

Institute for Public Policy Research



TRANSFORMED BY AI

HOW GENERATIVE ARTIFICIAL
INTELLIGENCE COULD AFFECT WORK
IN THE UK – AND HOW TO MANAGE IT

**Carsten Jung and
Bhargav Srinivasa
Desikan**

March 2024

ABOUT IPPR

IPPR, the Institute for Public Policy Research, is an independent charity working towards a fairer, greener, and more prosperous society. We are researchers, communicators, and policy experts creating tangible progressive change, and turning bold ideas into common sense realities. Working across the UK, IPPR, IPPR North, and IPPR Scotland are deeply connected to the people of our nations and regions, and the issues our communities face.

We have helped shape national conversations and progressive policy change for more than 30 years. From making the early case for the minimum wage and tackling regional inequality, to proposing a windfall tax on energy companies, IPPR's research and policy work has put forward practical solutions for the crises facing society.

IPPR
4th floor,
8 Storey's Gate
London
SW1P 3AY

T: +44 (0)20 7470 6100

E: info@ippr.org

www.ippr.org

Registered charity no: 800065 (England and Wales),
SC046557 (Scotland)

This paper was first published in March 2024. © IPPR 2024

The contents and opinions expressed in this paper are those of the authors only.

The progressive policy think tank



CONTENTS

Summary	5
1. Introduction: The moment we're in	9
1.1. The debate around automation and the labour market: will this time be different?.....	9
1.2. The policy debate around automation needs a step change	11
2. Why this time is different – generative AI is here	12
2.1. What can and can't existing generative AI do?	12
2.2. Assessing exposure to AI.....	13
3. Who is most exposed to generative AI?	15
3.1. Cognitive tasks are most exposed to generative AI	15
3.2. Back office jobs are most exposed to 'here and now' AI	17
4. Scenarios for how generative AI could affect work in the future	22
4.1. Whether AI will boost wages and growth will depend on how it is implemented	22
4.2. There might not be enough 'similar' jobs available for displaced workers	26
4.3. Industrial strategy can help create demand for new or less-exposed tasks	28
5. Towards a job-centric industrial strategy for AI	34
References	36
Appendix	38

ABOUT THE AUTHORS

Carsten Jung is senior economist at IPPR.

Bhargav Srinivasa Desikan is senior research fellow at IPPR.

ACKNOWLEDGEMENTS

The authors would like to thank Harry Quilter-Pinner, Melanie Wilkes, Carys Roberts and Professor Leigh Shaw-Taylor for their invaluable comments. And we would like to thank Abi Hynes and Richard Maclean for his copyediting and design. We would like to thank Rosie Day, Andy Dickerson, Gennaro Rossi and Rob Wilson for kindly sharing with us their crosswalk database.

Download

This document is available to download as a free PDF and in other formats at:

<http://www.ippr.org/articles/transformed-by-ai>

Citation

If you are using this document in your own writing, our preferred citation is:

Jung C and Srinivasa Desikan B (2024) *Transformed by AI: How generative artificial intelligence could affect work in the UK – and how to manage it*, IPPR. <http://www.ippr.org/articles/transformed-by-ai>

Permission to share

This document is published under a creative commons licence:

Attribution-NonCommercial-NoDerivs 2.0 UK

<http://creativecommons.org/licenses/by-nc-nd/2.0/uk/>

For commercial use, please contact info@ippr.org



SUMMARY

TECHNOLOGICAL CHANGE IS A DOUBLE EDGED SWORD

Technological change is a good thing. It has brought exponential gains to living standards and is the foundation of modern society. **Yet unmanaged technological change has always come with risks and disruptions.** For instance, the mechanisation of textile work in the late 18th century brought huge labour market disruption: it reversed women's employment participation. And the female employment share did not return to its late 18th century peak before the 1980s – 150 years later (Shaw-Taylor et al 2019).¹

Unmanaged deindustrialisation too – driven by globalisation and technological change – boosted growth but increased wage inequality in the United States and the UK, with economic scarring effects that last until this day. But equally, history offers examples **when technological change was well managed and helped to address societal challenges and enduringly boosted prosperity.** The Green Revolution in the mid-20th century is such a case when well-managed and policy-driven adoption of new innovations in agriculture increased crop resilience and avoided large scale food pressures around the world (Conway 1999).²

With another technological wave driven by generative AI on the horizon, these experiences show that policymakers should explore risks and benefits before deployment becomes widespread. **Generative AI refers to new computer software that can read and create text, software code and data. Cutting edge models have even shown ability reason and apply abstract concepts in a range of disciplines, often at undergraduate level.**³

Generative AI can be economically disruptive through its impact on wage inequality, wealth inequality and potential job displacement. In other words, there will be winners and losers. Policy can address this, on the one hand, through fiscal policy, eg by taxing the private gains made from technology deployment. And, on the other hand, it can support the creation of public goods - ie by boosting labour-augmenting automation and the creation socially beneficial work.

GENERATIVE AI IS SET TO TRANSFORM KNOWLEDGE WORK

The core of our findings is that the world of knowledge work will be transformed by generative AI and that we need to start preparing for this now. With existing technologies, but especially with those currently in the development phase, almost every aspect of knowledge work could in some form be aided by generative AI. Customers could habitually interact with AI-assisted advice systems, business inventory management could rely on AI engines and company's regular HR training could be delivered interactively by AI tutors. AI could even be used as input for creative and strategic products, such as generating a first draft of an article or a script – as it already increasingly is. But the exact shape of this will depend on design choices.

1 Little studied until recently, quantitative analysis of employment patterns during the Industrial Revolution in Britain show that female labour market participation in Britain fell by more than half from 25 per cent to 10 per cent between 1851 and 1901 (Horrell 2009). This is because it increased productivity but also concentrated work in factories, which were outside the purview of women at the time. See also Humphries and Schneider (2021).

2 Though the new methods adopted in the Green Revolution also created a new set of challenges, such as soil degradation, that had to be managed.

3 This has been reflected in a range of benchmarking exercises (eg applying legal knowledge to novel cases), where AI has been found to produce similar or even higher quality outputs than most humans, in a fraction of the time. In other domains (such as medical image recognition) it still falls short of human level output, though models are constantly advancing.

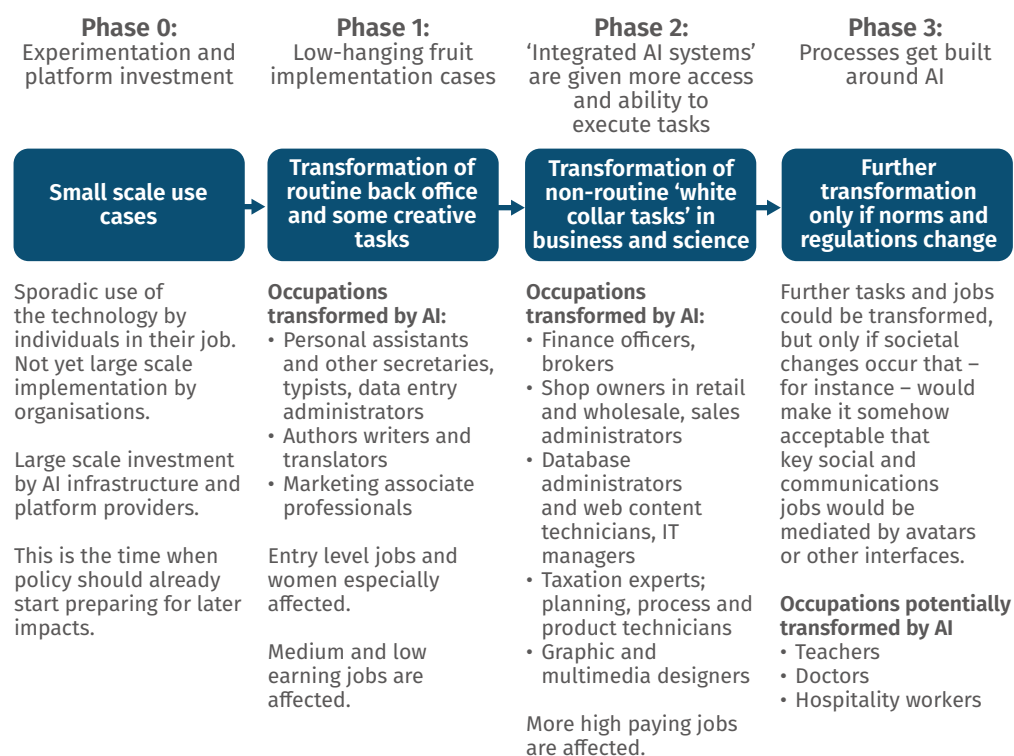
BACK OFFICE JOBS – WHICH ARE MORE LIKELY TO BE OCCUPIED BY WOMEN – ARE MOST EXPOSED IN THE FIRST PHASE OF GENERATIVE AI DEPLOYMENT

In a large scale assessment of 22,000 tasks in the UK economy, **we find that about 11 per cent of tasks are exposed to generative AI right now, and this could increase fivefold** if AI systems became more deeply integrated in organisational processes. We summarise our findings in figure S.1, showing a scenario for how this could play out over time, in four phases.

The phase which we have been in roughly since the advent of GPT4 in 2023, can be described as experimentation phase. Companies and governments are testing which types of tasks AI models can correctly perform, and how they can be introduced in workflows while ensuring quality and oversight. Meanwhile, large tech companies are investing hugely into building generative AI platforms, seeking to make it easier for organisations to integrate generative AI in their work (The Economist 2023a, 2023b).

Phase 1 refers to implementation in organisations that will likely target ‘low hanging fruit’ use cases. These are the cases where generative AI programmes are relatively easily plugged into existing IT processes, without many changes to workflows. **About 11 per cent of tasks would be heavily impacted** by this. **Back office jobs (such as personal assistants), entry level jobs and part time jobs will be most exposed in this first phase.** And we find that **women will be significantly more affected** (as they are more likely to work in the most exposed occupations, such as secretarial and administrative occupations). While the overall labour market impact in this phase could be limited, it could nonetheless be disruptive in these occupational groups. For example, in administrative occupations, about a third of jobs could be displaced.

FIGURE S.1: IN PHASE TWO OF GENERATIVE AI DEPLOYMENT, TWO-THIRDS OF UK JOBS COULD BE TRANSFORMED



Source: Authors' analysis

The second phase is one where generative AI becomes more deeply integrated with existing organisational processes. If organisations decide to integrate existing AI technology more deeply into their processes (which is not a given), we find that almost **five times more tasks – about 59 per cent of tasks – are exposed.** This means a large number of jobs could ‘feel’ its impact, and it will also increasingly affect high paying jobs too. Such integration into existing processes would, for instance, mean giving AI the ability to access proprietary data, providing inputs via apps or giving AI systems the ability to execute tasks (eg making orders or bookings). Whether this will materialise and who gains and who loses will depend on a number of policy and organisational factors. Crucially, it is likely that not all organisations will adopt the technology at similar rates, leading to inequalities.

WE NEED A JOB-CENTRIC INDUSTRIAL STRATEGY FOR AI IN ORDER TO REALISE THE BENEFITS OF GENERATIVE AI AND AVOID UNINTENDED COSTS

The aggregate effects of generative AI will depend on design, adoption and policy choices. ‘Augmentation’ means the technology is used to boost worker productivity to produce more or better output, while ‘displacement’ means it is used to lay off workers with less boost to output. We present three scenarios to illustrate this. **In our ‘full augmentation’ scenario we show that there can be zero job displacement.** If generative AI was widely integrated across the economy, we estimate it could provide an economic boost of 13 per cent of GDP. At the other extreme, in our ‘full displacement’ scenario, 8 million jobs could be lost with no GDP gains. In between those two scenarios falls **our central scenario where 4.4 million jobs disappear, but still with significant economic gains of about 6.4 per cent of GDP.** This shows that there is no one predetermined path for how AI implementation will play out.

We argue that the positive scenarios, with widespread gains to AI deployment, can only be realised through a wide range of targeted policy interventions, which we call a **job-centric industrial strategy for AI.** This includes a three pronged approach:

1. protect existing jobs and ensure gains for workers
2. boost creation of new tasks, jobs and support job transitions
3. address the fallout from lower labour demand (figure S.2).

A new centralised institution might be needed to coordinate these, helping ensure policies complement each other.

One policy worth considering is **‘ringfencing’ tasks from full automation,** thereby requiring a continued degree of human involvement. **This could be, for instance, ensuring that medical diagnoses are being overseen and delivered to patients by humans** or that much of early years education retains the ‘human touch’ (see chapter 5). This would be done through a combination of bottom-up classification of tasks and government policy incentives. It would require a significant level of cooperation across businesses, unions and government. Other policies include aligning regulatory and tax incentives so that they benefit job-augmentation over full displacement. Forward looking competition policy too will have to play a role in order to ensure a broad based job-centric adoption of AI. In chapter 3 we show quantitatively that **in some sectors work time reduction might be an important option to consider too.**

FIGURE S.2: POLICY PILLARS FOR A JOB-CENTRIC INDUSTRIAL STRATEGY FOR AI

Objective	Policy	Time to have an impact
1 Protect existing jobs and ensure gains for workers	Actions to augment jobs • Provide incentives ensuring that productivity gains from AI deployments benefit employment Medium term	Medium term
	Ringfence tasks for human involvement • Eg prevent process engineering aimed at job displacement Medium term	Medium term
	Change fiscal incentives • Change current tax advantages for automation, consider 'augmentation subsidies' Short term	Short term
2 Boost creation of new tasks, jobs and support job transitions	Incentivising creation of jobs with low exposure tasks • Subsidies or tax incentives for ringfenced and low-exposure tasks Medium term	Medium term
	Retraining and upskilling • Where jobs transitions constitute significant career changes, provide guidance and support Long term	Long term
	Competition policy • Use competition policies to support smaller firms with AI adoption Medium term	Medium term
3 Address the fallout from lower labour demand	Worktime reduction • Fiscal incentives or strengthen labour legislation to encourage working time reduction Short term	Short term
	Tax the gains from automation • Tax the gains from automation, eg via specific 'AI taxes' or wealth taxes Medium term	Medium term
	Social security reform, UBI • Expand social safety net to manage impact on incomes Medium term	Medium term

Source: Authors' analysis

Finally, our scenarios show that significant economic gains can be had from automation. But as even OpenAI CEO Sam Altman (2021) highlighted, there is no guarantee that these will be equally shared, leading him to call for a wealth tax paired with social security reform. In practice, the **tax response will likely have to be nuanced, balancing the incentive to innovate with the goal of internalising the unintended costs** from technology deployment. In line with taxation, social security assistance might also have to evolve to help attenuate any disruptive effects that might occur from fast or unexpectedly large labour market impacts of generative AI.

Our scenarios show that the potential range of impacts of generative AI (good and bad) is vast and that the way they play out hugely depends on policy choices. Given the speed and breadth of generative AI adoption that we show is possible across sectors, and given the long lags with which policies take effect, **we think it is urgent that policymakers start preparing a policy response now**. Our framework outlines the range of options. Fleshing these out in detail, in response to plausible near term scenarios, will be the key for reducing the risk of labour market disruption.

1. INTRODUCTION: THE MOMENT WE'RE IN

Spurred by breakthrough applications such as ChatGPT, and by the UK government's AI Safety Summit, the attention to artificial intelligence (AI) hugely increased over the course of 2023. It has revived previous debates around automation and its impact on the labour market. We have been here before.

1.1. THE DEBATE AROUND AUTOMATION AND THE LABOUR MARKET: WILL THIS TIME BE DIFFERENT?

There exists an extensive literature on the relationship between AI (or automation more broadly) and the labour market. It falls broadly into two strands.

The first strand argues that the labour market disruption from automation in the next years is likely going to be limited, and they urge the need to be cautious about 'futurology' predicting the demise of jobs. We call these the 'sceptics'. For instance, Bell (2024) points out that there have been many predictions about imminent technological unemployment over the last 14 years and, so far, these have not come to pass. Moreover, he points out that - if anything - many people work more, not less, hours than in the recent past (especially with higher earners working long hours). Labour markets in most advanced economies are tight by historic standards (with lots of unfilled vacancies) and he highlights the steadily increasing number of women entering the workforce over the last decades. Thus, there does not seem to be a shortage of demand for work.

We strongly agree with these points: so far, AI has not had significant effects on employment. That said, academic studies show that automation technologies (including software) had an effect on *wages*. For instance, Acemoglu and Restrepo (2022) find that that between 50 and 70 per cent of changes in the US wage structure in 1980-2016 can be attributed to wage impacts of automation technologies.

A second crucial point made by this strand of thinking is pointing to past technological transformations: they argue that, over the last half century, there were significant technological shifts and they have not led to a large increase in unemployment. For instance, Autor et al (2022) trace the creation of new jobs and tasks over time. They find that around 60 per cent of jobs in the US today did not exist in 1940. As gains from technology increased wages and wealth, demand for other goods and services grew, driving the creation of new jobs. They highlight that new jobs such as solar cell installers are being created all the time, and people are transitioning away from those that are made obsolete by technology.

From these type of historical findings, Autor (ibid) concludes that this time will not be different. He argues "AI is a tool, like a calculator or a chainsaw, and tools generally aren't substitutes for expertise but rather levers for its application." They can augment human expertise, he argues, and "it will further instantiate new human capabilities, new goods and services that create demand for expertise we have yet to foresee."

The other strand in this debate argues that – in the next decade or so – serious labour market disruptions from automation could take place. We call these the ‘warners’.

This strand argues that the technological change we are seeing unfold is different from what we have seen in the past. We are on the cusp of seeing a new type of AI that is more capable and widely applied than previous iterations. Through huge improvements in existing technical approaches, ‘generative AI’ can now produce original content that can be indistinguishable from that written by humans, solve complex problems, build websites, understand humour, produce highly detailed analyses of images and sound, and execute multi-step tasks online on behalf of users (Bubeck et al 2023). OpenAI’s GPT4 has scored better than nine out of 10 law students in the all-important Bar exam (Eloundou et al 2023)⁴. Past trends are therefore no guide for how the future will play out.

A growing literature is analysing how this could play out. McKinsey (2023) finds that “30 per cent of worked hours in the US could be automated” by generative AI. Elounou et al (2023) find that 19 per cent of the workforce have at least half of their tasks “impacted” by generative AI. IMF staff find that “in advanced economies, about 60 percent of jobs are exposed to AI, due to prevalence of cognitive-task-oriented jobs” – though not all of these are at risk of automation (Cazzaniga et al 2024). Korinek (2023) discusses how in the last five years AI technologies have become able to complete complex tasks that were thought to be squarely only doable by humans – such as reading comprehension and abstract reasoning, and that leading AI researchers continue to expect this frontier to shift further. Korinek thus highlights the need for *scenario planning* for the plausible (even if low probability) case that large numbers of jobs are significantly affected.

From a macroeconomic perspective, there are studies scrutinising how automation could impact the macroeconomy. Goldman Sachs (2023) estimates that widespread adoption of generative AI could increase the level of global GDP by 7 per cent compared to a counterfactual – a significant boost. But, through a macroeconomic model, Benzell et al (2018) show that the overall impact of automation also depends on how much inequality it creates. If too many people lose their job due to automation and distribution does not compensate the effect, this will drag down overall consumption and thus GDP. This is also discussed by Acemoglu and Restrepo (2021).

Our central take on these two debates is that, with the advent of generative AI, the game has changed. Rather than having to reason about possible future technical capabilities, a technology now exists that has now been proven to produce high quality outputs that are often indistinguishable from human ones, in a fraction of the time that a human would take, across a wide range of applications. It can hardly be overstated what an astonishing achievement this is. It was until recently thought to be feasible only in the distant future. And given many knowledge work processes are already digitalised, in many cases it does not require huge process changes or capital investments to introduce AI to these processes.

The speed of adoption is likely going to be faster than during past technological waves. Moreover, we see very rapid adoption of these technologies. ChatGPT reached 100 million users worldwide within only two months of its release (Reuters 2023). Steam engines took about 120 years to be adopted. Electricity took about 60 years to spread (Frey 2019). Conversely, generative AI, similar to ChatGPT might transform the realm of knowledge work within the matter of a few years – as many knowledge economy jobs are already digitised and the extra investment needed

⁴ Though some argue that, taking into account the type of training data used and test performed, the ranking should have been lower, but still higher than half of humans.

to use generative AI for them is relatively limited. Generative AI is already being widely deployed in the legal industry, journalism and content creation, finance and increasingly health care.⁵

At the same time, **we also agree with the ‘sceptics’ in that we do not expect there to be an immediate ‘job apocalypse’.** We highlight that generative AI deployment will likely take place in phases. The first one being experimentation and the second one having an overall small – but in some occupations significant – effect on employment. Only once generative AI gets more deeply integrated into organisations could a large number of jobs be transformed. Even so, we also highlight that design choices make a huge impact whether jobs are lost or not. And new tasks and new jobs can be created offering new opportunities for income, growth and social benefits that work brings. Exposure to AI does not mean certain automation, but this does not mean we can be complacent about its transformative potential.

1.2. THE POLICY DEBATE AROUND AUTOMATION NEEDS A STEP CHANGE

Finally, we think in both strands of the debate policy suggestions are often too narrow and need further fleshing out. This needs to urgently be addressed given the pace of generative AI deployment points towards a move to phase 1 and soon perhaps even phase 2 development in the coming years. Policy is not prepared for this.

Most commonly recommended policy suggestions often focus narrowly on retraining and work time reduction; which have a role to play but are unlikely to be sufficient on their own. Others assume that work will soon be automated and thus highlight universal basic income (UBI) as an immediate priority.

We stress that a much broader toolbox is needed in order to rise to the scale of the challenge. In particular we highlight the opportunity for policy to incentivise creation of new tasks and jobs that are less at risk of displacement. In short, we need a job-centric industrial strategy for AI.

5 See for instance Lexinexis (2023) and MDPI (2023).

2. WHY THIS TIME IS DIFFERENT – GENERATIVE AI IS HERE

In this section we introduce some key terms around generative AI and some of its capabilities. Readers who are familiar with these can consider skipping to chapter 3.

2.1. WHAT CAN AND CAN'T EXISTING GENERATIVE AI DO?

Since its release in February 2023 there have been many assessments of the capabilities of generative AI models. Notable findings are for instance:

- **Legal knowledge and reasoning.** GPT4 scored better than about 90 per cent than humans in the US Bar exam (Eloundou et al 2023). Martin et al (2024) found that Large Language Models (LLMs) such as GPT4 performed “on par with LPOs and Junior Lawyers, accurately determining legal issues within contracts”, while “a 99.6 per cent reduction in cost”. GPT4 also scored highly on standardised tests such as the SAT, GRE, and various AP exams (OpenAI 2023). Notably, these tests aim to check the ability to apply complex reasoning to *novel* tasks rather than simply checking whether models are able to ‘regurgitate’ answers from their training data. While there is still some controversy about this⁶, there is little doubt that generative AI models have significant reasoning abilities in certain tasks and contexts.
- **Medical knowledge and reasoning.** A comprehensive evaluation highlighted GPT-4’s proficiency in medical competency exams, where it exceeded the passing score on the United States Medical Licensing Examination (USMLE), without being explicitly trained on the exam (Nori et al 2023).
- **Assessing images and making inferences.** Another study assessed the ‘multimodal’ capabilities of generative AI, which allows analysing images. It was used for interpreting radiological images across various modalities, anatomical regions, and pathologies (Brin et al 2023). The findings revealed that while GPT-4V could correctly recognise imaging modalities in all cases, its performance in identifying pathologies and anatomical regions was inconsistent.⁷ This highlights that while it can play a role in assisting medics in certain cases, it will still need expert supervision in others.

Generative AI is constantly being improved. One focus of current developments is efforts to give models more abilities to access data and execute tasks. We call this ‘**integrated**’ generative AI. For instance, this could be accessing databases, executing actions online (such as making bookings), or liaising with other computer systems to achieve multi step actions. To some extent this is already being implemented (e.g. GPT4 can search the internet, and plug ins can assist in make travel bookings). Note that the access to existing systems can make existing generative AI systems significantly more powerful, even without the core models becoming any better than they are today.

6 Though some argue that many benchmarks do not successfully achieve this and thus leave the door open for ‘regurgitation’.

7 The overall accuracy for anatomical region identification was 69.2 per cent, but it varied significantly across different types of imaging.

There are a number of recent examples for this.

- **Multi step aid with data analysis and programming.** Agent-like AI tools are able to analyse data, modify it, create charts and draw inferences in multi step processes.
- **Multi step customer engagement.** There are various examples for where increasingly complex customer queries are being handled by agent-like, 'integrated' AIs. Klarna – a major fintech company specialising in online payment solutions – in February 2024 disclosed it had embedded GPT4 in customer management, with remarkable results. It said it was able to complete two-thirds of customer queries, “on par with humans on customer satisfaction” and with “higher accuracy” (Siemiatkowski 2024).
- **Large scale information processing.** Google’s just released Gemini 1.5 can now ingest large amounts of data in requests (equivalent in length to 20 novels). This allows it to provide answers instantaneously drawing on a huge amount of information that it was not trained on. Other developers are said to also be working on extending their context windows.

As we show below, these capabilities could bring large scale productivity benefits. Perhaps even more excitingly, they could allow us to do things that were hitherto thought impossible. In science, for instance, entirely new ways of doing drug discovery are already being developed (The Economist 2023b). Though as we highlight in Jung and Srinivasa Desikan (2023), many of the benefits will not automatically occur, or not to the extent that would be socially optimal, unless policy incentivises them.

2.2. ASSESSING EXPOSURE TO AI

To see which tasks and jobs will be affected by AI, we produce a metric that indicates how many tasks could be transformed by AI. We score each task with regards to whether a **human could perform it 50 per cent more quickly with the help of AI** (in line with some of the recent literature on this, eg Eloundou et al (2023) (see box 1). Note that this definition is not the same as full automation – often humans will be centrally involved in the tasks. This also makes many tasks more exposed, as involving generative AI does not require full blown digitisation. If we were to consider full automatability only, the degree of exposure would be significantly lower. However, as we show below, merely **speeding up the completion of tasks could nonetheless have large scale labour market implications**, as relative demand for human labour in certain jobs shifts. This definition also points to the potential opportunities for productivity increases and distributional implications. We did this both for two types of the phases of generative AI deployment:

- One is the **'here and now AI' exposure** – this is if existing generative AI such as GPT4 can already readily do the tasks involved.
- The other one is **'integrated AI' exposure**. This is the idea that generative AI is connected to other software systems, including databases and the ability to execute tasks (such as making bookings or orders).

In the coming sections we summarise our findings, construct scenarios and draw policy implications.

BOX 1: METHODOLOGY FOR ASSESSING EXPOSURE TO GENERATIVE AI

We calculated exposure to generative AI based on the tasks making up individual jobs. We identified tasks associated with a profession using a crosswalk of ONET with ONS LFS data. Following Eleonou et al (2023), we used the OpenAI API to pass each task and profession to the GPT4 model, querying it for generative AI capacity. We used the prompt “can the task ‘{task}’ by ‘{occupation}’ be done at least 50 per cent faster by a human using GPT4 or similar models?”. We also fed in a 'context prompt' that ensures there is clarity about generative AI's capabilities of text generation, data analysis, and usage of plug-ins. This resulted in all of the 22,000 tasks tagged with a “yes” or “no” based on GPT4's estimated ability to perform the task. Then, based on the number of exposed tasks, we calculate a (weighted) exposure metric for each occupation. This approach follows the widely cited study by Eloundou et al (2023). Our method differs from that employed by Felten et al (2023), or Frey and Osborne (2013), in that it calculates exposure by directly querying the AI model rather than rely on proxy metrics or a probabilistic approach.

We cross checked a sample of the model's assessment against our own scoring. We find an agreement of 87 per cent between the model and our own judgement for the 'here and now' AI and an agreement of 76 per cent for 'integrated' AI.

Overall, we found that while there was some uncertainty around how much speed or quality improvement generative AI could bring, it was often clear whether or not tasks could be significantly aided by generative AI. This points to our broader point that it is becoming increasingly clear where generative AI could have a significant impact, but the degree of it will depend hugely on design, adoption and policy choices. For more detail on this see the appendix.

3. WHO IS MOST EXPOSED TO GENERATIVE AI?

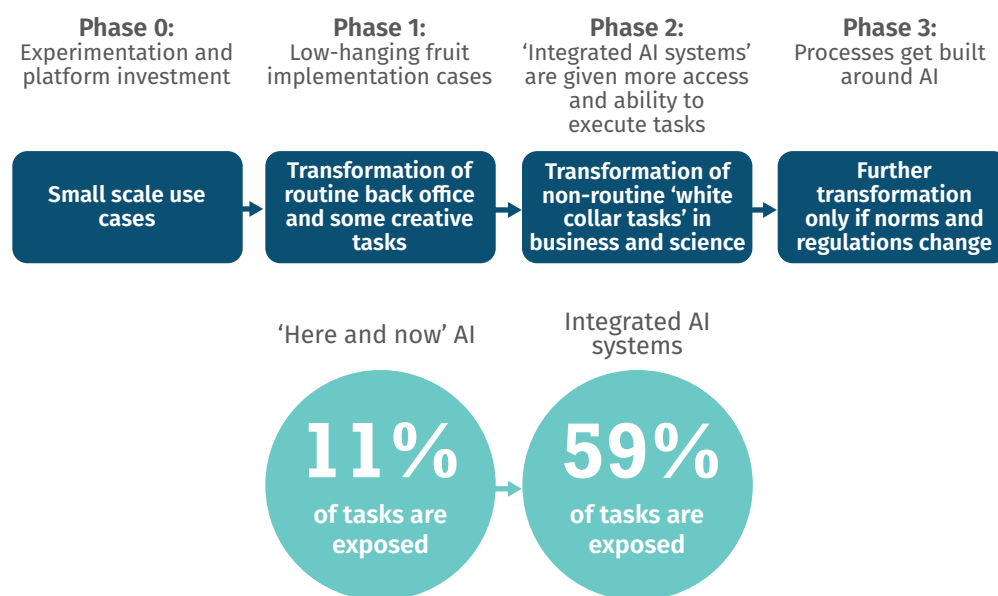
3.1. COGNITIVE TASKS ARE MOST EXPOSED TO GENERATIVE AI

The next phase of AI could see 11 per cent of tasks exposed

We find that **11 per cent of tasks** in the UK (scaled by hours worked) are exposed to 'here and now' generative AI – systems that are already widely available.

And, looking ahead to the near future, we find that about five times more jobs – **about 59 per cent of hours worked in the UK – are exposed when considering more integrated AI systems**, defined as those that have the ability to access proprietary data and the ability to execute tasks. **This is a significant impact**, similar in scale to how digitalisation has transformed most knowledge work since the 1990s. This size of exposure is similar to that found in a recent IMF study on the same topic (Cazzaniga et al 2024).⁸

FIGURE 3.1 A MUCH HIGHER SHARE OF TASKS IN THE UK LABOUR MARKET COULD BE DONE WITH GENERATIVE AI IF SOCIAL NORMS AND PROCESS WERE TO CHANGE



Source: Authors' analysis

None of these phases will occur naturally. But the capabilities – and commercial benefits – of generative will likely be a strong pull. Goldman Sachs suggests that only 5 per cent of chief executives expect AI to have a 'significant impact' on their business within one to two years, but 65 per cent think it will have an impact in the next three to five years (The Economist 2023b).

⁸ Though we use a different methodology from the IMF study.

We distinguish between three phases.

- Phase 1 implementation of 'low hanging fruit' cases: Tasks that can be done more quickly through generative AI, **without much process re-engineering** (eg extracting information from a database or a text which are already digitised could be done merely using a generative AI application instead of a human).
- Phase 2 refers to integrating AI systems with organisational processes that can be done through generative AI, but **require some process re-engineering**. For instance, an integrated generative AI could order supplies for a restaurant kitchen, but it would need employees to log stocks in an app. Or it could aid teachers with the grading of students' work, but it would require these to be uploaded (eg via pictures taken and submitted to an app) or submitted directly by students.
- Phase 3 refers to processes being rebuilt to enable the use of AI. This *includes* tasks that **could only be done if social norms** changed or that **would require significant regulatory change** (eg an AI system advising patients on treatment plans, or making personnel decisions at work).

This highlights potentially the most crucial point in our analysis: that we can actively change the way in which generative AI impacts the labour market, through making design choices. **We can ringfence certain types of tasks**, such as teaching, or policing, or therapy, or medical advice, to ensure that tasks retain the 'human touch' even if in theory they could be done via a computer. As we show below, a scenario where relatively more employment is maintained can result in more economic output, compared to a scenario where as much work as possible is automated.

There will also likely be two-speed adoption of the technology. Large firms were more likely to adopt AI, with an adoption rate of 68 per cent in 2020 – nearly twice the rate of medium sized companies (UK DCMS 2022). This could mean that larger companies gain a significant advantage over smaller ones.

In our assessment presented, unless otherwise stated, we only look at exposure that can be done in phase 2 and phase 3 described above, ie those without significant norm change, but perhaps with some adaptation. In the appendix we walk through some examples that highlight the principles behind which tasks are judged for their exposure to generative AI.

Cognitive tasks are most exposed to generative AI

There is a large diversity of tasks in the economy. To simplify this we break down all tasks performed in the UK economy into five intuitive groups, shown in figure 3.2. It shows for example that interpersonal and communications tasks (such as customer interactions in retail sales) make up the largest amount of tasks performed in the UK economy, constituting about 32 percent of overall hours of work in aggregate. Together, manual and social tasks account for about half of the hours worked in the UK economy.

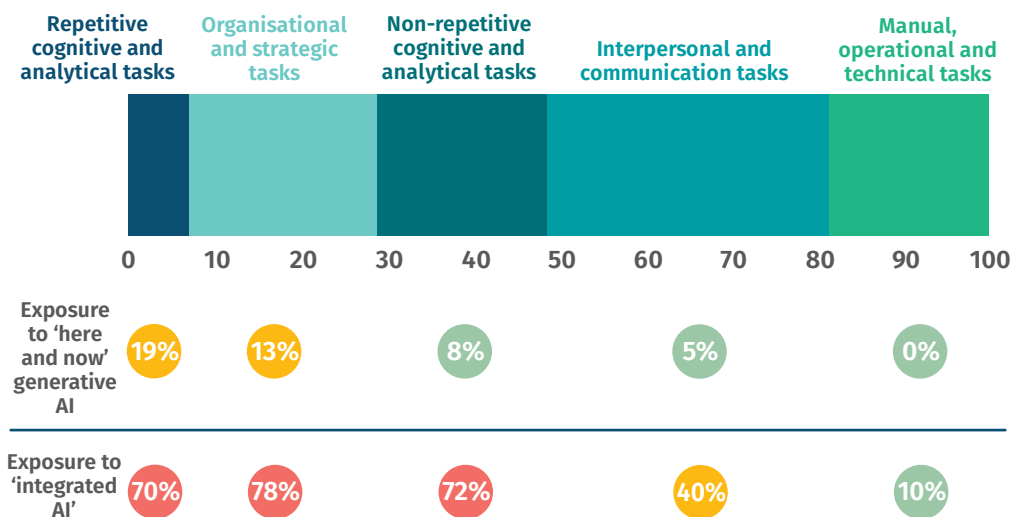
This matters for how AI will impact the economy. Below the bars in figure 3.2 we show how exposed to generative AI each task type is. It shows that **repetitive cognitive tasks and organisational and strategic tasks are the most exposed to 'here and now' AI**.⁹ But strikingly in the 'integrated' AI scenario, also non-repetitive cognitive and analytical tasks are highly exposed.

9 Organisational and strategic tasks encompass a broad range of activities aimed at managing resources, overseeing operations, and planning strategically to achieve business objectives. These tasks involve detailed planning, such as scheduling equipment inspections and managing inventories, to ensure operational efficiency and resource availability. Examples include developing operational standards and procedures for work units or departments, and managing finances like making bank deposits and performing bookkeeping duties.

It is worth noting that much of the literature on automation and jobs highlighted that it was usually routine jobs that were exposed to automation pressures since the 1970s (Autor 2022).¹⁰ Generative AI could clearly change this and thus have much more wide-ranging impacts on tasks across the knowledge economy.

FIGURE 3.2: ABOUT HALF OF THE TASKS IN THE UK ECONOMY HAVE LOW EXPOSURE TO EXISTING GENERATIVE AI TECHNOLOGIES BUT HIGHER EXPOSURES TO 'INTEGRATED' AI

Tasks as share of hours worked, and degree of exposure of these tasks go generative AI



Note: exposure means share of tasks in this category that AI could help conduct at least 50 per cent faster. High exposure is marked as red circles (defined as higher than 40 per cent of tasks), medium exposure is shown as yellow (between 11 and 40 per cent), and low exposure is shown as green, with 10 per cent or lower exposure.

Source: IPPR analysis of ONET (2023) and LFS (2023)

3.2. BACK OFFICE JOBS ARE MOST EXPOSED TO 'HERE AND NOW' AI

In this section we move from *tasks* to looking at jobs. Based on our bottom up analysis, we find that 6 per cent of jobs are automatable with 'here and now' AI. And two thirds of jobs are highly exposed to 'integrated' AI. In the former, the automated tasks are fairly concentrated in a relatively small number of occupations. In the latter, for integrated AI, the exposed jobs are more widely dispersed.

Figure 3.3 shows that occupational groups most exposed to 'here and now' AI are mostly **back office occupations**. The top five professions with the highest exposure are personal assistants and other secretaries (69 per cent exposure) human resources administrative occupations (68 per cent exposure), other researchers (65 per cent), marketing associate professionals (65 per cent), Authors, writers and translators (65 per cent). And we find that **women will be significantly more affected** (as they are more likely to hold the most exposed occupations, such as secretarial and administrative occupations) (see appendix).


Moreover, we find that for a given job, **entry level positions are more at risk than those of more experienced professionals**. For instance, administrative professionals

¹⁰ Autor (2022) states that “non-routine cognitive abstract-reasoning (expert judgment, creativity) and interpersonal (leadership, management) tasks have proven hard to automate because, simply put, we don't know ‘the rules.’” However one of the striking thing about generative AI has been that it was able to automate these tasks without every being explicitly programmed ‘the rules’, merely by ‘trained on’ gigantic datasets (Bubeck et al 2023).

with less training and experience have 14 percentage points higher risk of automatability than those with a higher level of experience.¹¹ Across all jobs, those with lower experience, on average, have 16 percentage points higher exposure.

FIGURE 3.3: 45 PER CENT OF TASKS IN THE 'SECRETARIAL AND RELATED OCCUPATIONS' ARE EXPOSED TO 'HERE AND NOW' GENERATIVE AI

Share of exposed tasks by occupation (25 percentile, median and 75th percentile)




Source: IPPR analysis of ONET (2023) and ONS (2023).

Figure 3.4 show the job exposures with regards to more integrated AI. Back office jobs are still at the top of the list with even higher exposure still. But other knowledge economy jobs are now also highly exposed, including those requiring more technical and scientific knowledge.

11 We measure experience via ONET's job zone qualifications. The statistics in this paragraph refer to a one unit difference in job zones.

FIGURE 3.4: IN PHASE 2, WHERE AI BECOMES MORE INTEGRATED IN WORKFLOWS, MORE THAN HALF OF OCCUPATIONS ARE HIGHLY EXPOSED



Source: IPPR analysis of ONET (2023) and ONS (2023).

There are already early signs that the occupations most at risk are seeing relative demand falling. We discuss some recent examples in box 2, from across the world.

BOX 2: EARLY EXAMPLES OF JOBS AFFECTED BY GENERATIVE AI IN THE US LABOUR MARKET

Recent employment trends in the labour market provide early supportive evidence for our derived ranking of exposure to generative AI, where we find that initial benefits and risks fall on back office jobs. Below we collect some examples from the US labour market, where employment decisions were directly related to generative AI developments.

For instance, some high profile AI-related job announcements in jobs in mid-2023. These were with regards to **human resource (HR) jobs**, which rank among the highest in our exposure metric. IBM’s decision to halt hiring for back-office jobs that AI could potentially replace is an early example of this (Bloomberg 2023). Moreover, the cases also suggest that where AI developments occur, companies tend to respond by reducing their hiring rather than outright laying off people. This is suggestive evidence toward our findings around entry level positions being relatively more at risk. The National Eating Disorders Association (NEDA) recently **replaced their entire human helpline workers with chatbots**. The move raised

questions about the balance between technological efficiency and the human touch, especially in sensitive sectors like mental health (NPR 2023).

The creative industry too has seen initial impacts of generative AI. A **film and TV industry** executive, in February 2024, halted a significant studio expansion after assessing the job displacement potential of AI, particularly influenced by OpenAI's Sora video generator (Guardian 2024). Early evidence suggests how copywriters, freelancers are feeling pressures due to advanced writing capabilities by generative AI (Washington Post 2023).

According to a consultancy report that tracks layoffs, AI contributed to 5 per cent of all layoffs across all sectors in the US in May 2023 (Challenger 2023). Note that this only consists of cases where the organisation declared it specifically as an AI related layoff, as opposed to other cases of organisational restructuring. Given the potential for bad press of announcing layoffs on account of AI, this 5 estimate could be a lower bound of all AI related layoffs.


There is also some early evidence of generative AI's potential to augment work. A study conducted by Harvard Business School in collaboration with Boston Consulting Group found that consultancy workers using generative AI became more efficient. Consultants using generative AI, specifically the large language model GPT-4, completed 12 per cent more tasks on average and **finished tasks 25 per cent faster** compared to those without AI access. Responses produced by consultants with access to generative AI were of higher quality, with a **40 per cent increase in quality** compared to those without AI support (BCG 2023).

'Here and now' AI will impact both medium and low wage jobs

Our findings indicate that both medium and some lower earning occupations are exposed to 'here and now' generative AI, with very few high paying jobs being exposed. This finding differs from the conclusions of Eloundou et al (2023) and Felten et al (2023) who find that higher earning jobs are relatively more at risk.

FIGURE 3.5: MEDIUM AND LOW EARNERS ARE EXPOSED TO CURRENT GENERATIVE AI

Share of tasks exposed to 'here and now' AI by occupation vs median hourly wage



Source: IPPR analysis of ONET (2023) and ONS (2023)

Yet, we argue that this is mainly related to the phase of deployment. In our 'integrated AI' scenario we observe a more pronounced increase in exposure for higher earners, bringing our results closer to those previously reported studies. Examples of jobs more exposed in that scenario are barristers, educational professionals as well as financial managers and directors (see figure A.1 in appendix).

4. SCENARIOS FOR HOW GENERATIVE AI COULD AFFECT WORK IN THE FUTURE

4.1. WHETHER AI WILL BOOST WAGES AND GROWTH WILL DEPEND ON HOW IT IS IMPLEMENTED

In order to think through policy implications of the arrival of generative AI, we play through some scenarios. We mostly focus on a central scenario that sees a combination of generative AI used for job augmentation and job displacement. But we also present a 'full augmentation' and a 'full displacement' scenario to illustrate underlying dynamics. In all scenarios, we assume generative AI is used wherever it can be used. In reality the speed and degree of adoption might vary significantly between, say, larger and smaller firms. Crucially, these are not predictions, but rather tools for thinking through how things could plausibly play out given a number of assumptions.


With 'here and now' AI, overall job disruption would be limited, but significant in some occupational groups

Figure 4.1 shows that 'here and now' would have a limited impact overall, but could have **quite disruptive impact in some occupational groups**.¹² For instance, in administrative occupations, about a third of all jobs could be displaced in our central scenario.

12 To arrive at this, we assume that those jobs that get automated have no increase in output and rather the same level of output is produced by fewer workers. For augmented jobs, we assume that all workers stay in place but that they use the AI technology in order to produce more output or (equivalently) output that has more value in proportion to productivity gains.

FIGURE 4.1: IN ADMINISTRATIVE OCCUPATIONS, ABOUT ONE=THIRD OF JOBS COULD BE DISPLACED

Number of jobs and exposure to augmentation or displacement




Source: IPPR analysis of LFS (2023) and ONET (2023)

At the same time, in our central scenario, **this loss of jobs would be paired with significant productivity gains**. If these were to accrue to employees, it could lead to large wage gains. Figure 4.2 shows illustrations of what this could mean. Secretarial and related occupations would see, on average, **productivity gains of 35 per cent**; customer service occupations would see productivity increases of 32 per cent and administrative occupations would see increases of 27 per cent. This degree of gain is not implausible in the light of early scientific studies on the productivity impact of generative AI. It is crucial that technology adoption is done in a way so that the fruits of productivity gains are enjoyed by workers. The TUC (2023) has outlined how this could be done. For example necessary consulting of workers during the introduction of AI tools and systems would allow for negotiation of wages to match productivity gains.

FIGURE 4.2: SECRETARIAL AND RELATED OCCUPATIONS COULD SEE THE LARGEST PRODUCTIVITY INCREASES IN THE CENTRAL SCENARIO.

Productivity gains through automation by occupation (%)



Note: these (2-digit) occupation level numbers are (4 digit) employment weighted, and take into account both direct productivity changes but also the change in employment in displaced occupations. The averages thus also reflect also compositional changes. Our methodology assumes that each task within a profession contributes proportionately to the value add of a job, which we, in turn, assume can be approximated via its wage level.

Source: IPPR simulations based on ONET (2023) and ONS (2023).

'Here and now' AI could lead to modest aggregate growth and employment impacts, but 'integrated' AI could be a game changer

We next consider what the aggregate implications of the above would be and also highlight how different types of scenarios could play out.

Table 4.1 highlights that there are stark differences in possible outcomes. In the augmentation and displacement scenario, significant productivity gains are possible but a significant change in employment occurs nonetheless.

The implications of different scenarios could be as follows.

- **The augmentation only scenario** highlights the possibility that workers with access to generative AI become more productive, being able to make either more or higher quality (and thus more valuable) outputs. All workers remain in place, producing more and directly boosting GDP. For instance, a marketing copy writer might produce 30 per cent more content. And a personal assistant might be able to support more people in the same amount of time. However,

these examples suggest that there might be limits to increased demand as well as to increased quality.¹³

- **The ‘displacement only’ scenario** is notable in that it illustrates that we could have an economy hugely impacted by generative AI, but without any *aggregate* output increase. In this scenario, exactly the same amount of output is produced, but fewer employees are required to achieve that – with those remaining made more productive by AI. So fewer people produce the same amount of output as before, with the displaced people finding no new work. This scenario is extreme, merely for illustrative purposes. It shows that large scale deployment of AI need not go hand in hand with improved aggregate performance. Displaced people could of course also move to new, other jobs. We discuss this in detail in the next two sections.
- **In scenarios with displacement**, those *staying* in work or the company owners are the biggest winners. As AI helps to produce more output with fewer workers, so the ones that remain become more productive – they are the ones who ‘reap’ the productivity gains. Of course, all the productivity gains might feed into to company profits and not in fact increase wages at all. Those who lose out in these scenarios are the people whose jobs are displaced, if there are no new jobs emerging.

TABLE 4.1: EMPLOYMENT AND PRODUCTIVITY SCENARIOS

	'Here and now' generative AI		'Integrated' AI	
	Change in employment due to displacement	Change in aggregate output due to augmentation	Change in employment due to displacement	Change in aggregate output due to augmentation
Augmentation only scenario	0	£92bn (4 per cent of GDP)	0	£305bn (13.4 per cent of GDP)
Displacement only scenario	-1,500,000	£ 0	-7.9 million	£ 0
Central scenario: Augmentation and displacement	-545,000	£64bn (3.1 per cent of GDP)	-4.4 million	£144bn (6.3 per cent of GDP)

Note: The augmentation and displacement scenario assumes that jobs with more than 40 per cent exposure are displaced, while those with less than 40 per cent exposure are augmented. **Central scenario:** assumes a mix of higher demand for goods produced by augmented jobs and constant demand for goods/services produced by displaced job. **Augmentation only scenario:** Assumes increase in demand (for quantity of quality) for goods produced by augmented occupations. **Displacement only scenario:** assumes no aggregate change in demand for goods and services. The change in aggregate output is calculated as follows. Given our exposure metric indicates whether a task can be done 50 per cent faster or more, we assume that this is exactly what happens. The share of a job’s tasks is exposed to generative AI is assumed to be done 50 per cent faster and that this increases productivity in proportion to the job’s wage level. The productivity of non-exposed tasks remains unchanged.

Source: IPPR analysis of ONET (2023) and ONS (2023)

13 For instance, there might be an upper limit for how many copywriting services are needed, so increased output (boosted by AI) might not necessarily be met by increased demand for marketing services. Benanov (2020) and Korineck (2023) make this point, highlighting that if demand for products stays consistent across most sectors, labour demand will drop.

What this shows is that there are different pathways for technology adoption. The augmentation scenario links with high productivity gains and low displacement of jobs. While on the other hand a full automation scenario would not increase overall output as the remaining workers would be more productive, but they would merely produce the same amount of output with fewer workers.

4.2. THERE MIGHT NOT BE ENOUGH ‘SIMILAR’ JOBS AVAILABLE FOR DISPLACED WORKERS

Many commentators highlight that historic episodes of economic upheaval through technological disruption have always led to new tasks and jobs being generated for displaced workers (Autor 2022). For instance, as technological change brings down prices of goods, people effectively have more income to spend on other things which increases demand for workers in the same jobs or elsewhere, including new hitherto non-existent tasks (Bessel 2018). This ‘income effect’ explains why unemployment today is low despite repeated waves of automation. A recent example of this is the digital and computing revolution. It put traditional typewriters out of jobs but created entirely new industries, fuelled by demand for digital tools.

But for three reasons, this should not give too much cause for comfort:

- First, historic episodes of automation took more time – multiple decades – to unfold than the changes we are currently seeing. New generations would do different jobs from their parents, rather than individuals retraining and going into new jobs (Broadberry et al 2015, Shaw-Taylor et al 2019). We simply have not seen large scale transitions at the scale and speed that are likely in train right now.
- Secondly, historically, there are important episodes when technological change did indeed cause large scale labour market dislocation. For example, Shaw-Taylor et al (2019) and Humphries and Schneider (2021) have argued that starting in the 18th century, labour market participation of women plummeted as a result of the advent of mechanisation of textile work. Women had been heavily involved in textile production when it took place in people’s homes. But with the advent of the mechanisation, production moved into factories which reduced the overall number of jobs in textile which were, due to social norms, less accessible to women (ibid).
- Third, job transitions – especially when they happen quickly – are of course dependent on people’s existing skills levels and the possibility for retraining. A large number of studies have explored how mismatches between skills and job demands can act as barriers to successful career transitions, affecting individuals’ career trajectories and overall well-being and that this is a particular problem in the UK (Sullivan and Al Ariss 2021; CIPD 2021). As we detail further on, transitioning is easier when there is a skill overlap between the AI exposed and AI resilient occupations.


To evaluate a scenario for this, **we analyse potential labour market transitions for the automated jobs in our central scenario.** We do this by developing a ‘job similarity metric’, which predicts to which jobs people could transition that are less at risk of automation. The similarity metric is built on two factors: how similar the types of tasks in two jobs are and how ‘senior’ the job is, such as to only suggest job transitions that are broadly at the same level of career progression. See box 3 for our methodology.

Based on this similarity metric, we can calculate how job transitions from one occupation to another. The result is an influx of workers from automated occupations to other non-automated ones. **We find that this influx into some non-displaced occupation would be large.**

Figure 4.3 shows how many employees could transition as a share of existing employees in each occupation. As would be expected, many job transitions based on job similarity are from back office jobs to other back office jobs that are less automatable.¹⁴ Some occupations – such as elementary administration occupations – **would receive inflows of more than 75 per cent of their existing workforce**. Some hospitality jobs receive a 50 per cent inflow. This could involve people moving and reskilling.

In our central scenario we assume that these ‘receiving’ jobs would not see an increase in demand. That is, the same amount of back office workers would be required across the economy. The implications of this could be twofold: First, it will exert **downward wage pressure in the jobs receiving an influx of displaced workers** and second, if the demand for these jobs stays unchanged, then new employees could be accommodated through work sharing and worktime reduction.

FIGURE 4.3: ELEMENTARY ADMINISTRATION OCCUPATIONS WOULD RECEIVE AN INFLOW OF OVER 75 PER CENT OF THEIR EXISTING LABOUR FORCE, HOSPITALITY JOBS ABOUT 50 PER CENT
Influx of displaced workers into non-displaced jobs as a percentage of existing workforce



Source: IPPR analysis of ONET (2023) and ONS (2023)

¹⁴ For instance, the most frequent job transitions are from ‘Administrative occupations’ into ‘Other managers and proprietors’. On a more granular level, this includes for example 4-digit transitions from ‘Other administrative occupations n.e.c.’ (for instance back office functions in retail), into ‘managers in transport and distribution’ (with an estimated 140k transitions). Another frequent transition is ‘customer service occupations n.e.c.’ into ‘Sales and retail assistants’ (120k transitions).

BOX 3: METHODOLOGY FOR JOB TRANSITIONS

To predict possible jobs transitions for displaced employees, we use so-called neural embeddings to create a network of 'skills similarity'. In other words, we look at which non-automatable tasks in displaced jobs are similar to non-automatable tasks in non-displaced jobs. Our approach to identifying similar jobs is by finding the similarity in the textual description of the tasks associated with the job. We use standard natural language processing methods to identify textual similarity between task descriptions (Srinivasa Desikan 2018).

4.3. INDUSTRIAL STRATEGY CAN HELP CREATE DEMAND FOR NEW OR LESS-EXPOSED TASKS

In the previous section we outlined how there are a vast range of possible outcomes resulting from adoption of generative AI. At one extreme we see huge increases in aggregate output and no jobs losses, while at the other extreme we could see no output increase at all and large scale labour market disruption. In this section we explore how industrial strategy policies can help bring us closer to the former scenario.


Industrial strategy can help shape the direction of growth of an economy, support critical sectors and, as a result, impact the composition of jobs in an economy (Alvis et al 2023). Technological change is not neutral, it can take place in various ways. Policy, in turn, can steer the direction of innovation. Most countries are already doing this with regards to climate change: policies have been put in place that are incentivising deployment of low-emission technologies. In Jung and Srinivasa Desikan (2023) we have argued that in order to steer the direction of new technologies, the UK needs an industrial strategy for AI. Economic incentives (subsidies, taxes, regulations) together with public digital infrastructure can achieve this.

A job-centric industrial strategy for AI would have multiple objectives. It would (1) help promote tasks and jobs that are less exposed to generative AI; (2) support workers to transition into such roles and (3) ensure that the fruits of automation are shared widely across the economy. We think continuing to enable people to work will be key due to the intrinsic value that is still widely given to work, both on an individual and societal level (Susskind 2020). Moreover, ensuring people have access to meaningful and well paying work can take some weight of income payments via the social insurance system.

A green industrial strategy could be a first step in this direction. In chapter 3 we showed that manual tasks are less exposed to 'integrated' generative AI. Figure 4.4 shows that **green jobs are less exposed to automation than non-green jobs**, as they have a relatively higher share of such tasks. Thus, delivering a green industrial strategy could have the positive side effect of also increasing the supply of non-automatable jobs. Note though that a green industrial strategy is unlikely to be enough on its own. In the UK it could create up to 725,000 jobs by 2030 (CCC 2023) which might be insufficient for the degree of impact of the 'integrated' AI as shown above. Of course not all displaced workers will have the appropriate skills to move into manual and technical jobs. We thus consider further options in the next section.

FIGURE 4.4: GREEN JOBS HAVE A RELATIVELY HIGHER SHARE OF MANUAL AND COMMUNICATIVE TASKS, WHICH HAVE LOWER EXPOSURE TO AUTOMATION

Task types, as a percentage share of all tasks



Note: The definition of 'green jobs' is based on an ONET (2023) classification.

Source: IPPR analysis of ONET (2023) and ONS (2023)

Thought experiment: An industrial strategy for the 'social' occupations

Beyond a green industrial strategy, we argue that an active labour market policy should go further in order to promote creation of jobs that more resilient to automation. One such example could be supporting what we call 'social occupations'¹⁵ for the purpose of this exercise.

There is a clear social need for this. Lots of evidence points to growing levels of social isolation in the UK. Existing social occupations – particularly social care and mental health services – are under-resourced and are having to ration the interpersonal time that is at the core of their work. There are lots of existing vacancies in these sectors now, but they are not being filled as the work is undervalued and low-paid (Patel et al 2023).

A social occupations focussed industrial strategy could shift that, looking at professionalising some social occupations, recognizing the skills and qualifications of workers, and building routes to progression.

Figure 4.5 illustrates this for five jobs, that currently have relatively high levels of employment and relatively high levels of interpersonal tasks. Two things stand out:

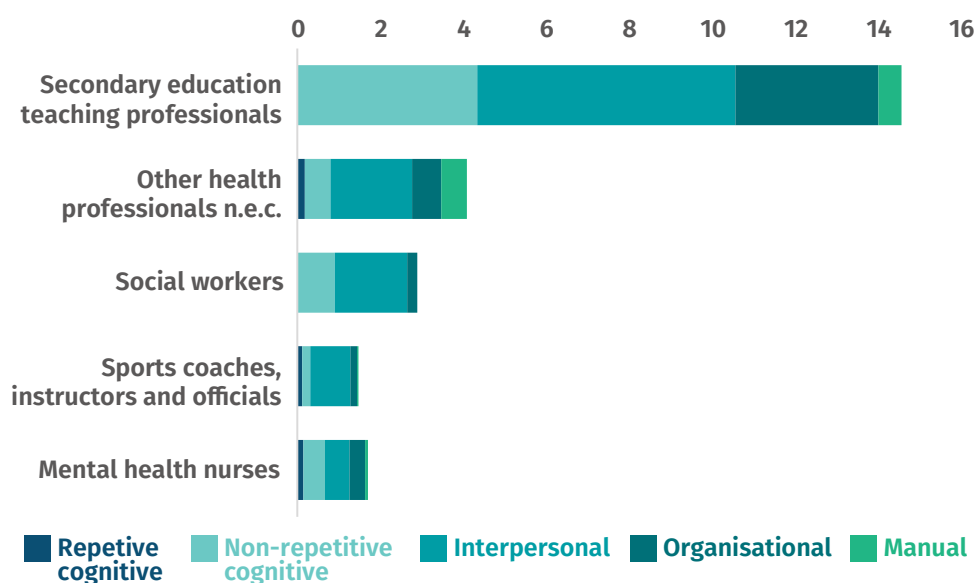
- **The share of interpersonal tasks could be further increased.** Even jobs considered largely 'social', will still have high shares of organisational or non-repetitive cognitive tasks. One could imagine that the latter two parts shrink in importance and the interpersonal part increases further still. This automation of administrative tasks could be a hopeful development in attempts to maximise NHS productivity, and in improving autonomy for clinicians in how they treat and support patients. In a previous IPPR report (Quilter-Pinner and Khan (2023)), further examples of such time saving approaches of automating repetitive tasks and increasing 'time to care' are developed.

¹⁵ Social occupations are defined as those which currently have at least 40 of their tasks (scaled by hours worked) allocated to interpersonal and communications tasks.

- **Sectors with high interpersonal task shares could be increased.** For instance, social care workers, mental health nurses, social workers and sports coaches have relatively high shares of interpersonal tasks, but employment in these sectors is very low, in spite of high social needs. Thus, similar to a green industrial strategy, there might be a case for supporting jobs that fulfil real social needs while also targeting tasks and jobs that are less exposed to automation.

FIGURE 4.5: AN INDUSTRIAL STRATEGY COULD FOCUS ON JOBS WITH A HIGH SHARE OF 'SOCIAL TASKS' AS THESE CAN BE LESS AT RISK TO FULL AUTOMATION

Total hours worked per week per task (in million hours per week)



Source: IPPR analysis of ONET (2023) and ONS (2023)

More broadly, ring-fencing of tasks – ie ensuring continued human involvement and work augmentation – should be at the heart of industrial strategy

Other tasks and occupations should be considered for a job-centric industrial strategy. Based on our review of tasks we suggest a few **illustrative types of tasks that would be worth considering for ringfencing:**

- **Tasks involving interpersonal relationships and trust.** Tasks that depend on building interpersonal trust and understanding, such as those in leadership, negotiation, or mentoring, benefit from human intuition, empathy, and social intelligence. Examples include: Effective leadership and mentorship, negotiation in complex situations (like diplomatic talks), or roles that require building deep trust and rapport (eg therapeutic professions) should minimise automation to retain these human qualities.
- **Tasks involving ethical sensitivity and moral judgments.** Tasks that require ethical sensitivity, moral judgments, or navigate complex social contexts should be approached with caution. Examples include: Social work or psychological counselling involves nuanced understanding of human emotions, ethical considerations, and cultural contexts that AI cannot fully replicate.
- **Tasks involving creative originality and artistic expression.** Generative AI can already generate art, music, and literature, tasks that value originality and unique artistic expression. But human touch in these areas might be considered crucial to add depth and authenticity. Protecting such work is

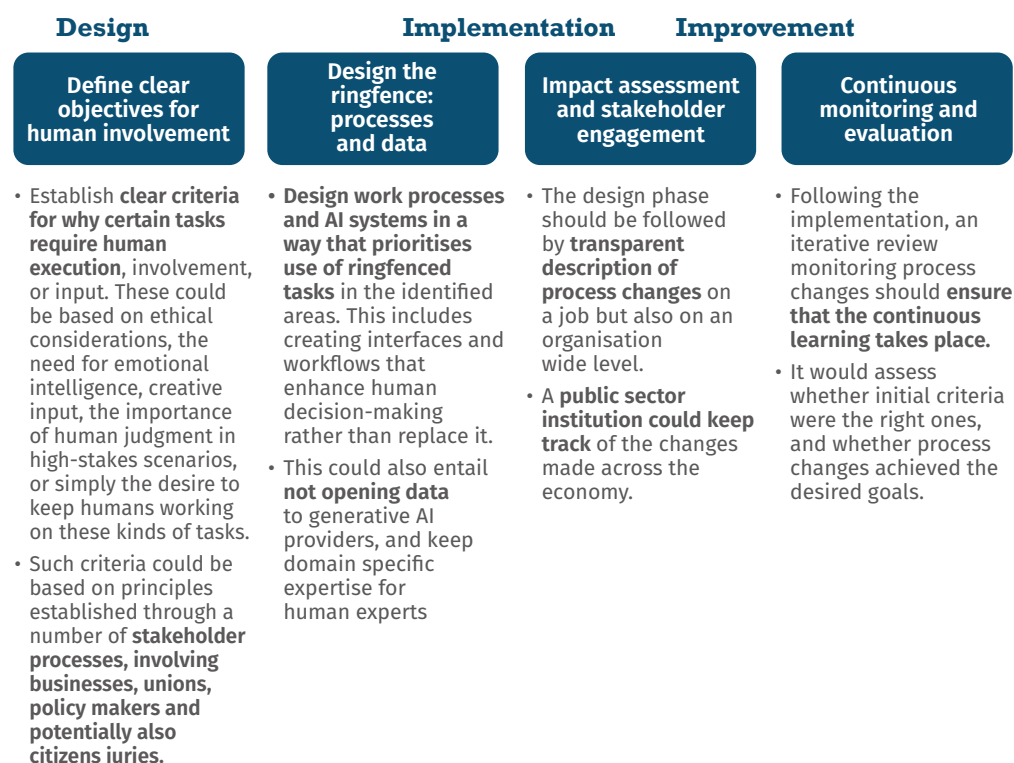
essential to ensure creative skills continue to be nurtured and valued among children and adults. Examples include: commissioning a painting, writing a novel, or composing music where the purpose is to convey deeply personal or innovative expressions should remain human-led to preserve originality and emotional depth

- **Tasks involving high-stakes decision making.** Tasks involving decisions with significant consequences for human life, health, and safety should be carefully shielded from full automation. While AI can support with data analysis and recommendations, the final judgment should remain with humans. Examples include: medical diagnosis and treatment planning, judicial decisions, or critical safety decisions in aviation and nuclear energy sectors require a level of accountability and ethical consideration beyond current AI capabilities.
- **Tasks involving complex problem solving in dynamic environments.** Tasks that involve solving novel problems in ever-changing environments, requiring adaptive thinking and innovation, should leverage human flexibility and creativity. Examples include: Strategic decision-making in businesses during uncertain times, research and development in novel areas, or emergency response during crises situations require a level of adaptability and innovation that people might not automate.

We think it can have social and economic benefits to **ringfence such tasks from full automation**, requiring a continued degree of human involvement. This would be done through a combination of bottom-up classification of tasks and government policy seeking to promote these tasks through the use of fiscal incentives, regulations, provision of digital infrastructure and other government policies (such as skills investment).

Ringfencing tasks might still involve deploying AI. But it suggests careful thinking about where human involvement is desirable and for what reasons – rather than rushing to automate all things that are technically possible to automate. Thus we suggest that, **for each of these tasks, a continuous design-implementation-improvement cycle should take place** – as shown in figure 4.6.

FIGURE 4.6: THE PROCESS OF RINGFENCING TASKS IN AN ORGANISATION, JOINTLY CONDUCTED BY A RANGE OF STAKEHOLDERS




Source: Authors' analysis

Ringfencing certain tasks from automation through generative AI, and others being naturally less exposed, would change the structure of the economy.

The result could be that non-ringfenced tasks deploy AI and their relative share of employment declines similar to, say, manufacturing since the 1970s. Meanwhile, the policy would incentivise the creation of *tasks* that have either low exposure to generative AI, or those that are ringfenced – ie are designed in such a way to be augmented rather than displaced. Following Baumol's (1965) famous 'cost disease' mechanism, this would have the following consequences: as technology gets better and automates more tasks, making products cheaper, people will spend relatively more of their money on services that technology is not making cheaper as quickly. This would include the social element of education, healthcare, or entertainment. These areas would not see costs fall as a result of generative AI in the same way as that back office work does, so they become a bigger part of the economy. Over the last century, this is exactly what's been happening: as agriculture and manufacturing became more automated and cheaper, people have been spending a bigger chunk of their income on services that have not seen those productivity boosts.

In a forward looking scenario, these jobs will be the ones that will employ relatively more people and thus play an important part for providing balanced demand (figure 4.7).

FIGURE 4.7: AUTOMATION-RESISTANT TASKS WOULD GROW AS A SHARE OF GDP



Source: Authors' analysis

5.

TOWARDS A JOB-CENTRIC INDUSTRIAL STRATEGY FOR AI

We argue that the more positive scenarios can only be realised through a wide range of targeted policy interventions, which we summarise as a **job-centric industrial strategy for AI**. In this section we introduce key pillars for achieving this, drawing on the analytical insights of this report. In future work we will seek to provide detailed policy guidance on each of its elements.

The three key pillars are:

1. protect existing jobs and ensure gains for workers
2. boost creation of new tasks, jobs and support job transitions
3. address the fallout from lower labour demand.

A new centralised institution would likely be needed to coordinate these, helping ensure policies complement each other.

PROTECTING EXISTING JOBS AND ENSURING GAINS FOR WORKERS

In the ‘scenarios’ section above, we have shown that the wage gains for workers could be huge – more than 30 per cent in some cases – but they could also be nil. Moreover, we have shown that **exposure to AI does not happen automatically**. It will depend on processes changing, AI platforms being built and implemented, and organisations integrating AI with existing databases and work processes. But organisations – incentivised by policy – **can decide not to go for full AI integration and maximum automation**. Instead, AI can be adopted in ways that do not lead to excessive displacement of workers. We have in detail discussed how ‘ringfencing’ of certain tasks can ensure that tasks in which we value human involvement are protected. Unions and policymakers, together with business, will play a crucial role in an iterative process to do this, which we detailed in section 6. Moreover, tax policy can help this process by not incentivising full automation over human labour, as is currently the case in the US and the UK (Acemoglu and Johnson 2023).

BOOST CREATION OF NEW TASKS, JOBS AND SUPPORT JOB TRANSITIONS

Policy can also be pro-active in helping creation of jobs that have low risk of automation. We have presented a thought experiment, outlining the idea of an industrial strategy that boost interpersonal jobs, while also ‘ringfencing’ certain social tasks. This could help address the big social needs for expanding the supply of currently under-resourced services, eg social care and mental health services. Incentivising such work, through subsidies or tax incentives, and by improving working conditions and pay will also have to be complemented by a retraining offer. This will be required because, as we have shown in section 5, these jobs are not necessarily ‘similar’ to the ones that are being displaced. One idea would be for a National Employment Service to be at the heart of supporting people with such labour market transitions (Wilkes and Parkes 2023). This would involve giving people both information about viable career transitions and point out routes for retraining.

ADDRESS THE FALLOUT FROM LOWER LABOUR DEMAND

Our scenarios have shown that there are potentially large gains from the deployment of generative AI – as much as 12 per cent boost to GDP in the 'integrated' AI scenario. But as OpenAI CEO Sam Altman (2021) has highlighted, there is no guarantee that these will be equally shared. He, as a result, called for a wealth tax. In practice, the tax response will likely have to be nuanced, balancing the incentive to innovate with the goal of internalising the unintended costs from technology deployment. In line with taxation, social security assistance might also have to evolve in order to help attenuate any disruptive effects that might occur from fast or unexpectedly large labour market impacts of generative AI.

FIGURE 5.1: POLICY PILLARS FOR A JOB-CENTRIC INDUSTRIAL STRATEGY FOR AI

Objective	Policy	Time to have an impact
1 Protect existing jobs and ensure gains for workers	Actions to augment jobs • Provide incentives ensuring that productivity gains from AI deployments benefit employment	Medium term
	Ringfence tasks for human involvement • Eg prevent process engineering aimed at job displacement	Medium term
	Change fiscal incentives • Change current tax advantages for automation, consider 'augmentation subsidies'	Short term
2 Boost creation of new tasks, jobs and support job transitions	Incentivising creation of jobs with low exposure tasks • Subsidies or tax incentives for ringfenced and low-exposure tasks	Medium term
	Retraining and upskilling • Where jobs transitions constitute significant career changes, provide guidance and support	Long term
	Competition policy • Use competition policies to support smaller firms with AI adoption	Medium term
3 Address the fallout from lower labour demand	Worktime reduction • Fiscal incentives or strengthen labour legislation to encourage working time reduction	Short term
	Tax the gains from automation • Tax the gains from automation, eg via specific 'AI taxes' or wealth taxes	Medium term
	Social security reform, UBI • Expand social safety net to manage impact on incomes	Medium term

Source: Authors' analysis

THE NEED FOR SPEED

Our scenarios show that the potential range of impacts of generative AI (good and bad) is vast and that the way they play out hugely depends on policy choices. Given the speed and breadth of generative AI adoption that we show is possible across sectors, and given the long lags with which policies take effect, **we think it is urgent that policy makers start preparing a policy response now.** Our proposed job-centric industrial strategy for AI outlines the range of options. Fleshing these out in detail, in response to plausible near term scenarios, will be the key for reducing the risk of labour market disruption.

REFERENCES

- Acemoglu D and Restrepo P (2019) 'Automation and New Tasks: How Technology Displaces and Reinstates Labour', *Journal of Economic Perspectives*. <https://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.33.2.3>
- Acemoglu D and Restrepo P (2022) 'Tasks, Automation, and the Rise in U.S. Wage Inequality', *Econometrica*. <https://onlinelibrary.wiley.com/doi/full/10.3982/ECTA19815>
- Acemoglu D and Johnson (2023) *Technology and Power*, MIT University Press
- Altman S, Anadkat S and Avila R (2023) *Gpt-4 technical report*, arXiv preprint arXiv:2303.08774. <https://arxiv.org/pdf/2303.08774.pdf>
- Autor D, Chin C, Salomons AM, Seegmiller B (2022) *New Frontiers: The Origins and Content of New Work, 1940–2018*, National Bureau of Economic Research, no w30389
- Alvis S, Dibb G, Gasperin S and Murphy L (2023) *Market making in practice*, IPPR. <http://www.ippr.org/research/publications/making-markets-in-practice>
- Altman S (2021) 'Moore's Law for Everything', blog. <https://moores.samaltman.com/>
- Barash Y, Konen E, Glicksberg BS, Nadkarni G and Klang E (2023) *Assessing GPT-4 Multimodal Performance in Radiological Image Analysis*, medRxiv, pp2023-11. <https://www.medrxiv.org/content/10.1101/2023.11.15.23298583v1>
- Baumol W (1965) 'On the performing arts: the anatomy of their economic problems', *The American Economic Review*. <https://www.jstor.org/stable/1816292>
- Bell T (2024) 'Britain isn't post work', blog. <https://www.resolutionfoundation.org/comment/britain-isnt-post-work/>
- Benzell SG, Kotlikoff LJ, LaGarda G and Ye VY (2021) *Simulating endogenous global automation*, National Bureau of Economic Research, no w29220. <https://www.nber.org/papers/w29220>
- Bessel J (2018) *AI and jobs: The role of demand*, NBER working paper. https://www.nber.org/system/files/working_papers/w24235/w24235.pdf
- Bloomberg (2023) 'IBM to Pause Hiring for Jobs That AI Could Do', news article. <https://www.bloomberg.com/news/articles/2023-05-01/ibm-to-pause-hiring-for-back-office-jobs-that-ai-could-kill>
- Brin D, Sorin V, Barash Y, Konen E, Glicksberg BS, Nadkarni G and Klang E (2023) 'Assessing GPT-4 Multimodal Performance in Radiological Image Analysis', *medRxiv*, pp2023-11. <https://www.medrxiv.org/content/10.1101/2023.11.15.23298583v1>
- Broadberry S, Campbell B, Klein A, Overton M, and van Leeuwen B (2010) *British economic growth: 1270-1870*. http://wrap.warwick.ac.uk/57581/1/WRAP_35.2010_broadberry.pdf
- Bubeck S, Chandrasekaran V, Eldan R, Gehrke J, Horvitz E, Kamar E, Lee P, Lee YT, Li Y, Lundberg S and Nori H (2023) *Sparks of artificial general intelligence: Early experiments with gpt-4*, arXiv preprint arXiv:2303.12712. <https://arxiv.org/abs/2303.12712>
- Cazzaniga M, Jaumotte F, Li L, Melina G, Panton AJ, Pizzinelli C, Rockall EJ and Tavares MM (2024) *Gen-AI: Artificial Intelligence and the Future of Work*, Staff Discussion Notes, 2024(001). <https://www.elibrary.imf.org/view/journals/006/2024/001/article-A000-en.xml>
- Challenger, Gray & Christmas (2023) 'May 2023 Layoffs Jump on Tech, Retail, Auto; YTD Hiring Lowest Since 2016', webpage. <https://www.challengergray.com/blog/may-2023-layoffs-jump-on-tech-retail-auto-ytd-hiring-lowest-since-2016/>
- Chartered Institute of Personnel and Development [CIPD] (2021) *People Profession 2021: UK and Ireland survey report*. https://www.cipd.org/globalassets/media/comms/news/al1people-profession-uk-ireland-2021-report-1_tcm18-102392.pdf
- Committee on Climate Change [CCC] (2023) *A Net Zero Work Force*. <https://www.theccc.org.uk/publication/a-net-zero-workforce/>
- Conway G (1999) *The doubly green revolution: food for all in the twenty-first century*, Cornell University Press. <https://www.jstor.org/stable/10.7591/j.ctvtv9340>

- 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*, arXiv preprint arXiv:2303.10130. <https://arxiv.org/pdf/2303.10130.pdf>
- Felten E, Raj M and Seamans R (2023) *How will Language Modelers like ChatGPT Affect Occupations and Industries?*, arXiv preprint arXiv:2303.01157. <https://arxiv.org/ftp/arxiv/papers/2303/2303.01157.pdf>
- Frey CB (2019) *The technology trap: Capital, labor, and power in the age of automation*, Princeton University Press. <https://press.princeton.edu/books/hardcover/9780691172798/the-technology-trap>
- Frey CB, Osborne MA (2017) 'The future of employment: How susceptible are jobs to computerisation?', *Technological Forecasting and Social Change*, 114, pp 254-280. <https://www.sciencedirect.com/science/article/abs/pii/S0040162516302244>
- Goldman Sachs (2023) 'Generative AI Could Raise Global GDP by 7 Percent', webpage. <https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html>
- Guardian (2024) 'Tyler Perry halts \$800m studio expansion after being shocked by AI', news article.
- Humphries J and Schneider B (2021) 'Gender equality, growth, and how a technological trap destroyed female work', *Economic History of Developing Regions*, 36(3), pp 428-438. <https://journals.co.za/doi/pdf/10.1080/20780389.2021.1929606>
- Jung C and Srinivasa Desikan B (2023) *Artificial intelligence for public value creation: Introducing three policy pillars for the UK AI summit*, IPPR. <http://www.ippr.org/research/publications/ai-for-public-value-creation>.
- Korinek A (2023) 'Scenario Planning for an A(G)I Future', *IMF Finance and Development Magazine*, December 2023. <https://www.imf.org/-/media/Files/Publications/FandD/Article/2023/December/30-33-korinek-final.ashx>
- Kung TH, Cheatham M, Medenilla A, Sillos C, De Leon L, Elepaño C, Madriaga M, Aggabao R, Diaz-Candido G, Maningo J and Tseng V (2023) 'Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models', *PLoS Digital Health*, 2(2), e0000198. <https://journals.plos.org/digitalhealth/article?id=10.1371/journal.pdig.0000198>
- Martin L, Nick W, Stephanie Y, Lizzie C, and Rivindu P (2024) *Better Call GPT, Comparing Large Language Models Against Lawyers*, arXiv preprint arXiv:2401.16212. <https://arxiv.org/html/2401.16212v1>
- McKinsey & Company (no date) *The Economic Potential of Generative AI: The Next Productivity Frontier*. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>
- National Public Radio [NPR] (2023) 'An eating disorders chatbot offered dieting advice, raising fears about AI in health', news article. <https://www.npr.org/sections/health-shots/2023/06/08/1180838096/an-eating-disorders-chatbot-offered-dieting-advice-raising-fears-about-ai-in-hea>
- Nori H, King N, McKinney SM, Carignan D, Horvitz E (2023) *Capabilities of gpt-4 on medical challenge problems*, arXiv preprint arXiv:2303.13375. <https://arxiv.org/pdf/2303.13375.pdf>
- ONET 28.2 Database [ONET] (2023) 'ONET Resource Center', National Center for O*NET Development, dataset. www.onetcenter.org/database.html.
- Office for National Statistics [ONS] (2023) 'Labour Force Survey Q1 2023', dataset.
- Parkes H and Wilkes M (2023) *Towards real social security: Embedding a long-term approach to universal credit*, IPPR. <http://www.ippr.org/research/publications/towards-real-social-security>
- Patel P, Quilter-Pinner H, Williamson A and Thomas C (2023) *For public health and public finances: Reforming health and social care*, IPPR. <http://www.ippr.org/research/publications/for-public-health-and-public-finances>
- Quilter-Pinner H and Khan H (2023) *Great government: Public service reform in the 2020s*, IPPR. <http://www.ippr.org/publications/great-government>
- Reuters (2023) 'ChatGPT sets record for fastest-growing user base: analyst note', news article. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>

- Shaw-Taylor L, Sudgen K, You X (2019) 'A preliminary estimate of the female occupational structure of England and Wales 1700–1911', *Cambridge Group for the History of Population and Social Structure Project Paper*, 33
- Siemiatkowski S (2024) 'tweet'. <https://twitter.com/klarnaseb/status/1762508581679640814>
- Srinivasa Desikan B (2018) *Natural Language Processing and Computational Linguistics: A practical guide to text analysis with Python, Gensim, spaCy, and Keras*, Packt Publishing Ltd. <https://www.packtpub.com/product/natural-language-processing-and-computational-linguistics/9781788838535>
- Sullivan SE and Al Ariss A (2021) 'Making sense of different perspectives on career transitions: A review and agenda for future research', *Human Resource Management Review*, 31(1), p100727. <https://www.sciencedirect.com/science/article/pii/S1053482218306788>
- Susskind D (2020) *A World Without Work*, Penguin. <https://www.penguin.co.uk/books/306864/a-world-without-work-by-susskind-daniel/9780141986807>
- The Economist (2023a) 'How businesses are experimenting with ChatGPT like services', news article. <https://www.economist.com/business/2023/04/19/how-businesses-are-experimenting-with-chatgpt-like-services>
- The Economist (2023b) 'What happened to the artificial intelligence investment boom?', news article. <https://www.economist.com/finance-and-economics/2024/01/07/what-happened-to-the-artificial-intelligence-investment-boom>
- Trade Union Congress [TUC] (2023) *Dignity At Work and The AI Revolution*. https://www.tuc.org.uk/sites/default/files/2021-03/The_AI_Revolution_2021_Manifesto_AW.pdf
- Sullivan SE and Al Ariss A (2021) 'Making sense of different perspectives on career transitions: A review and agenda for future research', *Human Resource Management Review*, 31(1), p 100727. <https://www.sciencedirect.com/science/article/pii/S1053482218306788>
- UK Department of Education (2023) *Impact of AI on UK Jobs and Training*. <https://www.gov.uk/government/publications/the-impact-of-ai-on-uk-jobs-and-training>.
- UK Department for Digital, Culture, Media, and Sport [DCMS] (2022) *AI Activity in UK Businesses*, report. https://assets.publishing.service.gov.uk/media/61d87355e90e07037668e1bd/AI_Activity_in_UK_Businesses_Report_Capital_Economics_and_DCMS_January_2022_Web_accessible_.pdf
- Washington Post (2023) 'ChatGPT took their jobs', news article. <https://www.washingtonpost.com/technology/2023/06/02/ai-taking-jobs/>
- Zhang P, and Kamel Boulos M N (2023) 'Generative AI in medicine and healthcare: Promises, opportunities and challenges', *Future Internet* 15, no 9. <https://www.mdpi.com/1999-5903/15/9/286>

APPENDIX

1. THE PROMPT FOR 'INTEGRATED' AI

As part of our approach to predicting generative AI's automating capacities for the near future, we used an adapted approach for our 'here and now' AI exposure metric. To offer a consistent approach we constructed a detail 'abilities' prompt, accompanying each prediction. With this prompt we obtained scorings similar to our personal annotation of task exposure. We 'remind' the model of its capabilities keeping in mind both its current skills, as well as research papers and current experiments predicting the capabilities of integrated AI. Notably, we also add strict instructions on cases to say no – where either significant process engineering or social context (authority, empathy, high stakes) are involved. This is in order to mimic the three phases of AI deployment which we highlight in the paper.

The prompt structure is as follows.

“You are to give the answer ONLY as a YES or NO, with no other explanation. Below is some information on the context of the situation and how to answer:

Consider a GPT4 level generative AI that is connected to existing software systems, with the ability to interact and prompt each other, capable of storing data to disk.

It can use the internet, code interpreters, access data and retrieve external information, analyse images and videos, and domain specific data, as well as other APIs and plug-ins.

However, say NO if the task involves exercising authority, usually requires significant amounts of empathy, or is associated with high stakes.

Also, answer with NO in case the task would require some significant process re-engineering first, such as installing cameras, or starting recording information, or re-organising the way the task is done.”

For robustness, we chose a sample of 250 tasks to match our annotation and the LLM annotations on which tasks it could automate and not. We found a high degree of agreement in our assessments and the LLM a Pearson correlation of 0.8. Indeed, the cases where we found disagreements were often 'borderline' situations such as research, analysis, and monitoring, where some process engineering and organising around LLMs would be required before being automated.

2. EXAMPLES OF TASKS AND HOW WE JUDGED THEIR EXPOSURE TO AUTOMATION

The way we judged tasks exposure to generative throughout was with regards to whether it could speed up processes by 50 per cent or more. Note that, as discussed above, this is not the same as automation but can still have significantly effect overall labour demand as the same amount of output can, in theory, produced by fewer workers.

A few principles through which we judged automation (and which we used to fine tune the GPT4 assessment was as follows):

- As highlighted in chapter 2, there already exists a wide range of cognitive and organisational tasks that generative AI can do accurately, largely on par with humans. In such cases we judged the tasks be able to be sped up by at least 50 per cent.
- Most advanced research and administrative tasks were not judged to be exposed to 'here and now' generative AI. But integrated AI – eg if connected to key databases - was judged to be able to significantly speed up research and work processes.
- Highly social tasks were considered not fundamentally exposed to AI, not to the degree that they could be done 50 per cent faster.
- Manual tasks that require physical presence were judged to not exposed to either current generative AI or agent-like AI.


TABLE A1:

Task	Can the task be done 50% more quickly by 'here and now' generative AI	Can the task be done 50% more quickly by 'integrated' AI	Notes
Marketing, sales and advertising directors: analyse marketing or sales trends to forecast future conditions.	Yes	Yes	Existing generative AI tools already have significant capabilities to speed up data analysis and interference generation capabilities.
Human resources and industrial relations officers: Monitor company or workforce adherence to labour agreements.	No	Yes	'Here and now' AI: A simple generative AI likely cannot significantly speed up this task. 'Integrated' AI: a more elaborate system, with access to databases can help make it faster by if connected to a database that captures labour agreements, other relevant documents and employees reporting infringements (eg via an app) and cross check them against existing practices.
Higher education teaching professionals: Initiate, facilitate, and moderate classroom discussions.	No	No	The task is social and while technically one could imagine it to be done via generative AI, in virtual classrooms, we scored it as not exposed. It would involve significantly re-designing how classroom discussions and teaching take place – something that we see would only happen in phase 3 of generative AI deployment.
Carpenters and joiners: Install structures or fixtures, such as windows, frames, floorings, trim, or hardware, using carpenters' hand or power tools.	No	No	Physical task cannot be done faster with generative AI, without physical infrastructure.
Secondary education teaching professionals: Prepare reports on students and activities as required by administration.	No	Yes	Here and now AI: would require significant administration and oversight by teachers. 'Integrated AI' AI: If current generative AI would get access to databases with student reports and activities, it could significantly speed up preparing reports.

Source: Authors' analysis

FIGURE A.1: INTEGRATED AI: EXPOSURE TO AUTOMATION IS FAIRLY EVENLY SPREAD ACROSS WAGE LEVELS


Task exposure by occupation vs median hourly wage



Source: IPPR analysis of ONET (2023) and ONS (2023)

FIGURE A.2: WOMEN ARE DISPROPORTIONATELY EXPOSED TO 'HERE AND NOW' AI

Task exposure by occupation vs share of women in the occupation



Source: IPPR analysis of ONET (2023) and ONS (2023)

Institute for Public Policy Research



GET IN TOUCH

For more information about the Institute for Public Policy Research, please go to www.ippr.org

You can also e-mail info@ippr.org or find us on X [@ippr](#)

Institute for Public Policy Research

Registered Charity no. 800065 (England & Wales),
SC046557 (Scotland), Company no, 2292601 (England & Wales)

The progressive policy think tank