How GenAI will change the world economy

Executive Summary

- **Generative AI has the potential to substantially improve the medium-term growth outlook for the economy. It arrives at a perfect moment – labour is going to be less supportive of growth as aging dynamics hit home and productivity is flagging.**
- **According to our modelling, the impact of GenAI on medium-term growth will be sizable. Using the US as a case study, we estimate GDP there could rise by between 1.8%-4% by 2032 depending on the speed of adoption. Other less flexible and knowledge-rich countries should also benefit, though to a lower extent.**
- **Gains from GenAI may be even greater if it not only helps to automate tasks but also leads to increased innovation. At this stage, however, we are sceptical of a permanent increase in productivity growth – we think the more significant productivity gains will come from the automation of a wide array of tasks.**
- **Although the economy wide-gains will arrive eventually, they won't be immediate. Developments in human capital and the infrastructure around GenAI will be needed to unlock the full benefits. And in line with the historical experience of reaping the benefits from innovation, that's likely to take years.**
- **In the meantime, higher productivity is likely to come at a cost. Labour market frictions mean that some workers will be displaced and will take time to find alternative employment.**
- **We think the impact of GenAI is likely to be a disinflationary force for the global economy in the medium term, as higher productivity is unlikely to be matched immediately by higher wage growth. In contrast, AI-related investment may raise demand, adding to cost pressures.**

After more than a decade of persistently weak productivity growth across advanced economies, there is hope that generative AI might help to arrest the slide. The timing could not be better. Workforce aging and the lack of reform to pension ages mean that labour supply is going to be less supportive for economic growth in the decades to come. We estimate that, relative to the pre-pandemic period, the contribution of labour to potential growth will halve over the next decade in advanced economies, and drop even more in emerging markets.

Up until now, labour force growth has been an important driver of wider economic growth, helping to offset a broader slowdown in productivity growth. That slowdown can be partly explained by the rise of the service sector, which is typically lower productivity and harder to automate (the so-called Baumol effect). Generative AI has the potential to change that.

But is this realistic? After all, many technological advancements have failed to deliver on their initial promise. In our view, while not all the hype should be believed, we think there are genuine reasons to be optimistic about the impact that GenAI will eventually have on the economy.

The current surge in US productivity

One reason why GenAI is raising expectations of a growth surge is that the US economy is currently enjoying a productivity boom. However, while it has coincided with the rollout of various AI technology, we would be wary of concluding that AI has been the key driver of the productivity improvement. The initial recovery from the pandemic led to intensive hiring, but as demand for labour has come into better balance with supply, productivity has begun to rise.

As we have pointed out previously, rising productivity is to be expected following a period of significant labour market churn (**Chart 2**). Unlike the stagnant pre-pandemic labour market, workers have the opportunity to move to roles and sectors that are able to better able to utilise their skills, while hybrid working may also have improved flexibility to a degree.

Chart 1: Labour market churn and increased business dynamism are leading to higher productivity

Source: Authors

What's more, productivity gains from technology occur when it reaches widespread adoption and filters down from a few frontier companies. So far, there is limited evidence that this broad adoption is happening among firms, even though take up by individuals has been rapid.

Quantifying the gains from GenAI

In theory, GenAI – or any technology for that matter – can have two types of impact on productivity. It can automate or augment tasks and improve the efficiency of workers, resulting in a temporary increase in productivity growth. It can, albeit rarely, also facilitate the production of new ideas and new ways of automating tasks – this implies a permanent increase in the growth rate of productivity and the economy.

Confusingly, temporary increases in productivity growth – improvements in efficiency – can be indistinguishable from permanent increases in real time. That can lead to overestimating the potential growth rate of the economy and associated policy mistakes – such as allowing the economy to run too hot for an extended period. This might have been one of the factors at play during the internet boom and in the run up to the global financial crisis.

Generative AI may eventually impact the economy in many ways, but the initial impact is likely to be through making workers more efficient. This is the effect that we have focused on to analyse the impact of GenAI on the economy.

Using machine learning techniques, we analysed over 18,000 tasks for their propensity to be automated by GenAI. Mapping these to occupations now and in the future, shows how the nature of demand for different types of jobs will change over time.

The first key finding of our study is that currently the majority of workplace tasks can't be automated or assisted by GenAI now or in the future (**Chart 3**). This is as much to do with the ability to harness GenAI as it is about the types of tasks undertaken today.

For example, it may be difficult to embed AI in the existing capital stock (e.g. driverless cars), or there may be a lengthy testing and learning phase before GenAI becomes helpful. A useful historical analogy is the introduction of the personal computer in the 1980s – the productivity boom associated with PCs did not really occur until the internet facilitated it in the late 1990s. In our view, by 2032 many of the structures around AI will be in place, allowing most tasks to be automated to some degree and around 12% of all tasks to be fully automated.

Chart 2: Potential tasks exposed to GenAI will grow over time

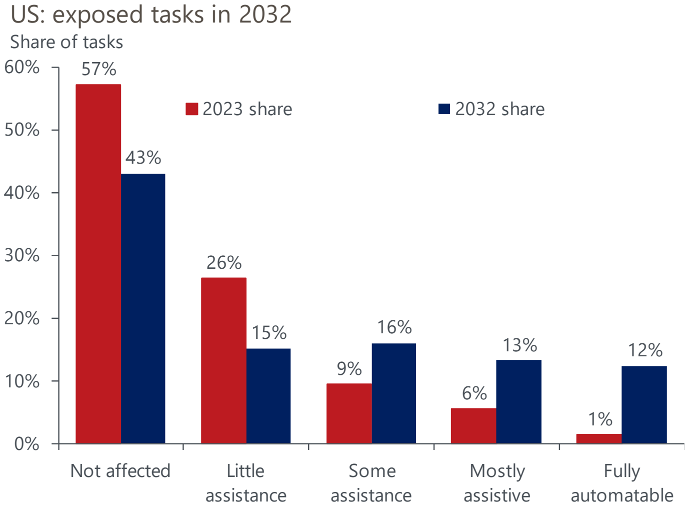

Source: Authors

A key factor that will determine how quickly these efficiency gains are realised is the GenAI adoption rate. We used three scenarios to reflect the inherent uncertainty in adoption speed, modelling adoption increasing to 12% in the 'low' scenario to 52% in a 'high' scenario (**Chart 4**).

The S-shaped profile of adoption rates reflects the likely evolution of the technology and consequently its applicability to different tasks. Indeed, it's clear from early evidence on company adoption that start-ups and very large firms that have been the early adopters, whereas mid-sized firms have been much slower. This may reflect the costs associated with early adoption – it takes time and resources to build up the human capital needed to unlock the benefits. This is a factor we have built into our modelling.

These adoption profiles were produced using a model calibrated on data that traces the relative importance of enterprise software as a share of US capital stock. The outcomes from this model depend on the historical reference period used – calibrating the model using data with an earlier starting year naturally creates a slower adoption profile.

Specifically, the adoption curves are based on the following three historical reference points:

- Low scenario (1987): the year in which Microsoft first launched Word and Excel, two tools that remain widely commercially adopted and were foundational to the spread of productivity gains from computing in office-based occupations.
- Medium scenario (1990): the year marking the introduction of HTML by Sir Tim Berners Lee, which greatly enhanced the commercial applicability of the internet.

- High scenario (1995): coinciding with the release of Windows 95, which integrated internet support into a mainstream operating system. By this stage, the internet had become much more widely accessible to the public.

Chart 3: Scenarios for GenAI adoption

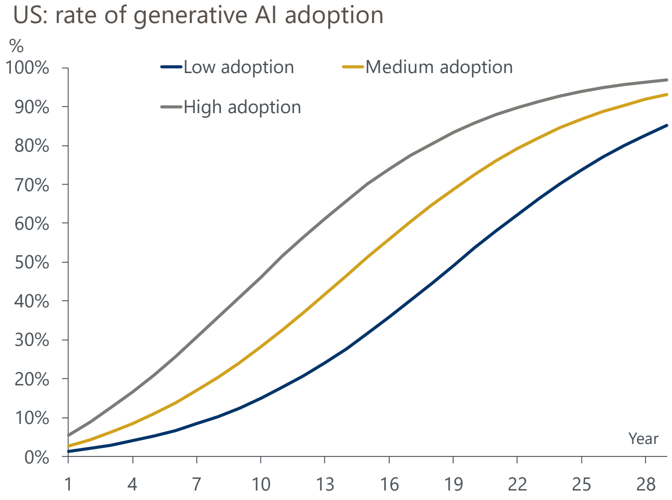

Source: Oxford Economics

Ultimately, the impact of GenAI adoption on productivity is striking. In our medium adoption scenario, the productivity of the workforce rises by over 10% (**Chart 5**). In historical terms this productivity surge would be similar in scale to the computer age – the 1990s to 2000s.

Chart 4: GenAI could raise labour productivity by 10% in a decade

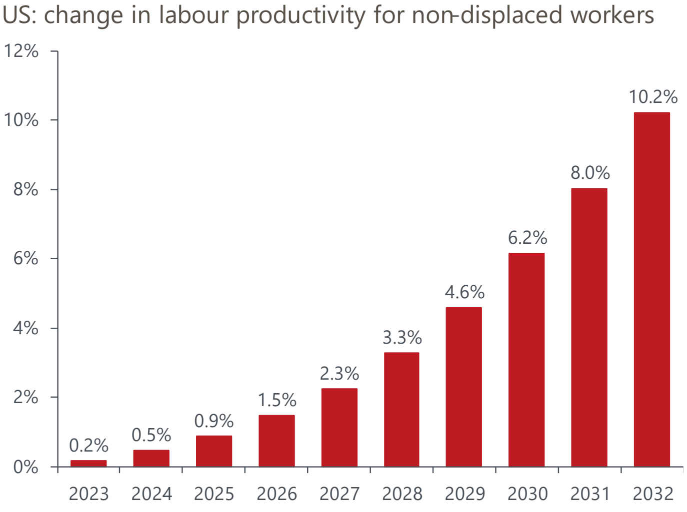

Source: Authors

Labour market frictions ensure some workers will be displaced

Translating the impacts into implications for the economy requires us to think of the dynamic impacts of AI. For example, at one extreme, it could be argued that the boost to productivity will fully translate into higher output, with no change in employment levels. On the other extreme, some have argued that it will result in the displacement of workers with no appreciable increase in the overall level of output.

Our view sits broadly in the middle. Some displacement of workers is inevitable and is a consistent feature of the introduction of new technologies. As job roles change and the demand for skills shifts to a different set of tasks, workers have to adjust. But this doesn't happen immediately. In a dynamic economy, there is always a loss of workers and a reduction in the labour force, at least temporarily.

Our modelling shows that computer and mathematical roles (e.g., roles involving heavy amounts of coding) are proportionately most at risk of displacement. In absolute terms though, office and administrative roles are likely to suffer the largest loss of demand (**Chart 6**).

Chart 5: Expected displacement by occupational group

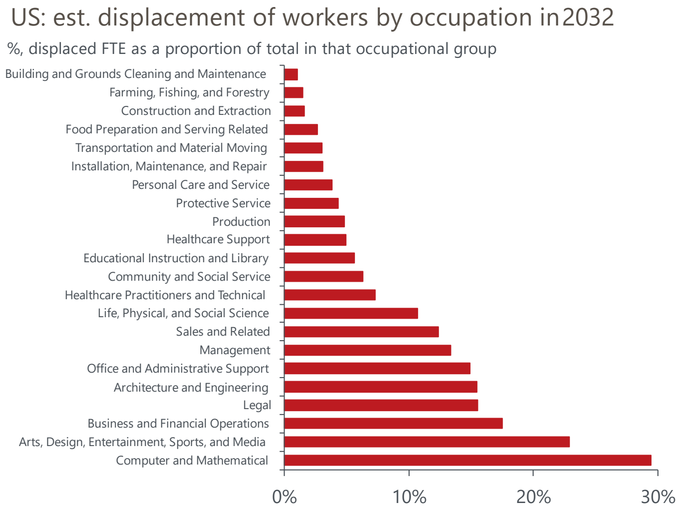

Source: Authors

Time out of the labour market for displaced workers is a function of how transferable their skills are. In other words, workers with highly transferable skills (e.g., computer and mathematical workers) may be highly exposed to the automation from GenAI, but they will also find it relatively easier to apply those skills in new roles. In contrast office and administrative support workers are highly exposed and will struggle comparatively to apply their skills to other jobs.

Chart 7 illustrates the degree of disruption to different occupational groups by plotting the exposure to automation against a 'frictional' score (ease of finding a new job). While higher paying roles tend to be more exposed to the effects of AI, they are also likely to experience lower hurdles to new employment.

Chart 6: Office and administrative workers face higher barriers to new jobs

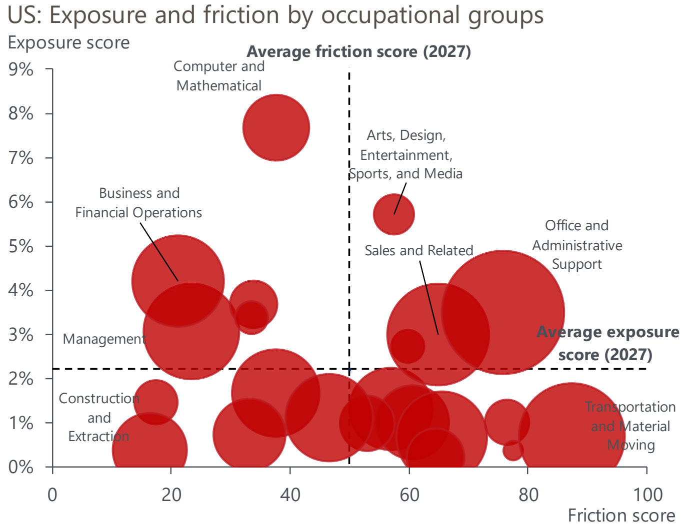

Source: Authors Note: size of bubble indicates relative size of employment

The implication is that the next decade may well feature a repeat of the hollowing out of mainly middle-skilled jobs, similar to that seen in the 2010s following the introduction of the internet. At that time, employment shares rose at the bottom- and top-end of the skills distribution, but workers in the middle found their jobs automated by the falling cost of computing.

This time, an aging workforce may help to mitigate some of the impact as workers leave the labour force. But another factor is the extent to which government policies are proactive in supporting the re-training of workers to aid their transition to different jobs.

Benefits set to outweigh the costs, especially for the wealthiest countries

Running all these inputs, including the productivity impacts, through our macroeconomic model, we estimate that generative AI will increase US GDP by 2.9% in the medium scenario by 2032 relative to our current baseline. This is significantly lower than the 10% increase in productivity we estimate for non-displaced workers, reflecting labour market frictions and the initial costs associated with applying generative AI. Depending on the speed of adoption, the gains could vary from a 1.8% boost to the level of US GDP to 4%.

We also think that thanks to an array of factors the US is likely to benefit more than other countries – adding further weight to the US growth exceptionalism narrative. Overall, high levels of human capital, a conducive regulatory environment, as well as a high share of service sector jobs that can be automated means that richer countries will benefit more (**Chart 8**).

Chart 7: Richer countries are better prepared to reap the benefits of GenAI

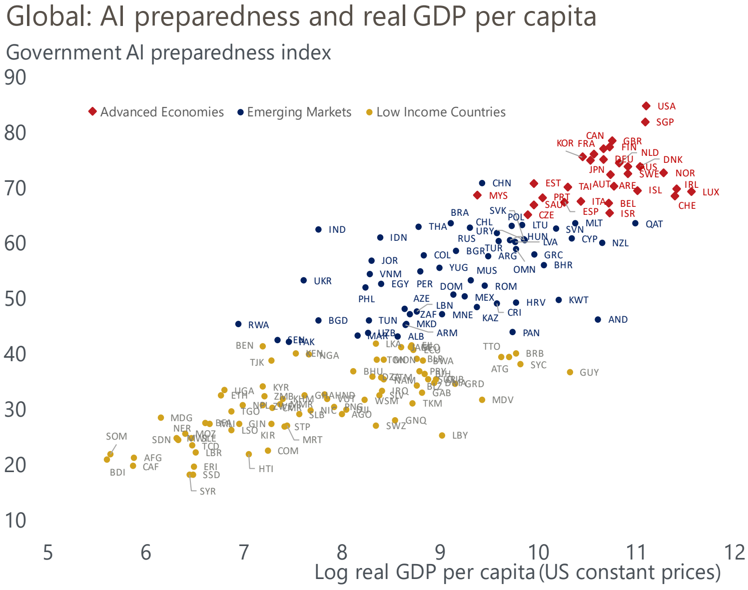

Source: Authors / Oxford Insights

GenAI is likely to have smaller impacts on economies that are more heavily reliant on industrial sectors where there is less scope for productivity gains, reflecting the more manual nature of tasks. That may mean that adoption is slower overall in many emerging or low income countries but it doesn't mean that adoption will be slower in some subsectors where generative AI might be particularly impactful. For example, Indian IT services and back-office support sectors.

More generally, emerging economies that are open to high levels of FDI might be faster adopters of GenAI, as they benefit from technology transfer from rich economies.

Much will also depend on the broader policy environment into which GenAI is being introduced. Compared to previous technologies, there seems to be already a much greater awareness of the broader societal risks that could be associated with the misuse of GenAI and that has led to a heightened level of scrutiny from policymakers.

We are already seeing differences emerging with the EU seemingly putting its energies into developing an all-encompassing top-down approach whilst other countries such as the UK opt for a sector-specific approach. In the UK, regulators have been charged with applying the general framework designed by the government which was underpinned by five core principles. We can, therefore, expect to see a sector-specific application of this framework allowing the various regulators some license to tailor guidelines to their respective domain. This contrasts quite sharply with the EU with the EU AI Act representing a piece of legislation that “seeks to create a far-reaching and comprehensive legal framework for the regulation of AI systems across the EU.”

In our view there are also clear dividing lines opening up between countries that have, to date, opted to develop a set of nonbinding principles (Japan, Singapore, UK, US) versus those that have implemented legislation (Korea, EU, Canada, China). Overall, it’s already quite evident that Europe is taking a more cautious regulatory approach as it seeks to balance risks and opportunities, and that may mean it is slower to adopt GenAI than, say, the US.

Impacts beyond GDP growth

Some of the most important indirect impacts of GenAI will be on inflation, interest rates, and asset prices.

A positive supply shock from higher productivity is likely to be disinflationary in the medium term (five to ten years). It could be argued that wages may rise in response to productivity gains, thereby cancelling out any impact on inflation. But in practise wages tend to be slow to respond for three reasons:

- The improvements in productivity are slow and difficult to articulate in real time (over and above trend improvements in productivity);
- Wages tend to be sticky, resulting in plenty of historical examples where real wages have lagged productivity; and
- The change in the labour market is taking place against a backdrop of displaced workers; therefore some slack emerging in the labour market would help to limit wage gains.

In contrast, in the near term (next five years), GenAI may add upward pressure on inflation. As above, the early stage of GenAI is likely to require some investment by firms before any benefits on the supply side can be realised. We are already seeing some evidence of bottlenecks in semiconductor markets, as well as demand far exceeding supply for data centres. If adoption is sufficiently widespread, this could add to inflationary pressure in advanced economies. What could be almost as important for broad demand in the economy is the ongoing rally in technology stocks and the associated wealth effects for households.

The implications for interest rates in the long term are more ambiguous based on our framework for neutral interest rates. Higher productivity – and therefore higher GDP growth – will push up on long-term equilibrium interest rates. But, at least to some extent, this is likely to be offset by lower longer-term inflation. Therefore, the implications for financial asset prices – outside of technology stocks – are highly debatable.

Still, our analysis does point to some implications for real assets. For example, a very high proportion of the job roles displaced by GenAI are office-based. We estimate an employment weighted AI exposure score of 18.5% for office workers compared with only 3.2% for the industrial property sector. This represents yet another structural headwind to a sector that is still getting to grips with the implications of hybrid working.

As with all slow-moving structural changes, the ripple effects beneath headline economic growth are likely to be the most profound and the most difficult to predict.