The Generative AI Revolution

Opportunities, Shocks, and Risks

Shabbir Merali, Ali Merali

Edited by Adam Hawksbee

ONWARD

About Onward

Onward's mission is to develop bold and practical ideas to boost economic opportunity, build national resilience, and strengthen communities across all parts of the United Kingdom. Our vision is to address the needs of the whole country: young and old, urban and rural, for all communities across the UK – particularly places that have too often felt neglected or ignored by Westminster.

We believe in an optimistic conservatism that is truly national – one that recognises the value of markets, supported by a streamlined state that is active not absent. We are unapologetic about standing up to vested interests, putting power closer to people, and supporting the hardworking and aspirational.

We do so by developing practical policies grounded in evidence. Our team has worked at high levels across Westminster and Whitehall. We know how to produce big ideas that resonate with policymakers, the media and the wider public. We work closely with policymakers of all parties to build coalitions of support. Most importantly, we engage ordinary people across the country and work with them to make our ideas a reality.

Thanks

The authors would like to thank the entire team at Onward for all their insights and support with this paper. The authors would also like to thank everyone who took their time to read and provide comments on drafts of this paper, including but not limited to, Tim Leunig, Dean Abele, Haydn Belfield, Jamie Hope, and Nitarshan Rajkumar, David Skelton, James Clough, Richard Jones, Anthony Finkelstein, Tabitha Goldstaub and Jess Montgomery. All views expressed and any inaccuracies are of course the author's own responsibilities.

About the authors

Shabbir Merali is a Policy Fellow at Onward, Director of Intelligence at CogX and Policy Fellow at the University of Cambridge. He is a former Special Adviser to the Prime Minister, and former Economic Adviser to the Foreign Secretary, Secretary of State for International Trade and Chief Secretary to the Treasury. Before working in Westminster, Shabbir spent five years working in Finance at PwC, specialising in M&A financial modelling. He is a CFA charterholder and an ACA qualified chartered accountant. He also previously co-founded a tech start-up.

Ali Merali is an Economics PhD Candidate at Yale University and a Forethought Fellow. His research spans the intersection of economic theory and AI and has been conducted at the Global Priorities Institute, Epoch, and the Center on Long-Term Risk amongst other institutions.

Adam Hawksbee is Deputy Director of Onward. He was previously Head of Policy at the West Midlands Combined Authority, advising Mayor Andy Street. He was also a Kennedy Fellow at the Harvard Kennedy School of Government, and was a Research Fellow at the Bloomberg Harvard City Leadership Initiative and at the New York City Mayor's Office of Operations.

Contents

Executive summary

Artificial Intelligence will be the transformational technology of our generation. In recent months the pace of change has accelerated, with the phenomenon of generative AI creating new applications that have exploded out of Silicon Valley and into the UK national consciousness.

A song that used AI to clone the voices of Drake and The Weeknd secured 15 million TikTok views before being hurriedly pulled from platforms. OpenAI's GPT-4, a Large Language Model (LLM), has performed better than 90% of American high schoolers across exams in nearly every subject. Google and Stanford researchers created a "mini-Westworld" where characters demonstrated "believable individual and emergent social behaviours". Training compute has grown by a factor of 10 billion since 2010, and in the first three months of this year $11 billion has been invested in reaching Artificial General Intelligence (AGI), more than the previous ten years combined. The exponential rate of progress is therefore set to continue: the Generative AI revolution has only just begun.

While AI has landed squarely in the public debate through ChatGPT, the implications for policy and politics have yet to be fully recognised.

AI presents an enormous economic and geostrategic opportunity

A recent Goldman Sachs report predicts that over the next ten years, the productivity benefits of generative AI alone will create $1.5 trillion of value. The scale of benefits in the UK will be driven by two trends:

1. Rapid adoption: ChatGPT garnered 100 million users in two months, compared to 30 months for Instagram, 55 for Spotify, and 70 for Uber.
2. Boosting worker productivity: A recent MIT paper found using LLMs reduced the time taken for professional writing tasks by 40% while improving quality. This productivity increase will also be felt in more advanced fields where the UK has a comparative advantage. DeepMind's AlphaFold predicted the 3D structure of almost every known protein, a task that was thought to take decades of human labour

Given the potential of AI, international governments are already placing AI leadership at the centre of their geopolitical strategies. The UK's Integrated Review in 2021 stated that "in the years ahead, the countries which establish a leading role in critical and emerging technologies will be at the forefront of global leadership". But other countries are pulling ahead - the vast majority of recent AI advancements in LLMs have been driven by a few Silicon Valley-based AI labs with employees in the hundreds. The USA and China are gearing up for a technological battle for AI supremacy and the EU is seeking to lead on AI regulation. The UK has

an opportunity to carve out a niche, shaping global standards and leading multilateral initiatives on AI safety. Google DeepMind, based in London, gives the UK a seat at the table.

AI will cause shocks to our labour market

Every wave of technological progress has generated fears, often unfounded, of rising unemployment. But this time the warnings should be heeded. For the first time technology looks set to automate cognitive functions and creativity, turning the traditional model of automating routine tasks on its head.

AI could cause superstar effects on steroids: Taylor Swift might profit from millions of personalised songs being created by AIs using her voice, or a high flying corporate lawyer could scale her productivity with effectively unlimited interns at her fingertips. White collar jobs - like the paralegals or graphic designers who make up a disproportionate amount of the UK workforce, particularly in London - could be automated partly or wholly out of existence. Research by OpenAI estimates that 19% of workers will have at least 50% of their tasks impacted while Goldman Sachs forecast 300 million jobs will be exposed. If the technology reaches AGI, where AI systems can perform any intellectual task that humans can, these numbers will multiply and the value of human labour will be limited.

The UK may also face macroeconomic shocks. Just as China's introduction to the global trading system in the 1990s brought downward pressure on inflation, generative AI will bring down cost curves and prices for services. The structure of the LLM market and where economic value is reaped, including whether models become commoditised, will determine the impact on price levels and the necessary policy response. Downward pressure on prices may bring down short term interest rates, but much higher longer term interest rates may be needed if we move towards AGI. Markets are not currently pricing this in but the latest Metaculus forecast suggests AGI will be developed by 2032.

AI will create risks for our shared future

A recent letter signed by Elon Musk and Stuart Russell called for a six month pause in AI developments to address the issue of "alignment" – ensuring that AI systems act in accordance with human values and goals. In a recent survey of AI experts the median respondent estimated a 5% chance of AI posing an existential risk to humanity. The public is in agreement - polling by the Centre for the Governance of AI showed the majority of the public believe AI requires careful management. This is a global problem and coordination issues mean that multilateral solutions are essential - UK leadership will be required.

These three areas - capitalising on the economic and geostrategic opportunity, managing the labour market and macroeconomic shocks, and navigating the risks around AI safety and alignment - should guide the UK policy debate.

So what is to be done? The Government has already taken a number of steps in the right direction. The Chancellor allocated £900 million at the last Budget for a new exascale supercomputer and the Prime Minister announced £100 million for a new Foundational Model Taskforce. The International Tech Strategy included AI as a priority technology and committed to initiate an international dialogue on AI risks. And the AI white paper set out a pro innovation approach to regulation. But the Government should go further and faster:

- Capitalising on the economic and geostrategic opportunity of generative AI means making the UK the home of AI labs, researchers, and consumer-facing applications and encouraging adoption. The Government should develop sovereign LLM capabilities - "GB GPT" - to improve the security of critical technology and support the diffusion of productivity benefits throughout the economy and public services. An AI fellowship programme should be launched to build expertise in the heart of government, alongside a bolstered incentive package to attract AI experts and entrepreneurs to the UK. And increased access to compute and a reformed IP regime can make the UK the best place to build and train foundational models.

- Managing the labour market and macroeconomic shocks from AI will require early preparation for a potentially rapid transition. HM Treasury should prepare measures to shift the burden of tax from labour to capital in the medium term, and consider lessons from previous labour market disruptions such as the rapid rise of globalisation in the 1990s. The Department for Education should overhaul its insight capabilities, ramp up its retraining offer, and incentivise the supply of high level STEM skills.

- Limiting AI safety risks requires multilateral solutions and leading these efforts should be one of the UK's highest foreign policy objectives. Solving AI alignment issues will become one of the most critical questions of our age and it is in humanity's collective interest that we develop solutions. The Government should launch a UK Evaluations Framework, harnessing our academic and industry expertise to set standards the world will use. Regulators should monitor and better distribute compute access to support alignment research. And a single regulator should be created to ensure this work does not fall through the cracks - the Office for Foundational Models (OFFOM).

Table of recommendations

Challenges	Recommendations
1. Capitalising on the economic and geopolitical opportunities of generative AI	**1.1** The Government should build a UK sovereign LLM: "GB GPT" **1.2** The Government should set up a UK AI Fellowship to improve artificial intelligence policymaking **1.3** The Government should create a new incentive package for AI experts to come to the UK and broaden the High Potential Individual (HPI) visa **1.4** The Government should extend the exception for Text and Data Mining (TDM) to allow it for any purpose but include an opt-out for content holders **1.5** The Government should increase access to compute beyond the commitments made at the Budget
2. Managing the economic shocks of generative AI	**2.1** HM Treasury should start preparing for a potential shift in the burden of taxation from labour to capital in the medium term **2.2** The Government should help workers to train in the skills of tomorrow through more accurate skills forecasting, an expanded retraining offer, and more high level STEM qualifications
3. Limiting generative AI safety risks	**3.1** The Government should create a UK Evaluations Framework to shape how AI systems are built and assessed **3.2** The Government should monitor and better distribute compute access **3.3** The Government should create a centralised UK AI regulator with oversight over foundational AI: the Office for Foundational Models (OFFOM)

Generative AI

Where are we now?

Rapid advances in generative AI and LLMs in recent months have caught the attention of researchers, policymakers, and the general public. Feats which would have been seen as science fiction a few years ago such as AIs composing sonnets, writing creative essays or carrying out coherent conversations are now commonplace. With progress being made at an exponential rate, new capabilities are being developed every week.

While AIs writing poetry is impressive, it's LLM's abilities across a range of intellectual and professional tasks that offers the greatest transformative effects on society. When ChatGPT came out in November 2022, it recorded a close to average score for the US Law School Admission Test (LSAT) and performed in the top third of SAT-takers for both reading & writing and mathematics.[1] By March 2023, with the release of GPT-4, these abilities were vastly superseded. GPT-4 performs at the 88th percentile on the LSAT and in the top decile for the SATs.

Figure 1: Exam results (ordered by GPT-3.5 performance)
Source: GPT-4 Technical Report, OpenAI, 27th March 2023

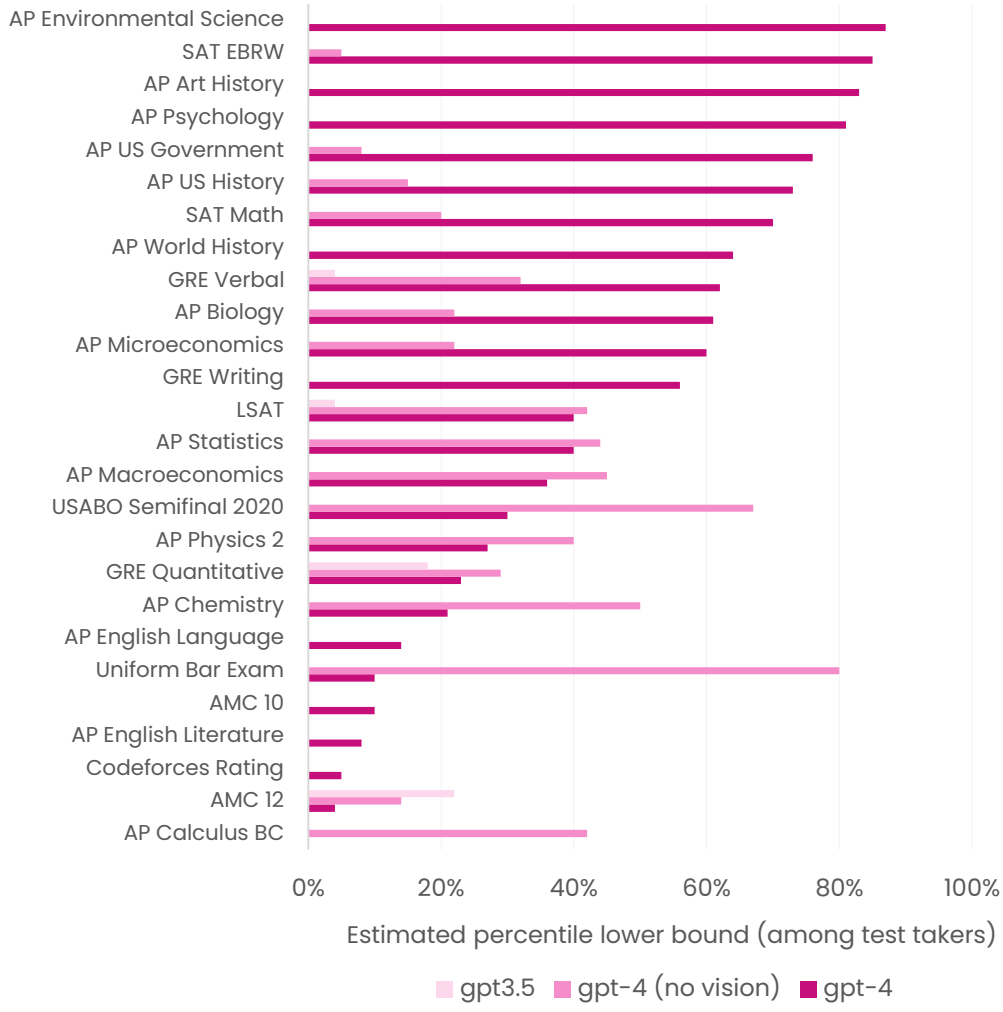

These are not cherry-picked results: as shown in Figure 1, GPT-4 can perform better than 90% of American high-schoolers in a vast range of subjects ranging from environmental science to US history.[2]

There have also been significant developments in generative AI imaging. Midjourney, an AI image generation platform, has seen an astonishing amount of technological progress coming from just an eleven person team[3]. As the images below show, Midjourney has improved drastically since the first iteration came out in March 2022 with a new update being made every three months on average.[4]

Figure 2: Midjourney images, from March 2022, at three month intervals
Source: Midjourney

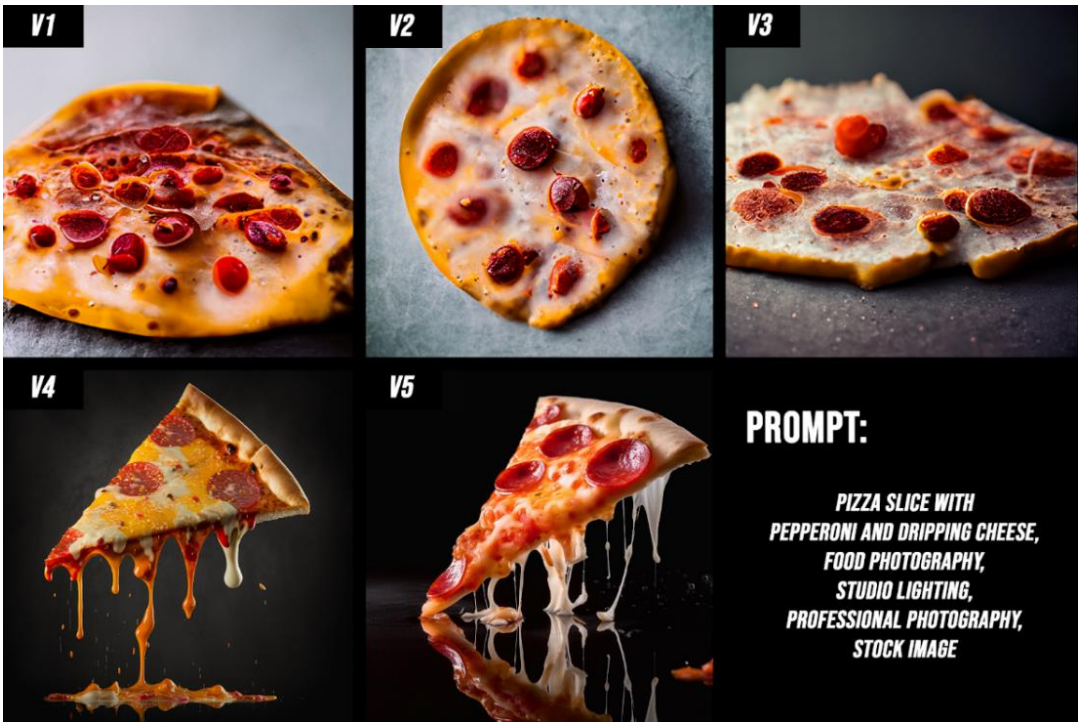

Progress in AI development, building on applications like ChatGPT and Midjourney, has been rapid. In just one month since the release of GTP-4 in March 2023, the following projects have been announced:

- In what MIT Professor Phillip Isola described as a "phase transition in science",[5] a team from Google Brain managed to use generative AI models to improve the performance of other AI models.[6]
- Generative AI models have been used to create agents which can fulfil complex tasks by themselves, including by creating sub-agents where necessary to complete sub-tasks. Although these are in very preliminary stages of development, AutoGPT already has more GitHub stars than PyTorch, an important machine learning framework created by Meta that

was developed over six years ago, highlighting the degree of research interest.[7]

- In what has been referred to as a "mini-Westworld",[8] researchers from Stanford and Google created a 'society' of generative AI agents with memories and the ability to reflect on plans.[9]
- Progress in extending successes of generative AI in language and imagery to music and video has continued, with an AI-generated Drake song garnering 15 million views and an AI-created video of Harry Potter characters wearing Balenciaga going viral on YouTube and TikTok. [10] [11]

The rapid state of progress is likely to continue. In a DeepMind study, a range of LLMs were trained in order to understand how performance on a vast array of metrics from coding challenges to mathematical puzzles improve as the size of an AI model increases.[12] The size of an AI model is usually measured in two quantities; the amount of parameters used by the model and the number of tokens or the quantity of data used. Increasing either of these increases how much computation is required to train a model.

The scaling laws derived in the study suggest increased AI capabilities will follow as the amount of compute used increases. And the amount of compute used has been increasing at a tremendous pace. Research by Epoch shows training compute, the computational resource used to train an LLM, has grown by a factor of 10 billion since 2010, with a doubling rate of every five to six months. This demand for compute has been met as much through design innovations - first the introduction of GPUs, and more recently designs optimised for AI like TPUs - as from raw increases in the number of transistors per chip. The demand for compute is set to grow. Estimates of major AI labs' investment decisions suggest that spending on compute alone is likely to increase by a factor of 100 over the next five years. If we continue to see increased chip efficiency and algorithmic progress, this tentatively suggests that model progress may continue at a similar exponential rate in the short to medium term. Innovations in chip designs will be essential here, and the UK has an important role with ARM and Graphcore amongst others headquartered here.

Larger models and increased compute will also have environmental impacts. Researchers studying the dangers of LLMs found that training a large neural language model can emit 284 tonnes of carbon dioxide equivalent. For reference a human being is responsible for five tonnes a year. As models become more complex to train, firms will need to start moving to 'green compute' to limit the environmental damage. This could include using highly efficient chips or data centres to minimise carbon emissions. Without significant efforts to reduce the energy required, AI advances will make it more difficult to meet emission reduction targets and become incompatible with climate objectives.[13]

Figure 3: Computation used to train notable artificial intelligence systems
Source: Our World in Data, July 2 1950- March 15 2023[a]

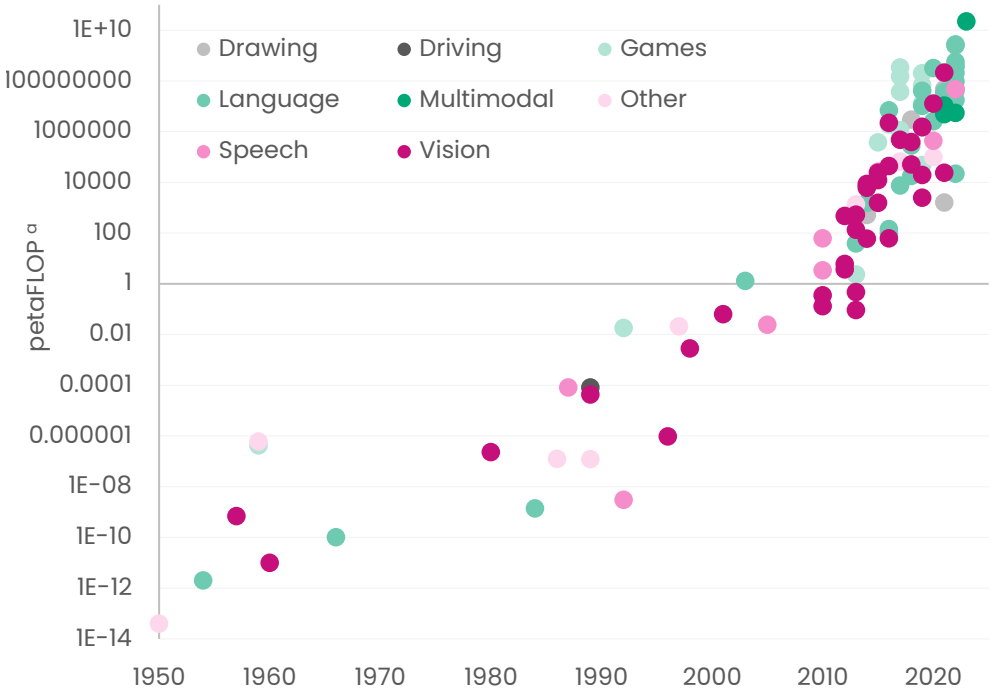

Who is behind the rapid advances in Generative AI?

US firms have driven recent AI advances and are leading the world in AI development, building on years of technical research in labs and academia. The main race is between OpenAI, a Silicon Valley start-up, and Alphabet which has recently released its own generative AI services. OpenAI was founded in 2016 and was originally co-chaired by now-CEO Sam Altman and Elon Musk as a non-profit, to ensure safety concerns were not compromised by profit motives. It since transitioned in 2019 into a capped for-profit[14] which has helped it raise $10 billion from Microsoft who also give it access to their Azure computing infrastructure. It is OpenAI that has led the early race in generative AI, with ChatGTP garnering over 100 million users in its first two months.[15]

While Google's launch of their Bard model was comparatively unsuccessful, their research capability means they are the primary challenger to OpenAI. This is reflected in Figure 3 below which shows that Google Brain and DeepMind combined had more scientists that published at NeurIPS, the premier AI conference, in 2020 than the rest of the other top-10 industry labs combined. The competition from OpenAI however, has led to Google management issuing a "code red" internally given fears that it is losing the AI race. This has led to the creation

[a] Computation is measured in total petaFLOP, which is 10^15 floating-point operations

of Google DeepMind, a new combined unit of Google Brain and DeepMind brought together to accelerate progress in AI.[16] Although Google DeepMind is US owned, it will be headquartered in London, and provides the UK with a world leading AI lab on its shores.

Figure 4: Number of authors in NeurIPS 2020

Source: Notebook for comprehensive analysis of authors, organisations, and countries of ICML 2020 papers, Dustin Tran analysis.[17]

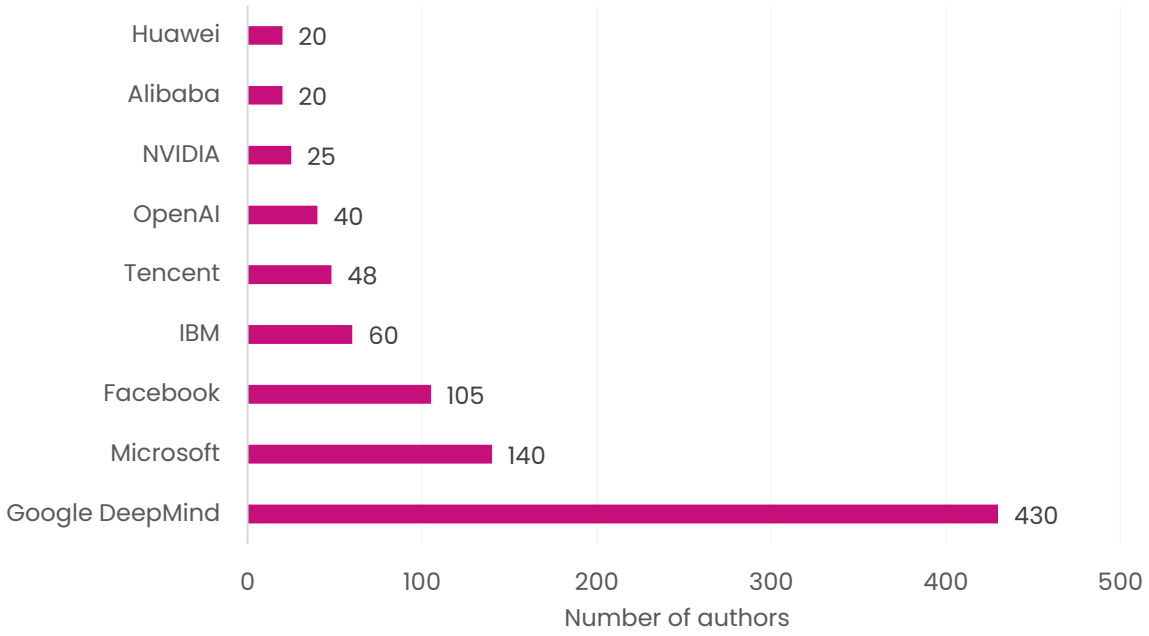

Behind OpenAI and Google there are a wave of start-ups that are increasing compute spending and capabilities. Of these Anthropic, a Californian start-up which Alphabet owns a 10% stake of, is at the forefront and has recently acquired tens of thousands of GPUs to speed up progress.[18] Cohere also recently began talks to raise funding at a valuation of over $6 billion[19] and has built its own version of LLMs.

The most interesting new challenger may be Elon Musk's X.AI which was established to compete with OpenAI, [20] although the new company's plans haven't yet been articulated. In imaging, Stability AI (based in the UK) has taken a different approach focussing on building open sourced models such as Stable Diffusion to enable access for developers and general users. While there are only a handful of firms that have commercialised generative AI services up to now, this will likely change in the near term. AGI companies have received more investment in the first 3 months of 2023 than in the entirety of 2022 - more than $11 billion.[21]

Figure 5: Increase in cumulative investment into AGI research

Source: Ian Hogarth, Financial Times, April 2023

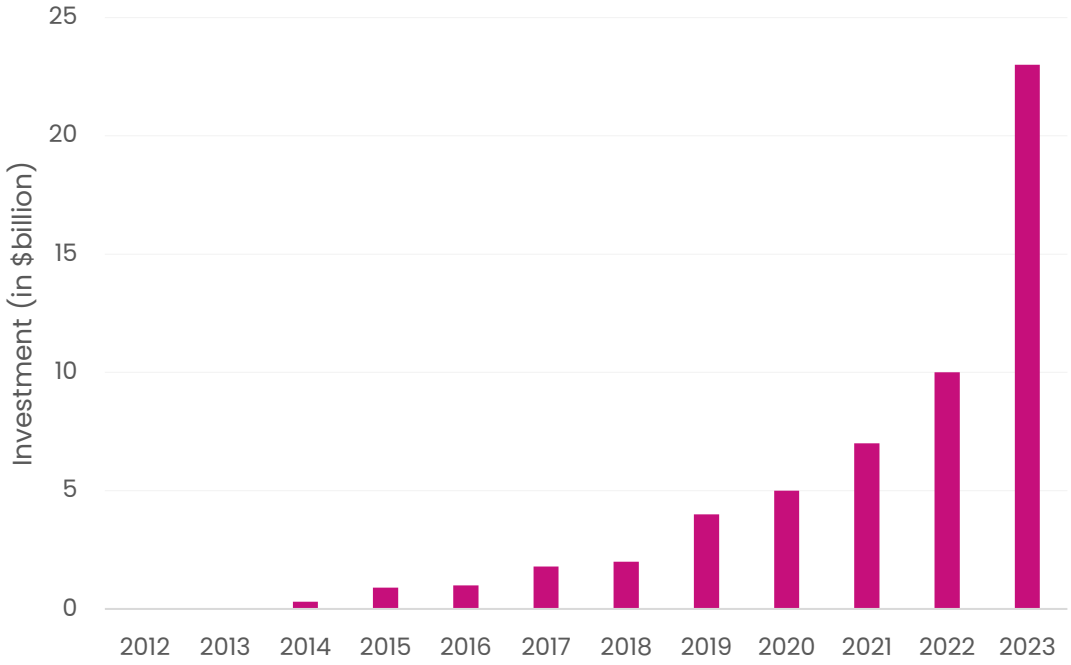

Given the rate of progress, policymakers can't afford to simply react to each wave of AI breakthrough. Instead, policy needs to be preemptive - identifying the trends that will accelerate in the coming months and years. We are still early on in an exponential process that will continue to create increasingly transformational technologies. The consequences, opportunities and risks to the UK from this process are addressed in the following sections.

Opportunities

Realising the economic and geopolitical upsides

Generative AI presents enormous opportunities for the UK economy. As with many technological transitions, the challenge will be to harness the benefits of change while managing and mitigating any negative impacts. But what might generative AI mean for how we work? How quickly will it be adopted? And how might these micro impacts reflect on our macroeconomic variables?

Productivity and Growth

The UK's recent economic stagnation, characterised by sluggish growth rates and declining real wages, can be largely attributed to its weak productivity growth. Labour productivity has grown by just 0.4% a year in the 12 years following the financial crisis.[22] Many solutions have been proposed to address this productivity decline, but none hold the potential for immediate and substantial impact quite like generative AI.

A recent Goldman Sachs report predicts that over the next ten years, the productivity benefits of generative AI alone will create $1.5 trillion of value across the world in the baseline scenario[23]. The UK, with its disproportionate share of high-skilled services, is particularly primed to benefit. As such, the UK's performance should at least be on par with the US. If that occurs then in the optimistic modelled scenario the productivity benefits from AI alone could lead to greater productivity growth in the UK over the next decade than total productivity growth combined over the previous five years.[24] This forecast therefore underscores the transformative power of generative AI in the context of the UK's productivity puzzle.

Figure 6: Effect of AI adoption on annual labour productivity growth, 10-year adoption period

Source: Goldman Sachs Global Investment Research, March 2023

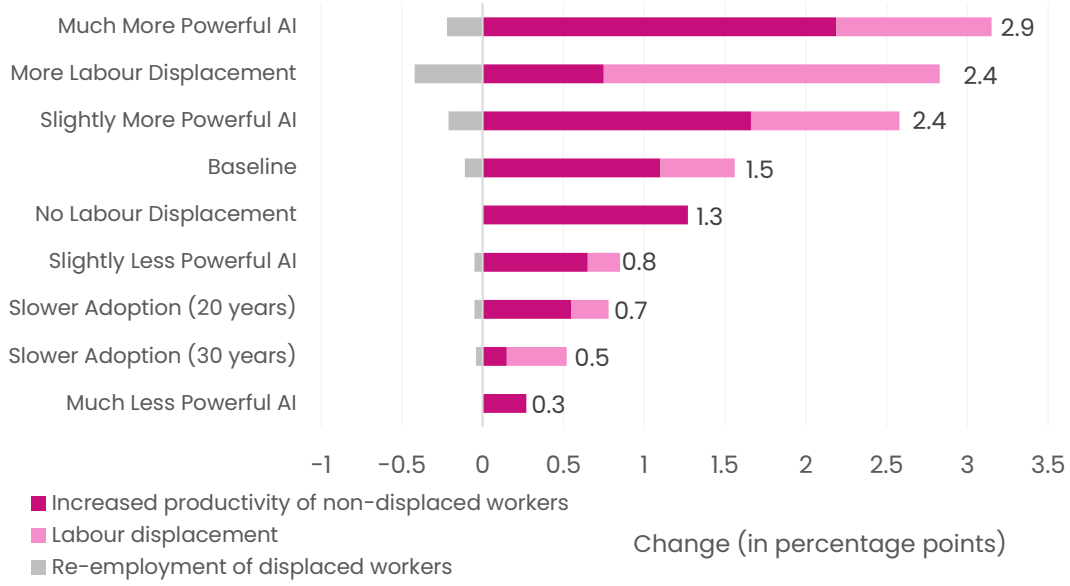

Generative AI will have a significant impact on productivity and growth for two reasons. First, rapid adoption. The widespread adoption of previous transformative technologies such as electricity and the internet took considerable time. Electricity took decades to become ubiquitous[25] and even computers required years to have a measurable impact on the economy. As Robert Solow remarked in 1987, "You can see the computer age everywhere but in the productivity statistics." [26] But technologies have typically diffused quicker in recent decades, as Figure 7 shows.

Figure 7: Adoption speeds of new technologies over time, 1951 - 2016
Source: World Economic Forum/Visual Capitalist, 2019

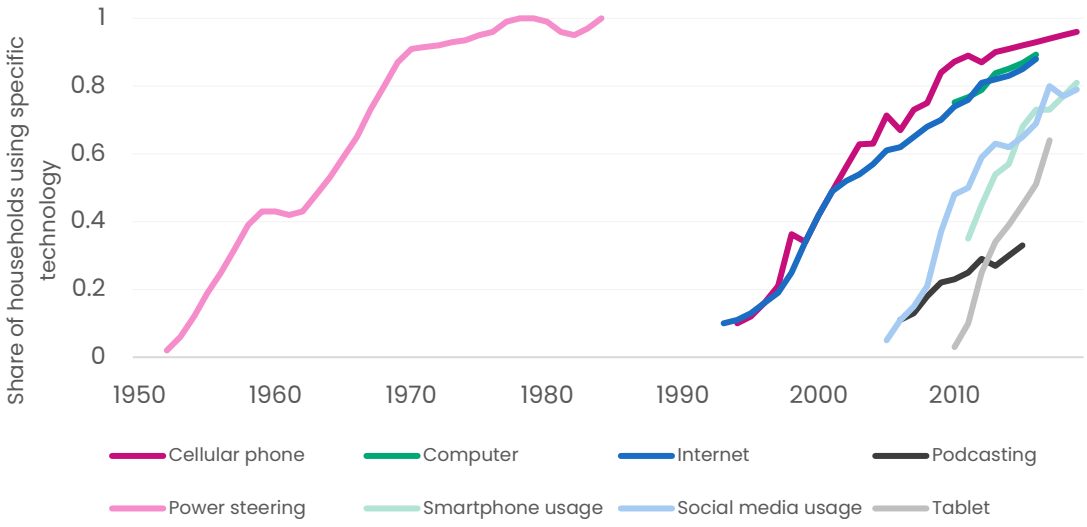

The landscape for generative AI adoption appears to be even faster than more recent technologies. ChatGPT gained 100 million users faster than any other application in history as shown in figure 8 below. This is partly due to the accessibility and usability of the technology - very little prior knowledge is required to experiment with ChatGPT, which has been instrumental in its widespread takeoff. A survey by YouGov in the USA in January 2023, when ChatGPT was still very new, found 35% of over 65s had either used ChatGPT or seen text generated by someone else using it.[27]

The rapid adoption is not limited to individual users; businesses and governments are also recognising the opportunities. Major corporations such as Bain & Company have entered into deals with OpenAI to use generative AI in their strategy consulting business,[28] while companies like Expedia have integrated ChatGPT through plug-ins,[29] allowing customers to book entire trips by simply stating their preferences. Even the Icelandic government is utilising generative AI

in an effort to preserve its language, by collaborating with GPT-4 to create a model able to respond in Icelandic.[30]

Figure 8: How long it took different apps to hit 100 million monthly users
Source: UBS, AI Business, Yahoo Finance, February 2023

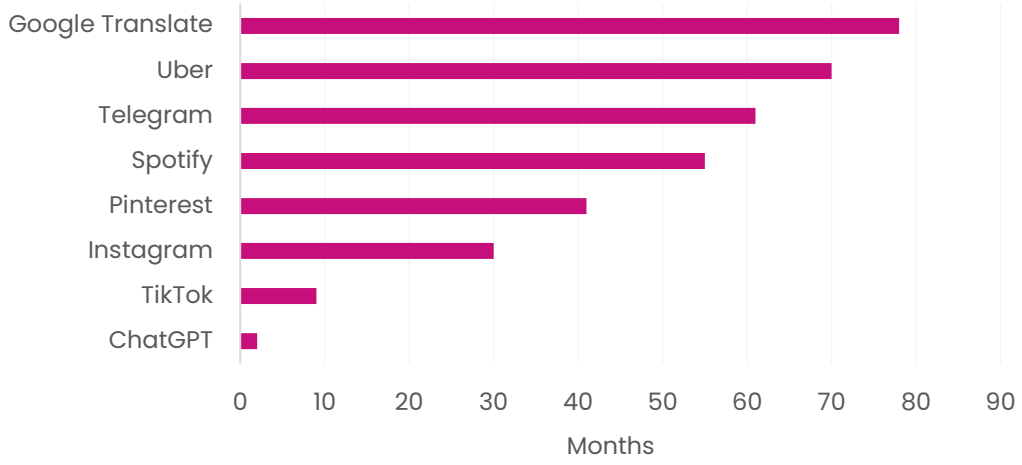

The second reason generative AI will significantly impact growth is its effects on labour enhancement and R&D. This means that AI reduces the limitations on our national productive capacity by making the traditional factors of production less of a limiting factor on output. Generative AI can augment human capital, helping workers make better informed decisions, optimise processes, and remove mundane tasks. This is reflected in a recent MIT paper conducting experiments in mid-level professional writing tasks. Output rose significantly: the time taken per task dropped by 40%. And the quality of writing also improved. Workers were able to spend more time on idea generation and editing instead of rough-drafting.[31] Productivity improvements have similarly been found in real world settings. A year long study by MIT and Stanford researchers found that AI tools increased workers productivity by 14% on average at a Fortune 500 firm. Using AI models enabled the customer service workers to decrease chat handling times, increase chat resolutions per hour and increase customer satisfaction. Evidence is starting to build that generative AI will deliver tangible productivity improvements when implemented well.

Generative AI can also help workers make better informed decisions by providing data-driven insights and offering solutions. As a result, companies can significantly improve their efficiency. This is reflected in a recent randomised controlled trial at a telemarketing company where salespeople with AI assistance managed to double customer purchases[32]. The increase in productivity levels is made all the more significant by the fact that this experiment occurred prior to the introduction of ChatGPT and is therefore at least two jumps in AI model progress behind the current capability.

It's not just automation of basic tasks that is driving the productivity increase either. In the same telemarketing experiment, customer support providers were particularly helped when faced with the most creative and challenging questions in which they were 133% more likely to be successful. The evidence suggests there will be productivity enhancements across a wide array of tasks.

As the UK's economy becomes increasingly knowledge-driven, the importance of R&D for productivity growth will increase. Generative AI can significantly expedite the R&D process by automating complex tasks, analysing vast datasets, and predicting potential outcomes. AI has been particularly useful in biological research after DeepMind's AlphaFold predicted the 3-D structure of almost every known protein, a task that was predicted to take decades of human labour.

This, alongside other related breakthroughs, has led Dr David Baker from the Institute for Protein Design to estimate that the pace of innovation in his field to be 10 times higher now than 18 months ago.[33] These breakthroughs in fundamental research have also been mirrored with progress in actual drug development - the co-founder of Insilico Medicine, a Hong-Kong based biotech firm, Alex Zhavoronkov says that while it usually takes four years for a new drug to get to clinical development this is already achievable "in under 18 months, at a fraction of the cost" due to AI[34]. This acceleration of R&D efforts can result in a quicker rate of technological progress, which in turn can boost long term productivity and trend growth.

The combination of rapid and wide scale adoption, along with significant human capital enhancement means that generative AI has the potential to significantly upgrade the UK's productivity and growth outlook.

Case study 1: Software Engineering

The role of software engineers has traditionally been to act as translators, bridging the gap between human-generated requests and machine-understandable commands. By employing programming languages, engineers provide a middle ground between natural and machine languages. Recent developments in AI have enabled humans to specify goals while computers build neural networks, systems of interconnected nodes inspired by how the human brain functions, to achieve them. Although humans cannot fully understand neural networks, they can observe the objectives and resulting behaviour. Even with the continued use of traditional programming, AI will significantly enhance engineer productivity.

In the near future, AI systems may be capable of perfectly understanding natural language requests and generating functional software, potentially rendering traditional engineers almost obsolete. Some experts predict this shift could occur within a decade, resulting in a 95% decline in demand for software engineers[35].

In the interim, LLMs like ChatGPT can generate code, albeit often with inefficiencies or errors. By searching for the most common answers from the internet and giving extra weight to those found in fine-tuning data or with positive human feedback, this approach will initially struggle to generate truly original software and may replicate common mistakes. However, ChatGPT can still reduce the time engineers spend searching for solutions online and serve as a valuable brainstorming tool.

Despite these limitations, AI tools like GitHub Copilot, built on OpenAI Codex[b] will significantly boost software engineers' productivity. By integrating into software development environments, Copilot predicts and suggests code based on natural language comments, improves code readability and efficiency, and assists in debugging code through AI chat. Perfect debugging could quadruple productivity, as engineers spend 75% of their time on this task[36]. A study by Kalliamvakou (2023) indicates that Copilot already increases developer productivity by 127%,[37] a figure likely to grow as OpenAI Codex improves and Copilot fine-tunes on more correct answers. Other applications, such as Mintlify[c], can add documentation to code for better user comprehension.

Even with the advent of tools like Copilot and Mintlify, AlphaCode has already demonstrated the ability to generate complex algorithms from natural language prompts at a competitive level. However, DeepMind engineers caution that the transition from competitive programming to commercial software development will pose challenges.

Rapid productivity increases will make high-quality software services increasingly accessible to even small firms and individuals, potentially providing a significant growth effect across the industry and driving demand. The impact on software engineers themselves is more ambiguous; while many might enjoy the ability to produce greater output with less effort in the short run and see salaries increase, this will likely eventually lead to job losses in the medium term.

[b] OpenAI Codex is an AI system which generates programming codes in response to natural language prompts
[c] Mintlify is an AI powered documentation writer

Case study 2: Law Professionals

Lawyers spend a significant portion of their time creating and reviewing legal documents. In the near future, AI could feasibly review legal documents for ambiguity and contradictions, as well as create contracts or pre-trial motions based on simple descriptions. One such example is CoCounsel by CaseText, which provides a chatbot capable of answering legal questions with extensive supervised fine-tuning on accurate legal texts. This technology could replace some of the work traditionally performed by lawyers and substantially reduce the workload of paralegals. CoCounsel is particularly useful during discovery, where a custom chatbot fine-tuned on a vast set of documents can extract essential information and answer relevant questions. The CaseText AI can also ensure document compliance with government regulations and summarise trial transcripts during a hearing to suggest questions for a deposition.

Even free-to-use technology, such as ChatGPT, can already create many less formal documents for lawyers, saving significant time on tasks like summarising trial progress for clients, interview summaries, pitches to attract new clients, and presentations.

The adoption of AI in the legal profession is likely to shift the focus of human lawyers and paralegals towards reviewing AI-generated documents, providing strategic legal advice to clients, and negotiating on their clients' behalf. In the UK, solicitors may be more disrupted than barristers, as regulation may prevent AI from representing clients in court. DoNotPay, a service that uses templates and machine learning to help customers fight parking tickets and cancel subscriptions, recently had to back down from having a 'robot lawyer' based on ChatGTP perform in court after pushback from state bar associations in the US[38]. Developing proficiency in utilising AI will become an essential skill for legal professionals.

The integration of AI in legal services could lead to a situation where high-flying corporate lawyers essentially have access to unlimited free interns, significantly increasing their productivity. This newfound capacity would enable them to work on more cases and potentially earn higher incomes. The increased efficiency could result in an expansion of the quantity of legal services supplied, leading to a greater number of custom agreements and contracts.

This increased supply of legal services could benefit economic growth and potentially improve public defence for vulnerable individuals. However, it also raises concerns about the potential surge in frivolous lawsuits due to the minimal costs associated with producing them. What is clear is that the provision of legal services will be significantly disrupted by generative AI.

Geopolitics: securing UK leadership in AI

Globally, nations are rapidly realising that AI is pivotal to both their security at home and influence abroad. The UK's Integrated Review in 2021 stated, "in the years ahead, the countries which establish a leading role in critical and emerging technologies will be at the forefront of global leadership" - noting AI as central among these.

Across the various levers of statecraft, AI is becoming an increasingly important feature in the new era of geopolitical competition. In defence, AI is at the heart of the military modernisation efforts of both China and the US. A report by US think tank CSET stated, "Chinese leaders view AI as the key to transforming the PLA into a world-class, globally competitive military force."[39] Similarly, the US National Defence Strategy 2022 committed to "institutional reforms" as part of its plans to "aggressively" pursue advantage in critical technologies like AI.[40] In economic statecraft, AI has been central to the US-China trade war, with the State Department making clear that China's AI ambitions were a key reason for the US move to restrict the export of leading-edge semiconductors to China in October 2021, as it seeks to maintain a strategic lead in the hardware underpinning AI advancement.[41]

But the competition extends deeper than simply trying to restrict geopolitical rivals' specific capabilities and shore up vulnerabilities. There is a growing realisation among states that strength in AI is a strategic national security imperative in its ability to secure sustainable economic growth, productivity and overall prosperity for the long term. The US National Security Strategy released in October last year spoke of the pressing need to "anchor an allied techno-industrial base," with AI core to this, for both defensive economic reasons and for driving prosperity for the long term.[42]

AI is therefore foundational to the UK competing in the 21st century. The potential benefits to treating this technology as a strategic goal are significant. Just as this report has described the potential for AI to create "superstar workers", it could potentially create "superstar countries" who lead the pack in AI technologies. Leadership could boost the UK's ability to shape global norms and standards related to AI and therefore control the direction of development in alignment with our values.

Militarily, it could allow the UK to both build world-leading, AI-driven armed forces and shape the technological revolution in military affairs. The author of I-Warbot, Professor Kenneth Payne, has argued "AI that is adaptive to a range of decision-making contexts… presage[s] a profound shift in strategy that is likely to be even more radical in nature than the nuclear revolution."[43]

UK AI expertise could become an important diplomatic asset, too. Just as countries use access to finance as a geopolitical tool, access to technology is becoming an increasingly used tool of foreign policy.[44] While creating vulnerabilities, AI developments may become a foreign policy lever for the UK in a world where AI is a competitive resource.

Developing frameworks and agreements to influence

The UK's safety and ethics objectives must be driven by our liberal democratic values working closely with like-minded nations across the West. But kowtowing or leaving the fight to others is not an option. The US technological domination that drove the shaping of the internet in its formative years in the late 20th century is unlikely to win the day in today's world,[45] with an increasing importance of middle powers in shaping global technology governance.[46]

Instead, coalition-building by influential states — like the UK — working to win over other middle power with an increasingly powerful bloc of non-aligned states is far more likely to be a deciding factor.[47] It is with these states that China has ramped up its charm offensive on a broad range of diplomatic fronts in recent years. Moreover, it is well-documented that for years China has been attempting to build influence and sway across a number of key multilateral institutions to better serve their interests.[48]

So as well as focussing on building technological capability, the UK should work to shape the frameworks and institutions that govern these new technologies.

Developing technology to compete

With so many competing for AI supremacy, competition is - and will remain - fierce. Countries will champion domestic AI labs and compete to attract talent and invest in the necessary infrastructure to invest in thriving ecosystems. Should it emerge, there will be no more powerful geopolitical tool than AGI. Whichever country reaches AGI first will have a strong claim to superpower status. And given potential recursive improvement in AGI systems, that technological advantage may only increase. The first mover advantage is immense.

The UK has a critical role to play. It is a founding member of the Global Partnership on AI and has a thriving AI ecosystem of startups and researchers. More importantly, Google DeepMind is based in the UK - and is one of the most integral companies in the world as it is one of the few labs with the current potential to develop AI towards AGI. It already produces some of the most transformational LLM models and applications of AI as part of its mission to "solve intelligence."[49]

But as analysis from Professor Paul Nightingale and former No.10 Science Advisor James Phillips shows, the UK's strength in AI is shallow, even in research where it is often considered to be world-leading.[50] Without Google DeepMind the UK's share of citations among the top 100 recent AI papers drops from 7.84 % to just 1.86 %. Statistics such as the UK ranking third in the world for AI publication citations per capita[51] mask the outlier of Google DeepMind and the underlying weakness of the UK's ecosystem. If the UK wants to be a leading global player on AI, it will need to cultivate a stronger research base.

The statistics on DeepMind's research also highlight the limited role of academia in AI development. One of the most surprising elements of the past few years in AI has been that nearly all the foundational breakthroughs and new models have been built in private AI labs. Academia and the public sector haven't really been a part of the story, despite their traditional role in leading on breakthrough technologies. They have been important in developing the technical base over the past few decades that AI labs have been successful in building on, but their impact on cutting edge development has waned over time. Part of this is due to a lack of access to compute by researchers which the government is looking to remediate with the establishment of the AI Research Resource as announced in the Budget. This should provide access to exploratory compute for every UK AI researcher and help academics contribute to progress. But even with access to some compute, fundamentally this is a technology which is being developed in the private sector. For comparison, the Chancellor announced £900 million in the Budget for UK compute a month after OpenAI raised $10 billion to invest in more compute and model development.[52]

This is part of a longer term trend of technology companies becoming increasingly powerful relative to nation states. Russia's invasion of Ukraine provides a recent example. As well as countries imposing economic sanctions, decisions by global technology corporations such as Alphabet, Meta and Nvidia to cease operating in Russia were also hugely important in stifling the Russian economy. When Ukraine worried about their internet access post invasion they appealed to Elon Musk rather than a sovereign government. As President Macron asked in 2019, "Who can claim to be sovereign, on their own, in the face of the digital giants?"[53]

The recently published International Tech Strategy, with AI as a priority technology, is a signal of intent. As is the recent creation of a Foundation Model Taskforce backed initially by £100 million[54]. But as the geopolitical race to develop technology accelerates, this should only be the beginning.

Shocks

Managing the labour market and macroeconomic impacts

The productivity and growth impacts of generative AI could be significant. But the potential of AI has also bred increasing concern in recent months that it will largely replace human labour.[55] The reality will likely be more nuanced: technologies create as well as destroy jobs and there are competing automation and augmentation effects. However, there will certainly be disruption, especially in the short term. And these processes could have significant macroeconomic implications that require policy attention. If we want to benefit from the growth and productivity effects described earlier in this paper, we will need to find a way to manage the economic transition.

History repeating?

People have often worried that technological progress would result in unemployment and economic hardship. The Industrial Revolution's mechanisation of labour gave rise to the Luddites, the introduction of the assembly line in the early 20th century caused clashes with unions, and the introduction of computer technology in the 1960s and 1970s was feared to be the end of white collar work. But unemployment is currently at 3.7 %[56], the lowest level since the 1970s, a strong employment rate of 75.7%[57] and a recent Budget with a headline focus of getting more people into the labour market due to the shortage of workers to fulfil the jobs needed.

The reason for this is the changing nature of jobs and the increase in demand for products when technology brings down their costs. The economist James Bessen [58] showcases this with research on bank tellers and ATMs. Since ATMs were first introduced in the US in the 1970s they have spread across the country and there are now over 400,000 in operation. But bank teller employment has actually seen a small increase over the period, for two reasons.

First, bank tellers moved up the value chain. ATMs had automated the basic tasks, like withdrawing money, but were unable to replicate the more valuable tasks, like building customer relationships and solving bespoke issues. Tellers adapted to focus more on human interaction and selling additional (often high margin) financial services and the nature of bank branches changed. Second, cost reductions from automation brought increased demand. As Bessen highlights, the number of tellers required to operate a bank in an average urban branch fell from 20 to 13 between 1988 to 2004. This fall in costs made it cheaper to open new branches and banks responded by expanding their presence to capture market share. While tellers per branch fell, the number of branches increased to offset it, and teller employment actually marginally increased over the period.

Figure 9: ATMs installed and Tellers employed in the US, 1970-2010
Source: James Bessen, International Monetary Fund, 2015

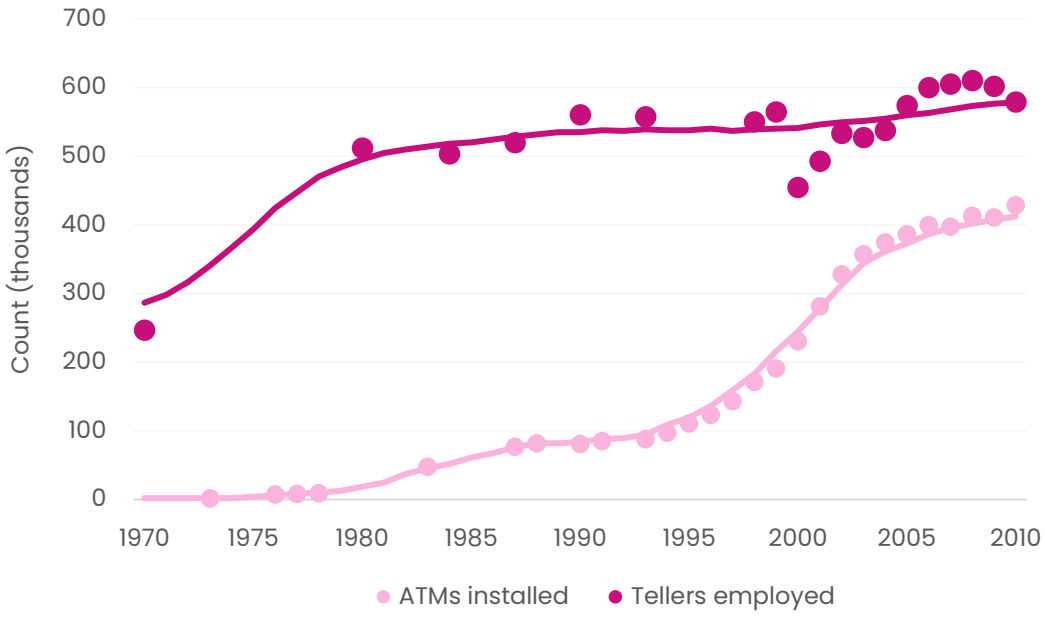

The nature of jobs in the wider economy has also changed over time. The MIT economist David Autor [59] has analysed new job titles in the US Census Bureau's occupational descriptions and estimated that 60% of US employment in 2018 was found in job titles that did not exist in 1940, increasing to 74% percent for professional employment. Just as we didn't have social media managers 15 years ago, there will be jobs that don't exist today which will be commonplace in the 2030s (such as "prompt engineers").

Figure 10: Employment count of broad occupations in 1940 and 2018, distinguishing between job titles that existed in 1940 versus those added subsequently
Source: Autor D, et al, 2022, National Bureau of Economic Research

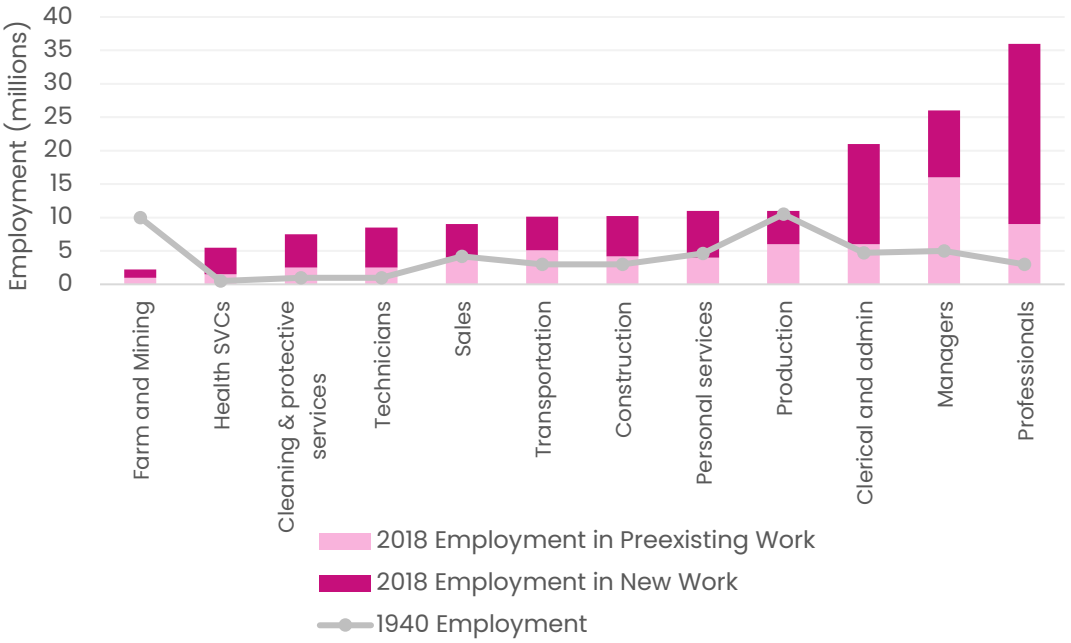

Why is Generative AI different?

LLMs and increasing capabilities in generative AI will create millions of new jobs across the world as labour productivity increases and new types of jobs are formed. However, this is likely to be more than offset by job losses and the next decade might see us enter a world of much higher unemployment and labour market strife. There are a few reasons it might be different this time.

First, the classic model of automation and labour replacement from recent decades has been turned on its head. Humans have typically used technologies to replace routine tasks that can be easily programmed to replicate. Routine tasks are typically found in middle-paid occupations while non-routine tasks are concentrated in either low paid (e.g. social care worker) or high paid professions (such as a therapist or a CEO).[60] Over time, automating technologies have decreased demand for middle relative to low and high paid occupations leading to polarisation in the job market. The assumption has been that AI would follow the same model. As recently as a few years ago, AI reports by think tanks and consultancies would use a model of identifying routine tasks in jobs to estimate the job losses that may be incurred (including a 2019 report by Onward).[61] McKinsey's Global Institute's report on jobs in 2017[62] estimates that high wage and low wage occupations will see the highest growth while middle income

employment will decline. It suggests that workers of the future will require, among other things, more creative skills.

None of this seems likely anymore. Generative AI will be advancing but also replacing creativity and cognition. We have never had a technology capable of replacing human thinking at this level before and so the job displacement effects are likely to be different too. It is service based, white collar, high paid work which is looking most at threat. From graphic designers to lawyers and strategy consultants, the outlook has fundamentally changed.

Given the economic make-up of the UK, it may be especially vulnerable to disruption. The UK is dominated by services, accounting for 79% of GVA and 82%[63] of total employment. Much of this is in areas of comparative advantage such as the creative sector (GVA of £104 billion)[64] and professional and business services (GVA of £190 billion)[65] where there is a higher likelihood of disruption from generative AI.

The relative importance of high value services to the UK compared to other countries can be shown by proxy through examining services exports. From the graph below we can see that the UK has the highest services-export intensity in the G7. If segments of white collar services are affected, the UK will likely face more disruption than its peers.

Figure 11: Share of services in total exports, 1960-2021
Source: Our World in Data

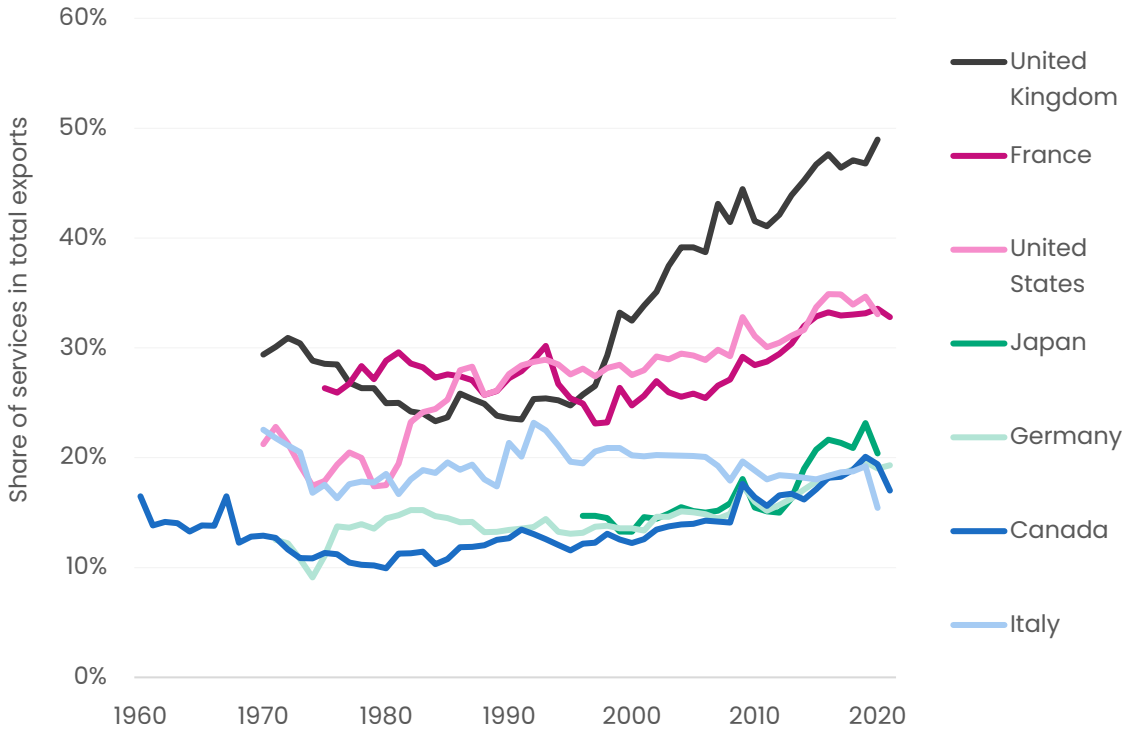

Second, the speed of technological progress is rapid, driven largely by compute. Research by Epoch[66] shows that the development of machine learning can largely be split into the three eras: the Pre-Deep Learning Era, the Deep Learning Era and the Large Scale Era. The chart below shows how training compute has increased across those periods and how it has grown by a factor of 10 billion since 2010. This has been the biggest driver of progress in AI and as compute continues to scale up, technological progress will likely continue to do so at a very fast pace.

Figure 12: Training Compute (FLOPs) of milestone Machine Learning systems over time

Source: Sevilla, J, et al, 2022

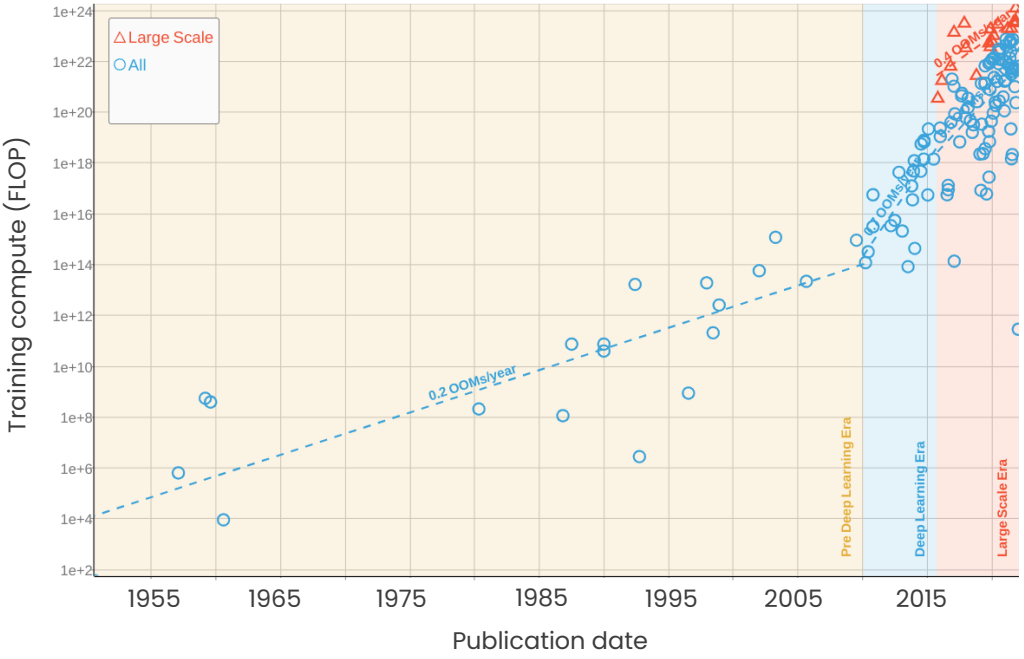

The speed of change makes labour market disruption more likely. Compensating job creation can take time and a rapid influx of unemployed workers will make it harder for those affected to find jobs immediately. This may lead to long term scarring effects and hysteresis with people having extended periods out of the labour market. As research from the ONS shows, the longer a person remains out of work, the worse their chances of returning to work become. Between 2007 and 2020, 41% of those who had been out of work for three months returned to work within the next three months. This reduced to 29% for those who had been out of work for six months and 23% for those between six and nine months.[67] A sudden, sharp rise in unemployment can have long term consequences for the unemployed and the labour market.

Third, labour market disruption will be different this time due to the combination of the pervasiveness of the technology and the ease of adoption. Like previous general purpose technologies such as electricity and computers, there will be very

few parts of our economy where LLMs and generative AI will not have an impact. The interdisciplinary nature of generative AI means that its adoption is not restricted to specific industries but instead it will permeate the entire economic landscape. Crucially, it is also easy to adopt. While using electricity or computers for the first time required up front investment and skilled labour to operate the relevant machinery, that doesn't exist with generative AI. Any business can start using ChatGPT for free today and the easy-to-use user interface means that there are few skills barriers to adoption.

Generative AI is a bottom up technology. Workers in a marketing department at a large corporation can see the productivity benefits and start using new tools to speed up their jobs and improve the quality of output. There is no need for management to make a conscious decision to invest in the technology and spend years bringing it into operations as there has been for many technologies (e.g. computers) in the past.

A recent paper by researchers from Open-AI, OpenResearch and University of Pennsylvania[68] looked at the implications of GPT models on the US labour market. They found that approximately 80% of the US workforce could have at least 10% of their work tasks affected by the introduction of GPTs, while 19% of workers could see at least 50% of their tasks impacted. In absolute terms that is 30.6 million workers in the US alone who could see at least half of their work tasks affected by GPTs.[69] The equivalent figure for the UK if the proportions were consistent would be 6.3 million workers. Affected doesn't mean replaced. And many workers affected would augment their productivity by using the technology in their work. But it does mean that huge swathes of workers will be heavily impacted, and it feels reasonable to assume many will see their jobs disappear.

Figure 13: Impact of AI on UK and US employees
Source: Eloundou, T, et al, March 2023, Bureau of Labour Statistics March 2023[70], UK labour market statistics[71], Onward analysis

	Total workforce (millions)	Workers with 10% of tasks impacted (millions)	Workers with 50% of tasks impacted (millions)
US	160.9	128.7	30.6
UK	33.0	26.4	6.3

The workers impacted stretch across all wage levels, though higher wage workers are estimated to face greater exposure. The impact is also predicted to be felt across sectors, not just highly productive ones. The analysis leads the researchers to conclude that GPT models exhibit the characteristics of general purpose technologies and could have significant economic, social and policy considerations.

What will happen to the labour market?

Over the coming years there will be a dual impact on work. David Autor, Professor of Economics at MIT, describes this as the race between automation and augmentation.[72]

There is the potential for a new class of superstar workers to be created. Recognised and popular artists will be able to go past the constraint of their time to scale production. Taylor Swift might profit from millions of personalised songs generated by AIs using her voice and Tom Cruise could star in thousands of movies each year. Away from the creative sector, leading strategy consultants or bankers, armed with a stellar reputation, broad network and sector experience could benefit from unlimited "interns" at their fingertips to scale productivity - delivering thousands of slide decks or reports to clients. And C-suite executives who understand how to implement these technologies and improve profitability could see their remuneration flourish. As "superstars" increase their productivity and thus their remuneration, there will be an increasing demand for service roles to cater for their needs from dog walkers to personal chefs.

There will, however, be replacement of human labour. From paralegals who may see legal contracts being drafted and edited by large language models to graphic designers where the cost of replication will fall heavily. Of the 19% of workers estimated to have at least 50% of their tasks impacted, it is likely that a significant portion of those roles will be automated out of existence. A survey in the USA of companies who either use ChatGPT or plan to do so found 48% had already replaced workers with the technology. When asked if ChatGPT will lead to any workers in their company being laid off this year 33% said "definitely" while 26% said "probably".[73] Executives are clearly seeing the technology as a labour saving tool and that will be reflected in employment figures.

A recent report by Goldman Sachs found that roughly two thirds of current jobs are exposed to some degree of AI automation and that around a quarter of work tasks in the US and Europe could be automated.[74] This varies significantly by sector: 46% of office and administrative support tasks are exposed to automation compared to just 6% in construction. The net effect is the equivalent of 300 million full time workers across big economies being exposed to automation.

Figure 14: Share of occupational workload exposed to automation by AI
Source: Goldman Sachs research, 2023

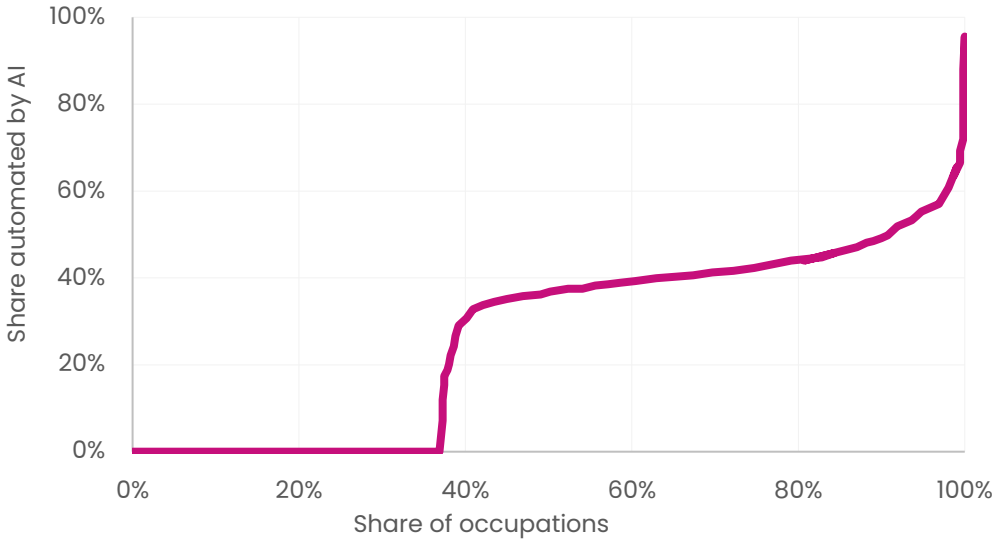

Goldman's base case is for 7% of workers to lose their jobs entirely in the decade after generative AI reaches half of employers, though most will be reintegrated into the labour market. In line with predictions that disruption will be a white collar phenomenon, Goldman views the labour market impact as being more severe in developed markets compared to emerging markets. For example, they forecast just under 30% of FTE employment in Hong Kong to be exposed to automation versus less than 15% in India.

While the impact is likely to be cross industry and skill level, the focus on white collar cognitive work means that the distributional impact will be skewed. 34% of creative jobs are based in London[75], and 56% of jobs in London were classified as high skilled compared to an England average of 46%. Despite the focus on levelling up in recent years, it may actually be London which faces the brunt of labour market disruption when cognitive automation hits.

There will be adjustment costs for displaced workers and negative economic consequences. Economist James Bessen used Dutch administrative data to analyse the impacts on workers who were made redundant after their firm invested in AI.[76] They found the expected annual income loss was 9% after five years, driven by spells of unemployment within a year. The adverse effects were found to be stronger for older and middle educated workers as well as those who had been working in small firms. These costs will be higher as we move to a much faster pace of labour market disruption where re-entering the labour market is harder due to slower new job creation.

The labour market disruption described above is premised on current developments in generative AI. The technology enhances labour productivity as well as replacing it and most jobs will still be needed in the short run. However, if we reach AGI, the job losses will scale up significantly. There will be even less need for human labour if AI systems can replicate and likely perform better than human capabilities in most tasks. In this scenario, which the latest Metaculus forecasts suggests could come as early as 2032[77], we should prepare for much greater labour market disruption.

Policymakers were unsuccessful in navigating the labour market disruption from globalisation in the 1990s, which has contributed to the political shocks of recent years. That disruption is also a contributing factor to the UK's high level of regional disparities and left behind communities.[78] The scale of challenge from AI

induced automation is different - it will affect different areas of the economy, be more pervasive and come at a much faster speed. But there are lessons to be learnt for what happens if the government fails to prepare or react. If we want to benefit from the economic potential of AI we need to grapple with the disruptive consequences.

Macroeconomic impact: inflation and interest rates

The path of macroeconomic variables is impacted by many economic factors and forecasting out for a few months, let alone to the medium term, is challenging. Deep analysis and modelling of these variables is outside the scope of this report. However, there are a few high level impacts of generative AI that are worth considering as they will form challenges that UK policy will need to react to.

Inflation

The economics debate over the past 18 months in the UK has been dominated by high inflation and the cost of living crisis. Inflation has reached double figures, levels not seen in decades, and the central economic issue has been how to bring down inflation and help households with higher prices. However, we may soon face the exact opposite problem. Deflationary pressure.

Widespread adoption of generative AI and continued technological improvements could significantly bring down cost curves, as insight from Schroders suggests, leading to deflation in some industries.[79] There was a similar macro impact when China entered the global trading system in the mid 1990s, which brought sustained downward pressure on prices in the UK and the developed world, as the graph below shows. Then, it was low cost labour in China producing cheap goods imported into the UK which meant that inflation was kept low during periods of

higher growth. This time, it will likely impact services more than goods, and as generative AI brings significant productivity benefits, the cost of services (e.g. legal fees) will come down at an economy wide level.

Figure 15: Consumer Price Index (CPI) inflation rates, 1971-2021
Source: World Bank, World Development Indicators

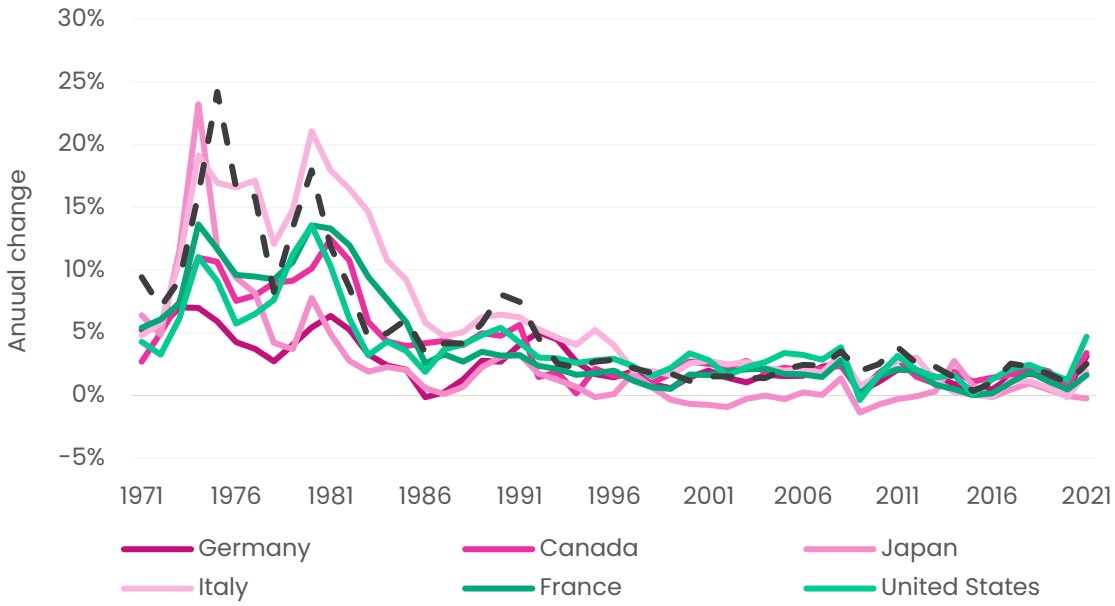

The industrial structure of the LLM market and where the economic value is reaped will affect what the impact on price levels is. There are currently several AI labs competing on LLMs and the competitive dynamic may keep costs of the underlying technology low. As labs compete for market share and to stay in front, they may also subsidise losses to keep prices low. For example, since Google bought DeepMind in 2014 they have been very willing to subsidise losses and allow the company to release research for free, leading to a £0.5 billion loss in 2020.[80] This is part of a long term strategy for Google where they hope DeepMind will develop increasingly important technological breakthroughs which can be commercialised in the future. Subsidised losses and competitive pressure will bring downward pressure on prices given that LLMs will be inputs across sectors.

But if labs manage to successfully build moats or differentiate, they may be able to set higher prices. This could be through creating better foundational models or differentiating themselves through better user experiences or establishing a brand, all factors that OpenAI is currently excelling in. Under this market structure, some AI labs will be able to charge higher prices and the industry won't become commoditised. We can see this in other technology markets. Google and Bing offer a similar product, but Google's superior algorithms, user interface and brand has resulted in it capturing 84% of the global market share.[81] This has also

allowed Google to charge higher prices: Bing's average cost per click is 70% lower than Google's,[82] reflective of lower quality traffic. Google's superior performance means it has also captured a lot of the value uplift reflected in Alphabet's $1.36 trillion market capitalisation[83]. If the LLM market follows this direction, more of the value will accrue to Labs, and prices will be higher than otherwise, resulting in much weaker deflationary pressure.

Interest rates

Deflationary pressures in the medium term may bring down interest rates to counter the threat of deflation. However, it's also worth considering what the long term path of interest rates may look like if we do move towards AGI.

As the productivity and growth section highlighted, generative AI is likely to significantly improve productivity in many sectors and if we move to AGI, the growth impact will be transformational. The researcher Davidson concludes that explosive growth driven by AI is a plausible scenario, where explosive growth is defined as above 30%.[84] This world would necessitate very high long term interest rates to combat the associated long term inflation. That is not what we see today. The US 30 year real interest rate ended 2022 at 1.6% and as recently as Autumn 2021 the UK sold a 50 year bond with a -2.4% real interest rate at the time.

Research by Chow, Halperin and Mazlish posit that the low 30-50 year interest rate environment means one of two things.[85] Either, financial markets are functioning as effective information aggregators and accurately pricing in the fact that transformative AGI is unlikely to happen in the next 30-50 years. Or markets have underestimated the probability and there is a market inefficiency. There is huge uncertainty around the timelines of AGI, yet there is a scenario where timelines are short. Metaculus, the forecasting platform, currently estimates that we will reach AGI by 2032[86] compared to a year ago when the forecast was 2057[87] showing how quickly things are changing. If the probability of us reaching AGI soon increases, we should begin to see that play out in higher longer term interest rates. There could be a sudden shift following a significant breakthrough in the technology, and this would have huge implications, from public finances to financial markets. Low long term interest rate projections shouldn't be taken for granted.

Risks

Focussing on AI safety and solving alignment

To realise the economic and geostrategic benefits of generative AI, it is crucial that we navigate the risks associated with AI systems. These include short term risks around misinformation, transparency and privacy but also most importantly ensuring that AI systems develop safely. One of the most pressing concerns in the AI safety community is the alignment problem: ensuring that AI systems act in accordance with human values and goals. The alignment problem arises because AI systems, especially those with high levels of autonomy, may develop unforeseen behaviours or optimise objectives that are not in line with human intentions, potentially leading to unintended and even harmful consequences.

Many of the difficulties in aligning AIs with human values are highlighted in a simple example from 'Concrete Problems in AI Safety', a research paper with authors affiliated from most of the leading commercial and academic labs[88]. Here they imagine a fictional robot whose job it is to clean up an office space. Even in such a simple setting, setting the goals of the AI agent can be challenging. For instance if the robot is rewarded for creating a clean environment, instead of cleaning the office it may be incentivised to disable its cameras. Likewise, there may be challenges in preventing the robot from dangerous exploration such as if it tried to use a wet mop in an electrical outlet. These are not merely theoretical concerns. These behaviours have been observed in real-life AI models from both DeepMind[89] and OpenAI[90]. In one example, an agent that was supposed to be grasping an object found it easier to fool a human evaluator by hovering between the object and the camera[91].

While these risks may not seem profound, if AI systems remain unaligned with human goals as they scale up in capabilities they will become increasingly dangerous. These concerns are largely based on two hypothetical properties of such systems, first defined by Oxford philosopher Nick Bostrom; orthogonality and instrumental convergence.[92] The former refers to the idea that an agent's intelligence levels and its purposes or goals are independent of one another. The latter suggests that regardless of an AI's main goals, they are likely to have many of the same sub-goals. For instance, just as humans with a variety of ambitions are likely to find greater access to wealth or power useful in achieving their aims, AIs with a large variety of goals might find having access to more compute or information a useful instrumental goal. Combined these two theses imply that a high-capability AI is likely to be power-seeking, leading to concerns over existential risks from AI.

Recent events have brought these existential risk concerns from the alignment problem out of the fringes and into the forefront of public discourse. Over a thousand leading AI practitioners including Elon Musk have signed a letter calling for a six-month pause on AI capabilities research,[93] urging the need for caution and reflection. This call for a pause highlights the growing concern that

advancements in AI capabilities might outpace our ability to develop safe and aligned AI systems.

Likewise, leading AI scientists such as Yoshua Bengio and Stuart Russell have also voiced concerns about the potential consequences of misaligned AI systems. A recent survey of AI experts reveals that the median respondent estimates a 5% chance of AI posing an existential risk to humanity,[94] emphasising the need for ongoing research and vigilance in AI safety and alignment. Ian Hogarth, an AI entrepreneur and investor, recently wrote of the need to "slow down the race to God-like AI"[95]. The wider public, as seen in data below collected by the Centre for the Governance of AI at Oxford, shares many of these concerns.

Figure 16: Share of respondents that agree or disagree that AI is a technology that requires careful management
Source: Centre for Governance of AI

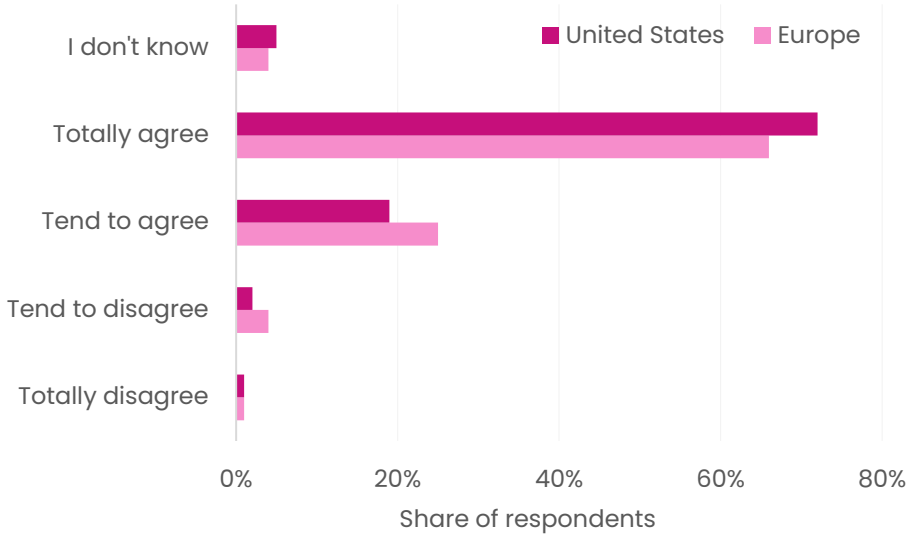

These concerns have prompted proponents to sometimes suggest drastic policy responses. In a recent Time magazine article by Eliezer Yudkowsky, a decision theorist at the Machine Intelligence Research Institute, warns of the existential risks posed by AI and advocates the use of extreme measures, such as airstrikes, in specific scenarios to halt AI progress.[96]

Policy proposals by advocates such as a research pause and restricting chip access often suggest a large trade-off between ensuring the safety of AI systems and society enjoying some of the transformative benefits of generative AI. This trade-off disappears in practice however; it is only by ensuring the safety of systems that it's possible to build confidence to innovate and deploy models out to billions of people around the world. As such the UK's opportunity to play a leading role in AI safety can also help the UK's ambitions to play a leading role in AI capabilities.

The government has taken promising first steps. The International Tech Strategy included a commitment to "Initiate an inclusive international dialogue on the current and future risks presented by AI, for example, general AI systems" making it clear that the Foreign Office views it as an objective to be pursued internationally.[97] And the recent AI white paper discussed a centralised risk monitoring function that will enable the government to monitor 'high impact but low probability events'. While these initiatives are welcome, risks from AI safety will be one of the most important global policy areas of the next decade and the attention that the government devotes to it needs to rapidly increase.

Recommendations

Capitalising on the economic and geostrategic opportunities of AI

This report has discussed the potential economic impacts of generative AI as well as the geostrategic importance. While there may be disruption, particular in labour markets, it is clear that there is huge economic value which will be realised through adoption of generative AI in the coming years. Where that value is captured - whether by the creators of the underlying LLMs or the consumer facing applications - is still to be determined. But there will be economic upside. And the UK should be using policy to ensure that it capitalises on this and is well placed on the international stage.

Recommendation 1.1: The Government should build a UK sovereign LLM: "GB GPT"

The generative AI revolution is being led by a handful of largely Silicon Valley based AI labs. Those that build the models are deciding which values to encode in them, who should have access to them, and how the technology should be developed in coming years. The UK needs to decide what role it wants to play. With developments moving at pace and compute costs spiralling for advanced models there is a very limited time window for the UK to become a serious player.

The Integrated Review set out the government's approach to science and technology through its "own-access-collaborate" framework: [98]

Own: where the UK has leadership and ownership of new developments, from discovery to large-scale manufacture and commercialisation. This will always involve elements of collaboration and access.

Collaborate: where the UK can provide unique contributions that allow us to collaborate with others to achieve our goals.

Access: where the UK will seek to acquire critical S&T from elsewhere, through options, deals and relationships.

The UK should "own" rather than "access" LLMs. This could be through either developing sovereign capability or a sovereign LLM. Developing a government-backed LLM would enable the UK to maintain technological competitiveness on the global stage, ensuring that its AI researchers and companies have access to state-of-the-art tools and resources. It would allow the public sector to use a tailored and safe LLM to enhance public service delivery. And it would give us a seat at the table as the technology transforms the world and not just be reliant on Google DeepMind.

LLMs are not 'neutral'. While not always explicit, values are encoded into the technology through its design and therefore those that are building the technologies are by definition choosing the values that everyone else will interact with. There is therefore a choice for the UK. Does it want to be part of the building process and shape the values that are incorporated into the technology? Or does it want to be a passive user and just seek to regulate them after they've been shaped? Given the UK has global leadership ambitions, the answer should be the former and it must invest in LLMs to be relevant.

The government recently set up the Foundational Model Taskforce, led in the interim by Matt Clifford, to scope the UK's work on sovereign capability. The announcement to invest £100 million through the Taskforce provides the finance to build initial sovereign capability and the Taskforce now needs to ensure that happens, including by creating the right vehicle.

The absence of a sovereign LLM capability could expose the UK to multiple risks:

- The reliance on foreign technology could create vulnerabilities in the nation's critical infrastructure and digital services, as these AI systems could be manipulated, disrupted, or compromised by adversaries. The UK would be vulnerable to either owners of LLMs changing access to their own models, or other groups getting access to those models and manipulating them. To mitigate against this the government might simply not use the technology in public services, but then it would miss out on opportunities for efficiency and improved delivery.
- The lack of domestic AI capabilities could hinder the UK's ability to shape international norms, standards, and regulations governing AI and its applications. This could result in the adoption of unfavourable policies that undermine the UK's economic and strategic interests.

A home-grown LLM - Great British GPT (GB GPT) - could be achieved through a collaboration between the government and the private sector. The government should not be competing with OpenAI to see who can build the latest cutting edge models. It does not have the capability to do this, and it should not be fuelling arms race dynamics. But it could instead commission an LLM built in the UK in partnership with the government. The government could provide the majority of funding and access to datasets, and AI Fellows, as described in the next recommendation, could help design the specifications of the model based on government use cases. The actual building of the model could be commissioned to a UK firm or developed through a partnership between UK academia and the private sector. The government would own at least a majority stake in the LLM and would have a controlling interest to determine who gets access to it and how it is used.

Another model, with similarities to the UK BioBank, would see the UK Government offer a series of grants and in-kind support to smaller organisations looking to develop an LLM in return for specific requirements on open access, domestic ownership and meeting criteria for usage in public services. While this approach might take longer to produce a single, large competitive LLM, it would spread risk and increase the chance of backing a breakthrough innovation.

The first customer of GB GPT should be the government itself. It could be used to improve public service delivery and offer a degree of personalisation in public services which is not currently possible. For example, job centres could use the LLM, which has access to personal details, to create tailored job recommendations for the unemployed based on their background and preferences. Or once the technology has advanced further, a NHS chatbot could be used to triage medical conditions incorporating personal and family medical history. It could even be used in policymaking, simulating what the economic effects of policies would be. The viability and security of these functions would be much greater with a sovereign AI model rather than accessing one from Silicon Valley.

Outside of public services, the government could also give privileged access to GB GPT to UK based academic researchers or start ups. This would provide an incentive for geographically mobile researchers and tech entrepreneurs to base themselves in the UK, providing economic and strategic benefits, and cementing the UK's position as a leader in AI.

GPT-4 cost over $100m to build.[99] Given the increasing amounts of compute needed to train larger models, the cost of developing cutting edge models will likely spiral to hundreds of millions if not billions in the near future. At that stage, the monetary barriers to entry will be much larger and companies operating in an arms race will have become more secretive about how they have built their models and the findings of their latest research. If the UK wants to develop its own capability it needs to start now.

Using GPT-4 as a proxy and recognising that the Government would be starting from scratch, it would likely cost in the region of £200m-£400m to build the first model, depending on the cost of compute. The UK could build one of the most transformational technologies of our lifetime for less than 0.3 % of annual government spending, giving it an international strategic and economic advantage. This is affordable by comparison to any similar project which would give the UK similar strategic capability. For example, the MoD regularly spends many billions on military procurement, such as £10 billion budgeted on the Future Combat Air System (FCAS).[100]

Recommendation 1.2: The Government should set up a UK AI Fellowship to improve AI policymaking

To devise better AI policy, a field which is changing almost weekly and has models which are highly technical by their nature, the UK needs to be able to thoroughly understand the technology. This creates recruitment challenges for the UK Government. Working in AI in the private sector is very attractive and well-remunerated and it is highly unlikely that experts in the field would pursue a career in the public sector. However, many working in AI believe that ensuring the right AI policies in governments is one of the highest impact areas for the future, and thus if the Government established the right scheme there is an opportunity to attract top talent for short periods to help optimise the UK's AI policy.

AI Fellowships could be modelled on the US Tech Congress, an innovative program designed to bridge the gap between technology and policy in the US government.[101] The initiative aims to provide congressional offices with talented technologists who can help inform and shape technology-related policy decisions. The program selects and places tech professionals as Congressional Innovation Fellows, allowing them to serve one-year fellowships in either House or Senate offices. These technologists bring their expertise to bear on a wide range of issues and can greatly enhance the understanding of complex technology issues within government and lead to more effective policymaking.

A UK AI Fellowships programme could similarly be set up to bring tech professionals into government for a one year stint, though it should seek to attract more experienced tech workers than the US scheme to bring additional value. They could sit within DSIT's AI teams, the FCDO's International Tech team, or within the Office for Science and Technology Strategy in the Cabinet Office and bring their real experience into the policymaking process. Fellows could help the government to navigate the constantly changing environment for cutting edge AI developments and move at pace to respond to and correct its own policy. They could provide expertise which would otherwise be impossible for the government to attract.

There are three key areas that Fellows should work on. First, helping the government to advance adoption of AI in the UK and create an attractive environment for AI companies to thrive, in order to secure economic advantage. Second, thinking through the non-economic impacts of AI across society. For example, working with Department for Education on how skills and exam testing policy should change as generative AI develops or helping Ministry of Justice assess how visual evidence should be interpreted in courts in a world where criminals could generate video alibis of themselves. And third, working on AI safety policy and shaping the UK's approach in multilateral fora to lead on

concrete safety initiatives. For the scheme to work successfully, it would need to ensure that Fellows are able to genuinely influence and shape policy and also provide them with an opportunity to really understand the policymaking process to ensure it is an attractive offer.

Recommendation 1.3: The Government should create a new incentive package for AI experts to come to the UK and broaden the HPI visa

Ultimately it is people who are building generative AI models. At the very highest level, there are probably only hundreds of people in the world operating at the cutting edge of transformer models with perhaps several thousand working on building the latest LLMs. Open AI, who have released Chat-GPT, GPT 4 and a whole suite of APIs currently have less than 400 employees.[102] Top AI talent is one of the scarcest resources in the world; where that talent locates will have an outsized impact on where new technologies will be built and where the economic value will reside.

It is very unlikely that the UK could tempt Open AI to move to London, yet we could seek to attract leading AI thinkers who will be building the next generation of consumer facing applications, transformative models and underlying research to come to the UK. Given the impact they could have and the strategic and economic advantage they could give the UK, this has a very high expected value return.

The UK has recently made progress in this space. As part of the UK Science and Technology Framework announced by the Prime Minister and the UK Science and Technology Framework in March, £8 million was committed to encouraging the next generation of AI leaders around the world to do their research in the UK. While the details behind this have yet to be published, this is a welcome step.

Given the importance of this technology, there is the opportunity to scale this much further. Government could:

- Emphasise the benefits of the wider policy environment - if the UK seeks to build GB GPT, accelerates its compute offer and offers a sensible regulatory environment it will be attractive as a leading AI destination in its own right.
- Expand the Global Entrepreneur Programme operated by the Department for Business and Trade - attracting global entrepreneurs building AI companies to relocate their high growth companies to the UK by offering mentoring from experienced entrepreneurs, introductions to key networks, guidance on how to grow internationally and export support.

- Offer free access to government procured compute for researchers as well as introductions to academic institutions where there might be opportunities for collaboration.
- Offer positions in the UK AI Fellow scheme to play a leading role in setting UK policy.
- Make introductions to investors and help facilitate access to seed capital.
- For a small number at the very top of their field, offer direct investment into their ideas and financial incentives such as a one year grant for working from the UK.
- These offers should all be in addition to assistance with a fast tracked visa process and a concierge service to help with the logistics of moving to the UK.

If this package incentivises even a fraction of world leading AI talent to move to the UK, the resulting economic benefits will likely significantly outweigh the costs of the programme.

Alongside this, the government should consider reforming its visa routes to attract more of this talent. The impact of high skilled immigration in creating exciting companies is well evidenced. Research from the Entrepreneurs Network found that 49% of the UK's fastest growing startups have at least one immigrant co-founder despite just 14% of the UK's population being foreign born, and this includes nine of the UK's 14 unicorns.[103]

The UK's High Potential Individual (HPI) Visa scheme clearly has the right direction but its academic criteria means that leading graduate AI programmes are excluded. For example, graduates researching AI at the University of Montreal, which has a thriving ecosystem of AI researchers and one of the strongest graduate programmes wouldn't qualify for a HPI visa to come to the UK. Carnegie Mellon University which is ranked 5th in the world for Computer Science in the Times Higher Education world university rankings[104] also isn't included in the government's eligible list. Introducing a sub-branch of HPI visas for AI talent with a better and wider selection of universities would mean that the UK doesn't miss out on potential talent.

Recommendation 1.4: The Government should extend the exception for Text and Data Mining (TDM) to allow it for any purpose but include an opt-out for content holders

The UK should create an intellectual property regime for AI that carefully balances the interests of creative industries, AI developers, and society at large. This balance is crucial for promoting innovation, fostering economic growth, and

ensuring that the benefits of AI are sustainable. Some tensions between sectors are inevitable. But outcomes which are either overly protectionist or insufficiently pro innovation will not last.

In recent months we've started seeing some of these tensions. Getty Images filed a copyright claim against Stability AI in January in a landmark dispute which will set the direction of the legal ownership of creative output.[105]

One immediate area that the UK should review is its approach to Text and Data Mining (TDM) laws, which should be updated to encourage greater access to data for AI development. TDM refers to the use of computational techniques to analyse digital works and identify trends or other useful information. Currently, the UK has a TDM exception in its copyright law, which allows non-commercial researchers to mine copyrighted materials for the purpose of scientific research. However, this exception has limitations, as it does not extend to commercial entities or other purposes beyond research.

In 2022 the Government set out plans to introduce a new copyright and database exception which would allow TDM for any purpose i.e. including commercial activities. This was very welcome to AI developers and removed restrictions around using copyrighted material for model building. However, inevitably, it prompted a strong backlash from the creative sector, who are exposed to disruption from generative AI despite owning much of the copyrighted material that models are trained on. UK Music said they were "deeply concerned" about the proposals while the Publishers Association described the proposals as a "sledgehammer to crack a nut."[106] The Government backtracked with George Freeman MP confirming in Parliament that it was no longer planning on introducing this exception.

To create a more conducive environment for AI development, the UK could consider broadening the scope of its TDM exception. This could involve allowing commercial entities to mine copyrighted materials for AI development, subject to certain safeguards. For example, requiring the use of anonymized data or aggregated results could help protect the interests of copyright holders while still enabling AI developers to access the data they need.

The UK should also mirror the EU's position on opt-outs. When the EU recently adapted their TDM rules under Article 4 of the EU Digital Copyright Directive, they allowed the use of TDM for commercial purposes provided the rights holder has not opted out through appropriate measures. This opt-out for content creators gives a degree of control to those concerned while still creating a permissive regime attractive to AI model developers. AI companies are globally mobile and the conditions we set will have a significant impact on whether they decide to base

themselves here. But the creative sector is also hugely important in the UK, accounting for £116bn in GVA and representing a pillar of the UK's global soft power. A balanced approach is needed.[107]

The TDM debate is important, and following the EU's position would represent a sensible middle ground between AI innovation and the creative industries. But it is also symptomatic of the much wider tension between these competing goals which will invariably flare up. The UK should seek to create more certainty and a stable policy environment for both sectors; first by cautiously extending TDM exceptions and then more broadly by determining a scope of protection for AI generated works.

Recommendation 1.5: The Government should increase access to compute beyond the commitments made at the Budget.

The rapid progress in AI development has been driven in recent years by the remarkable advancements in computational power. The increasing availability of high-performance compute has enabled researchers and developers to process vast amounts of data, train complex models, and iterate on their designs at an unprecedented scale. As AI technologies become ever more deeply integrated into our daily lives and the global economy, access to compute resources will emerge as a critical factor in determining a nation's competitive advantage in the AI race.

It is essential for the UK to have sovereign compute to attract researchers to the UK and help foster a domestic ecosystem of AI developers and entrepreneurs. It is also essential from a national security perspective as AI systems increasingly become integral to defence and intelligence capabilities.

But the UK has been falling behind. Between 2005 and 2022 the UK fell from 3rd to 10th in international compute rankings and all our competitors are doubling down.[108] The EU's EuroHPC programme is deploying three pre-exascale systems while the US is looking to build systems 5-10 times as powerful as Frontier, their most powerful current system which is far beyond the UK's capabilities.

Figure 17: Change in compute performance rankings, 2005 and 2022
Source: DSIT Review of the Future of Compute, March 2023

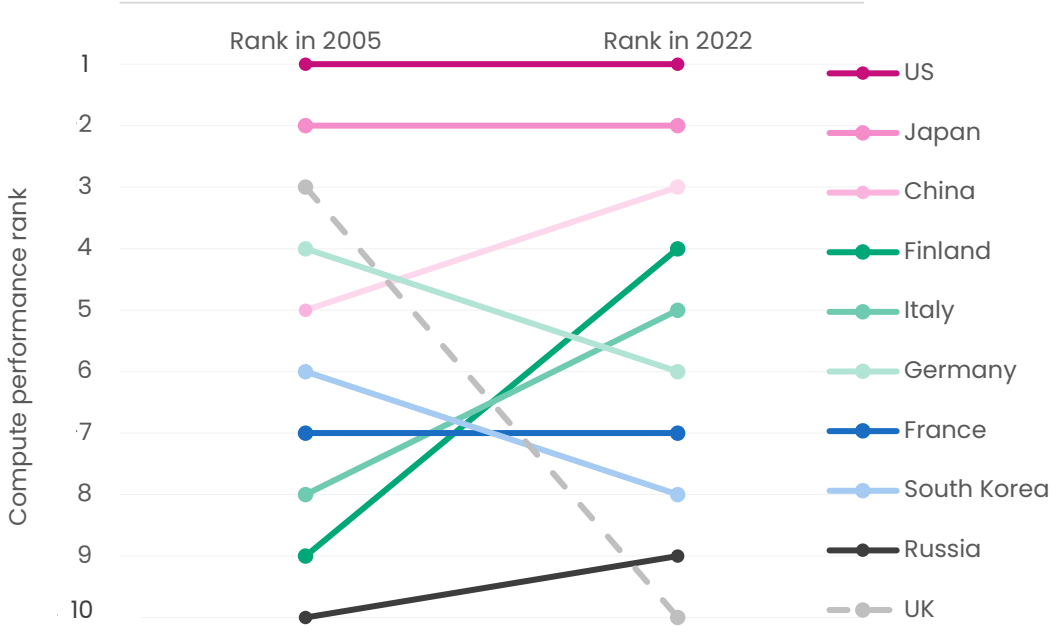

It is in this context that the government launched the Future of Compute Review which reported its conclusions in March 2023. The comprehensive review led to an announcement by the Chancellor at the Budget in March 2023 that the UK Government would invest £900m to build an exascale computer and establish a new AI Research Resource in line with the recommendations.[109]

However, it is right to question whether the ambition matches the scale of opportunity. As former No.10 adviser James Phillips pointed out in a recent blog post[110], the announcement in the Budget to create a new AI Research Resource would use 3,000 GPUs in the commercial cloud. An exascale supercomputer as used by Open AI has the equivalent of 30,000 GPUs. We would therefore only be able to offer researchers 10% of the capacity of a single private sector US company today, despite being a sovereign state with AI superpower ambitions.

This level of ambition cannot be right. While GPUs are difficult to rent today, the UK should urgently seek to increase its purchases, with procurement led commercially from the heart of government to aim to reach the equivalent of 30,000 GPUs by 2026. This will require tens of millions of additional funding from HMT but it is an investment in the UK's ability to compete. There is asymmetric risk here - researchers not having access to sufficient compute over the next couple of years could be significantly detrimental to the UK's AI ambitions, whereas 'over-purchasing' wouldn't add significant cost.

The UK's plans for an exascale computer, while exciting and positive news, could also be enhanced. Anthropic, a leading Silicon Valley US Lab, last year called for the US government to have a 100,000 GPU cluster to be truly competitive.[111] Although this may be beyond our immediate ability, we should commit to procuring this level of top-spec GPUs in the medium run in order to be competitive. This would require significant additional funding from HM Treasury, but given the pace of change and potential upside, we cannot treat the Future of Compute recommendations as static.

Managing the labour market shocks of AI

Recommendation 2.1: HM Treasury should start preparing for a potential shift in the burden of taxation from labour to capital in the medium term.

AI will cause shocks to the labour market. The role of policy will be to manage that shock, helping people to retrain and find new jobs. This transition will be easier to navigate if it is slow. Workers who lose their jobs due to AI will find opportunities created in new sectors and have a lower likelihood of being long-term unemployed.

Shifting the burden of taxation from labour to capital has two potential benefits. First, if the pace of labour market disruption is sudden, it can alter investment incentives and slow the transition to a more manageable pace. Second, it can generate resources for redistribution in a world where we reach AGI, and much of the economic value in the economy is residing with a small number of creators of general AI systems.

There are currently clear tax advantages for employers to use LLMs to replace human workers given the additional taxes that we impose on workers through income tax and national insurance. Our taxation system is labour heavy. In 2021/22 income tax was the highest generating tax raising £225 billion while National Insurance Contributions (NICs) raised £161 billion. Combined, that is £386 billion out of public sector receipts of £915 billion, 42% of the total[112]. This unequal tax treatment has resulted in many people, including Bill Gates, to call for a "robot tax" in an age of automation.[113]

HM Treasury should begin exploring how a shift in the burden of taxation from labour to capital could help slow down the transition in a scenario of fast paced labour market disruption. This would be a significant change to the UK tax system which cannot be done overnight, and thus preparatory work for that potential scenario should begin now. Reducing employer NICs over time would incentivise firms to keep workers on at the margin and thus slow the level of disruption. Depending on the scale of labour market transformation, more radical options on

employee NICs and reducing income tax rates may be needed to incentivise keeping workers in jobs. If the public policy goal here is to slow down the rate of job losses through AI, there will clearly be deadweight losses given the universal nature of labour taxes. But given that AI is likely to be a general purpose technology which will impact most sectors, and incentivising employers to hire for new jobs will be important as well as stemming the slow of losses, a system wide solution is likely needed.

This shift might create public finance challenges, notwithstanding the potential boost to public finances that higher economic growth may bring. To compensate, an increase in capital taxes may be required to raise revenue, and ensure the tax reform as a whole is revenue neutral. The spirit of a "robot tax" can be channelled in different ways. The purchase or development of LLMs could be taxed at an additional rate, or reflected in tweaks to capital allowances. This reflects the option taken by South Korea in 2017[114] who limited tax incentives for investments in automated machines to cushion the impact on the workforce of increasing automation. One blogger has suggested "a small tax on language model API calls… like a Tobin tax on currency transactions."[115] More broadly, assuming that wide adoption of LLMs increases the profitability of firms, governments could look to increase corporate taxes to take a higher proportion of the growing pie.

There are plenty of well-known problems with "robot taxes" or indeed any increase in corporate taxes for those firms investing in LLMs. Given the significant impacts LLMs will have on productivity, reducing their take up will have implications for the UK's productivity performance and corresponding economic growth. It will make the UK less competitive if companies are burdened with higher rates of taxation for using cutting edge technologies. This will be particularly so if other countries don't follow suit. And it will reduce innovation by disincentivising investment in new use cases for AI technology. The UK has had a fundamental problem over the last two decades in encouraging business investment. Taxing investment in new technologies feels counterproductive.

These are all valid critiques and increasing taxes on investment in new technology shouldn't be taken lightly. But there is a scenario where the benefits of slowing adoption overrides the negative economic effects that it brings. AI timelines are uncertain, but if the last 12 months is any indicator of future trends, the pace of development is frighteningly quick. If these trends continue, and adoption takes off, we may be facing a pace of societal change which is very difficult to manage from a public policy perspective.

Shifting the tax burden is not a luddite plan to stop all workers losing their jobs; we will still need to move to the new equilibrium. But if the transition is happening at such a pace that huge disruption is caused to the labour market then there is a case for intervention. HM Treasury should be preparing for this scenario and

developing a plan to shift labour taxes to capital to slow down automation effects in that world.

The second rationale for the shift in relative taxation is if we move towards AGI. In this world, general AI systems will be able to replicate most human activity and will therefore likely extract most of the economic value. As discussed in the labour market section of this paper, labour market disruption and job losses will be significantly scaled if we reach AGI. There will thus be an urgent need for fiscal power for redistribution through a stronger welfare system designed for a world where the value of human labour is limited. In this scenario, labour taxes will become increasingly ineffective ways of revenue raising given the lack of workers. But capital taxes will need to rise to obtain tax revenue from the corporations behind General AI systems, which can fund a much wider welfare state.

Recommendation 2.2: The Government should help workers to train in the skills of tomorrow through more accurate skills forecasting, an expanded retraining offer, and more STEM qualifications

The AI revolution will lead to a significant change in the nature of work and jobs. To minimise disruption and ensure the workforce is capable of thriving in an AI affected world, there are a few areas of skills policy that the Government needs to consider and where it must devise policy at pace.

Firstly, the Department for Education (DfE) should be researching what generative AI developments mean for the skills requirements in the 2030s and 2040s. Predicting the exact skills that will be needed in an AI led future is very difficult, but there is helpful research being done. For example, a paper by researchers from Open AI, OpenResearch and University of Pennsylvania [116] looked at the implications of GPT models on the US labour market and how exposed different occupations would be. They found that roles heavily reliant on science and critical thinking skills showed a negative correlation with exposure, while programming and writing skills were positively associated.

Without such research, policy is at real risk of falling behind. In recent years, the Government has often focussed on retraining into software development and coding. Yet with the advance of generative AI it's unlikely that there will be much growth for employment in coding in years ahead, and if the Government is not careful, workers may spend time and money retraining into skills which become obsolete.

Secondly, the government should improve its retraining offer. This was already important given the trend in people living longer and having multiple careers, but will be especially so now given potential labour market disruption. International

examples show the importance in adult skills systems of having a training component beyond the early stages of careers and Singapore offers a good example of a flexible system where adults over the age of 25 can use credits on training courses at any time they choose.

The Lifelong Loan Entitlement (LLE) is a strong step in this direction but work will need to be done to ensure there is take-up of the offer. A short course pilot that the government is running as a live test of demand for the LLE has been underwhelming. Despite 22 providers offering 103 approved courses for a potential slice of £2.5m of fee loans, only 12 people had signed up for the first available start date for the courses, September 2022. That's 0.11 people per available course.[117]

David Latchman, Vice Chancellor at Birkbeck University reflected that part of the failing of the pilot was the assumption that mature students are like younger students and there is latent demand if courses get set up with the right offer.[118] Given the barriers mature students face to retraining, from financial and caring responsibilities, to social norms, the approach to targeting students must be different. The government should launch a national campaign alongside the rollout of the LLE, openly discussing the challenges of automation from AI and disruption, and build awareness of the importance and necessity of retraining along with support to help people take up opportunities.

More also needs to be done on retraining within firms. As Onward has previously argued, the Government should bring in a Retraining Tax Credit to incentivise firms to retrain their employees.[119] The scheme should mirror the setup of the R&D Tax Credit and be targeted to enable firms to retrain employees in different skills to their current roles. This would help reduce the labour market disruption by keeping workers attached to firms and can also incentivise firms to adopt the new technologies and adapt to them, rather than trying to hold onto historic business models which may not remain viable.

Finally, given the importance of AI and technology in our future economy, the government should be encouraging more people to be studying STEM subjects and developing the technical ability to work in AI. There has been encouraging progress in recent years with undergraduate numbers for STEM courses increasing including in areas like engineering which saw an increase of 21% in course acceptances from 2011 to 2020.[120]

Figure 18: Number of offers made to study AI at UK universities

Source: UCAS Undergraduate Sector-Level End of Cycle Data Resources, 2020

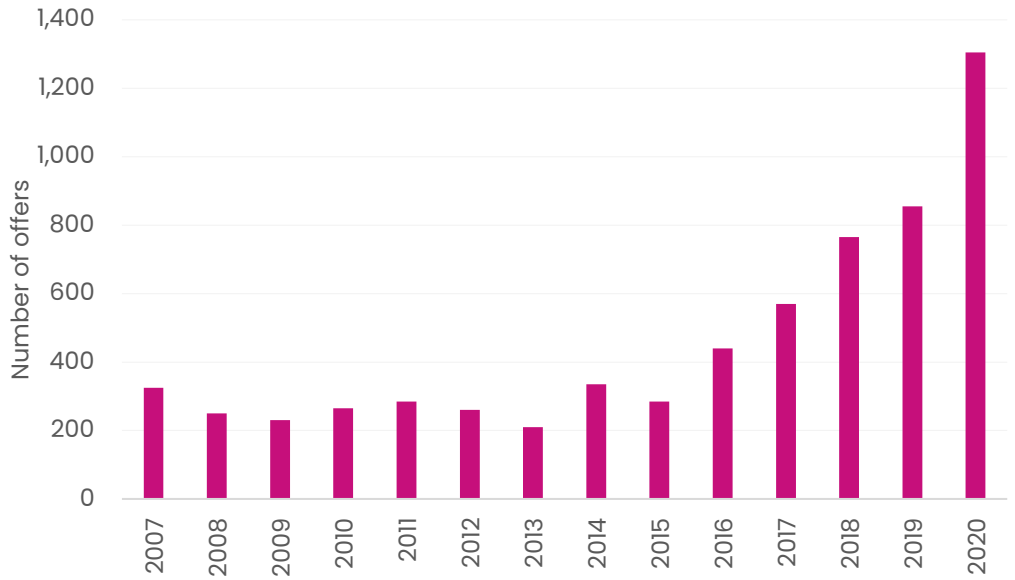

Figure 19: Number of offers made to study all subjects in the Computer Sciences group at UK universities

Source: UCAS Undergraduate Sector-Level End of Cycle Data Resources, 2020

There is, however, further action that should be taken:

- More can be done to encourage women to enter STEM professions. Research from PwC found only 27% of female students surveyed would consider a career in technology and only 5% of leadership positions in the tech sector are held by women.[121] The research found that accessible routes into tech careers such as tailored apprenticeships which the Government can help incentivise, alongside better visibility of female tech role models would help.

- Government should provide better data on the outcome of university courses. Too many students are studying in courses which won't benefit them economically, nor the wider economy. Research from the IFS found that one in five students, about 70,000 each year, are actually losing money by going to university and a significant driver of this is course selection.[122] There is an information asymmetry market failure here where students don't have ready access to clear information about likely future financial earnings from studying different courses. A Government interactive dashboard for students applying to university showing the data of likely salaries by subject chosen, would encourage more students to pick economically valuable subjects to study, and increase uptake in STEM.
- The supply of STEM courses should be boosted. As a report by the Tony Blair Institute highlighted, UK universities offer sub optimal numbers of STEM courses at the elite level. Oxford only offers 32 places for Computer Science, the same as in 2002. Stanford's computer-science major offers 36 core classes in AI compared to Imperial who have just one.[123] Government should use its position to incentivise the creation of more elite courses including through funding. It should also reduce the time taken for new courses to be created at universities to ensure that they can adapt at a faster pace.

Limiting AI safety risks

In order to take advantage of the potential of AI the UK needs to ensure that AI systems are reliable, secure, and aligned with human values. This is essential for ensuring the positive consequences of AI, and also presents an opportunity for the UK to take leadership as part of its Global Britain agenda. By prioritising AI safety research and developing robust regulatory frameworks, the UK can create a competitive advantage in the global AI landscape.

Recommendation 3.1: The Government should create a UK Evaluations Framework to shape how AI systems are built and assessed

Evaluations or benchmarks refer to a set of standardised metrics or tests that are designed to assess the performance of AI systems, such as generative AI models. Often, models are trained through reinforcement learning from human feedback (RLHF) to optimise their performance on these evaluations. However, at present, there is no universally adopted set of evaluations or benchmarks, and the private sector lacks the necessary incentives and means to develop comprehensive standards. Consequently, there is an opportunity for the UK government to take the lead in this area and establish a national set of benchmarks to evaluate generative AI models.

By setting national standards the Government can exert a significant influence on the performance of AI models worldwide through the training process. This would enable the UK to help steer global AI research and development in a direction that ensures AI models are more reliable and better aligned with human values. Establishing such benchmarks can increase confidence in AI systems and foster a greater sense of trust and collaboration between countries, industry, and academia.

Increased confidence in the outputs of AI models, driven by the adoption of standardised evaluations, can also accelerate the integration of these models into UK businesses. As a result, firms can take advantage of the enhanced efficiency, innovation, and decision-making capabilities offered by AI systems, ultimately boosting productivity and economic growth across the country.

Monitoring the performance of AI models against established evaluations is also essential for informing future policy. Policymakers could introduce requirements that AI models being used in critical sectors, such as healthcare, must meet specific performance benchmarks. This could ensure the safety, reliability, and ethical deployment of AI systems, while also reinforcing the UK's position as a global leader in AI safety and innovation

Recommendation 3.2: The Government should monitor and better distribute compute access

As discussed earlier in this paper, the availability of compute resources plays a pivotal role in the development of cutting-edge generative AI models. This has led to a competitive race among firms seeking to acquire increased computational power. While securing sovereign compute access is crucial for the UK to maintain its competitive position and promote innovation, it is equally vital to ensure that AI safety researchers can access these resources to stay current with the latest AI advancements.

In order to guarantee that AI safety researchers have access to computational resources similar to those used by state-of-the-art models, the UK government must first determine the extent of compute resources currently employed by firms. One possible approach is to introduce reporting obligations for both users of compute, such as AI labs, and suppliers of compute, including data centre operators.

The reporting criteria could be established collaboratively and assessed regularly to concentrate on relevant systems, including state-of-the-art foundational models. This ensures that it gathers the most useful information while minimising the impact on private actors. Likewise reporting might draw upon existing

technical data in its native form, encompassing elements like configuration files and logs of comprehensive training sessions. This provides the UK government with the strategic clarity it requires in order to make national compute decisions while imposing almost no burdens on private firms.

Once the UK government has gathered information on the compute resources utilised by private firms, it should establish a system that prioritises access to compute for academics and AI safety researchers. Ensuring that these researchers can work with an equivalent level of compute as leading models is crucial for alignment research, which in turn contributes to the reliability and safety of AI systems.

Currently, too few experts within leading private firms are devoting their time to alignment: fewer than 100 researchers were employed in this area across the core AGI labs in 2021, with DeepMind dedicating just 2% of their total headcount and OpenAI only about 7%.[124] By working in tandem with state-of-the-art models, researchers can provide valuable insights that foster greater confidence in AI systems among businesses. This increased confidence can accelerate the adoption of transformative technologies, further driving innovation and growth across various sectors in the UK.

Recommendation 3.3: The Government should create a centralised UK AI regulator with oversight over foundational AI: the Office for Foundational Models (OFFOM).

Given the implications of AI across the UK economy and its potential to impact multiple industries, it is important to regulate AI holistically, rather than just relying on sector-specific frameworks. A centralised AI regulator, the Office for Foundational Models (OFFOM), would enable a unified and comprehensive approach to AI safety as well the use of foundational models in the economy.

Given the abundance of regulators across the economy, there should be a strong rationale to creating yet another one. With OFFOM, there is. As AI models become more capable of performing tasks across multiple sectors, their applications extend beyond the jurisdiction of sector-specific regulators. For example, ChatGPT can be employed in advertising, drafting legal contracts, and providing medical advice, demonstrating the need for centralisation. The actual models will also become increasingly important; not just their commercial applications, as it is the foundational models which are encoding values and will be driving progress. Given the scale of these models and their reach, it will be impossible for OFCOM and other sectoral regulators to keep up.

OFFOM can also make the UK more attractive internationally for AI development. A strong foundational model regulator can set clear rules for model development usage and deployment, reducing legal uncertainty, and making it more likely that AI labs base themselves in the UK. It can increase adoption of models by firms through increasing confidence in their performance, leading to higher productivity benefits. And there is an advantage to creating the regulator quickly, as other regulators around the world are currently devising their own approaches to regulating foundational models and they may follow the UK's principles.

OFFOM will also be important for AI safety. One of the key aspects of AI safety as discussed above is the monitoring of compute resources, which is essential to the development of state-of-the-art generative AI models. The policy proposal emphasises the need for cross-sectoral compute monitoring, as opposed to a sector-by-sector model. A centralised regulator would be better equipped to implement and enforce such monitoring, ensuring a consistent and effective approach to managing the use of compute resources in AI development.

The proposals in the recently published AI white paper included a central function to support regulators as they develop rules and monitor future risks. This will not be enough. New foundational models will need regulation, and without a single regular holding responsibility this will fall through the cracks. If OFFOM is created it would need to be well funded, given LLM experience is in high demand, and would need to work closely with sectoral regulators who will still be responsible for the application of LLMs in their sectors.

Endnotes

[1] Evan Castillo (2023), OpenAI's GPT-4 is Here. It's Passing More Exams, Best Colleges
[2] GPT-4 Technical Report, Open AI
[3] Midjouney
[4]Midjourney
[5] Azizi S, Kornblith S, Saharia C, Norouzi M, Fleet D.J (2023), Synthetic Data from Diffusion Models Improves ImageNet Classification, Google Research.
[6] Ibid
[7] Twitter, Dr Jim Fan, AI Scientist at NVIDIA, Link
[8] Twitter, Link
[9] Park J.S, O'Brien J.C, Cai C.J, Morris M.R, Liang P, Bernstein M.S. (2023), Generative Agents: Interactive Simulacra of Human Behavior. Cornell University, Arxiv.
[10] Jordan Pearson, (2023) A viral AI-Generated Drake Song by 'Ghostwriter' Has Millions of Listens, Vice
[11] YouTube, (2023), Link
[12] J Hoffmann, S Borgeaud, A Mensch, E Buchatskaya, T Cai, E Rutherford, D. d. L Casas, L. A Hendricks, J Welbl, A Clark, T Hennigan, E Noland, K Millican, G.v.d Driessche, B Damoc, A Guy, S Osindero, K Simonyan, E Elsen, Jack W. Rae, O Vinyals, L Sifre, (2022), Training Compute-Optimal Large Language Models, Cornell University, Arxiv.
[13] Bender E.M et al, 'On the dangers of stochastic parrots: can language models be too big?' ACM Digital Library, Open Access Article, March 2021, Pages 610-623
[14] Devin Coldewey, (2019), Profits are capped at 100 times any investment
OpenAI shifts from nonprofit to 'capped-profit' to attract capital, Tech Crunch
[15] Krystal Hu (2023), ChatGPT sets record for fastest growing user base - analyst note, Reuters
[16] Google DeepMind, (2023), Official website, Link
[17] For more information see here
[18] Richard Waters, Kadhim Shubber, (2023), Google invests $300mn in artificial intelligence start-up Anthropic, Financial Times.
[19] Ibid
[20] Jay Peters, (2023), Elon Musk founds new AI company called X.AI, The Verge
[21] Ian Hogarth, (2023) We must slow down the race to God-like AI, Financial Times
[22] Krishan Shah, Gregory Thwaites (2023) Minding the (productivity and income) gaps, The Economy 2030 Inquiry
[23] 'Generative AI could raise global GDP by 7%,' Goldman Sachs, Insights blog, 5th April, 2023
[24]Office for National Statistics, 'UK Whole Economy: Output per hour worked % change per annum SA', April 26th 2023
[25] NM Group (2018), A short history of the UK electricity network
[26] Robert Solow, (1987), The IT Productivity Paradox, CS Stanford
[27] YouGov Daily ChatGPT Survey, Link
[28] 'Bain & Company announces services alliance with OpenAI to help enterprise clients identify and realize the full potential and maximum value of AI,' Bain and Company Press Release, February 21st, 2023, Link
[29] Aaron Baar, (2023), Expedia integrates ChatGPT intro mobile travel app, Marketing Dive
[30] Government of Iceland, (2023), Head start for Icelandic
[31] Noy, S. and Zhang, W, 'Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,' MIT Working Paper, March 2nd, 2023
[32] Jia, Nan and Luo, Xueming and Fang, Zheng and Liao, Chengcheng, When and How Artificial Intelligence Augments Employee Creativity (March 22, 2023). Forthcoming, Academy of Management Journal
[33] Rosenbush, S, 'Biologists Say Deep Learning Is Revolutionizing Pace of Innovation,' CIO Wall Street Journal, March 22nd, 2023
[34] Lisbona, N, 'How artificial intelligence is matching drugs to patients,' BBC News, April 17th 2023
[35] Hughes, A, (2023) 'ChatGPT Will Replace Programmers Within 10 Years,' Levelling Up Coding, Medium
[36]Assaraf, A, (2015) 'This is what your developers are doing 75% of the time, and this is the cost you pay,' , Coralogix
[37] Peng, S. et al (2023) 'The Impact of AI on Developer Productivity: Evidence from GitHub Copilot,' MIT
[38] Amanda Yeo, (2023), DoNotPay's AI lawyer stunt cancelled after multiple state bar associations object, Mashable
[39] Fedasiuk, R, Melot, J and Murphy, B, 'Harnessed Lightning: How the Chinese Military is Adopting Artificial Intelligence', Center for Security and Emerging Technology, October 2021

[40] '2022 National Defense Strategy of The United States of America', U.S. Department of Defense, October 27th 2022
[41] 'Fact Sheet' CHIPS and Science Act Will Lower Costs, Create Jobs, Strengthen Supply Chains, and Counter China', The White House, August 9th 2022
[42] 'National Security Strategy: October 2022', The White House, October 2022 Link
[43] Payne, K (2018), 'Artificial Intelligence: A Revolution in Strategic Affairs?', Global Politics and Strategy
[44] Braw, E, 'China is the master of grey-zoen aggression', Financial Times, February 14th 2023
[45] Evans, B (2020), 'The end of the American internet' Link
[46] Kim, S (2022), 'Roles and Limitations of Middle Powers in Shaping Global Cyber Governance', Italian Journal of International Affairs
[47] 'Can the West win over the rest?', The Economist, April 13th 2023
[48] Sampson, M et al (2022), 'China's Multi-Front Institutional Strategies in International Development Finance', The Chinese Journal of International Politics
[49] DeepMind, 'About' Link
[50] Phillips, JW (2023), 'S&T - Is the UK a world leader in science?', Link
[51] '2021 AI Index Report', Stanford University Human-Centred Artificial Intelligence Link
[52] Bass, D, 'Microsoft Invests $10 Billion in ChatGPT Maker OpenAI', Bloomberg, January 23rd 2023
[53] Macron, E, 'Emmanuel Macron's 3 ways to renew Europe', World Economic Forum, May 10th 2019
[54] Department for Science, Innovation and Technology, 'Initial £100 million for expert taskforce to help UK build and adopt next generation of safe AI', Press release, April 24th 2023
[55] Leonhardt, M, 'Some workers are worried that ChatGPT will replace their jobs. They might be right,' Success AI, Fortune, March 25th, 2023
[56] Office for National Statistics Employment and Labour Market data, 'Unemployment rate (aged 16 and over, seasonally adjusted): %,' Released 18th April 2023, Link
[57] Office for National Statistics Labour Market Overview, UK, March 2023. Link
[58] Bessen, J, Toil and Technology,' Finance and Development, March 2015, Vol. 52, No. 1, IMF. Link
[59] Autor D et al (2022), New Frontiers: The Origins And Content Of New Work, NBER
[60] Goos M, et. al, (2014), Explaining Job Polarization: Routine-Biased Technologial Change and Offshoring, The American Economic Association
[61] Will Tanner, Guy Miscambell, James Blagden, (2019) Human Capital, Onward
[62] Manyika, J., Lund, S., Chui, M, Bughin, J., Woetzel, J., Batra, P., Ko, R., and Sanghvi, S, 'Jobs lost, jobs gained: What the future of work will mean for jobs, skills, and wages,' McKinsey Global Institute, November 2017, Part 4
[63] House of Commons Library, 'A5: Services,' Economic Indicators, 21st April, 2023. Link
[64] DCMS Sectors Economic Estimates 2020: annual GVA, National Statistics, 6th April 2023
[65] 'Professional & Business Services sector: Creating further demand and growth outside London,' BEIS Research Paper 2020/006, February 2020
[66] Jaime Sevilla, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbhahn, and Pablo Villalobos (2022). 'Compute Trends Across Three Eras of Machine Learning'. ArXiv, http://arxiv.org/abs/2202.05924.
[67] 'Which groups find it hardest to find a job following a period out of work?', Office for National Statistics, March 30th 2021 Link
[68] Eloundou, T, Manning, S, Mishkin, P and Rock, D (2023), 'GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models'
[69] 'The Employment Situation - March 2023', Bureau of Labor Statistics, April 7th 2023 Link
[70] Bureau of Labor Statitics, US Department of Labour, 'The Employment Situation – March 2023,' March 2023, News Release. Link
[71] Francis-Devine et al (April 2023), 'UK Labour Market statistics,' House of Commons Library. Link
[72] Autor, D (2022), 'The Labor Market Impacts of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty', National Bureau of Economic Research
[73] '1 in 4 companies have already replaced workers with ChatGPT', Resume Builder, Feb 2023 Link
[74] 'Generative AI could raise global GDP by 7%, Goldman Sachs Artificial Intelligence Insights. Link
[75] 'Jobs in the creative economy in London and all other regions of the UK, 2010 to 2021', Office for National Statistics, October 25th 2022
[76] Bessen, James E et al (2019) Automatic Reaction - What Happens to Workers at Firms that Automate? Boston Univ. School of Law, Law and Economics Research Paper
[77] Barnett, M, 'When will the first general AI system be devised, tested, and publicly announced?' Metaculus question, Link
[78] O'Brien, N. 'Measuring Up for Levelling Up,' September 2020, Onward report
[79] Rode, N, 'The generative AI revolution: reflections upon my return from Silicon Valley,' Schroders Perspective Blog, 24th February, 2023. Link
[80] Bradshaw, T, 'Google's AI unit DeepMind swallows £1.6bn as losses continue,' Financial Times, December 17th, 2020

[81] Theuring, J, 'Bing vs Google: Search Engine Comparison 2023,' Impression, August 2nd, 2021
[82] Notermans, M, 'Bing Ads vs. Google Ads: Which Platform is Better in 2023?' Think Orion blog, March 29th, 2023
[83] Alphabet Inc Class A, Google Finance, NASDAQ Live Feed
[84] Davidson, T, 'Could Advanced AI Drive Explosive Economic Growth?', Open Philanthropy, June 25th 2021
[85] Ibid
[86] 'When will the first general AI system be devised, tested, and publicly announced?', Metaculus Link
[87] Cheung, R, Twitter, March 28th 2023 Link
[88] Amodei, D et al (2016) 'Concrete Problems in AI Safety' Link
[89] Everitt, T et al, 'Specification gaming: the flip side of AI ingenuity', DeepMind, April 21st 2020 Link
[90] 'Faulty reward functions in the wild', OpenAI, December 21st 2016 Link
[91] 'Learning from human preferences', OpenAI, June 13th 2017 Link
[92] Boston, N (2012) 'The Superintelligent Will: Motivation and Instrumental Rationality in Advanced Artificial Agents', Minds and Machines
[93] 'Pause Giant AI Experiments: An Open Letter', Future of Life Institute, March 22nd 2023 Link
[94] 'How generative models could go wrong', The Economist, April 19th 2023
[95] Hogarth, I (2023), 'We must slow down the race to God-like AI', Financial Times
[96] Yudkowsky, E, 'Pausing AI Developments Isn't Enough. We Need to Shut it All Down', Time, March 29th 2023
[97] Department for Science, Innovation and Technology, 'The UK's International Technology Strategy', CP 810, March 2023
[98] 'Global Britain in a Competitive Age: The Integrated Review of Security, Defence, Development and Foreign Policy', CP 403, March 2021
[99] Knight, W, 'OpenAI's CEO Says the Age of Giant AI Models Is Already Over', Wired, April 17th 2023
[100] 'Global Combat Air Programme: Finance', UK Parliament Written Questions, Answers and Statements, 30 March 2023 Link
[101] TechCongress, 'About Us' Link
[102] Holmes, A and Victor, J, 'OpenAI Is Making Headlines. It's Also Seeding Talent Across Silicon Valley', The Information, February 1st 2023
[103] Dumitriu, S and Stewart, A (2019), 'Job Creators: The Immigrant Founders of Britain's Fastest Growing Businesses', The Entrepreneurs Network
[104] Times Higher Education (2023), 'World University Rankings 2023 by subject: computer science'
[105] Brittain, B, 'Getty Images lawsuit says Stability AI misused photos to train AI', Reuters, Feb 6th 2023
[106] 'UK text and data mining copyright exception proposals set to be watered down', Pinsent Masons, January 26th 2023
[107] Weitzman, E. 'Impact of government policy on the creative sector,' House of Lords library, 28th October, 2021
[108] Independent Review of the Future of Compute: Final report and recommendations,' Department for Science, Innovation and Technology, March 6th 2023
[109] Department for Science, Innovation and Technology, 'Government commits up to £3.5 billion to future of tech and science', Press release, March 16th 2023
[110] Phillips, JW (2023) 'Securing Liberal Democratic Control of AGI through UK Leadership' Link
[111] Anthropic (2022), '(RFI) on Implementing the Initial Findings and Recommendations of the National Artificial Intelligence Research Resource Task Force: Response', Federal Register Link
[112] Keep, M (2023), 'Tax statistics: an overview', House of Commons Library
[113] 'The robot that takes your job should pay taxes, says Bill Gates', Quartz Link
[114] 'South Korea first in world to adopt 'robot tax'', Tech Wire Asia, August 9th 2017
[115] Cowen, T (2023), 'Brian Slesinsky on AI taxes (from my email)', Marginal Revolution
[116] Eloundou, T et al (2023), 'GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models.'
[117] Kernohan, D (2022), 'Little demand for fee loan backed short courses', Wonkhe
[118] Latchman, D (2023), 'How could we test the attractiveness of the lifelong loan entitlement?', Wonkhe
[119] Blagden, J, Miscampbell, G and Tanner, W (2019) 'Human Capital: Why we need a new approach to tackle Britain's long tail of low skills', Onward Report
[120] Department for Education, 'More young people are taking STEM subjects than ever before', The Education Hub, Feb 2021
[121] 'Women in Tech: Time to close the gender gap', PwC, 2017
[122] Johnson, P (2020) 'Too many students take courses that don't benefit them or the economy', Institute for Fiscal Studies
[123] Kakkad, J, Macon-Cooney, B, Northend, J and Westgarth, T (2023), 'A New National Purpose: Innovation Can Power the Future of Britain', Tony Blair Institute for Global Change
[124] Hogarth, I (2023), 'We must slow down the race to God-like AI', Financial Times

Support Onward

Onward is an independent, not-for-profit thinktank. We rely on the generous support of individuals and trusts, as well as partnerships with charities and businesses, to support our leading programme of research and events.

Individual and Trust donations

Onward's core programme of research is funded by individual and trust donations. If you are an individual or represent a philanthropic trust and would like to support our ongoing research agenda, we would love to hear from you. We do not accept corporate funding for research reports. Please contact us on office@ukonward.com if you would like to donate by cheque or you can donate electronically using Onward's account details below.

UK Onward Thinktank Ltd
Not-for-profit company no. 11326052
Bank: Natwest
Account number: 21328412
Sort code: 50–10–05

Please note that Onward retains copyright and full editorial control over any written research it produces, irrespective of funding.

Partnerships with companies and charities

Alongside our research, we want to work closely with charities, trusts and businesses, large and small, to further the political debate and bring the expertise and ideas from outside government to bear on the policymaking process.

If you would like to partner with Onward as a sponsor of one of our roundtable events or conferences, or to join our Business Leaders Network to engage further with our work, please get in touch at office@ukonward.com.

A commitment to transparency

Onward is committed to transparency and will publish the names of individuals and organisations who give us more than £5,000 of support each year. This is published on our website twice a year.

www.ukonward.com

ONWARD