VAN LANSCHOT KEMPEN

INVESTMENT MANAGEMENT

Generative artificial intelligence:

Boon, or threat?

MARCH 2024
FOR PROFESSIONAL INVESTORS ONLY

AI Executive Summary

1) Generative Artificial Intelligence (GenAI) is a form of machine learning that learns from patterns in training data to generate new text, video, images, code or music that can be indistinguishable from what humans can create. Moreover, its appeal is that it can potentially be performed faster, cheaper and with fewer errors (over time).

2) Large Language Models (LLMs) are a subset of GenAI that utilise advanced algorithms to understand, generate, and manipulate human language. These models have the potential to take over routine tasks in the service sector (e.g., software engineering, marketing & sales, legal and financial services) and make processes more efficient. Examples of applications include data analysis, content creation and summarisation, product design, and customer service. This provides companies with two incentives to invest in LLMs: augmenting higher-wage, white-collar workers, and alleviate some of the pressure coming from tight labour markets due to adverse demographics.

3) For a technology to materially boost labour productivity growth – a key driver to long-term economic growth – it needs to be widely adopted. Adoption Curve Theory helps explain why some technologies are widely adopted, while others fail to gain traction. When applied to LLMs, the theory's three main factors (affordability, attractiveness and accessibility) suggest that this technology is still in the 'early adoption' phase.

4) Potential macro-economic impact: we expect that over the next few years, employees could increasingly be assisted by conversational bots, which could have a slight positive impact on labour productivity growth. In the long term, LLMs can have a strong productivity effect, depending on how technology, regulations, and costs develop. We anticipate that LLMs will have a slight deflationary effect in the coming years, primarily driven by the prospect of cost inflation. We expect a more diffuse effect from demand inflation, depending on how it affects income distribution and the labour market. Historically, the labour market has experienced positive effects from technological innovation, with increasing labour demand offsetting concerns about job losses.

5) GenAI brings risks in terms of security and privacy, which necessitates regulation (which could impact its economic potential). Forecasting and mitigating these threats is key. Some risks are already visible (e.g., hallucinations, data privacy and disinformation), while others may only manifest themselves longer term (e.g., job displacement and autonomous weapon systems).

6) Investment implications: shares of several global tech companies (e.g., the 'Magnificent 7') have greatly benefited since the beginning of 2023 due to the attention and potential of GenAI. One risk of investing in these growth companies is that market valuations often overshoot, resulting in certain stocks pricing-in GenAI potential that may never materialise. We see opportunities for investors in specific service sectors, although there are significant differences between and within sectors, and it will take time for the integration of new technology to lead to higher productivity growth, and higher profit margins.

7) GenAI's environmental footprint is considerable, and growing: it uses enormous amounts of energy (carbon emissions) for computation and storage purposes, and millions of gallons of fresh water to cool data centre equipment.

Contents

Artificial intelligence is at the top of everyone's agenda regarding future technology, with its myriad possibilities and implications. In this paper we focus on generative artificial intelligence and its subset 'large language models'. These models have the potential to take over routine tasks and could potentially transform the service sectors in the long term. We weigh up the benefits in efficiency and productivity against the risks to safety and security; where we are in terms of oversight and regulation; and what it all means to the investor.

What is generative artificial intelligence?

Simply put, artificial intelligence involves computer systems that can perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making and translation between languages.

In this paper, we specifically look at generative artificial intelligence ('GenAI') – a form of machine learning that learns from patterns in training data to generate new text, video, images, code or music that can be indistinguishable from what humans can create. Moreover, its appeal is that it can potentially perform it faster, cheaper and with fewer errors.

Within the broad spectrum of GenAI, Large Language Models ('LLMs') is a subset that has recently surged in popularity (chart 1). LLMs utilise advanced algorithms to understand, generate, and manipulate human language. Examples of applications include data analysis, content creation and summarisation, product design, and customer service. This provides companies with two incentives to invest in LLMs: augmenting higher-wage, white-collar workers, and alleviate some of the pressure coming from tight labour markets due to adverse demographics.

Chart 1 - Mentions of GenAI in company calls rose from 200 to 1300 over last year

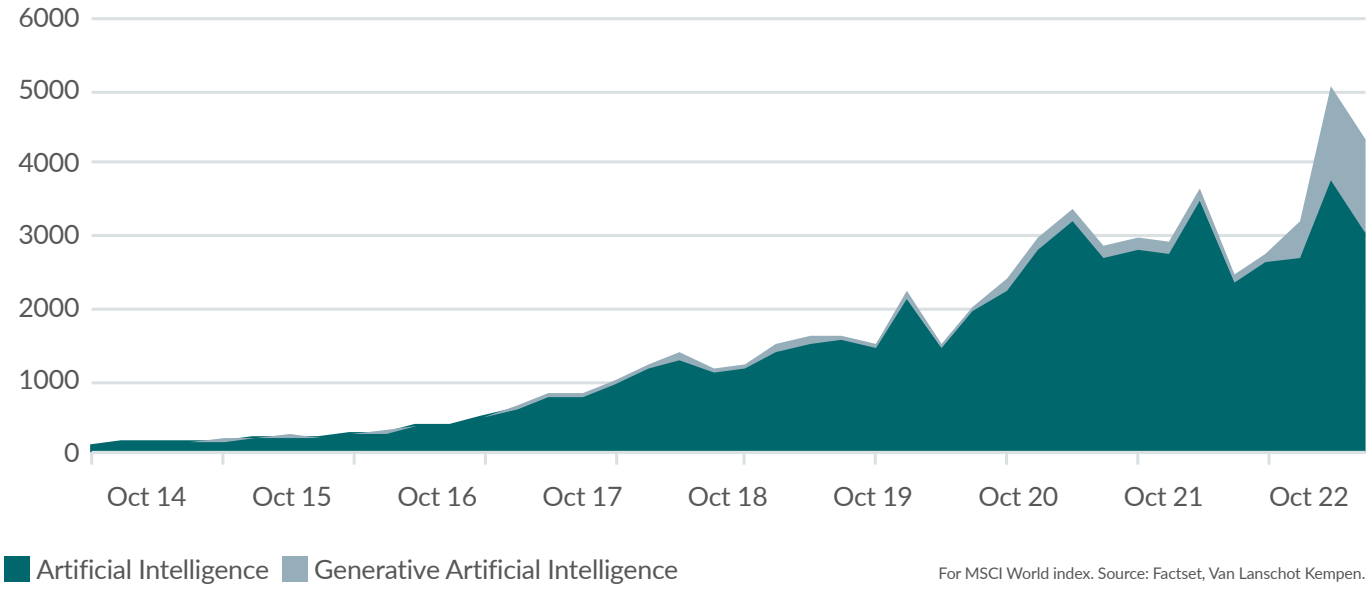

For MSCI World index. Source: Factset, Van Lanschot Kempen.

Service sectors will benefit most

Although most technological breakthroughs in the past few decades have focused on improving manufacturing sector productivity, the impact of GenAI is expected to be mainly on service sectors. And, as developed markets have a relatively higher share of service sectors, these will likely be more affected.

This also implies that GenAI mostly augments higher-wage, white-collar jobs (chart 2), which provides an incentive for companies to invest in LLMs to boost labour productivity. Another incentive is scarcity of labour due to declining demographics.

As with many earlier emerging technologies, initial use cases only scratch the surface of the more pronounced, long-term shifts that might be ushered in. LLMs have many potential applications, extending beyond text-to-speech solutions and image creation. This versatility makes it a potentially powerful tool in various fields (e.g., software engineering, customer services, marketing and sales, legal services and financial services), with several use cases: from analysing data, to creating content, designing products, and customer service tasks or simply summarising and transcribing meeting recordings (chart 3).

Law firms already use LLMs to help create the first draft of legal documents, while the healthcare industry has started using LLMs for medical-imaging analysis and conversations with virtual physicians.

Chart 2 - LLMs potential usage by income group

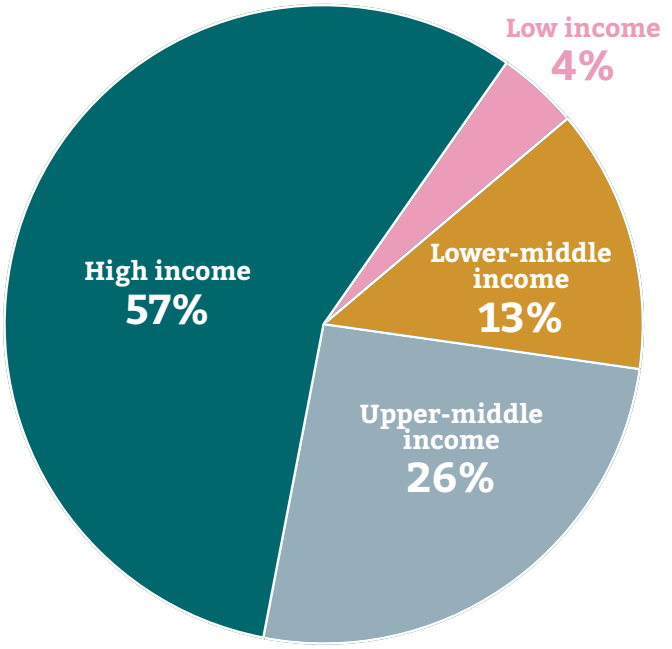

Source: International Labour Organization, Van Lanschot Kempen.

Moreover, costs and complexity rise substantially where information needs to be accurate and relevant. In our view, accuracy, delivering on its promise and data privacy remain work-in-progress for the foreseeable future.

Chart 3 - LLMs automation potential by occupation

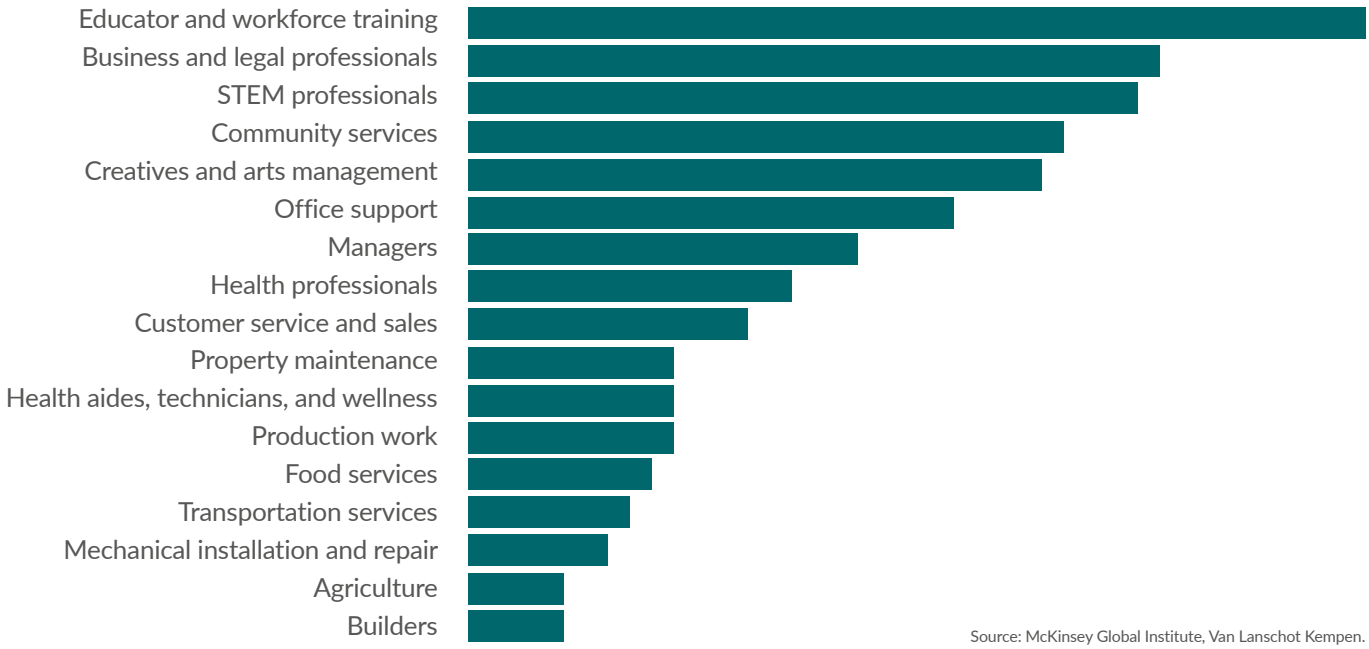

Source: McKinsey Global Institute, Van Lanschot Kempen.

The potential adoption rate

Labour productivity growth is a key driver to long-term economic growth. For a technology to materially boost labour productivity growth at a broader macro-economic level, it needs to be widely adopted (households and organisations) and users able to perform additional value-enhancing tasks.

And this is a challenge. For instance, while the launch of OpenAI's ChatGPT-4 in March 2023 sparked widespread interest in LLMs, most users still only use LLMs at home on their private laptop. Adoption at the organization level is more complex as organisations are not just the aggregate of its individuals, but also have their own system with procedures and norms.

By applying the Adoption Curve Theory to LLMs, we can take a deeper dive into the potential adoption rate as well as possible economic impact of LLMs

Adoption Curve Theory

Every new technology has an adoption curve, which is the cumulative rate at which a population adopts it over time. Sociologist Everett Rogers popularised adoption curves in his 1962 book 'Diffusion of Innovations'.

Rogers argued that there are 5 different recipients with dissimilar needs and wants regarding new technologies: innovators, early adopters, early majority, late majority and laggards (chart 4) Each recipient has its own characteristics and adopts new technology at a different rate. The theory holds that 'crossing the chasm' is the most critical moment in the adoption process: enthusiasts have accepted the new technology, but pragmatists still need to be convinced. These pragmatists are rational and interested in new technologies, but require proof of their advantages and potential impact, and are not convinced by a hype.

Chart 4 - Adoption curve theory

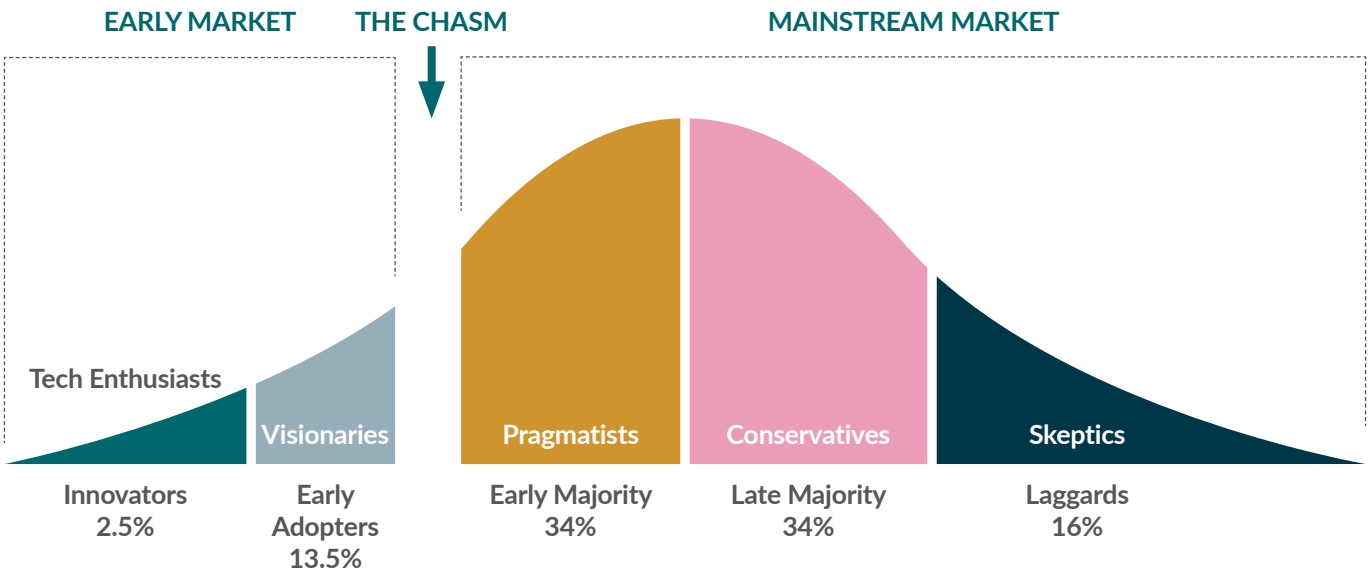

Historic adoption rates

Over the past 15 years, the adoption rate of various US consumer technologies has accelerated (steeper curves). It took the landline telephone 50 years from its introduction to become universally adopted. More modern inventions, like HDTVs and smartphones, have taken half as much time to reach similar adoption levels (chart 5). Some of these technologies have already lost their appeal.

Chart 5 - Technology adoption rates in US households, 1900 - 2020

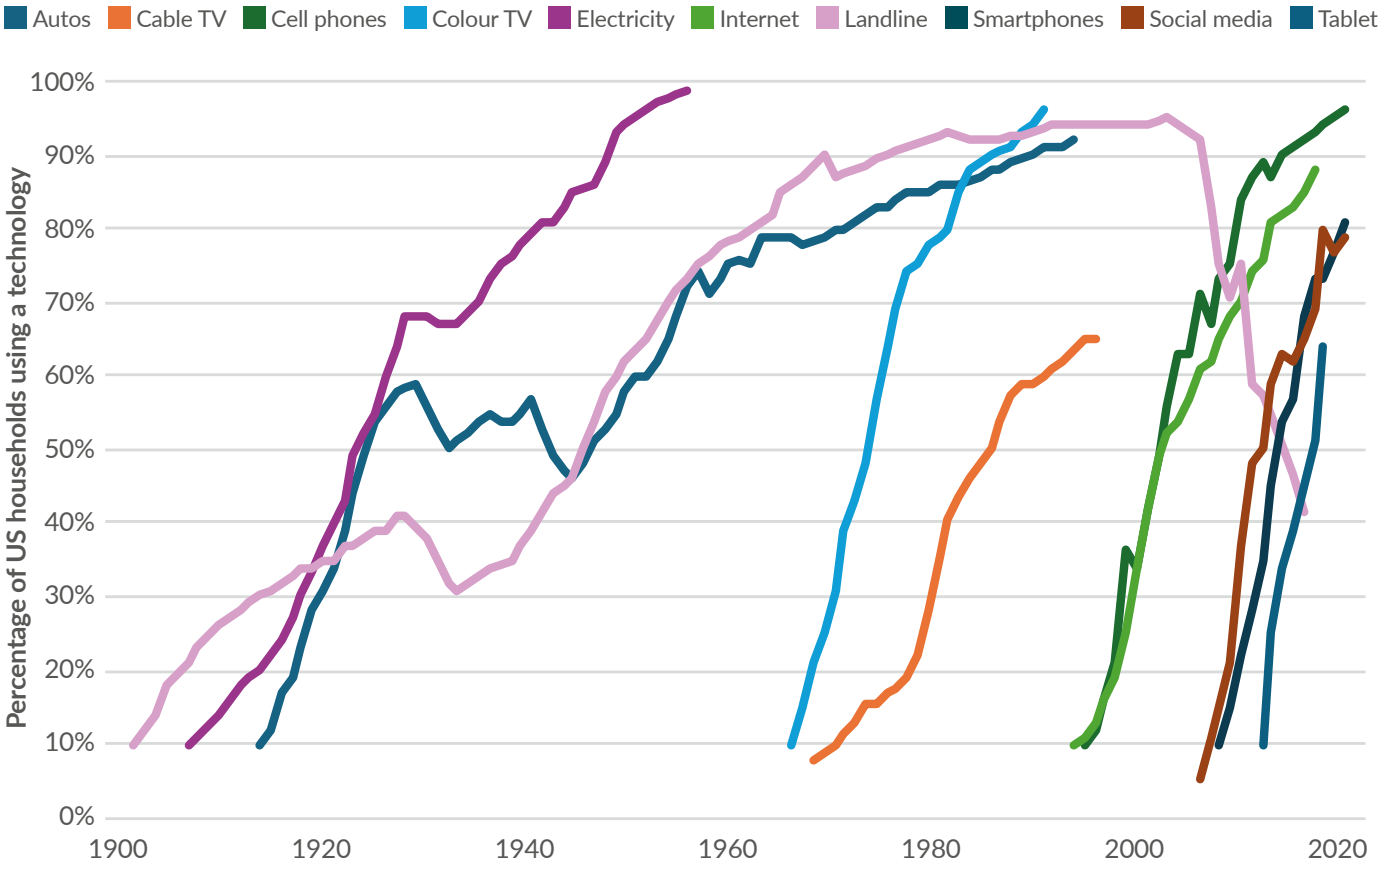

Source: Comin and Hobijn (2004), ourworldindata.org, Van Lanschot Kempen.

Factors that play a crucial role in the speed and breadth of adoption are:

- **Affordability:** If a technology becomes (relatively) affordable, it gains wider acceptance and adoption. This can be stimulated through technological advances/lower relative cost (e.g. microchips), and policy support (e.g. subsidies, while taxing alternatives).
- **Attractiveness.** When a technology offers superior performance or additional benefits (e.g. simplicity and cost advantages), it becomes more appealing to consumers and firms.
- **Accessibility.** The availability of supportive infrastructure (e.g. information exchange through the Internet) and the ability to try out (e.g. free trial mobile phone apps) help to keep entry barriers low. Adoption barriers are, for example, regulation and values and norms of societies.

Applying Adoption Curve Theory to LLMs

We can apply these three key points to LLMs:

- **Affordability:** LLMs require strong data and technology infrastructure, and considerable investment in employee up-skilling. At the same time, some of the (basic) tooling is available through the Internet for free. However, as computational power requirements increase with each improved version of LLMs (Chart 6) and if the value proposition of LLMs becomes clearer, we expect that technology firms will start charging (higher) fees to users.
- **Attractiveness:** The technology has already delivered evidence for several use cases in particular domains, showing the potential to improve individuals' productivity as several routine tasks can (largely) be delegated to LLMs such that the employee can take on more/other productive tasks. As a sign of its appeal, ChatGPT gained 1 million users in just 5 days (Chart 7), and currently has around 100 million active users;
- **Accessibility:** Its infrastructure is already widely available (largely free of charge) to end users in most developed countries through the Internet. Behind the scenes, however, the capital investment required from developers, in terms of computing power and data, is substantial. Overall, it is fair to assume that the diffusion of LLMs in developed countries (which comprise most of global GDP) can be smooth.

Chart 6 - Significant improvement in exam results taken by GPT-4 version

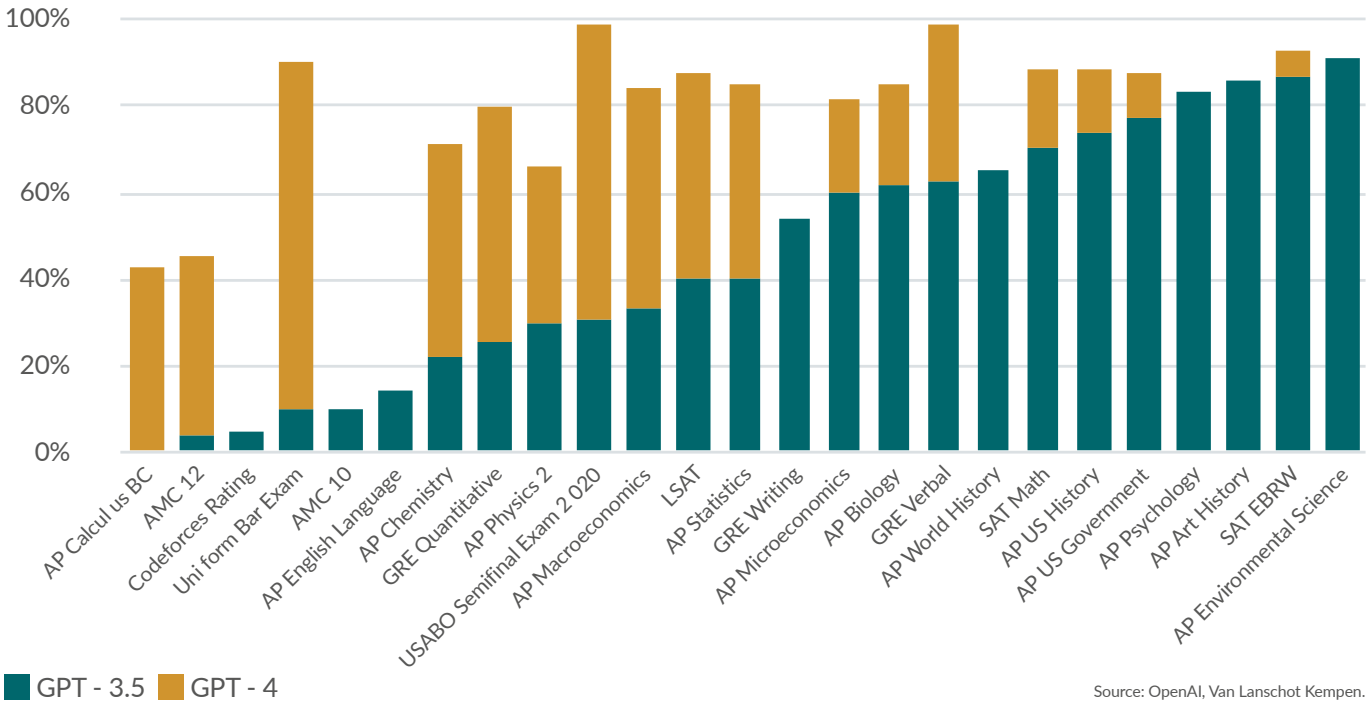

Source: OpenAI, Van Lanschot Kempen.

Chart 7 - Time taken to acquire the first 1 million users for each app

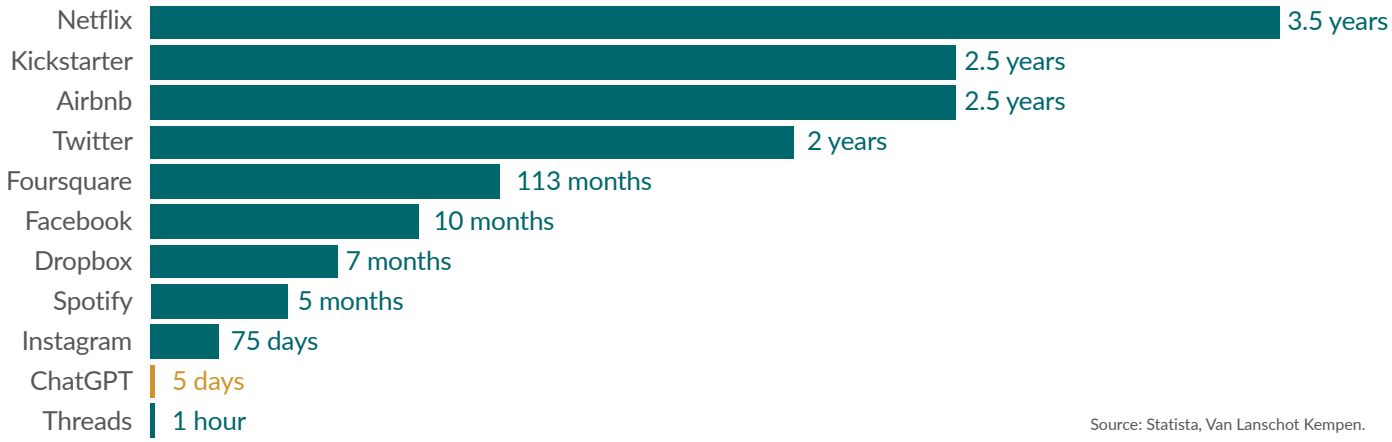

Source: Statista, Van Lanschot Kempen.

Still in the 'early adoption' phase

Despite evidence of several LLM use cases in particular domains and for individual companies, LLM technology has only recently been released to the wider public. As LLMs are still mainly used at home by enthusiasts but not yet broadly adopted by organisations, we would call this the 'early adoption' phase. This makes it premature to gauge if, and by when, this technology will start to have a material impact on broader economies.

In the short term, companies will first need to map where business processes can be assisted by LLMs in an efficient and effective way, and consequently encourage / authorise end users to integrate these into key processes. We expect that affordability, attractiveness and accessibility will keep their appeal, while the technology itself and use cases continue to evolve.

Preconditions for companies to deploy LLMs at scale include allowing employees to experiment, having appropriate data management and technological infrastructure (e.g., integration with key systems), and having clear governance to prevent data leakage.

Users of LLMs need to have sufficient knowledge about how to recognise when a model's output is incorrect. If engineers cannot reliably identify and fix a tool's inevitable errors, or if a model produces low-quality output too often, organisations could lose conviction in the technology.

Over the long term, LLMs could be something that is baked into everyone's job in some form or other.

Chart 8 - Our estimate of the current adoption curve stage for LLMs

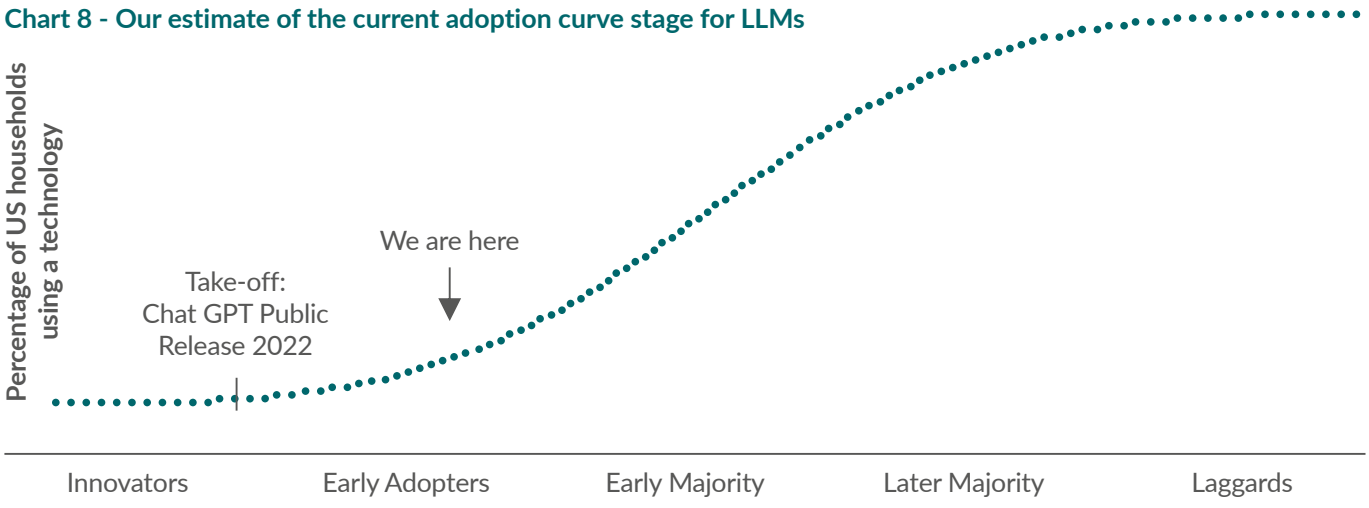

Source: Diffusion of Innovations, Van Lanschot Kempen.

Users of GenAI need to have sufficient knowledge about how to recognise when a model's output is incorrect. If engineers cannot reliably identify and fix a tool's inevitable errors, or if a model produces low-quality output too often, organisations could lose conviction in the technology.

Slight positive economic growth effect

Any economic impact of GenAI on the global economy is expected mainly through labour productivity growth. In recent times, labour productivity growth in developed markets has decelerated, and sits below 1% in most countries (chart 9). Reasons include difficulty in improving productivity in service sectors, already high education levels, the after effects of the financial crisis and that automation generally has focused on low- and middle-wage workers leading to relatively lower cost savings.

Furthermore, in assessing a new technology's impact, we need to distinguish between augmentation and substitution. In an augmentation scenario, adoption of a new technology helps employees to become more productive, while substitution displaces employees.

Although most research on GenAI point to an augmentation effect, some studies[1] find that for some occupations (e.g., communication, supervision and documentation) there is relatively higher potential for substitution. Technological innovation has always created anxiety among workers (e.g., fear of robots replacing humans). Historically, however, innovation has raised labour demand in other sectors and created new jobs, leading to a positive net effect.

For GenAI, several recent academic studies found evidence that assistance of a conversational bot can boost productivity. NBER[2] found that the productivity of customer support agents improved by 14%, with this number rising to 30% for new workers. Another study by Harvard Business School[3] found that BCG consultants completed 12% more tasks and were 25% quicker when assisted. In law, GPT-4 can already pass the Uniform Bar Examination[4].

Chart 9 - Labour productivity growth has declined for decades

Source: OESO, Van Lanschot Kempen.

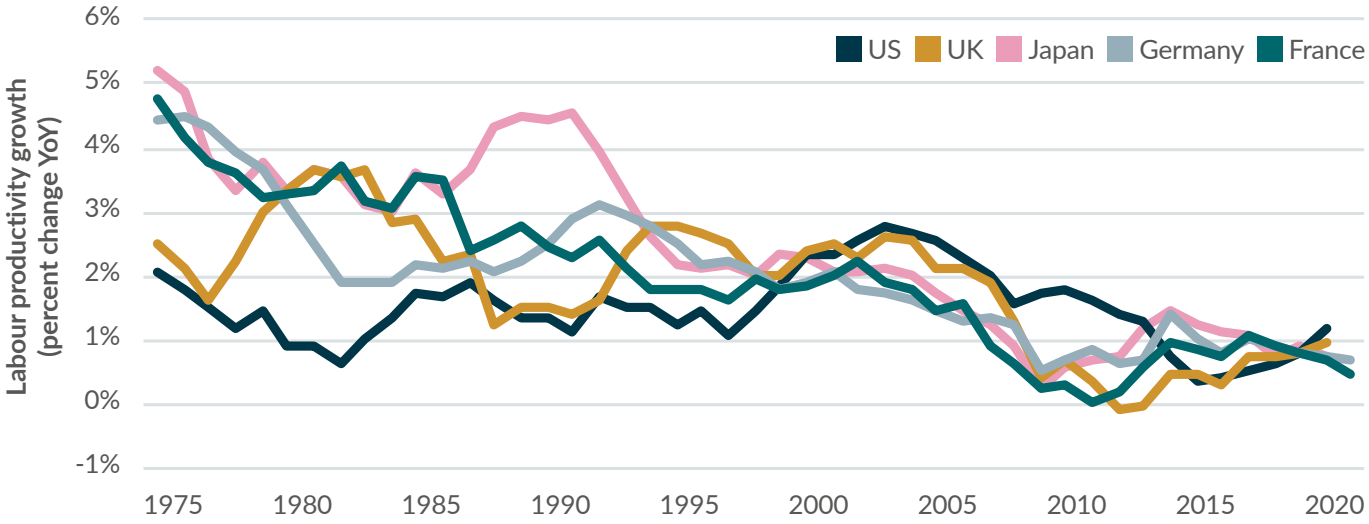

We believe that over the next few years employees may increasingly be assisted by conversational bots that help optimise their productivity across a wide array of tasks. Although it is early days and there is considerable uncertainty around its adoption (e.g. how the technology, regulation and costs evolve and if it ultimately manages to 'cross the chasm'), we expect that GenAI could give a boost to labour productivity growth over the long term (chart 10).

Chart 10 - Potential impact on labour productivity growth could be substantial

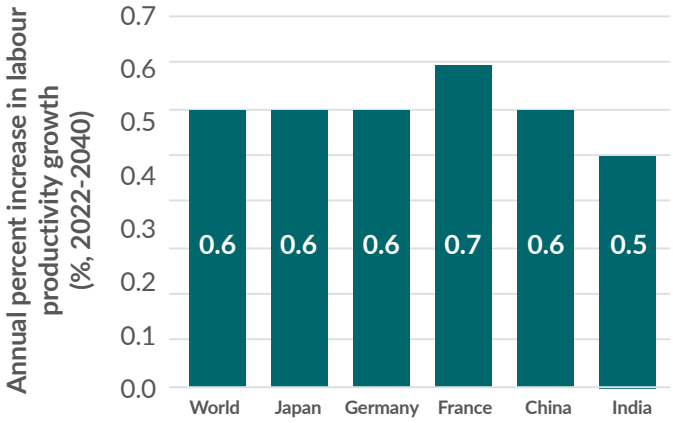

Source: Conference Board Total Economy Database, McKinsey and Van Lanschot Kempen.

[1] ILO 'Generative AI and jobs: A global analysis of potential effects on job quantity and quality', Working Paper (2023).
[2] Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. National Bureau of Economic Research (2023).
[3] Dell'Acqua, Fabrizio, et al. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Harvard Business School Technology and Operations Mgt. Unit (2023).
[4] Katz, Bommarito, Gao, Arredondo (2023). 'GPT-4 Passes the Bar Exam'. Available at SSRN.

A mild deflationary impact

Dependent upon the degree of augmentation and/or substitution, the impact of GenAI on inflation can be analysed through a **cost-push** and/or **demand-pull** lens.

Cost-push inflation

If an employee is more productive due to augmentation, his/her absolute compensation (wage) could be expected to rise. However, since 1979 this relationship has diluted: productivity has grown 3.7x faster than pay, as policymakers tolerated higher unemployment levels to deal with higher inflation, and the globalisation trend led companies to outsource work to lower-wage emerging markets (chart 11).

Chart 11 - Since 1979, productivity significantly outgrew wages

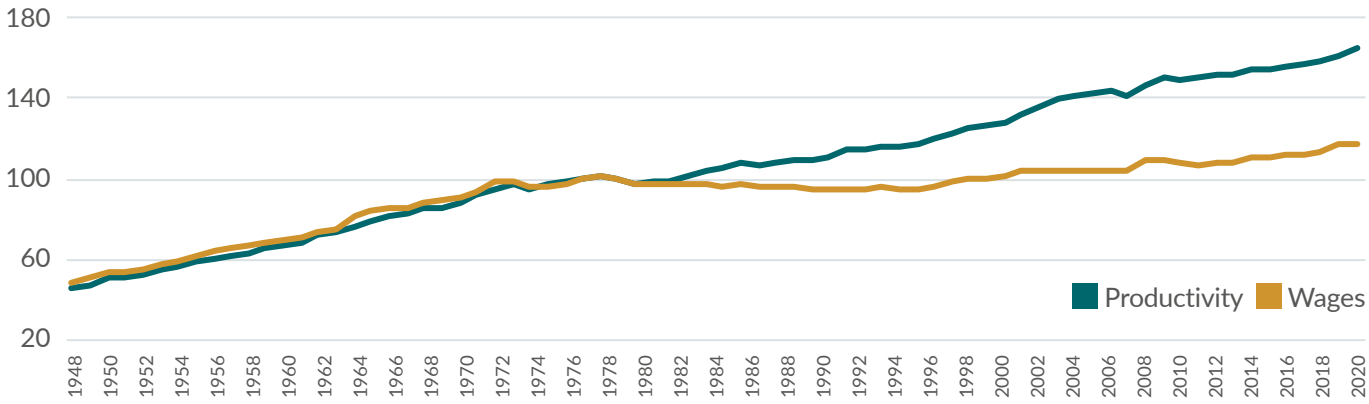

US data, ratio of net productivity and production compensation of the total economy. Source: Economic Policy Institute, Van Lanschot Kempen.

However, this may reverse amid more onshoring due to rising geopolitical tensions, and as ageing global populations lead to structurally tighter labour markets. An increase in productivity through successful GenAI adoption would not necessarily result in higher wages. Firms can produce more at similar costs, reducing the average selling price, which has a deflationary effect. Similarly, in case of substitution, until GenAI can match or exceed human abilities, we see the cost of GenAI capped at the marginal cost of substitution by labour – companies will only apply GenAI for substitution if it helps to bring costs down, and therefore it will likely be a deflationary force.

Chart 12 - Goods productivity has been relatively less compensated

US economy data. Gross wages and social contributions as a percentage of gross value added. Source: OECD, Van Lanschot Kempen.

Demand-pull inflation

If successful GenAI adoption leads to improved productivity, this could lead to higher output, which in turn increases society's ability to spend, fostering inflation. However, this depends on how it affects income distribution (in)equality. The more (all) workers benefit, the higher the risk that inflation prevails. Whereas, in an unlikely scenario where the broader unemployment rates would rise as employees are substituted, the impact is likely to be deflationary.

Overall, we expect a mild deflationary effect from the cost-push side, while we see a more mixed effect from the demand-pull side, dependent upon how it affects income distribution (in)equality. A deflationary effect is welcome in a global context that has mostly inflationary forces such as ageing of the labour force, the energy transition and higher defence spending.

Possible risks

There continues to be much talk of the potential risks surrounding GenAI both in the near term and longer term.

<table>
<tr><th>Near-term risks</th><th>Long-term risks</th></tr>
<tr><td>Hallucinations and misuse: generating incorrect, made-up content that can sound plausible and convincing: this could lead to scams (e.g. ‘deepfake’) and erosion of trust in media and governments and increased polarisation.</td><td rowspan="2">Job displacement could lead to social unrest: either by redundancy of work positions or by workers’ inability to use the augmentative capabilities of newly developed tools and respective productivity gains. Europe generally has stronger worker protection laws and higher rates of union membership, which makes this threat less likely.</td></tr>
<tr><td>Data privacy and copyright infringements: revealing private, sensitive data without consent, or plagiarism.</td></tr>
<tr><td>Generalisation of content: a computer system makes discriminative assumptions about a population (Friedman and Nissenbaum, 1996). Think of a credit-scoring algorithm that associates one’s last name to criminal activities.</td><td rowspan="2">For AI besides LLMs, it could spur development of autonomous weapon systems (e.g. drone swarms, tanks with autonomous iteration and targeting); this is expected to materially alter modern warfare.</td></tr>
<tr><td>Environmental footprint: GenAI uses enormous amounts of energy (carbon emissions) for computation and storage purposes, and millions of gallons of fresh water to cool data centre equipment.</td></tr>
<tr><td>Too much reliance on the accuracy of GenAI: Humans ‘falling asleep at the wheel’, not spotting imprecisions[3].</td><td rowspan="2">Systems could become more intelligent than humans: we need to ensure that AI objectives are aligned with ours. This threat is probably still decades away and unlikely to materialise by simply scaling up GPT.</td></tr>
<tr><td>Financial markets: it is difficult to forecast the future direction of new technologies; companies identified as GenAI global leaders (e.g. the ‘Magnificent 7’) may be pricing in GenAI potential that may never materialise.</td></tr>
</table>

To mitigate these risks, many countries are developing regulatory frameworks, although at different speeds and stringency levels, and, so far, with limited coordinated effort. We believe that governments, in the end, will be able to control the development of GenAI and ensure that the positive impact of the technology will outweigh its risks to society.

[3] Dell’Acqua, Fabrizio. Falling asleep at the wheel: Human/AI collaboration in a field experiment on HR recruiters. Working paper (2022).

An uneven regulatory landscape

In May 2023, hundreds of leading figures working for AI influential companies such as OpenAI, Microsoft, Anthropic, and Google's DeepMind released a statement that "mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war."

In November 2023, during the AI Safety Summit, several countries issued the Bletchley Declaration which called for international co-operation to effectively handle the challenges and risks of AI. Overall, regulators have expressed their goal of developing appropriate policies; however, they are not all acting at the same pace.

European Union: first to act, stricter policy

AI Act: In April 2023, the European Commission proposed the first EU regulatory framework for AI. It suggests that different AI systems should be analysed and classified according to the risk they pose to users. Then, the different risk levels would be translated into different layers of regulation. The AI Act also proposes to make model creators liable for how the models are used – picture the inventor of the telephone being responsible for what is discussed during a call. This is causing dissatisfaction among EU businesses, as they regard that this legislation would harm their competitive capabilities via disproportionate compliance costs.

In November 2023, France, Germany and Italy came out with a joint paper supporting "mandatory self-regulation through codes of conduct" and opposing untested norms. The AI Act was approved in December 2023, but it includes a grace period of about two years before becoming law.

United States: early stage, taking a broad view

Although the White House has released several guiding documents on the risks of AI, the US have let the industry **self-regulate** up to now. Legislators are undertaking a comprehensive study to first identify which aspects of the technology may require new regulation and which can be addressed by current laws.

In July 2023, Microsoft, OpenAI, Google, Amazon and Meta came together to sign a series of voluntary commitments for the development of AI. They agreed to emphasise 'safety, security and trust'. In November 2023, the US announced the creation of an AI Safety Institute .

With its abundant human and capital resources, the US exerts significant influence on global technological development, and so their regulatory framework, or lack of it, will be a key determinant for the industry.

China: in-between regulation and competitiveness

Chinese regulatory development sits in between their EU and US counterparties. The country has already issued 24 guidelines that aim to provide a framework for the ethical and responsible development and use of AI. These require, for instance, that developers register their services and conduct a security check before bringing products to the market. The country also wants to control the output of the models.

Once the Chinese agencies opt to implement effective regulations, their strong governmental control should streamline the process of imposing them. However, lawmakers in Beijing must perform a **balancing act:** ensuring enough state control over the technology without harming firms' ability to compete on a global scale – especially if US companies are left to regulate themselves without oversight.

Investment opportunities

New technologies stand to benefit different industries at different times. GenAI first helped software developers in 2021 with a code-writing assistant. In 2022, GenAI started assisting consumers with producing words and pictures (ChatGPT, DALL-E). As investors grew more excited about the technology's prospects, technology firms and data centres benefited in 2023 as key suppliers of (inputs for) GenAI and/or the supporting infrastructure. This has been reflected in the share price growth of these 'Magnificent 7' companies since early 2023 (chart 13). Nvidia, a producer of graphics processing units (GPUs) that power GenAI applications emerged as the clearest winner so far.

Chart 13 - Strong share price growth for the 'Magnificent 7' companies since 2023

Source: Refinitiv, Van Lanschot Kempen.

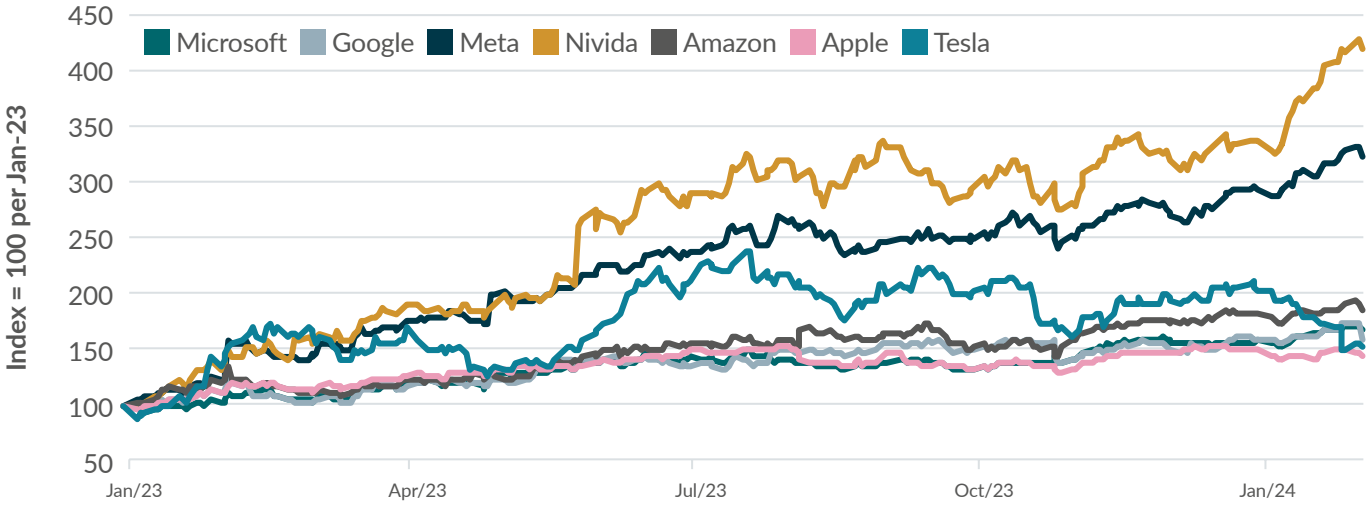

A risk of investing in these types of growth companies is that it is difficult to predict the future of a new technology like GenAI, while financial markets often overestimate the potential of new technologies. As a result, certain stocks may price in GenAI potential that may never materialise. For 2024, we expect that, following these tech giants, a broader group of companies in the service sector, after a year of experimenting, will start deploying GenAI on a larger scale. For these companies, GenAI can increase productivity growth and profit margins, while eventually reducing their cost base. Although it looks promising that the top 3 sectors in broad stock indices for developed countries (e.g. MSCI World) are at the forefront of GenAI implementation, we expect it will take time to integrate the key benefits of this new technology into value chains, and we expect to see significant differences between and within (sub)sectors.

Chart 14 - Valuations lofty for 'Magnificent 7'

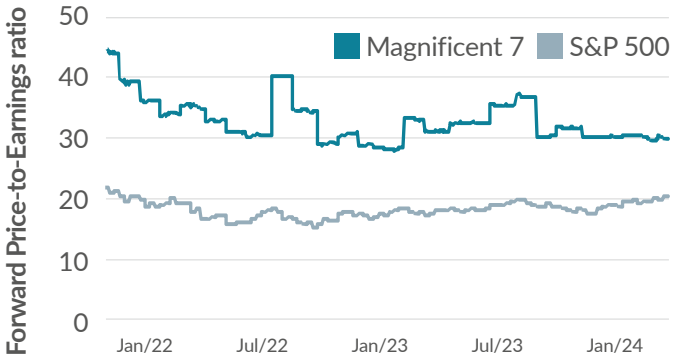

Chart 15 - Sector composition MSCI World

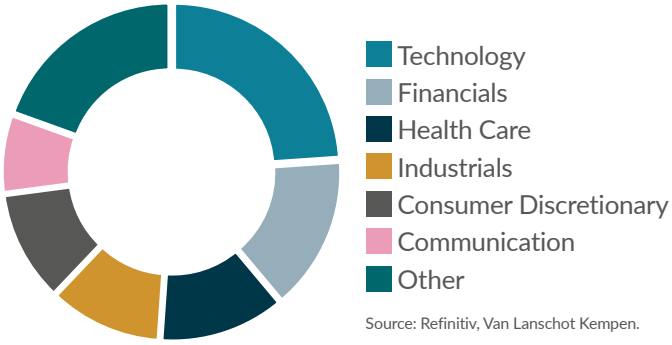

Source: Refinitiv, Van Lanschot Kempen.

To successfully implement GenAI, companies need to establish the right data and technology infrastructure, set up a minimal governance structure regarding this technology, and work on employee understanding and empowerment (taking into account the 5 different recipients). Similar to previous technology waves, such as the rise of smartphones and the internet, smaller companies usually adapt the fastest, as they are often more agile than large companies and see new technologies as a means to gain an edge. This could provide opportunities for (bottom-up) investors in small-cap companies.

Authors:

Pieter Heijboer
Head of Investment Strategy & Research

Daniel Costa
Analyst Client Portfolio Management

The Netherlands
Beethovenstraat 300
1077 WZ Amsterdam
+31 (0)20 348 8700

United Kingdom
20 Gracechurch Street
London, EC3V 0BG
+44 (0)20 3636 9400

France
16 Cours Albert 1er
75008 Paris
+33 (0)18 375 6273