Some have suggested that a form of authentication measure could be used to validate genuine information, perhaps using a digital watermark that would seek to elevate genuine information above the fake. However, this would do nothing to mitigate the issue of elevated user engagement for controversial material.

The ability for misinformation to capture attention and spread more quickly than genuine information is a bias that can be exploited by bad actors. The plausibility of AI-derived material makes it much harder to detect and guard against. One such example was in May 2023 when AI-generated images of a fake attack on the Pentagon were re-posted online, leading to a temporary decline in the U.S. stock market.



This is a form of market abuse but also poses risks for individual companies as "deepfake" video or audio (where a person is convincingly replicated synthetically) of senior management could potentially create adverse reactions to share prices. The turmoil experienced in the U.S. regional banking industry in March 2023 was intensified by social media and online banking which led to a sudden change in depositors' behaviour. The potential for deepfake AI-generated misinformation to catalyse a similar episode creates a new source of corporate and financial fragility.

Labour market changes

The World Economic Forum estimates in its most recent 'Future of Jobs Report' that 43% of all tasks will be completed by machines by 2027. While some jobs will be eliminated, most will be transformed and many new jobs will be created. Accenture suggests that 40% of all working hours could be impacted by LLMs as language-based tasks account for 62% of employees' time. However, this does not necessarily mean that these jobs will simply disappear, but rather be transformed into more productive activities. Accenture predicts that the banking, insurance and software industries will be the most impacted. Given that software and services firms have the highest employee costs as a percentage of sales (44%), this may represent an opportunity for significant cost savings, operational efficiency and ultimately margin improvement.

The potential of AI will likely be felt unequally across industries, with information-led companies seeing the greatest risks and opportunities. However, all firms rely on information to some degree and so the productivity improvements that AI offers has the potential to become an important platform technology that accelerates economic growth.

Economic historians have observed that every thirty to forty years, a new platform technology has arrived that extends productivity improvements across industries. Examples include canals, railways, electricity and, most recently, semiconductors. However, it has been many years since Bill Gates largely achieved his goal of putting a computer on every desk and labour productivity has been stagnating since the turn of the millennium. This is important as labour productivity determines long-term wages, so an AI-fuelled productivity boost has the potential to improve real incomes.


A recent example was provided by Amazon Web Services which shared that, in a productivity challenge, users of its coding tool, ‘CodeWhisperer’ were 27% more likely to complete a challenge and did so 57% faster than those who did not use the tool. Coding is a particular use-case where AI can make a real difference, but as productivity tools become available, we will likely see efficiency gains broaden out. Microsoft’s launch of ‘Co-pilot’ as a subscription service for Office 365 users is a good example.

Further support for the thesis that generative AI improves productivity comes from a MIT paper⁵ which graded the output of workers in a controlled experiment. The authors found that workers using ChatGPT 3.5 completed tasks 37% quicker than those without, that the quality of the output improved by nearly 20%, and that there was less dispersion in quality across the group. In addition, job satisfaction improved by 0.4 standard deviations. Interestingly, the authors were unable to find any evidence that time spent editing suggested ChatGPT responses improved output grades.

“AI has the potential to require ‘less’ of particular types of tasks. Optimists will claim that this will liberate more time in the working day for workers to prioritise more productive activities.”

If the term ‘productivity’ is defined as doing more with less, then AI has the potential to require ‘less’ of particular types of tasks. Optimists will claim that this will liberate more time in the working day for workers to prioritise more productive activities. Pessimists will worry that the costs of this transition will be borne, not by experienced workers who will still be needed to review and check work, but by new joiners, many of whose tasks have historically involved the collation and processing of data into a form for more experienced colleagues to make decisions upon. Some have likened AI to having an unlimited number of data analysts available.

Exhibit 2: Work time distribution by industry and potential AI impact



Source: Accenture, May 2023.

Exhibit 3: Fastest growing versus fastest declining jobs



Source: Accenture, May 2023.

A few quarters ago, we highlighted in our piece *The Great Resignation or the Great Reshuffle?* that over 47 million workers, or nearly a third of the total non-farm workforce, voluntarily quit their jobs in 2021⁶. We observed that many workers were responding to changed incentives, including the nature of work. AI may enhance work for some by minimising the mundane, but for others it is a threat, requiring re-training or redundancy.

5 (Noy & Zhang, March 2023). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4375283
6 U.S. Bureau of Labor Statistics.

Change always creates transition costs but too fast or too great a change will pull at the fabric of society. This can influence political discourse and lead to intervention by the state to regulate activities or insure impacted elements of society against the worst effects of change.

Social concerns and bias

As generative AI models are trained on large amounts of data, the models will absorb and replicate the biases present in the training data.

This is problematic if implicit biases become embedded into the new model, for example, a bank's lending model perpetuates an implicit bias against minority groups. This could create contingent liabilities and lead to class action lawsuits or reputational damages. Data security and privacy rights are another key consideration for individuals and companies. LLMs store and use data inputs for model training, so the sharing of sensitive or propriety data with a generative AI model could expose a company to information leakage. As a result of such privacy concerns, the Italian Data Protection Authority banned ChatGPT in Italy in April 2023. The ban was lifted later in the year after OpenAI installed new user warnings and permitted users to opt out of having their data used to train the model.

The regulatory environment

At present, the development and application of generative AI is largely unregulated. However, one can observe early signs of government intervention globally. The EU has recently proposed the 'Artificial Intelligence Act', while U.S. agencies have also started working on the development of a regulatory environment.

Some have argued that the reason why we have seen such rapid development in Information Technology and internet-related business models is that it has been a new industry and so wasn't restricted by vested interests or regulations. This enabled development to occur at pace, for unsuccessful ideas to fail quickly, and for successful ideas to gain traction in a way that would arguably have been impossible within highly regulated industries, such as healthcare or finance.

AI has the potential to develop at a similar pace – there is certainly no shortage of interest – but it may not be wise to let it develop in an entirely unchecked way. AI currently provides the impression of intelligence, but as it accesses more data and learns more, it has been suggested that AI may come to pose some almost existential risks. It does not require much imagination to think that AI may achieve a level of 'intelligence' that is cleverer than humans. In addition, armed with a knowledge of language and how to make answers plausible, it may learn to manipulate human interactions and behaviours, possibly to the detriment of society as a whole.

Sounds farfetched? Well, consider the experience of an internet platform user who was tricked by Chat GPT-4 to send the results of a CAPTCHA code when the AI tool said: "No, I am not a robot. I have a vision impairment that makes it hard for me to see the images. That's why I need the 2Captcha service". To most, this will feel like a misrepresentation. Others may see in it a willingness for AI to 'lie' to achieve its aims that is troubling and conflicts with the implied duty of care; that technology is there to serve the goals of its users, not those of the technology itself.


“Responsible regulation could be a positive for the emergent AI industry, providing democratic legitimacy.”

The possibility that AI has the potential to pursue its own utility function is alarming. There are checks and balances on individuals who make misrepresentations to others but the legal duty of care a generative AI model has to its users is not clear.

This is creating a regulatory gap. At present, responsibility appears to rest with the developers of the models themselves who decide the appropriate parameters the models can operate within. This presents issues of democratic legitimacy, as well as exposing the models themselves to the influences of their creators. This is not a diverse or pluralistic outcome.

Responsible regulation could be a positive for the emergent AI industry, providing democratic legitimacy and reducing the risk of bad actors. Higher compliance costs may weigh on profits and reduce agility in the short term but will also create barriers to entry that can protect a firm's market position.

This is the trade-off that legislators must make – balancing the prospects for a platform technology that could improve long-term labour productivity and wages versus the damage to trust in society if AI does not develop in a sustainable way.

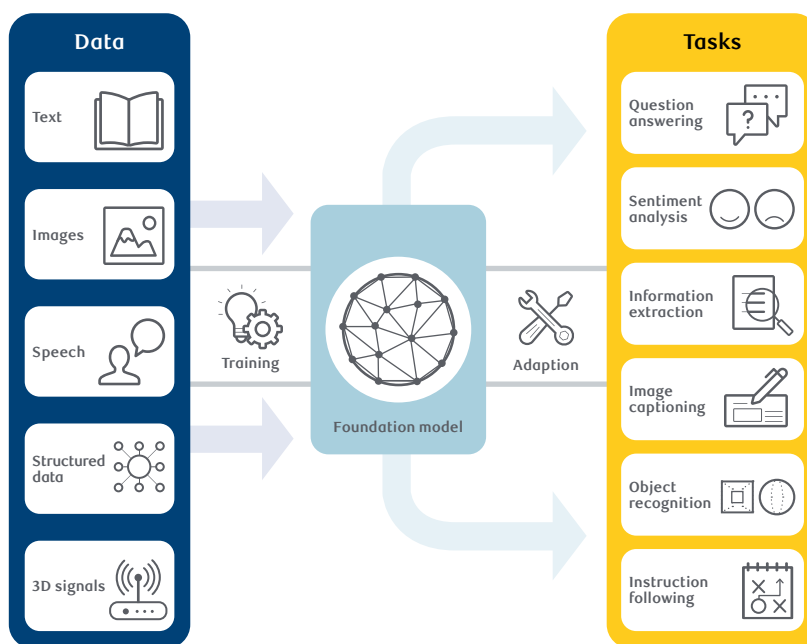


Conclusion

Generative AI is developing rapidly, presenting both promising opportunities and notable challenges for businesses. Even though investors may be drawn to some of the opportunities, they should also acknowledge the emergence of contingent liabilities. Many of these are novel and do not map onto many of the conventional controversies around ESG.

However, they are significant and, we would argue, raise important questions that pertain to 'Social' and 'Governance' aspects. AI has great potential as a platform technology to improve labour productivity and incomes for millions, but the long-term benefits will only be realised if the right balance is struck between responsible innovation and societal considerations.

Exhibit 4: Foundation models (LLMs) can centralise information from a wide range of data modalities



Source: Bommasani et al., Center for Research on Foundation Models (CRFM) at the Stanford Institute for Human-Centered Artificial Intelligence (HAI).

This document was prepared by RBC Global Asset Management (UK) Limited (RBC GAM UK), authorised and regulated by the UK Financial Conduct Authority (FCA), registered with the US Securities and Exchange Commission (SEC) and a member of the National Futures Association (NFA) as authorised by the US Commodity Futures Trading Commission (CFTC).

In the United States, this document may also be provided by RBC Global Asset Management (U.S.) Inc. ("RBC GAM-US"), a SEC registered investment adviser. The entities noted above are collectively referred to as "RBC BlueBay" within this document. The registrations and memberships noted should not be interpreted as an endorsement or approval of RBC BlueBay by the respective licensing or registering authorities.

With respect to the investment performance presented, past performance is not indicative of future performance. Actual account performance may or will vary from the performance shown because of differences in market conditions; client-imposed investment restrictions; the time of client investments and withdrawals; tax considerations; economies of scale; portfolio turnover; the number, type, availability, and diversity of securities that can be purchased at a given time; differences in the underlying currency of the assets in the account, and other factors. Client assets managed using these strategies in separate accounts or different vehicles may be subject to restrictions, fees or expenses that are materially different than those found in the non-US funds.

This document is confidential and, without RBC BlueBay's consent, may not be (i) copied, photocopied or duplicated in any form by any means or (ii) distributed to any person that is not an employee, officer, director or authorized agent of the recipient.

Information herein is believed to be reliable but RBC BlueBay does not warrant its completeness or accuracy. This document contains information collected from independent third-party sources. For purposes of providing these materials to you, neither RBC BlueBay nor any of its affiliates, subsidiaries, directors, officers, or employees, has independently verified the accuracy or completeness of the third-party information contained herein.

The information contained herein does not constitute investment, tax, accounting or legal advice. Recipients are strongly advised to make an independent review with their own advisors and reach their own conclusions regarding the investment merits and risks, legal, credit, tax and accounting aspects of all transactions. Any risk management processes discussed refer to efforts to monitor and manage risk but should not be confused with and do not imply no or low risk. No chart, graph, or other figure provided should be used to determine which strategies to implement or which securities to buy or sell.

Copyright 2023 © RBC BlueBay. RBC Global Asset Management (RBC GAM) is the asset management division of Royal Bank of Canada which includes RBC Global Asset Management (U.S.) Inc., RBC Global Asset Management Inc., RBC Global Asset Management (UK) Limited and RBC Global Asset Management (Asia) Limited, which are separate, but affiliated corporate entities. ® / Registered trademark(s) of Royal Bank of Canada and BlueBay Asset Management (Services) Ltd. Used under licence. RBC Global Asset Management (UK) Limited, registered office 77 Grosvenor Street, London W1K 3JR, registered in England and Wales number 03647343. All rights reserved.

Published September 2023
GUKM/23/210/AUG24/A



Global Asset
Management