TONY BLAIR INSTITUTE FOR GLOBAL CHANGE

The Impact of AI on the Labour Market: Technical Annex

The nature of AI's integration into different sectors of the economy and different jobs remains highly uncertain. We are still in the foothills of AI adoption, so robust empirical studies assessing AI's impact on specific jobs and tasks remains limited.

Any estimates of AI's potential impact must therefore largely be based on a forward-looking assessment of how AI could affect the future of work rather than how it already has affected work. All studies that provide a macroeconomic-level assessment of AI's impact on the labour market thus rely on an ex-ante assessment of the technology's potential at a point in time and rest on specific assumptions that are open to debate. This study is no different.

This paper builds on existing TBI analysis that has estimated the potential time savings for public-sector workers from using artificial-intelligence tools.[1] Our methodology starts, as is now common in the literature on technology and employment, with the premise that jobs consist of several distinct tasks. It is not "jobs" that are replaced by technology but tasks: new technologies might perform some of these tasks, but not others. As with other similar studies, we obtain a list of tasks associated with each occupation from the O*NET database.[2]

The main aim of our analysis is to identify what types of tasks could be performed by AI and how much worker time could be saved as a result. Here, there are at least three broad

[1] Owing to the use of a very similar method, the material included in this section of the technical annex closely aligns with that previously included in our paper, *The Potential Impact of AI on the Public Sector Workforce.*
[2] https://www.onetcenter.org/database.html

approaches that have already been trialled in the literature – each with advantages and drawbacks:

- **Method 1: Broad categorisation of tasks based on individual human judgement:** This method, used by economists Briggs and Kodnani (2023) in a Goldman Sachs research paper, involves identifying broad clusters of "work activity" that could potentially be performed by AI.[3] In the O*NET database there are 39 categories of work activity (for example, "getting information" or "monitoring and controlling resources"). The authors assume 13 of these work activities could be performed by generative AI based on their own judgement and a reading of the existing evidence on AI. Within each work activity there are also seven levels of complexity identified in the O*NET database, and the authors assume that AI can perform up to a difficulty level of four across each of the 13 categories. Then, because the study is only focused on generative AI, they assume that any occupation that involves a significant share of workers' time spent outdoors or performing physical tasks cannot be automated by AI at all. Any task that meets all the above criteria is then assumed to be fully automatable and hence could lead to time savings of up to 100 per cent. The authors' headline result is that 25 per cent of UK workforce time is exposed to automation from generative AI.

  This approach is transparent and defensible, but as with all methods it does have some drawbacks. First, its categorisation of tasks is broad, so it cannot account for how AI's capabilities to perform individual tasks might vary across different professions. For example, it may be easier for AI to perform writing tasks in professions where large amounts of machine-readable data are already available – such as the legal and financial professions – than in other settings where data are more siloed or expensive to digitise. Another drawback is that the analysis relies on fixed thresholds related to the difficulty of tasks that AI can perform, but this is somewhat arbitrary and may overestimate AI's abilities to perform some basic tasks and underestimate its ability to perform more complex tasks more efficiently. Finally, this method implicitly assumes that just because AI could perform 100 per cent of a task, all of that task is at risk of automation. However, in some professions this is unlikely to be the case, either because it is more efficient for a human to work in conjunction with an AI tool on a particular task, or it is socially desirable to keep a human in the loop. For example, for tasks that have a large impact on other human lives – such as judgments in criminal courts – generative AI may be able to produce

[3] https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html

judgments based on the submissions of both parties in a case, but it is unlikely that society would find this acceptable.

- **Method 2: Broad categorisation of tasks based on the wisdom of the crowd:** This approach, deployed by Felten, Raj and Seamans (2021), creates an AI-exposure index by profession.[4] They use information from the Electronic Frontier Foundation's (EFF) 2017 AI Progress Measurement study,[5] which identifies ten broad types of activity where AI has already been shown to be capable of performing particular tasks (for example, generating images, reading comprehension and so on). The authors then map those AI capabilities onto the 52 "occupational abilities" contained in the O*NET database based on a crowd-sourced survey of responses from 2,000 gig workers on the Amazon Mechanical Turk platform. The resulting AI Occupational Exposure index identifies which jobs are most exposed to disruption from AI.

  This method has been deployed by other researchers, and by the UK government[6] and the International Monetary Fund (IMF).[7] The latter has built on Felten et al's study by adding an "AI complementarity" component, which seeks to screen out tasks from the exposure index that would be socially unacceptable for AI to perform (such as decisions by judges). The IMF's headline result is that almost 70 per cent of the UK workforce has some exposure to generative AI.

  This approach is novel in that it is using the wisdom of a crowd (rather than experts) to link AI capabilities with particular tasks, and the IMF's extension is innovative in that it accounts for the social acceptability of adopting AI in different settings. But this method also has its drawbacks. First, the assessment of AI's capabilities is static, based on the EFF's 2017 study, so if there is a new breakthrough in AI's capabilities (such as its recent ability to generate video content) then these may not be captured. Second, as with Method 1, the broad categorisation approach could mask some differences in AI's ability to perform similar tasks across different professions. Third, it is not obvious that asking a large number of non-experts will definitively give an answer that is closer to the truth than asking a small number of experts or relying on a large language model's (LLM) training data – knowledge about task content is not widely dispersed. Finally, while this approach helpfully identifies exposure, it does not provide a specific assessment of how much time

---

[4] https://oar.princeton.edu/handle/88435/pr11551
[5] https://www.eff.org/ai/metrics
[6] https://www.gov.uk/government/publications/the-impact-of-ai-on-uk-jobs-and-training
[7] https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2024/01/14/Gen-AI-Artificial-Intelligence-and-the-Future-of-Work-542379

could be saved by adopting AI in each profession. This time-savings component is essential in our analysis to identify the potential efficiency gains of adopting AI.

- **Method 3: Granular categorisation of tasks using LLMs cross-checked against human judgement:** This method, pioneered by Eloundou, Manning, Mishkin and Rock (2023),[8] involves using AI itself to help categorise whether a particular task can be performed by generative AI or not. The authors initially categorise a wide range of detailed work activities from the O*NET database using human annotators familiar with the capabilities of LLMs. They categorise activities into three groups: 1) no exposure to AI; 2) direct exposure to generative AI (where it could help complete the task at least 50 per cent faster); and 3) exposure to generative AI when paired with other software (where the pairing can lead to time savings of at least 50 per cent). They then provide GPT-4 with a rubric of prompts to perform the same classification exercise and find a high degree of alignment between the human and AI-generated assessments, which suggests GPT-4 is a reasonable proxy for human judgement in this case. Overall, the authors find that LLMs could help complete 47 to 56 per cent of all worker tasks in the United States significantly faster.

  Other studies, including a recent report by the Institute for Public Policy Research (IPPR),[9] have applied Eloundou et al's approach to the UK and found that up to 59 per cent of workforce tasks could be affected by generative AI in the coming years. IPPR's analysis is very similar to Eloundou et al's, but relies more directly on ChatGPT's results, which are then cross-checked against the authors' judgements for a sample of 250 tasks.

  These studies are novel in that they are using ChatGPT as a form of the wisdom of the crowds – on the basis that the LLM was trained on a large volume of data that reflects the accumulated knowledge of the world's population and provides a probabilistic assessment of the most likely outcome based on those data. This approach also has an advantage in that it provides a more granular assessment of AI's capabilities across different tasks and professions and is more specific about identifying how much time could be saved from adopting AI. However, as with the other studies, these advantages also come with drawbacks. Eloundou et al acknowledge that both methods they use to categorise tasks are flawed – the human annotators used in their study are knowledgeable in the capabilities of LLMs but not in how they could be applied to specific professions. Meanwhile, ChatGPT's results, even though they match well with the human annotators, do give some

[8] https://arxiv.org/pdf/2303.10130
[9] https://ippr-org.files.svdcdn.com/production/Downloads/Transformed_by_AI_March24_2024-03-27-121003_kxis.pdf

contradictory assessments for some tasks. In addition, this approach does not distinguish between tasks that can be performed by AI in theory and those that should be performed by AI in practice; it includes tasks even if there are strong ethical reasons not to deploy AI in such a setting (as Method 2 adjusts for).

Clearly there are trade-offs between the different methods. None is perfect. Greater reliance on human judgement can limit the analysis to a broader categorisation of tasks with less specificity over time savings. On the other hand, pursuing a more detailed categorisation typically involves relying more on AI to support the assessment.

Our approach builds on this existing body of work in several ways. First, the scope of our study is broader; while the previous studies focused on generative AI, we attempt to assess the impact of all types of AI, including AI-enabled hardware that can perform physical tasks. Second, we provide a more granular assessment of AI's potential to save time within the workforce. Third, we utilise ChatGPT to help perform our analysis but do so by using it to create a system of filters that we then use to refine its results according to our own assessment.

We began our analysis with a training data set of around 200 work tasks from the O*NET database. We then used these data to iteratively develop a rubric of prompts that pushed OpenAI's GPT-4 Turbo model to produce results that accorded closely with our own judgement of AI's capabilities. Our own judgements were informed by the empirical studies of AI mentioned above, the latest research on AI's existing capabilities, conversations with AI-technology experts, and cross-checking the results from GPT-4 for particular tasks against those of each member of our own research team to benefit from the wisdom of crowds.

This process resulted in our using the following sequence of decision-tree prompts to GPT-4 to help categorise each task (which we have simplified here for brevity into short-form questions):

- **Is the task fully cognitive or does it require some sensory input or manual input to perform?** This is to identify whether a task can solely be performed by AI software (for example, generative AI or machine learning) or whether it requires complementary hardware as well (for example, a headset, microphone or more sophisticated AI-enabled hardware such as a drone or autonomous vehicle). This distinction is important as the base responses from ChatGPT tend to overestimate AI's current capabilities to perform physical tasks.

- **a. For cognitive tasks: Does the task require a high degree of human empathy, or does it have a significant impact on peoples' lives if errors are made?** This is to provide an assessment not just of whether AI can perform a task, but whether it is socially optimal to do so – like the IMF's "AI complementarity" component outlined in Method 2.

b. For tasks involving manual input: Is the task repetitive and performed in a stable environment, or is it a complex physical task involving high degrees of autonomy in a changing environment? This is to correct the bias in ChatGPT's assessment of physical tasks, whereby it assumes that AI-enabled hardware can already perform a range of complex physical tasks (such as autonomous driving), whereas in practice most AI-enabled hardware is mainly used in controlled environments to perform repetitive tasks (for example, collecting stock in warehouses and distribution centres).

- **Given these previous characteristics, could AI perform the task?** As part of the prompt for this question, we provide GPT-4 with a strong prior (based on the previous answers) that certain tasks – such as complex physical tasks or cognitive tasks involving high degrees of human empathy – cannot be done by AI. The model should only overwrite this prior if it has a very high degree of confidence in its answer (that is, a high probability of being true).

- **If the task can be done by AI:**

    a. **Is it more efficient to be done solely by AI or in conjunction with a human?**

    b. **Would it be cost effective to use AI?**

    c. **Would it be socially desirable for the task to be performed by AI?**

    d. **What type of AI would be required to perform the task?**

    e. **How much time could AI save relative to a human solely performing the task?**

The first four prompts are designed to provide filters for the data set enabling us to exclude certain tasks that would not be performed by AI in practice, even if they could theoretically be performed by AI. The final prompt then creates an estimate of the potential time saving from AI for each task.

Once we had refined our prompts on the training data set of 200 tasks, we then deployed the same rubric to GPT-4 Turbo to categorise the remaining nearly 20,000 tasks in the O*NET database – effectively giving GPT-4 a steer as to how to efficiently apply our judgements across a much larger data set, but still allowing it to apply the rich information in its training data to the nuances of each individual task.

Once we had this information at the task level, we then aggregated our results up to the occupation level by using information on the importance of each task in a profession to create an importance-weighted average. This gives us a potential time saving for each occupation in the O*NET database at the US Standard Occupational Classification (SOC) code level.

To apply this information to the UK Labour Force Survey (LFS) data, we then used a crosswalk to match US SOC codes to UK SOC codes so that we can calculate the potential saving for each occupation as defined in the UK SOC.

As part of the model-refinement process, we also cross-checked our results at various points to ensure they accord with frontier studies that assess AI's potential impact on work. For example:

- **At the task level:** Dell'Acqua et al (2023)[10] conducted a field study of 758 consultants from the Boston Consulting Group (BCG) to test the ability of AI to perform a range of consultancy tasks. They found that AI helped perform these tasks 25 per cent faster on average. Our analysis of the "management analyst" and "project-management specialist" professions, both of which involve tasks closely related to those in the study, estimated that the use of AI would introduce a time saving of 24.7 per cent and 19.5 per cent respectively – very close to the empirical results from the study. Moreover, both these figures are conservative when compared with other studies on the ability of AI to improve the writing speed of business professionals. For example, Noy and Zhang (2023) show that professionals who use ChatGPT to help with writing tasks can save about 40 per cent of their time.[11]

  Peng et al (2023) conducted a separate study focused on computer programmers and found that access to GitHub Copilot, an AI pair programmer, can save 55.8 per cent of time for some coding tasks.[12] Our estimates, which apply to a wider range of coding and software-based tasks, are more conservative but in a similar ballpark – indicating an average saving of 39 per cent across these tasks, or 29 per cent for all tasks associated with computer programmers.

[10] https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdfhttps://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321
[11] https://www.science.org/doi/10.1126/science.adh2586
[12] https://arxiv.org/abs/2302.06590

- **At the occupation level:** Recent evidence from the UK Department for Education (DfE) found that AI can save teachers at least 4 per cent of their time, while a new study by Oak National Academy, a provider of digital teaching resources, suggests a time saving of up to 8 per cent.[13, 14] Again, this range closely matches our own estimate of time saved by primary and secondary school teachers through the use of AI (6.3 per cent and 7.6 per cent respectively). And again, our results are more conservative than some other studies, for example a McKinsey study suggested time savings of 20 to 40 per cent were possible for teachers.[15]

- **At the economy-wide level:** We can also compare our results at an aggregate level with those of the macroeconomic studies mentioned earlier. Our overall potential time saving of 25 per cent across the whole economy closely matches Briggs and Kodnani's estimate[16] highlighted in Method 1. Our results on the share of tasks impacted by AI are also in a similar ballpark to other studies. We estimate 50 per cent of employment-weighted tasks across the whole economy are potentially exposed to AI – in line with Eloundou et al's 47 to 56 per cent range for the United States,[17] but slightly less than IPPR's 59 per cent figure for the UK.[18] Since our study includes a broader range of AI tech than these studies, which focus only on generative AI, this suggests our estimates are generally more conservative than other studies.

These robustness checks provide reassurance that the numbers produced in this paper are consistent with other expert judgements and the emerging real-world evidence. However, as noted earlier, the figures rely on a forward-looking assessment of AI's potential, so both higher and lower numbers are possible. These figures should thus be treated as indicative of the scale of potential gains that could emerge, rather than a precise point forecast of what will happen.

Scenario Analysis: Key Assumptions and Supporting Evidence

Having calculated the potential time saving from the use of AI tools, to build our scenarios of how AI might impact the UK labour force and economy more widely we needed to make a

---

[13] A recent DfE (2024) report indicated that using generative AI tools can save multiple hours of teachers' time per week. Given that teachers work 50 hours a week on average, based on DfE statistics, and "multiple" implies at least two hours, this equates to at least a 4 per cent time saving. Meanwhile, Oak National Academy has shown that teachers can save at least four hours a week by deploying AI tools, which equates to an 8 per cent time saving.
[14] b3816c4fd6b7e92d301bf034753f465be334bb7c.pdf (thenational.academy)
[15] https://www.mckinsey.com/industries/education/our-insights/how-artificial-intelligence-will-impact-k-12-teachers
[16] https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html
[17] https://arxiv.org/pdf/2303.10130
[18] https://ippr-org.files.svdcdn.com/production/Downloads/Transformed_by_AI_March24_2024-03-27-121003_kxis.pdf

range of assumptions around the speed and distribution of the rollout of the various types of AI tool across the economy. To design our assumptions, we drew upon the extensive evidence and literature regarding previous patterns of technological diffusion as well as expert insights around the fixed and continuous costs related to various forms of AI-based technologies.

ASSUMPTION 1: WHAT SHARE OF THE POTENTIAL TIME SAVINGS FROM AI WILL BE REALISED IN PRACTICE?

It is possible that AI could deliver even greater time savings than we have estimated. Our assessment is based on the current capabilities of AI, but AI technology is advancing all the time. For example, further advances in driverless vehicles, video creation and humanoid robots may well lead to AI being able to assist with other tasks too. Greater time savings might therefore be possible

Equally, it is possible that not all the time savings identified in our analysis may be realised in practice. Firms may not find it worthwhile to use AI tools to partially replace human workers as it may not be cost effective in all cases. Svanberg et al (2023) show that, for AI tools that require high-cost computer-vision technology, if bespoke AI solutions have to be developed by each firm, only a minority of time savings – 23 per cent – are likely to be economically viable to realise, barring significant reductions in the cost of the technology.[19] If only larger firms (those with more than 500 employees) are able to economically exploit AI tools which fall into the categories of “bespoke AI systems” or “high-cost equipment and robotic AI”, the overall time saving falls from 23.8 per cent of total time worked by private-sector workers to 15.8 per cent. Our “breeze” scenario incorporates this reduction in the amount of workers’ time that will ultimately be saved by using AI tools.

[19] https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4700751

Figure A1: Contribution to overall time savings by firm size and technology type

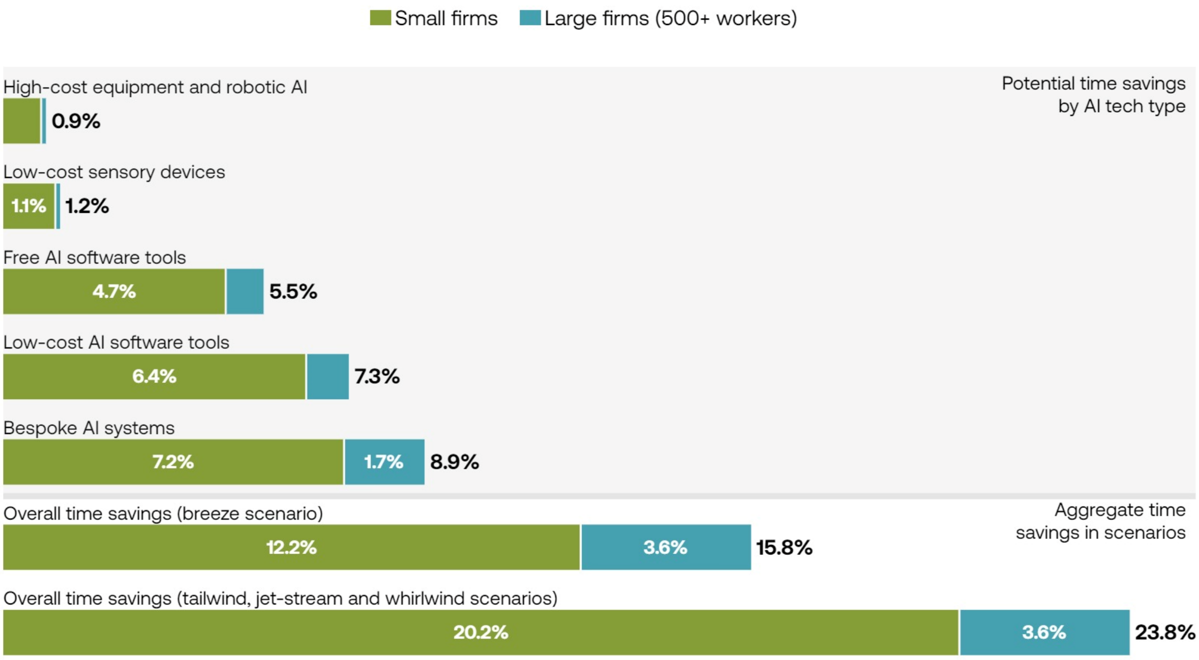

Source: TBI analysis using LFS and O*NET data

Svanberg et al also shows that tech adoption is likely to be more prevalent if it is possible to develop solutions at an industry or sector level, for example, if specialist AI startups can develop applications that can then be applied to individual firms at minimal cost. In this case, almost all the time savings can be realised. We consider this latter case to be more realistic given recent trends in technological diffusion, such as the market for cloud computing. We therefore proceed with the assumption that all the time savings discussed above are achievable in our “tailwind”, “jet-stream” and “whirlwind” scenarios while acknowledging there are both upside and downside risks to this assumption.

ASSUMPTION 2: HOW QUICKLY WILL AI BE ROLLED OUT ACROSS THE ECONOMY?

Predicting when AI will start to have a significant impact on the economy is inherently difficult. Nevertheless, some bold forecasters have started to incorporate the potential impact of AI into their forecasts for economic growth. For example, recent Goldman Sachs forecasts estimate that AI is unlikely to have a significant effect on the economy until at least 2027 but has the potential to significantly boost growth in the decade after that. [20]

[20] https://www.goldmansachs.com/insights/articles/ai-may-start-to-boost-us-gdp-in-2027

Other studies that have examined the potential for AI to boost growth have not put a timescale on when this might occur.[21]

We are of course already beginning to see the impacts of AI and its effects in the workplace as some experimental studies have already demonstrated its effectiveness.[22] Many workers are also already using “free” AI tools such as ChatGPT, which in March generated 1.8 billion visits.[23] However, these visitor numbers are volatile, which likely reflects the fact we are still in the experimental phase of AI adoption. Reaching the predicted time savings reported in our paper and others will require a critical mass of users of ChatGPT and other tools, to move beyond experimentation to regular use and deployment within their work processes. This process of experimentation eventually leading to deployment is reflected in the history of technological adoption which tells us that it can take a long time for the impact of new inventions to be fully felt and to appear in productivity and growth statistics.[24] Firms need to first experiment to work out how to best reorganise their production to take advantage of new technologies.

That said, there are good reasons to believe that the AI rollout will be relatively rapid. Much of the digital infrastructure that firms require to exploit AI is already in place. And we can observe that successive waves of innovation have been rolled out across the economy more quickly. The steam engine took more than a century for its effects to be fully felt, electricity more than 60 years and the internet between 30 and 40 years.[25] In the 21st century, innovations in production and globalisation have enabled new technologies to diffuse at even faster rates with smartphones, containing hundreds of thousands of times more computing power than the Apollo 11 mission,[26] taking only eight years from first introduction to reach the point where the majority of the US population owned one.[27] Early evidence suggests that the rollout of AI may be similarly rapid, with ChatGPT the fastest-growing consumer application in history, reaching an estimated 100 million monthly active users within just two months of launching.[28] Moreover, AI use is not limited to consumers - many of the world’s largest non-tech companies including Novo Nordisk, Eli Lilly and JPMorgan Chase are among the biggest investors in AI-era technology within their sectors. For example, JPMorgan Chase invests $15 billion a year in technology[29] and now employs more than 2,000 AI experts.[30] Given these promising early indicators, it seems reasonable

[21] www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2024/01/14/Gen-AI-Artificial-Intelligence-and-the-Future-of-Work-542379.
[22] https://www.science.org/doi/10.1126/science.adh2586
[23] https://explodingtopics.com/blog/chatgpt-users
[24] https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html
[25] https://econpapers.repec.org/bookchap/oxpobooks/9780199290895.htm
[26] https://www.realclearscience.com/articles/2019/07/02/your_mobile_phone_vs_apollo_11s_guidance_computer_111026.html
[27] https://news.microsoft.com/on-the-issues/2019/07/16/telephones-television-adoption-broadband/
[28] https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/
[29] https://www.jpmorgan.com/technology
[30] https://www.jpmorganchase.com/ir/annual-report/2023/ar-ceo-letters

to expect AI’s impact will continue to grow and should begin to show up in macroeconomic statistics by the end of this decade.

While we are still in the nascent stages of AI adoption, the available evidence suggests that the speed of implementation will almost certainly vary by firm size. Recent survey evidence from British firms points to marked differences in the speed of digital-tool adoption by smaller SMEs compared with even slightly larger firms.[31] Larger firms benefit from economies of scale that make investments with high upfront fixed costs more cost effective for them. It is also easier for large firms to take risks and invest in new forms of technology such as AI – if investments fail to be productive these firms have a larger cushion to fall back on. For small firms, upfront costs involved in deploying new technologies will represent a larger portion of revenues, potentially making investments, especially in higher-cost forms of technology, prohibitively expensive.

One paper that has explored this dichotomy is Svanberg et al (2024), which looks at models of the potential adoption of computer-vision technology (a high-cost hardware enabled by AI) and in particular the differences between large and small firms. The paper argues that it is likely to take 20 years or more for this kind of high-cost AI-enabled hardware to develop to the point where it is economically viable to be implemented in smaller private-sector companies.[32]

While at first the costs of these forms of tools are expected to be prohibitively high for all but the largest firms, over time costs will fall, or specialist AI providers will develop “AI-as-a-service” options allowing smaller firms to access these tools too. Under either scenario of fast platformisation or rapidly reducing costs, Svanberg et al demonstrate the possibility for even the rollout of complex AI hardware to be complete within 20 years. But they also show that rollout could be faster or slower than this if costs decline faster or slower, or if platformisation is faster or slower than this. If costs fall very quickly or there is 20 per cent compound growth in platformisation, they show that close to full coverage of AI tools could be possible within 10 years. On the other hand, if both are slower, it could take 30 years or more for all possible applications of AI-enabled computer vision to be economically viable.

Bringing this all together, in all our scenarios we assume that it takes longer for smaller firms to adopt all kinds of AI than larger firms (those with 500 or more employees). However, we assume that this lag varies by technology type.

---

[31] https://www.xero.com/uk/media-releases/digital-drag-report-uk/

[32] Maja Svanberg et al, "Beyond AI Exposure: Which Tasks Are Cost-Effective to Automate with Computer Vision?", *SSRN*, (2024), doi:10.2139/ssrn.4700751.

- For “free” AI, off-the-shelf AI software tools and AI-enabled low-cost sensory devices we assume smaller firms only experience a small lag in AI adoption relative to larger firms, reflecting their historically more cautious approach of waiting for the benefits of various tools to be proven by larger firms before they take steps to adopt them. Once benefits are proven, however, we assume that the speed of adoption between the different firm sizes will be only slightly slower within small firms given the low barriers to entry (low fixed costs) involved in adopting technologies of this type. In our tailwind scenario, the rollout of these technologies is expected to reach 60 per cent across firm types by 2040 for most forms of low-cost AI, with the rollout of “free” AI expected to have reached more than 75 per cent by this point. By 2050 the adoption of both these kinds of low-cost tools is essentially complete.

- For more costly AI tools – including bespoke AI solutions and high-cost AI equipment – we anticipate a much larger difference in rollout speed between large and small firms. In our tailwind scenario only 65 per cent of small firms are fully utilising these high-cost forms of AI by 2050 compared with more than 95 per cent for larger firms. By contrast, in our jet-stream and whirlwind scenarios the rollout happens more quickly, in line with Svanberg et al’s most optimistic case; 90 per cent of small firms use the tools by 2050 while 90 per cent of large firms achieve the same feat by 2040. By contrast in our breeze scenario, rollout of these high-cost technologies is assumed to be zero in 2050 among smaller firms as we assume they are too costly to implement. This compares with 80 per cent deployment for larger firms by 2050, which is in line with the more pessimistic scenarios presented by Svanberg et al (2024).

Figure A2: Expected speed of AI rollout for each scenario – all firms

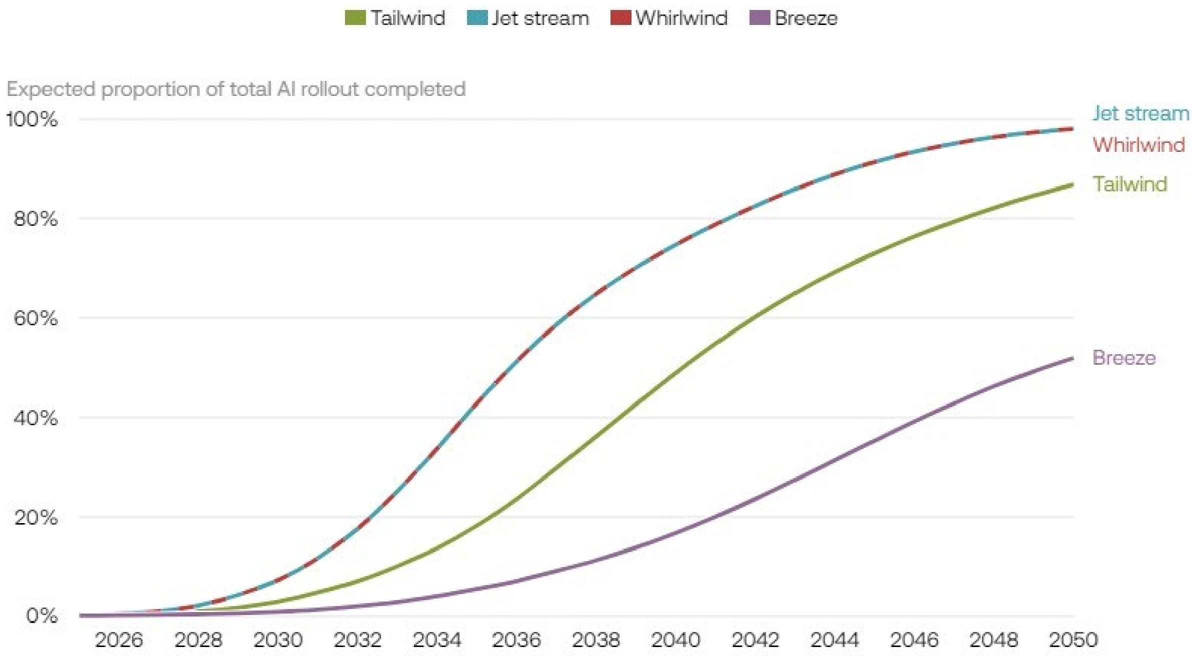

Source: TBI analysis

ASSUMPTION 3: HOW MANY WORKERS WILL BE MADE REDUNDANT DUE TO AI ADOPTION?

There is no occupation in our data that can be fully replaced by AI – human input will still be required for some tasks in all occupations. But it is nonetheless likely that firms will use at least some of the time savings generated by the use of AI tools to reduce employment. For example, if a firm has two workers of a particular occupation who can use AI to perform their jobs in half the time, the firm could retain just one AI-enabled worker to produce the same amount as two workers without AI. Using this logic, if firms were to use all the time savings to reduce headcount, nearly six million private-sector jobs would be at risk.

In practice, however, we expect the number of redundancies to be far fewer than this. AI will make workers more productive, making them more valuable to their employers and giving opportunities for growth by retaining them. Furthermore, this productivity boost will raise demand for goods and services across the economy, increasing the demand for labour. Finally, firms may choose to retain workers to help develop new products, some of which may themselves be enabled by AI.

To ascertain a credible estimate of the extent to which firms will use time savings to shed workers, we take several approaches. First, we examine the approach taken by other studies examining the labour-market impacts of AI. These other studies tend to estimate that between one-quarter and one-third of worker time savings will be used to reduce worker numbers. Briggs and Kodnani (2023) estimate that 7 percentage points of their estimated 25 per cent time saving will be used to reduce labour input. This is based on their estimate of the share of occupations where time savings are greater than 50 per cent.[33] The equivalent using our time-saving estimates, which are typically smaller (as we do not assume that AI can reduce the time taken by 100 per cent in all instances), would be around 15 per cent. A recent report by IPPR uses a central scenario where around a third of potential time savings are used for job displacement rather than labour augmentation.[34] Similarly to Briggs and Kodnani,[35] this is based on occupations with a greater than 40 per cent exposure to AI leading to job displacement and those with a less than 40 per cent exposure leading to augmentation only. Applying a similar threshold in our analysis would imply that around 25 per cent of time savings led to job displacement.

Another way of ascertaining likely firm responses is to examine how firms respond to reductions in demand for goods and services during recessions. Previous research has found that firms do not reduce employment one-for-one in response to a reduction in demand: instead, they tend to hoard labour to avoid firing-and-hiring costs when the economy recovers. Arpaia and Curci (2010)[36] find that in large European economies, employment has only fallen by between a quarter and a third as much as GDP at the start of recent recessions. Although this is clearly a different situation to the one we are studying here – a temporary reduction in demand rather than a permanent increase in labour productivity – it is nonetheless informative about likely responses and confirms the reasonableness of these previous estimates.

It is also possible that firms' responses will vary between firms of different sizes. Smaller firms may not have sufficient time savings to make one person redundant, so they may choose to expand production rather than reduce labour input. By contrast, larger firms with more staff of a particular occupation would find it easier to make redundancies. If large firms (those with more than 500 employees) were to use any time savings to reduce headcount (that is, a 100-per-cent shedding for large firms), and small firms (with fewer than 500 employees) used time savings to increase output, only around a quarter of time savings would lead to job displacement (Table 1).

---

[33] Joseph Briggs and Devesh Kodnani, *The Potentially Large Effects of Artificial Intelligence on Economic Growth*, Goldman Sachs (2023) www.ansa.it/documents/1680080409454_ert.pdf.

[34] https://www.ippr.org/articles/transformed-by-ai

[35] https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html

[36] Arpaia, Alfonso and Curci, Nicola, *EU Labour Market Behaviour during the Great Recession*, European Commission (2010) https://data.europa.eu/doi/10.2765/39957

Table 1: Summary of possible labour-shedding assumptions

Possible assumptions	Share of exposed labour displaced (Using distribution and levels of exposure in our model)	Estimated total number of workers displaced
a) Full shedding: Firms choose to use all AI-related time savings to reduce the size of their workforce	100 per cent	6 million
b) Shed only those workers where at least half of their time can be saved by AI (Briggs and Kodnani, 2023)[37]	15 per cent	840,000
c) Shed only those workers where at least 40 per cent of their time can be saved by AI (IPPR, 2024)[38]	25 per cent	1.5 million
d) Firms shed labour in line with how they normally respond to economic downturns	25–30 per cent	1.5–1.8 million workers
e) If large firms shed workers but small firms retain them	25 per cent	1.5 million
Scenario assumptions		
Tailwind and jet stream	25 per cent	1.5 million
Breeze	25 per cent (but applied to the partial level of AI adoption assumed in this scenario)	1 million
Whirlwind	50 per cent	3 million

Taking all of this evidence together, the most likely outcome appears to be around one-quarter of time savings will be used to reduce headcount and this is the baseline assumption we use in our tailwind, jet-stream and breeze scenarios. By contrast, in the whirlwind scenario we assume half of all time savings will be used to reduce headcount – an extreme assumption based on the evidence above, but one that is designed to illustrate the implications of AI causing greater labour-market disruption.

Overall, these figures imply 1.5 million people will be made redundant in the tailwind and jet-stream scenarios, 1 million in the breeze scenario, and 3 million in the whirlwind scenario.

[37] https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html
[38] https://ippr-org.files.svdcdn.com/production/Downloads/Transformed_by_AI_March24_2024-03-27-121003_kxis.pdf

Crucially, these redundancies will not all occur at once. Rather, they will occur gradually in line with the pace that AI is rolled out, with the peak impact coming at the time the rollout is fastest.[39] Combined with our timing assumptions outlined above, this means that in our central tailwind scenario redundancies peak at just under 100,000 a year, with a range from 60,000 (breeze scenario) to 270,000 (whirlwind). To put these figures in context, there are around 25 million private-sector workers, and there have been on average 450,000 redundancies a year across the whole economy over the last decade,[40] so even our most pessimistic scenario involves less than doubling the number of redundancies compared to a typical year and an annual risk of redundancy of around 1 per cent.

Figure A3: Annual number of redundancies in each of the four scenarios

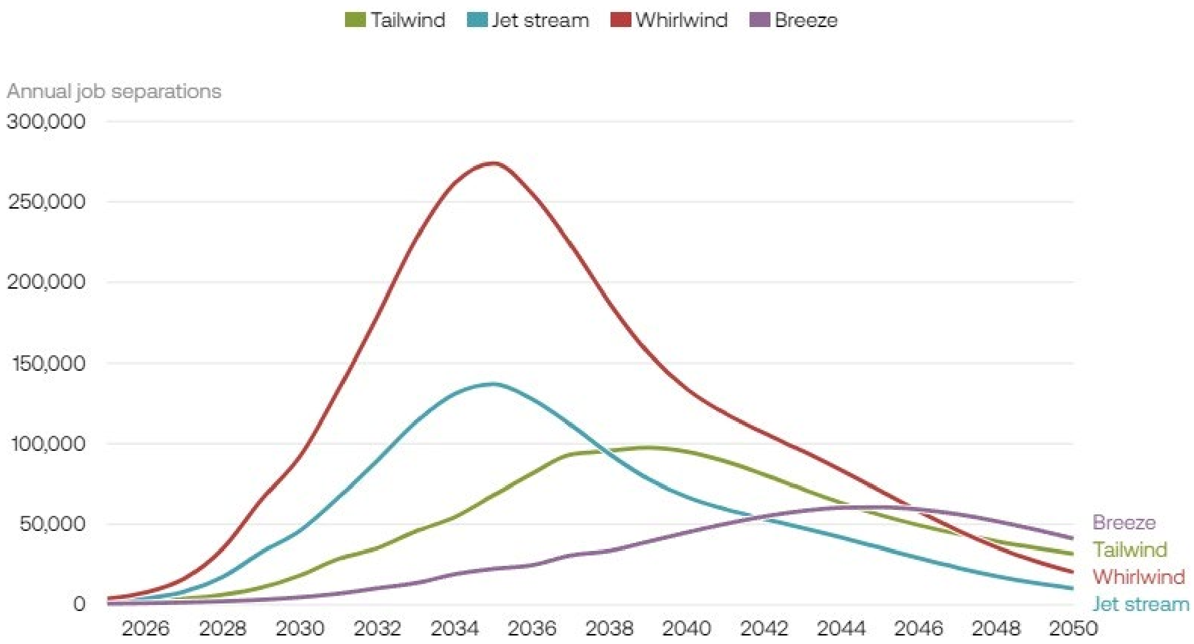

Source: TBI analysis using LFS and O*NET data

ASSUMPTION 4: HOW QUICKLY WILL DISLOCATED WORKERS BE RE-EMPLOYED?

Historically, technological unemployment has tended to be offset over time by increases in demand for labour from two sources. First, higher productivity leads to higher demand for goods and services and hence higher demand for labour. Of course, some of this higher demand will be used to soak up the increased production of firms that do not use all the time savings from the use of AI tools to reduce their labour input. But the remainder would

[39] That is to say, the number of redundancies in a given year is in proportion to the share of the rollout that occurs in that year.
[40] Source: ONS redundancy data using LFS.

lead to higher labour demand. Following Morén and Wändal (2019),[41] in our analysis a 1 per cent increase in productivity leads to a 0.57 per cent increase in demand for labour – roughly in line with the labour share of income. In all our scenarios, this occurs with a one-year lag between higher productivity and higher labour demand: unemployment typically lags changes in aggregate demand as firms are reluctant to hire more workers until they are convinced that the rise in demand for their product is not just temporary.

Increased labour demand will also come about through new tasks facilitated by AI itself. Observing past patterns in the data – specifically, comparing US productivity growth over time from the US Bureau of Labor Statistics and estimates of new task creation from Acemoglu and Restrepo (2019) – we find that in the past, increases in labour demand from the introduction of new tasks tend to occur around a decade after productivity increases, and that in the past a 1 per cent increase in productivity growth is associated with a 2.5 per cent increase in demand for labour through new task creation.[42]

But there is good reason to think that AI will start to generate new tasks for workers to do more quickly than this. In line with the accelerating speed of rollout of general-purpose technologies discussed above, we reduce this to five years in our tailwind scenario. However, it may occur even more quickly. There is evidence that AI is already boosting labour demand; firms that have invested in AI have seen higher employment growth[43] and Amazon has increased employment even as it has adopted AI. AI has generated new roles for workers in the repair and maintenance of robots and for "quarterbacks" who control the movement of robots on the warehouse floor. Of course, this is not evidence that AI has boosted employment overall – firms that have been early adopters of AI may simply be taking market share from those who have not, and it is in any case too early to give a definitive answer to this question – but this evidence is at least suggestive that the overall impacts of AI on unemployment may be smaller than those we show in our tailwind scenario. In our jet-stream scenario we therefore reduce the lag between productivity growth and new task creation to three years. Similarly, to illustrate what could happen if the lag is not smaller than during previous waves of technological improvement, in our breeze and whirlwind scenarios we keep the lag at ten years.

These re-employment effects reduce the stock of unemployment that builds up through job separations. In each scenario, the impact on unemployment goes negative once these job-creation effects kick in, in other words, AI could become a net job creator. However, given constraints on the number of potential workers available, we do not assume that overall

---

[41] Victoria Morén and Elias Wändal, "The Employment Elasticity of Economic Growth" (2019). https://gupea.ub.gu.se/bitstream/handle/2077/61745/gupea_2077_61745_1.pdf
[42] https://www.aeaweb.org/articles?id=10.1257/jep.33.2.3
[43] https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3651052

employment increases, but rather that wages and the share of national income received by workers increase once this point is reached.[44]

Overall, then, in all four scenarios unemployment rises initially as redundancies occur when AI is rolled out but then it falls back again as demand for workers to perform new tasks facilitated by AI starts to be created. Figure A4 shows how this occurs in the tailwind scenario, while Figure A5 compares unemployment effects across all four scenarios. The peak impact on unemployment will thus depend on the extent to which firms respond to time savings by shedding staff and how quickly new task creation occurs. The whirlwind scenario, with much higher levels of job shedding and a longer lag before new task creation starts to raise demand for workers, thus involves the highest unemployment levels.

Figure A4: Unemployment impacts of AI, tailwind scenario

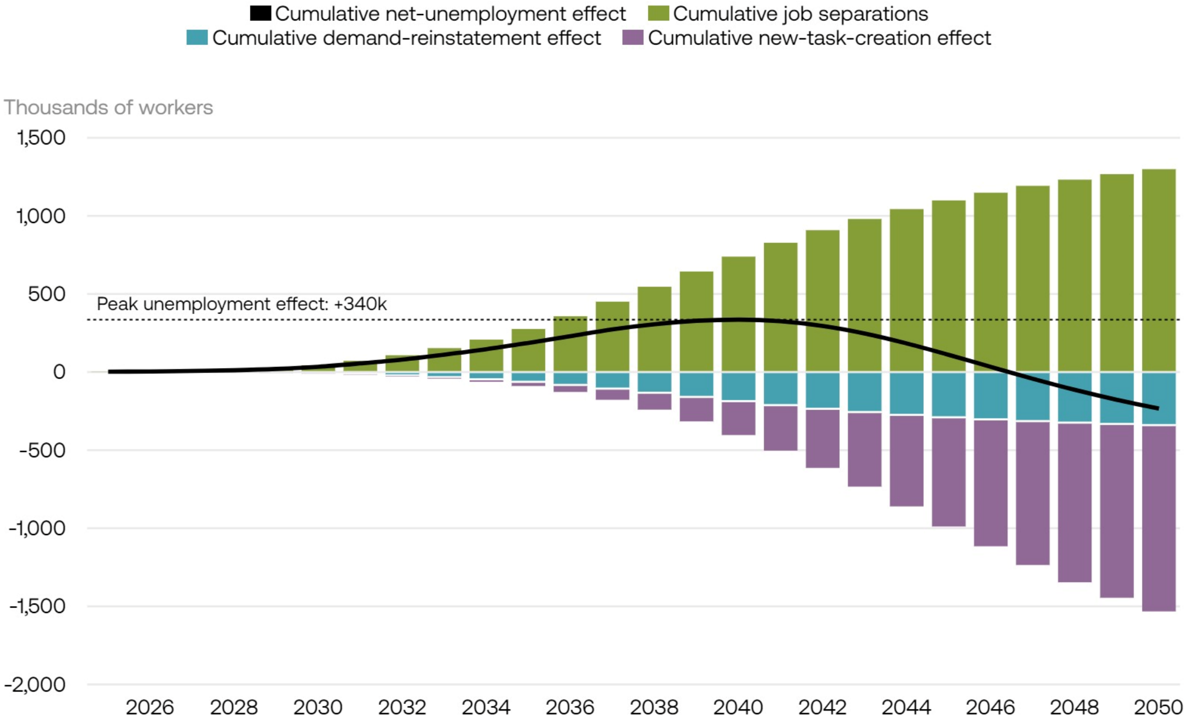

Source: TBI analysis using LFS and O*NET data

[44] Note that job separations are still occurring at this point, so in principle there could still be additional "frictional" unemployment as displaced workers take time to find new jobs. We implicitly assume that any increases in frictional unemployment caused by ongoing job separations are offset by reductions in frictional unemployment from better job matching (see section on Job Matching and Labour-Market Efficiency).

Figure A5: Impact of AI on unemployment across all four scenarios

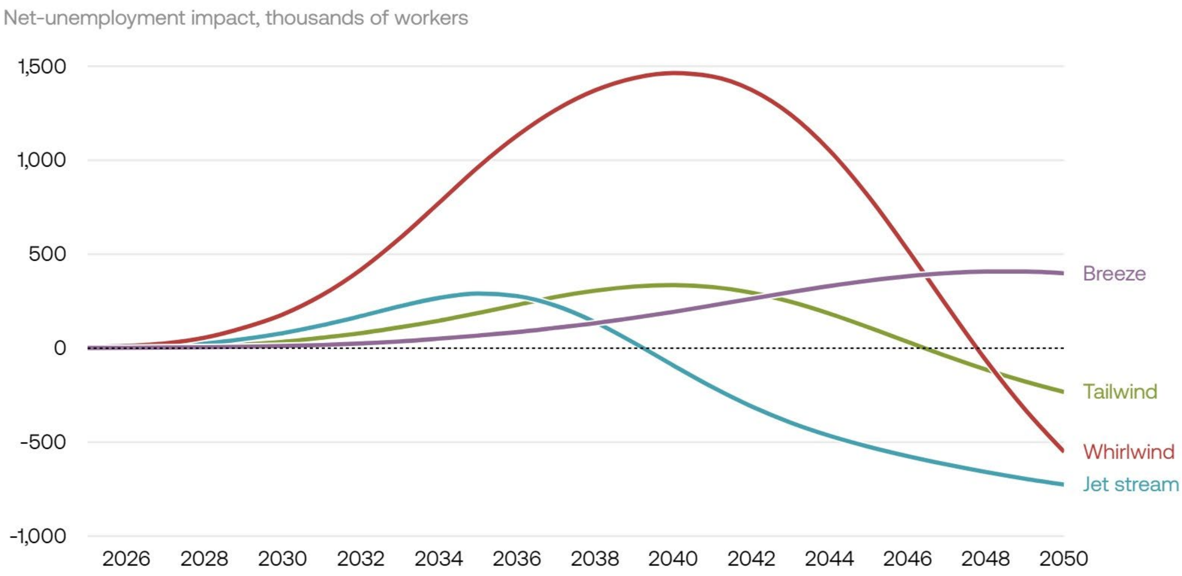

Source: TBI analysis using LFS and O*NET data

How Will AI-Related Time Savings Affect GDP?

AI-related time savings are likely to affect GDP through three main channels: higher productivity from the use of AI tools; higher investment in machinery which will also boost output; and changes in employment levels discussed above.

PRODUCTIVITY EFFECTS

We follow Acemoglu (2024),[45] who shows that the impact of AI on total factor productivity can be estimated using the following formula:

$$\Delta \log TFP = labour\ share\ \times share\ of\ cost\ saved\ \times percentage\ cost\ saving$$

That is to say, the increase in productivity from AI depends on the share of GDP that is affected by automation (the labour share times the earnings-weighted proportion of workers'

[45] Daron Acemoglu, "The Simple Macroeconomics of AI", NBER Working Paper 32487 (2024) https://www.nber.org/papers/w32487

time saved) multiplied by the cost saving that automation yields. We have calculated the time savings already; the figure for cost saving in the private sector as a share of employment costs across the whole economy (including the public sector) is 19.3 per cent in our tailwind, jet-stream and whirlwind scenarios and 12.5 per cent in our breeze scenario. The productivity effect thus chiefly depends on how much overall cost reduction can be achieved by replacing some of workers' time by using AI tools. The intuition for this is that productivity improvements come about by replacing labour with cheaper capital. We detail below the specific cost assumptions we have made for each of the five AI technology types and how these vary by firm size and across different scenarios. Overall, we estimate substantial cost savings – ranging from 60–80 per cent in aggregate across the scenarios. This suggests that AI is unlikely to be another "so-so" technology that only offers marginal cost savings which leads to jobs being lost but does not increase productivity.

COST SAVINGS

The cost-saving estimates utilised in this report are based on expert best judgement and real-world evidence. In our central tailwind scenario, we have used relatively conservative cost assumptions across each of the five AI technologies – detailed below – which we then vary for the alternative scenarios.

"Free" AI Tools

While labelled as "free", we conservatively estimate that implementing "free" AI tools will still involve certain costs, such as training employees and integrating these tools into existing workflows. These costs could be relatively low if no or minimal training is required, and employees can integrate the tools organically. However, costs might be significantly higher if substantial investments are needed to develop the technical capabilities to use these tools effectively, such as data upgrades or extensive training.

Evidence from a study of BCG workers found only very minor differences in task completion and quality for those who received some training in how to use ChatGPT before being assigned a task versus those who did not.[46] This finding suggests that for most uses of free AI tools it would be appropriate to assign only minimal costs to adoption. In our tailwind scenario we assume a relatively conservative cost-savings level of 95 per cent. Based on the average wage rate of workers exposed to this technology type (and the estimated levels

---

[46] https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf

of time saving: 5.5 per cent), this implies that the cost of deploying this technology is just over £100 per worker on average, or enough to cover roughly an-hour-and-a-half of training for every exposed worker or a full day of training each year for around every 1 in 5 workers.[47] We model that cost savings could rise to 98 per cent in our whirlwind and jet-stream scenarios, where AI is expected to be more user-friendly allowing for seamless integration and lower investment in training. Conversely, in our breeze scenario, cost savings fall to 80 per cent, reflecting greater challenges in implementing even free forms of AI.

We do not anticipate significant differences in cost savings from free AI tools between large and small firms. While large firms may benefit from better technological preparedness, they could be hindered by more rigid structures and heightened privacy concerns, which might limit the use of these tools – issues that smaller firms may encounter to a lesser extent.

Low-Cost AI Software Tools

Various relatively low-cost AI tools are already available on the market, such as Copilot Pro for general tasks (£228 per user per year),[48] Jasper.ai for copywriting (around £290 per user per year),[49] Calendly for scheduling (£92 per user per year),[50] and Otter.ai for meeting transcription (around £190 per user per year).[51] We estimate that these tools could save 7.3 per cent of time across the private sector. Given average wage rates for private-sector workers exposed to this form of AI tool, this time saving equates to approximately £2,700 per worker annually. In our tailwind scenario, we assume that using these tools could lead to cost savings of 85 per cent, translating to a spend of just under £400 per user. This budget would comfortably cover several AI subscriptions per worker exposed to these technologies, along with any necessary training or integration costs. In our breeze scenario we use a lower cost-saving estimate of 70 per cent, reflecting the possibility that firms may need to invest in additional or more expensive AI tools to achieve the desired time savings for more complex tasks. In our jet-stream and whirlwind scenarios, we project cost savings of around 90 per cent, considering the potential for AI tool prices to decrease due to increased competition and further technological advancements, enabling firms to achieve the same time savings at lower costs. We also assume that there are no significant

[47] Training costs per day estimated using figures from the 2022 Employer Skills Survey https://www.gov.uk/government/statistics/employer-skills-survey-2022
[48] £19 per month at time of writing: https://www.microsoft.com/en-gb/store/b/copilotpro?msockid=152984e0b14a60d0384590d4b02961c3
[49] https://www.jasper.ai/pricing
[50] https://calendly.com/pricing
[51] Based on a business subscription at $20 per month: https://otter.ai/pricing

differences in cost savings between large and small firms when adopting these tools as they are freely purchasable on the open market.

Bespoke AI Tools

Bespoke AI tools and systems, such as internally trained chatbots, AI customer-service agents and financial monitoring systems, inherently require access to significant amounts of firm-specific data and significant investment in model building and training. Large financial institutions with their substantial internal data and assets have been early adopters of these technologies. For instance, Bank of America developed its internal chatbot, Erica,[52] to assist customers with banking needs, while BlackRock created its own AI to assist with investment decisions, which it refers to as a Thematic Robot tool.[53] Other examples include retail giants like Walmart and consulting firms like McKinsey, as outlined in the main text.

In our analysis of implementing AI within the public sector, we drew on expert advice which suggested that the cost of developing each bespoke AI tool would be £2.5 million (with yearly running costs of £250,000 for compute and fine-tuning) and assumed that each task identified would require a custom-built solution. Despite these high costs, the very large size of the public-sector workforce meant that when considering the large value of time saved by these tools, projected cost savings reached 98 per cent. We anticipate that introducing bespoke AI tools in the private sector will be somewhat more expensive due to the inability to achieve the same economies of scale as in the public sector.

We follow the model proposed by Svanberg et al (2024) to develop our costing assumptions. This paper proposed two potential scenarios: one in which all AI tools must be developed in-house at prohibitive cost levels for small firms and another within which AI-as-a-service firms develop, allowing AI tools to proliferate at lower cost and be adopted even by smaller firms.

Within our least optimistic scenario, breeze, we assume that this proliferation does not occur and all costs of tool development remain internalised. This assumption means that for small firms we assume no cost savings from the use of these tools and subsequently they are not adopted. For large firms we model cost savings of 55 per cent which, given wage rates, would be enough for roughly one entirely new tool every five years at each large firm.

[52] https://promotions.bankofamerica.com/consumer/ericabyyourside
[53] https://www.blackrock.com/us/individual/insights/ai-investing

Across our other scenarios we assume that platformisation occurs, allowing for cost-savings to be unlocked by small firms and reducing the costs faced by large firms as AI-provider firms do the heavy lifting of tool development, building off standard models and competing on price. We essentially model cost-saving levels as sensitivities around the entirely self-built tool scenario, modelling for large firms' cost savings of 70 per cent in our tailwind scenario and 85 per cent in our whirlwind and jet-stream scenarios. As these tools impose large, fixed costs, invariant to the number of workers affected, we model cost savings as lower for small firms even under this AI-as-a-service assumption with cost savings expected to reach 50 per cent in our tailwind scenario and 65 per cent in the more optimistic whirlwind and jet-stream scenarios.

Low-Cost Sensory Devices and AI

Low-cost sensory devices like microphones, cameras and radio-frequency identification (RFID) tags can be paired with AI to achieve significant time savings across various tasks. These tools can often be utilised in conjunction with free AI software or affordable subscription-based models, such as transcription services. Many firms may already possess these sensory devices, but they remain accessible at a low cost – for example, standard headsets can retail for as little as £7,[54] while RFID tags can cost just 3 pence each.[55] Even slightly more expensive sensory devices, like smartphones, can be purchased for around £100.[56] Even above this price point, they represent a cost-effective investment for companies implementing AI tools. For instance, Walmart now provides smartphones worth more than £300 to at least half of their store associates, enabling access to AI-powered apps on the go.[57] We generally assume that the costs of providing and maintaining these forms of hardware are estimated to be minimal and hence the cost savings from these kinds of AI-enabled hardware are expected to be comparable to those achieved through software alone.

Low-cost hardware can be paired with various forms of AI software but for the majority of tasks, our analysis suggests the software is likely to be of the free or low-cost subscription type. As a result, we estimate that the costs associated with these AI-enhanced tools will fall between those predicted for "free" and low-cost subscription forms of AI, leading to an estimated cost saving of 90 per cent in our tailwind scenario for both large and small firms.

---

[54] https://www.cromwell.co.uk/shop/office-supplies/headsets/24-1512-hp512-economy-stereo-headset-boom-microphone/p/CTL8030015N
[55] Multiple tags would likely need to be purchased to build an effective inventory system or similar, but this is not a major constraint given their exceptionally low unit cost. https://www.mckinsey.com/industries/retail/our-insights/rfids-renaissance-in-retail
[56] https://www.amazon.co.uk/Samsung-Galaxy-A15-Black-128GB/dp/B0CR77BGWB/ref=asc_df_B0CR77BGWB/
[57] https://edition.cnn.com/2021/06/03/business/walmart-employees-samsung-phone/index.html

In our breeze scenario cost savings are lower at 75 per cent, potentially due to a greater reliance on paid software or higher-than-expected hardware costs. Conversely, in our jet-stream and whirlwind scenarios, we estimate cost savings of up to 95 per cent, reflecting lower additional capital investment and/or more affordable AI software to complement these tools. We assume no significant difference in cost-effectiveness between small and large firms, as these costs are generally per user and do not substantially benefit from economies of scale.

AI-Enabled High-Cost Equipment and Robotics

The expected cost savings from high-cost equipment and robotics are, given current market conditions, generally expected to be much lower compared with other forms of AI. Nevertheless, these technologies have already begun to make an impact, particularly among large firms. For example, Amazon's deployment of robots in its warehouses reportedly reduced operating costs by 22 per cent as early as 2016.[58] Given the rapid advancements in robotics, these savings may have since increased further, with reports on Amazon's "Digit" robots suggesting that operating costs have been declining from just under £10 per hour per robot when introduced towards a predicted level of less than £2.50.[59]

Drawing on this evidence from Amazon, we estimate in our central scenario that robotics and other AI-enabled high-cost equipment could lead to cost savings of 25 per cent for large firms. For smaller firms, where economies of scale are less pronounced, we assume cost savings of just 10 per cent – because of small firms needing to buy the technology from larger firms (at a premium) and applying it to a smaller workforce. As with bespoke AI tools, we assume that in the breeze scenario small firms will see no time savings from these forms of AI at all. This is because all development costs would be internalised, making the investment prohibitively expensive, in line with the findings of Svanberg et al (2024).[60]

However, in our jet-stream and whirlwind scenarios, AI-as-a-service providers emerge, driving down costs through competition and making these technologies more accessible, even to small firms. In these scenarios, cost savings reach 40 per cent for large firms and 25 per cent for small firms.

---

[58] https://www.inc.com/betsy-mikel/amazons-secret-weapon-could-save-their-warehouses-25-billion.html
[59] https://earnyourleisure.com/news/business/amazons-humanoid-warehouse-robots-will-eventually-cost-only-3-per-hour-to-operate/
[60] https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4700751

INVESTMENT IMPACTS

Increases in productivity from using AI tools also impact the return on capital. Higher returns should lead to greater investment, which leads to further increases in the capital stock and even larger increases in GDP. In all our scenarios, the increase in the capital stock is sufficient to get the rate of return on capital down to its original level.[61] As Acemoglu and Restrepo (2018) argue, this is what one would expect in a standard economic framework with a representative household with exponential discounting and time-separable preferences.[62]

EMPLOYMENT LEVELS

The impact of AI on unemployment and hence employment levels is discussed in the previous section. We account for any increase in unemployment reducing GDP in line with the reduction in labour input. We again use a Cobb-Douglas production function with a labour share of 59 per cent.

OVERALL GDP IMPACTS

Average cost savings are around 70 per cent in our tailwind scenario based on a weighted average of savings across the five technology types, around 80 per cent in our jet-stream and whirlwind scenarios and 60 per cent in our breeze scenario. Using the formula from the start of this section and assumptions about the degree of time savings realised in 2050 for each scenario, this implies an overall impact on productivity of around 8 per cent in our tailwind scenario, 10 per cent in our jet-stream and whirlwind scenarios and 4 per cent in our breeze scenario.

These productivity gains then increase the capital stock by the same amount, which further raises GDP by the capital share (41 per cent) times these figures. The overall increase in GDP in 2050 is thus around 11 per cent in our tailwind scenario, around 14 per cent in our jet-stream and whirlwind scenarios and around 5 per cent in our breeze scenario.

---

[61] For all these calculations we use a standard Cobb-Douglas production function with labour share of 59 per cent.
[62] https://www.nber.org/system/files/working_papers/w24196/w24196.pdf

FOLLOW US

facebook.com/instituteglobal

twitter.com/instituteGC

instagram.com/institutegc

GENERAL ENQUIRIES

info@institute.global

Copyright © October 24 by the Tony Blair Institute for Global Change

All rights reserved. Citation, reproduction and or translation of this publication, in whole or in part, for educational or other non-commercial purposes is authorised provided the source is fully acknowledged. Tony Blair Institute, trading as Tony Blair Institute for Global Change, is a company limited by guarantee registered in England and Wales (registered company number: 10505963) whose registered office is One Bartholomew Close, London, EC1A 7BL.