THE ECONOMIC IMPACT OF GENERATIVE AI IN THE US

METHODOLOGICAL APPENDIX

DECEMBER 2023

ABOUT OXFORD ECONOMICS

Oxford Economics was founded in 1981 as a commercial venture with Oxford University's business college to provide economic forecasting and modelling to UK companies and financial institutions expanding abroad. Since then, we have become one of the world's foremost independent global advisory firms, providing reports, forecasts and analytical tools on more than 200 countries, 100 industries, and 8,000 cities and regions. Our best-in-class global economic and industry models and analytical tools give us an unparalleled ability to forecast external market trends and assess their economic, social and business impact.

Headquartered in Oxford, England, with regional centres in New York, London, Frankfurt, and Singapore, Oxford Economics has offices across the globe in Belfast, Boston, Cape Town, Chicago, Dubai, Dublin, Hong Kong, Los Angeles, Mexico City, Milan, Paris, Philadelphia, Stockholm, Sydney, Tokyo, and Toronto. We employ 600 staff, including more than 350 professional economists, industry experts, and business editors—one of the largest teams of macroeconomists and thought leadership specialists. Our global team is highly skilled in a full range of research techniques and thought leadership capabilities from econometric modelling, scenario framing, and economic impact analysis to market surveys, case studies, expert panels, and web analytics.

Oxford Economics is a key adviser to corporate, financial and government decision-makers and thought leaders. Our worldwide client base now comprises over 2,000 international organisations, including leading multinational companies and financial institutions; key government bodies and trade associations; and top universities, consultancies, and think tanks.

December 2023

The modelling and results presented here are based on information provided by third parties, upon which Oxford Economics has relied in producing its report and forecasts in good faith. Any subsequent revision or update of those data will affect the assessments and projections shown.

To discuss the report further please contact:

Vasilis Douzenis: vdouzenis@oxfordeconomics.com

Oxford Economics

4 Millbank, London SW1P 3JA, UK

Tel: +44 203 910 8061

TABLE OF CONTENTS

1. INTRODUCTION AND OVERVIEW

Oxford Economics was commissioned by Cognizant to assess and forecast the economic impact of generative artificial intelligence (Generative AI) technology on the United States over the next 10 years. This document provides a comprehensive description of our methodological approach and an explanation of the associated rationale.

The project was executed through five main phases of work that sequentially enabled us to develop assumed inputs for Oxford's Global Economic Model (GEM). These inputs are designed to reflect how we anticipate that the implementation of Generative AI technologies by businesses in the United States over the forecast horizon will influence structural drivers of economic activity, notably total factor productivity (TFP) growth. To reflect the uncertainty inherent in such a process we have run three scenarios to develop a range of outcomes.

Fig. 1. Phase-by-phase project methodology

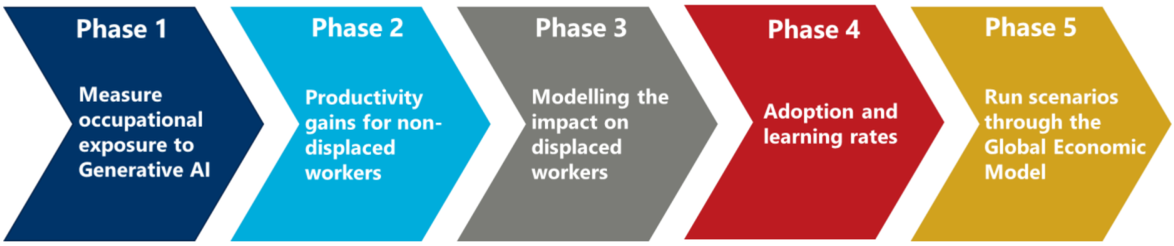

The remainder of this document outlines the process we followed in detail.

Section two discusses how we assessed the potential impact of Generative AI on automating workplace tasks and the impact this will have on productivity and employment. This was informed by a machine learning model was built to assign an automatability score to all workplace tasks and the scores were then aggregated to obtain theoretical occupational exposure scores for all US jobs.

Section three describes how we formed assumptions for the rate of adoption of Generative AI by US companies in each scenario. The adoption rates were used to adjust the theoretical occupational exposure scores to projected exposure scores in each year of the forecast period. These scores were used to estimate the proportion of the workforce that would be displaced in each occupation.

Section four outlines how we investigated the implications of task automation for the productivity of workers who remained in their jobs and, separately, for those who were displaced. Our research sought to account for the transitional costs that will be created by displaced workers taking time to find a new job, the extent to which any workers might permanently exit the labour market and due to the integration of new technology by non-displaced workers.

Finally, as described in section five, the aggregated impacts of Generative AI on productivity and employment were modelled as shocks applied to Oxford's Global Economic Model (GEM) to estimate the contribution of the technology to US GDP growth above Oxford's baseline forecasts.

While modelling labour market adjustments in response to the displacement caused by Generative AI was outside this study's scope, we conducted quantitative analysis to understand the degree of difficulty displaced workers may face when trying to regain employment. This is covered in section six.

2. MEASURING OCCUPATIONAL EXPOSURE TO GENERATIVE AI

2.1 O*NET DATABASE

The US O*NET program is a comprehensive system for collecting and disseminating information on occupational and worker characteristics, sponsored by the US Department of Labor and Employment and Training Administration. At the center of the program is the O*NET database, which contains detailed information on the skills and requirements for around 1,000 occupations. All of these data are structured according to the O*NET "content model", which defines the key features of an occupation as a standardized set of variables called "descriptors". These descriptors are designed to provide an exhaustive list of the characteristics and requirements for any given occupation. Information is organized across six "domains" which comprise both worker-oriented characteristics and job-oriented characteristics (Fig. 2).

Much of the information is collected via self-reported assessments by existing employees using standardized questionnaire surveys, and is supplemented by professional assessments by job evaluation analysts. These data collection methods are undertaken on an ongoing basis, enabling information to be updated regularly. Indeed, the O*NET database now spans over 20 years.

Fig. 2. O*NET content model[1]

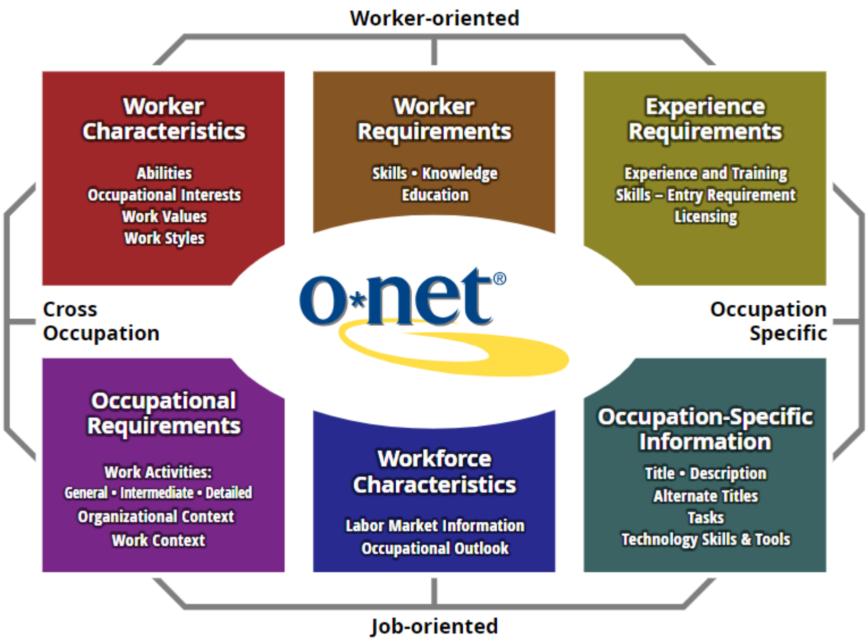

For the purposes of the machine-learning model developed in this project, we used data from the "Tasks" domain. This domain presents a list of specific detailed tasks that are relevant for each

[1] O*NET, "The O*NET Content Model", accessed November 2023

occupation covered in the O*NET database. Occupations in the US economy collectively undertake over 18,000 unique tasks. The number of relevant tasks for each occupation range from 4 to 40. Each task is assigned an importance score which indicates its significance to the occupation, ranging from "Not important" (1) to "Extremely important" (5).

For example, a Marketing Manager has 20 relevant tasks. The most important of which (with an importance score of 4.3) is to "identify, develop, or evaluate marketing strategy, based on knowledge of establishment objectives, market characteristics, and cost and markup factors."

2.2 CLASSIFYING OCCUPATIONAL TASKS

The approach requires us to classify over 18,000 O*NET occupational tasks into categories that reflect their maximum potential for 'automation' through Generative AI.

We seek to classify each task in O*NET based on its potential for automation using the current capability of Generative AI (2023) and its potential capability in 10 years' time (2032). To this end, we defined a set of five distinct categories each reflecting a different potential exposure to Generative AI.

Fig. 3. Generating exposure buckets to classify tasks according to their potential for automation.

Maximum potential for automation	Potential exposure to automation through Generative AI
Not automatable	0%
Little assistance	25%
Some assistance	50%
Mostly assistive	75%
Fully automatable	100%

Since it would have been too time-consuming to manually label 18,000 tasks according to each exposure category, we instead developed a machine learning classification model that focuses on natural language classification.

These types of estimators are supervised machine learning models and require labelled inputs to learn semantic patterns in textual data. A sample of 1,400 tasks was manually labelled and used to train the model. To ensure the sample contained a strong mix of sampled tasks across a range of occupations, the sample was stratified across 2-digit occupational categories.[2]

[2] The 2018 Standard Occupational Classification (SOC) system is a statistical standard used by federal agencies to classify workers into occupational categories using a hierarchical structure. There are 867 detailed occupations (6-digit SOC) which can be combined to for 23 major groups (2-digit SOC). These major groups include categories such as "Management occupations", "Legal occupations" and "Sales and related" occupations.

We recognise that prior belief around the potential for automation through Generative AI will vary across individual. To help reduce the impact of uncertainty on our classifications we employed a two-step approach:

1. Consensus labelling by members of the Oxford Economics' team
2. Leveraging domain expertise with Cognizant's members

The first consensus approach involved two economists at Oxford Economics simultaneously labelling all 1,400 sampled tasks according to their potential for automation through Generative AI. Each economist labelled against a set of guidelines that were derived through discussions with Cognizant. In cases where the label selected by each economist differed, a final 'consensus' decision was reached following internal discussions.

These initial classifications were then reviewed and adjusted by Cognizant's subject matter experts (SMEs). Given their domain expertise, we gave priority to changes made by SME members.

Based on the final set of 1,400 sampled occupational tasks, around 58% were classified as having no potential for automation through Generative AI today, a share that was expected to fall to 54% by 2032.. Of the tasks that are predicted to have at least some degree of automation today, our analysis points to a general shift in the distribution towards greater automation potential by 2032.

Fig. 4. Share of sampled tasks classified into its potential for automation.

Maximum potential for automation	Share of classification 2023	Share of classification 2032
Not automatable	58%	54%
Little assistance	23%	7%
Some assistance	9%	13%
Mostly assistive	6%	13%
Fully automatable	3%	15%

These findings suggest that tasks that have no automation potential today are likely to remain that way in the next decade. However, if a task has some degree of automation today, then it's likely to be 'more automatable' in the future.

2.3 DEVELOPING A MACHINE LEARNING MODEL

Next, we used a machine learning model to train a classifier on our 1,400 sample points which was then exposed to the full sample of 18,000 tasks to produce an estimate of its potential for automation.

To understand how the potential for the technology evolves over time, we ran two separate pieces of analysis. The first reflects the current stock of Generative AI technology as of 2023, and the second reflects its projected capability in 2032.

In each case, we developed the following machine learning classification models:

(1) Artificial neural network
(2) Logistic regression
(3) Random forest
(4) LSTM
(5) Count vectorization
(6) TF-IDF
(7) Average of all models

To understand how well each model generalises to unseen data, we randomly set aside 20 observations from each automation class from the training set to be used as a testing set. This is data the model will not see during training, and therefore provides a good indication as to how well the model can predict tasks it has not yet seen.

Each model was then assessed according to how well it predicted class labels from the withheld testing set. As part of this evaluation, we considered in how 'far off' each prediction is from its true label, i.e., the degree of miss-classification, in addition to its point accuracy i.e., the share of tasks that were classified correctly. For instance, a model that predicts a 'not automatable' label as 'little assistance' is better model than one that predicts 'not automatable' as 'fully automatable'.

The table below summarises which model had the highest out-of-sample predictive accuracy. Around 80% of samples were classified as being up to one classification degree away from the true label, whilst around 95% were up to two misclassification labels out.

Fig. 5. Calculating out of sample model accuracy

Chosen model	Point accuracy	Accuracy within 1 degree of misclassification	Accuracy within 2 degrees of misclassification
2023: Average of all models	57%	81%	95%
2032: TF-IDF	58%	80%	94%

2.4 EXPOSING THE MACHINE LEARNING MODEL TO THE FULL SAMPLE OF TASKS

Next, we applied our preferred models to the full range of 18,000 tasks to estimate their expected potential for automation today, and in 2032.

We find that the class distribution for all 18,000 tasks broadly resembles the input sample distribution in both 2023, and 2032. As in the case of the input sample, we find that in 2032 there is rightwards shift in the degree of potential automation across our tasks.

Fig. 6. Share of task predictions by potential for automation

Potential for automation	Share of prediction 2023	Share of predictions 2032
Not automatable	57%	43%
Little assistance	26%	15%
Some assistance	9%	16%
Mostly assistive	6%	13%
Fully automatable	1%	12%

The final set of classifications were then sense checked against our understanding of the potential for Generative AI based on the sample input. In those instances where the model implied that a task would become 'less automatable' over time, we constrained the 2032 classification to its 2023 label.

Fig. 7. Example tasks with their potential for automation in 2023 and 2032

Task	Potential for automation 2023	Potential for automation 2032
Prepare budgets and submit estimates for program costs as part of campaign plan development.	Little assistance	Mostly assistive
Climb ladders to position and set up vehicle-mounted derricks.	Not automatable	Not automatable
Order stock, and price and shelve incoming goods.	Little assistance	Fully automatable
Conduct or direct investigations or hearings to resolve complaints or violations of laws or testify at such hearings.	Not automatable	Little assistance
Review financial statements, sales or activity reports, or other performance data to measure productivity or goal achievement or to identify areas needing cost reduction or program improvement.	Little assistance	Fully automatable

2.5 THE EVOLUTION OF GENERATIVE AI TECHNOLOGY

For tasks that increase in their potential exposure to automation between 2023 and 2032, we assume that the time-path of this evolution is linear over time.

In other words, we assume that any increases in automation does not all occur all at once in 2032, but instead increase evenly over time.

2.6 CALCULATING OCCUPATION EXPOSURE SCORES

The output of the machine-learning model described previously is a theoretical maximum automatability score from Generative AI for every unique task performed in the US economy (over 18,000 unique tasks). We used the task-composition of each occupation to aggregate automation scores for tasks into an exposure score by occupation. These occupation exposure scores provide a measure of the susceptibility of each occupation to automation by Generative AI.

For each occupation, the relative weight assigned to each task was calculated using the share of importance scores for all relevant tasks for that occupation. Task-based automation scores were multiplied by these weights to produce a weighted average score, known as the occupation exposure score. See Fig. 8 below for an example calculation of occupation exposure score for Travel Agents in 2032. This process was repeated for each year from 2023-2032 with task importance weights held constant over time for each occupation.

Fig. 8. Example calculation of occupation exposure score (Travel Agent)

Task	IM Score	Weight	Automatability label - 2032	Automatability score - 2032
Collect payment for transportation and accommodations from customer.	4.60	13%	Some Assistance	38%
Plan, describe, arrange, and sell itinerary tour packages and promotional travel incentives offered by various travel carriers.	4.57	13%	Mostly Assistive	63%
Converse with customer to determine destination, mode of transportation, travel dates, financial considerations, and accommodations required.	4.57	13%	Mostly Assistive	63%
Compute cost of travel and accommodations, using calculator, computer, carrier tariff books, and hotel rate books, or quote package tour's costs.	4.57	13%	Mostly Assistive	63%
Record and maintain information on clients, vendors, and travel packages.	4.49	13%	Mostly Assistive	63%
Book transportation and hotel reservations, using computer or telephone.	4.42	12%	Little Assistance	13%
Print or request transportation carrier tickets, using computer printer system or system link to travel carrier.	4.30	12%	Mostly Assistive	63%
Provide customer with brochures and publications containing travel information, such as local customs, points of interest, or foreign country regulations.	4.04	11%	Mostly Assistive	63%
Exposure score				**53%**

The occupation exposure score represents the theoretical maximum proportionate decrease in man hours required for an occupation to produce the same level of output as today due to the use of Generative AI. For example, the occupation exposure score of 53% for Travel Agents in 2032 means that we estimate Travel Agents will need to work 53% fewer hours due to Generative AI to produce the same level of output as today. Note: this is based on universal adoption across the entire economy and should therefore be considered the theoretical maximum impact rather than a prediction of what will happen.

To calculate the overall national (theoretical maximum) exposure score, the occupation exposure scores were weighted by 2022 employment shares sourced from the Bureau of Labor Statistics (BLS) Occupational Employment and Wage Statistics (OEWS).[3]

[3] Bureau of Labor Statistics, "Occupational Employment and Wage Statistics", accessed September 2023

3. ADOPTION RATES

3.1 HISTORICAL CONTEXT

The rate at which firms adopt Generative AI affects the timing and extent to which different occupations are exposed to automation. Historical precedents from transformative technologies such as steam engines, electricity, automobiles, and computers provide insightful context for understanding adoption dynamics.

The adoption patterns of different major technologies have varied. For instance, the adoption of electricity in manufacturing took several decades in the early 20th century. The transition was not just about replacing steam engines with electric motors; it involved rethinking and redesigning entire production processes to optimize the new power source's potential.

Similarly, the adoption of computers and internet technology during the late 20th century showed a staggered pattern. Initial adoption was slow due to high costs and limited understanding of the technology's potential. However, as the technology matured and its benefits became clearer, adoption rates accelerated.

While Generative AI presents significant opportunities for enhancing productivity, its adoption will likely follow a similar pattern of gradual integration and optimization. Our approach considers the technological capabilities but also the socio-economic and organizational factors influencing the integration of Generative AI into productive activities.

3.2 MODELLING ADOPTION RATES

To profile the speed at which Generative AI technology will spread across businesses in the US we have applied a Bass Diffusion Model. This framework is commonly used to predict the speed and shape of adoption for new technologies and products that enter the market.

The Bass Model distinguishes between two types of adopters:

- Innovators: Early adopters who are less influenced by others and are quick to try new technologies.
- Imitators: Those who adopt technology influenced by the number of previous adopters.

Although the Bass Model provides a standardised framework for our purpose, we developed tailored assumptions for the coefficients that describe the innovator and imitator rates We estimated a regression model applied to the historical path of enterprise software within the US's capital stock. We believe that the integration of enterprise software into the capital stock represents the best available proxy to use as a benchmark for Generative AI. Both have very widespread applicability across different sectors and occupations and, therefore, would be reasonably categorized as general purpose technologies (GPTs). More practically, a very extended time series, dating back to 1960, describing the evolution and relative importance of software is published by the US statistical authorities.

3.3 DEALING WITH UNCERTAINTY: THREE ADOPTION SCENARIOS

At this point, the speed at which Generative AI technologies diffuse and are applied by businesses is highly uncertain. To account for this uncertainty, we have designed three alternative scenarios for adoption. These are all premised on the assumption that the current rate of adoption in 2022 is zero which we believe is broadly representative of the state of the market.

Each scenario was developed by altering the starting point of the previously referred to regression model. We selected three start years, 1987, 1990, and 1995 with these years chosen to mark significant developments and milestones in the software industry, which reflect varying stages of technological maturity and industry readiness.

3.3.1 1987: Introduction of Key Office Software

In 1987, Microsoft introduced Excel and Word for Windows. These tools were significant advancements in office productivity software, offering improved user interfaces and functionalities over previous versions and competitors. Excel became a vital tool for data management and analysis, while Word set new standards in word processing.

Software in this era was characterized by relatively limited functionality and user interfaces that were less intuitive compared to later generations.

3.3.2 1990: Birth of the World Wide Web

The year 1990 was marked by Tim Berners-Lee's development of HTML at CERN, laying the foundation for the World Wide Web. This innovation was pivotal in transitioning the internet from a largely academic and military network to a platform with broader applications and accessibility.

This phase saw software becoming more user-friendly and versatile, with increasing integration into various business functions beyond basic computing needs.

3.3.3 1995: Widespread Internet Adoption and Enhanced Web Capabilities

By 1995, the internet was increasingly accessible to the general public, and web technologies were rapidly evolving. Key developments included the introduction of JavaScript, enhancing web interactivity, and the release of Windows 95, which integrated internet support into a mainstream operating system.

Software during this time became integral to business operations, with functionalities expanding into web-based applications, e-commerce, and enterprise resource planning (ERP) systems.

By selecting these specific years as start points, each scenario captures the unique characteristics and adoption patterns of different software generations. The chosen years represent distinct events in the evolution of software technology - from basic computing in 1987, to the onset of the Internet era in 1990, and the rapid expansion of Internet-based business applications by 1995.

3.4 COMPANY SIZE ADJUSTMENT

We refine the estimation of Generative AI adoption rates by incorporating adjustments based on company size using insights from the Annual Business Survey published by National Center for Science and Engineering Statistics.[4]

The survey provides a breakdown of AI adoption across different company sizes. It categorizes companies based on their employment size and details the extent of AI use, from no use to high use. This cross-sectional data reveals significant variations in AI adoption rates across different company sizes. We have used data on AI as a proxy due to the lack of data for Generative AI.

These differences were used to adjust Bass Model parameters for small and medium enterprises (SMEs) and large companies respectively and, therefore, develop alternative forecast adoption curves for businesses depending on their size.

[4] National Center for Science and Engineering Statistics, "Table 85: Use of artificial intelligence as a production technology for goods and services, by company size: 2016–18", accessed October 2023

4. PRODUCTIVITY IMPACTS

4.1 PRODUCTIVITY IMPACTS FOR NON-DISPLACED WORKERS

To calculate predicted occupation exposure to Generative AI, the theoretical maximum occupation exposure scores (discussed in section **Error! Reference source not found.**) were scaled down by the predicted industry adoption rates (discussed in section **Error! Reference source not found.**).

For example, a Marketing Manager has a theoretical maximum exposure score of 35.5% in 2027. This position is filled in many industries including Business Services and Real Estate. By 2027, 21.3% of businesses in the Business Services sector are predicted to have adopted Generative AI whereas in the Real Estate sector adoption is only predicted to be 9.8%. This difference stems from variation in business structure, reflecting the respective sector's share of employment in large firms versus SMEs. The predicted occupation exposure for Marketing Managers employed in the Business Services sector is therefore 7.6% (35.5%*21.3%) compared to 3.5% (35.5%*9.8%) in the Real Estate sector. This process was repeated for every occupation-industry pair.

Exposure scores represent the percentage reduction in man hours required to produce the same level of output. Fewer workers required to produce a given level of output implies that the productivity of workers not displaced by Generative AI must have risen. To calculate the productivity gain for non-displaced workers we assumed full productivity pass through such that the productivity impact on non-displaced workers can be calculated as 1/(1-exposure).[5] For example, non-displaced Marketing Managers in the Business Services sector would be expected to be 8.2% (1/(1-7.6%)) more productive due to the use of Generative AI. This process was repeated for every occupation-industry pair.[6]

This formula was applied to the national predicted exposure score to estimate the overall national predicted increase in labour productivity each year from 2023 to 2032 due to Generative AI.

4.2 PRODUCTIVITY AND UNEMPLOYMENT IMPACTS FOR DISPLACED WORKERS

4.2.1 Unemployment assumptions

Fears that previous technological advances which automate parts of workers' jobs would lead to permanent increases in unemployment have largely proven to be unfounded.[7] However, the implementation of Generative AI in the workplace would have a disruptive effect and it is reasonable to expect that there will be some type of transitional costs.

[5] Workers may take time to adapt to the use of Generative AI in the workplace and therefore the time saved may not be entirely spent on productive tasks resulting in a productivity pass-through less than 100%. However, after this initial "learning phase" we would expect productivity pass-through to reach 100%.

[6] The productivity increase occurs with a lag due to the J-curve effect which reflects a mix of learning costs required to adjust to the new technology, onboarding costs and figuring out how to redeploy spare time. This is discussed in section 4.3.

[7] Kerstin Hötte et al (2022), "Technology and jobs: A systematic literature review", accessed October 2023

Our review of the Current Population Survey (CPS) of displaced workers from the US Census Bureau data identified two labour market effects that we have built into our modelling framework.[8] We specified three scenarios to reflect the inherent uncertainty of labour market impacts associated with technological advancements.

Workers leaving the labour market

Workers who get displaced by Generative AI may struggle to regain employment. The difficulties of finding a new position will likely increase as their unemployment duration increases due to skills atrophy. Workers in such a situation may continue to search for work and be unemployed or end their job search and leave the labour market. As noted earlier in the section, major technological advances have not resulted in permanent material increases in the unemployment rate. To account for this possibility, however, we have relaxed the assumption of full re-employment in some scenarios. Based on US Census Bureau data on displaced workers, we found that 11% of workers left the labour market after being displaced.[9] We assumed this is the case in our central and high scenarios but full re-employment of workers in our low scenario.

Frictional unemployment

We expect that workers who get displaced by Generative AI will take time to adapt to the changing labour market and find a new position. Based on the US Census data, we assumed that it will take, on average, 38 or 39 weeks for displaced workers to regain employment, depending on the scenario.[10] This was calculated in two steps. Firstly, the average time to find a new role by aggregate occupation was calculated from the CPS data. 72% of individuals found a role within 1.5 years on average after being displaced.[11] For those taking longer, we estimated that it will take an additional year to be re-employed based on the literature and assumptions.[12,13] In the central and high scenarios, the duration is shorter due to our assumption that individuals expected to take the longest to secure new employment are likely to exit the labour market. As a result, those finding new roles are expected to do so slightly faster on average.

Lower productivity in new role

Studies have shown that workers who have been displaced from their previous jobs might struggle to regain employment in a similar capacity and instead move into a new role which is not fully suited to

[8] United States Census Bureau, "Displaced Worker", accessed October 2023. The information from the 2020 supplement was used given that the 2022 data encompasses the period affected by the COVID-19 pandemic and may therefore not be representative of normal economic conditions.

[9] Based on those who responded that they were displaced due to insufficient work or that their position or shift was abolished as these categories are more aligned to the impact new technology could have.

[10] Average calculated for those who were displaced due to insufficient work or their position or shift was abolished.

[11] The CPS displaced workers supplement asks individuals in January 2020 whether they have been displaced within the previous 3 years and if they found a job after this. As this spans a 3-year period, we assume it represents the labour market outcome of an average displaced worker 1.5 years after being displaced.

[12] Fredrick Andersson et al (2014), "Job displacement and the duration of Joblessness", accessed October 2023

[13] Estimates were calculated based on the aggregate change found in Andersson (2014) from 1.5 years to 2 years and an assumption that everyone found a role within 2.5 years of being displaced.

their skillset.[14] Therefore, the productivity in their new role would likely be lower. It is commonly assumed that wages can serve as an indicative measure of productivity and research suggests that workers displaced from a full-time job who subsequently find a new full-time job earned 11% less than they would have if have had they not been displaced.[15] Therefore, those who have been displaced were considered 11% less productive than non-displaced workers in the same occupation and industry.

4.2.2 Summary of assumptions

Fig. 9 below provides a summary of the unemployment and productivity assumptions for displaced workers that are used in each of the three scenarios.

Fig. 9. Summary of assumptions for displaced workers

Scenario	Frictional unemployment[16]	Workers leaving the labour market	Lower productivity in new role
Low	39 weeks	Full re-employment	11% lower productivity
Central	38 weeks	11% workers leave the labour market	
High			

4.2.3 Approach

For each occupation-industry pair, we assumed that their 2022 wage (sourced from BLS OEWS) is reflective of their productivity. A productivity time series from 2023-2032 was calculated for non-displaced workers in each occupation-industry pair by growing 2022 productivity according to predicted exposure scores.

The number of workers displaced per year for each occupation-industry pair is calculated using their respective exposure scores. For each annual cohort of displaced workers, we produced a productivity time series applying the assumptions outlined in section 4.2.1. For each occupation-industry pair we combined the productivity time series of each annual cohort of displaced workers to produce an overall average productivity time series of displaced workers in that role. These were weighted by displaced worker shares to estimate a national productivity time series of displaced workers.

[14] For a summary of the literature, see Glenda Quintini and Danielle Venn (2013), "Back to work: Re-employment, Earnings, and kill use after job displacement", accessed October 2023.

[15] Henry S. Farber (2011), "Job loss in the great recession: Historical perspective from the displaced workers survey", accessed October 2023

[16] Difference in duration of frictional unemployment duration between low scenario and central and high scenarios arises because in the latter scenarios the workers who permanently leave the labour market are those who would otherwise be expected to take the longest to regain employment. As a result, those that do regain employment are expected to do so slightly quicker than in the low scenario which assumes full re-employment of all displaced workers.

4.3 OVERALL DIRECT PRODUCTIVITY IMPACTS

The overall direct impact of Generative AI on labour productivity is a combination of the productivity impact on non-displaced workers and that on displaced workers. The national productivity time series for displaced and non-displaced workers were weighted by employment shares to produce an aggregative productivity time series. The ratio of this aggregate productivity time series to the non-displaced productivity time series was used to scale the labour productivity impact for non-displaced workers in section **Error! Reference source not found.**, to produce a macroeconomic labour productivity impact.

This was converted into a total factor productivity (TFP) effect using an assumption that labour's share of output is 65% and capital's share is 35%.[17,18]

Adjusting for productivity mismeasurement (J-curve)

Breakthrough technologies often require complementary investments across human and capital processes. Brynjolfsson et al (2021)[19] show that these types of investment are often poorly measured, despite creating a valuable asset for a company.

During the initial stages of technological adoption, firms will incur labour and capital costs that do not immediately translate into increased output, reducing firm productivity. Over time, this investment begins to pay off and productivity begins to increase. This productivity profile is termed the 'J-curve' in the literature.

To capture this effect in our model, we included an estimate of the J-curve for the US economy as provided by Brynjolfsson et al (2021). Its impact is to reduce our 'maximum theoretical' productivity figure to account for the initial investment and learning costs associated with new technologies.

[17] Total factor productivity represents the portion of output growth that cannot be attributed to growth in capital or labour inputs.

[18] The Federal Reserve Bank of New York, "The FRBNY DSGE model", 2013, accessed October 2023

[19] Erik Brynjolfsson et al (2021), "The productivity J-curve: how intangibles complement general purpose technologies", accessed September 2023

TFP mismeasurement adjustment (%)

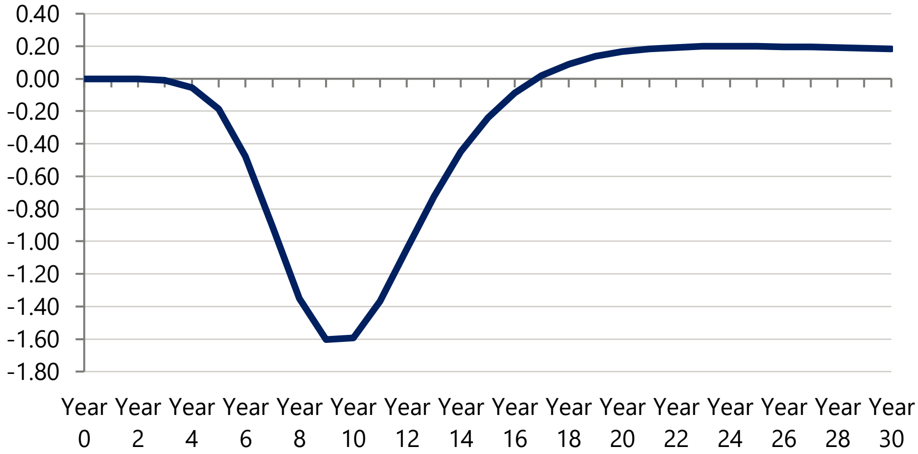

Source: Brynjolfsson et al (2021), Oxford Economics

5. MODELLING THE IMPACT ON GDP

Finally, our estimates of Generative AI's TFP and labour market participation impacts, discussed in previous chapters, were used to 'shock' Oxford's GEM and to develop alternative macroeconomic scenarios. The GDP impact of Generative AI on the US economy was captured as the difference in GDP forecasts between the 'shocked' scenarios and Oxford's baseline GDP forecast.

The GEM provides a rigorous and consistent structure for forecasting and scenario analysis. The model has consistently evolved over the past three decades and is the most widely used commercial International Macro Model.

The model's equation structure has been developed to achieve a balance between empirical validity (the extent to which the equations predict the historic path of each indicator) and theoretical coherence. The latter helps to ensure that when using the model for scenarios analysis, the results are more interpretable and intuitive.

The GEM takes the form of an error correction model, a form of multiple time series model that estimates the speed at which a dependent variable returns to its equilibrium profile after a shock to one or more independent variables. This is useful as it estimates both the short and long run effects of variables on the given variable in question. The model exhibits 'Keynesian' features in the short run: factor prices are sticky and output is determined by aggregate demand. In the long-run, its properties are neoclassical, such that prices adjust fully and the equilibrium is determined by supply factors – productivity, labour and capital – and attempts to raise growth by boosting demand only lead to higher prices.

The model's baseline forecast reflects our global outlook. It covers 85 economies in detail and six regional blocks interlinked through trade, prices, exchange rates, and interest rates.

The GEM also balances a trade-off between detail and tractability. In general, our approach has been to aggregate where it is not clear that disaggregation improves the quality of forecasts or analysis or serves particular users' needs. Many financial flows have been aggregated, and government accounting conventions have been standardised at a relatively high level of aggregation. But we continue to disaggregate series such as personal income components, investment, and the energy sector. Greater granularity of some series makes for more accurate forecasts and provides clients a comprehensive outlook for forecast components. In total, the GEM, which is predominantly quarterly, has tens of thousands of variables.

In the standard mode, the GEM assumes adaptive rather than forward looking expectations because we believe that introducing expectation on the basis of economic theory is more advantageous than using the forward-looking assumption ubiquitously. Where appropriate, the model does introduce expectations implicitly and explicitly, therefore accounting for how and the extent to which agents respond to information about changes in fundamentals.

Oxford Economics continually monitors and tests the model's performance. In addition to centralised testing of the model, internal usage of the model means that every month the structure is looked at and assessed by Oxford Economics' own forecasts and our wide range of clients.

6. FRICTION ANALYSIS

Generative AI is expected to disrupt the labour market by displacing some workers from their roles due to its capability to automate (fully or partially) a portion of the workplace tasks. However, this will likely be offset by expected growth in demand for other roles, including those which currently do not exist. We did not explicitly model the way the labour market may adjust in response to this disruption but conducted some quantitative analysis to understand the degree of difficulty displaced workers may face when trying to regain employment.

This analysis relied on the O*NET database which contains standardised scores regarding the characteristics and requirements of all occupations in the US economy (see content model in Fig. 2).

6.1 PAIRWISE FRICTION SCORES

We measured the difficulty a displaced worker may face trying to meet the requirements of other roles, such as in terms of required skills or knowledge, using "friction" scores. For each occupation, the "friction" they face in meeting the demands of another occupation was calculated as the average of the following three dimensions:

- ***Distance*** = measures the overall similarity between two occupations and therefore how suited an individual is to another role.
- ***Total improvement*** = measures the extent of improvement/upskilling needed to meet the demands of another role.
- ***Job zone difference*** = measures the difference in overall preparation needed for their current job compared to another role.

These are discussed in further detail below.

6.1.1 "Distance" between requirements and characteristics of the occupations

For this analysis we used O*NET data regarding Abilities, Knowledge, Skills and Work Activities. For each of these categories, there is a consistent list of descriptors which are scored according to the Level and Importance for every occupation.[20] The former reflects the proficiency required for the occupation (scored from 0–7) while the latter reflects the degree of importance to the occupation (scored from 1-5). We converted all scores onto a 0-100 scale.

For each of the O*NET categories mentioned above we calculated the importance-weighted average difference in level scores between each pair of occupations. For example, the "distance" between the skillset of occupation 1 to that required in occupation 2 can be summarised as follows:

[20] In the O*NET database there are 52 unique abilities, 33 unique knowledge categories, 35 unique skills.

$$Skill\ Distance_{1,2} = \sum_{s=1}^{s=35} \sqrt{(LV_{s,2} - LV_{s,1})^2} * \frac{IM_{s,2}}{\sum_{s=1}^{s=35} IM_{s,2}}\ where\ IM_{s,2} \geq 50$$ [21] [22]

Where LV = level score, IM = importance score, s= skill

This process was repeated for each of the four O*NET categories mentioned above (Abilities, Knowledge, Skills and Work Activities). The overall distance between two occupations is the average of the distance scores for each of those four O*NET categories.

6.1.2 Total improvement required to meet the demands of another occupation

The distance score above is a measure of job dissimilarity in which being over-skilled or under-skilled are treated in the same way. We therefore calculated a second score, called "total improvement" which measures the degree to which an individual needs to upgrade their skillset to meet the requirements of other roles.

For each of the four O*NET categories, we calculated the sum of pairwise level score differences in which the current occupation is under-skilled relative to the requirements of the potential new occupation. For example, the improvement in skills of occupation 1 needed to meet the demands of occupation 2 can be summarised as follows:

$$Skill\ Improvement_{1,2} = \sum_{s=1}^{s=35} (LV_{s,2} - LV_{s,1})\ where\ IM_{s,2} \geq 50\ and\ LV_{s,2} > LV_{s,1}$$

Where LV = level score, IM = importance score, s= skill

This process was repeated for each of the four O*NET categories. The overall improvement required to meet the demands of the other role was given by the sum of the improvement scores for each of the four O*NET categories.

6.1.3 Difference in job zone scores

The O*NET database also contains a variable called "Job Zones" which groups occupations into one of five zones according to the degree of preparation needed to perform the role. Preparation is assessed according to how much education, related experience and on-the-job training people need to be able to do the work required in an occupation.

For each pair of occupations, we calculated the difference in the job zone of the current occupation to that of the other occupation.

6.2 AGGREGATING PAIRWISE SCORES INTO OCCUPATION SCORES

The approach outlined in 6.1 produced friction scores for each pair of occupations, for example the frictions a Marketing Assistant may face when moving role to become a Finance Analyst. This section

[21] We only consider descriptors which are deemed important (importance score>=50) for the occupation that is being moved into.

[22] Positive differences (i.e. being over-skilled) and negative differences (i.e. being under-skilled) are treated equally.

describes how pairwise friction scores were aggregated to produce an average friction score for each occupation. This score provides an overall measure of the difficulty that a given occupation may encounter when attempting to move into any other role.

For each occupation-industry pair, the friction scores to other roles were weighted according to the relative frequency with which people have been recorded moving into those roles. These insights were calculated using longitudinal extracts from the US Census Bureau's Current Population Survey in which individuals' employment status, occupation and industry are recorded over time. Job moves which have occurred more frequently were therefore weighted more heavily when aggregating the friction scores that a given occupation-industry pair faces when moving into any other role.

During this aggregation process we only permitted an occupation to move into other roles which have a lower exposure score to Generative AI. The rationale for this restriction is that a worker displaced from a given role would be unlikely to move into another role which has displaced an even greater share of workers. Instead, it is expected that a displaced worker would likely move into a role that has been less disrupted by Generative AI and is therefore more likely to have job openings that need to be filled.

The aggregation step above produces friction scores for each occupation-industry pair. The next step is to aggregate scores across industries to produce friction scores for each occupation. The weights applied to each industry, for a given occupation, reflect the proportion of workers displaced in each industry. At the end of this aggregation process, we produced three friction scores ("distance", "total improvement" and "job zone difference") for each occupation.

Fig. 10 below presents summary statistics of the scores for the three methods.

Fig. 10. Summary statistics of friction scores

Statistic	Distance	Total improvement	Job zone difference
Mean	11.46	487.70	-1.89
Standard deviation	2.67	210.51	17.33
Minimum	0	0	-44.19
Maximum	27.74	1,640.57	46.47

6.3 COMBINING INTO A SINGLE FRICTION SCORE

The three friction scores provided different yet complementary insights into the challenges that workers may face when trying to move roles after being displaced by Generative AI. We combined these three scores into one single friction score that provided an overall measure of the expected difficulty to successfully move roles. Given the very different distribution of friction scores for the three dimensions (as illustrated in Fig. 10), we first standardised scores onto the same scale before combining into a single score.

Friction scores for each of the three dimensions were first converted in z-scores by subtracting the mean and dividing by the standard deviation. Z-scores were then converted onto a 0-100 scale by multiplying the z-score by 25 and adding 50. Where this formula produced values less than 0 or more than 100, the values were constrained to 0 or 100 respectively. The overall friction score was given by the average of these 0-100 scores across the three dimensions.

This approach provided some useful features when interpretating scores. A value of 50 is equal to the average, while a score of 25 or 75 can be interpreted a 1 standard deviation below or above the average. Scores of 0 or 100 can be interpreted as 2 or more standard deviations below or above the average.

OXFORD ECONOMICS

Global headquarters
Oxford Economics Ltd
Abbey House
121 St Aldates
Oxford, OX1 1HB
UK
Tel: +44 (0)1865 268900

London
4 Millbank
London, SW1P 3JA
UK
Tel: +44 (0)203 910 8000

Frankfurt
Marienstr. 15
60329 Frankfurt am Main
Germany
Tel: +49 69 96 758 658

New York
5 Hanover Square, 8th Floor
New York, NY 10004
USA
Tel: +1 (646) 786 1879

Singapore
6 Battery Road
#38-05
Singapore 049909
Tel: +65 6850 0110

Europe, Middle East and Africa

Oxford
London
Belfast
Dublin
Frankfurt
Paris
Milan
Stockholm
Cape Town
Dubai

Americas

New York
Philadelphia
Boston
Chicago
Los Angeles
Toronto
Mexico City

Asia Pacific

Singapore
Hong Kong
Tokyo
Sydney

Email:
mailbox@oxfordeconomics.com

Website:
www.oxfordeconomics.com

Further contact details:
www.oxfordeconomics.com/about-us/worldwide-offices