ROBECO
The Investment Engineers

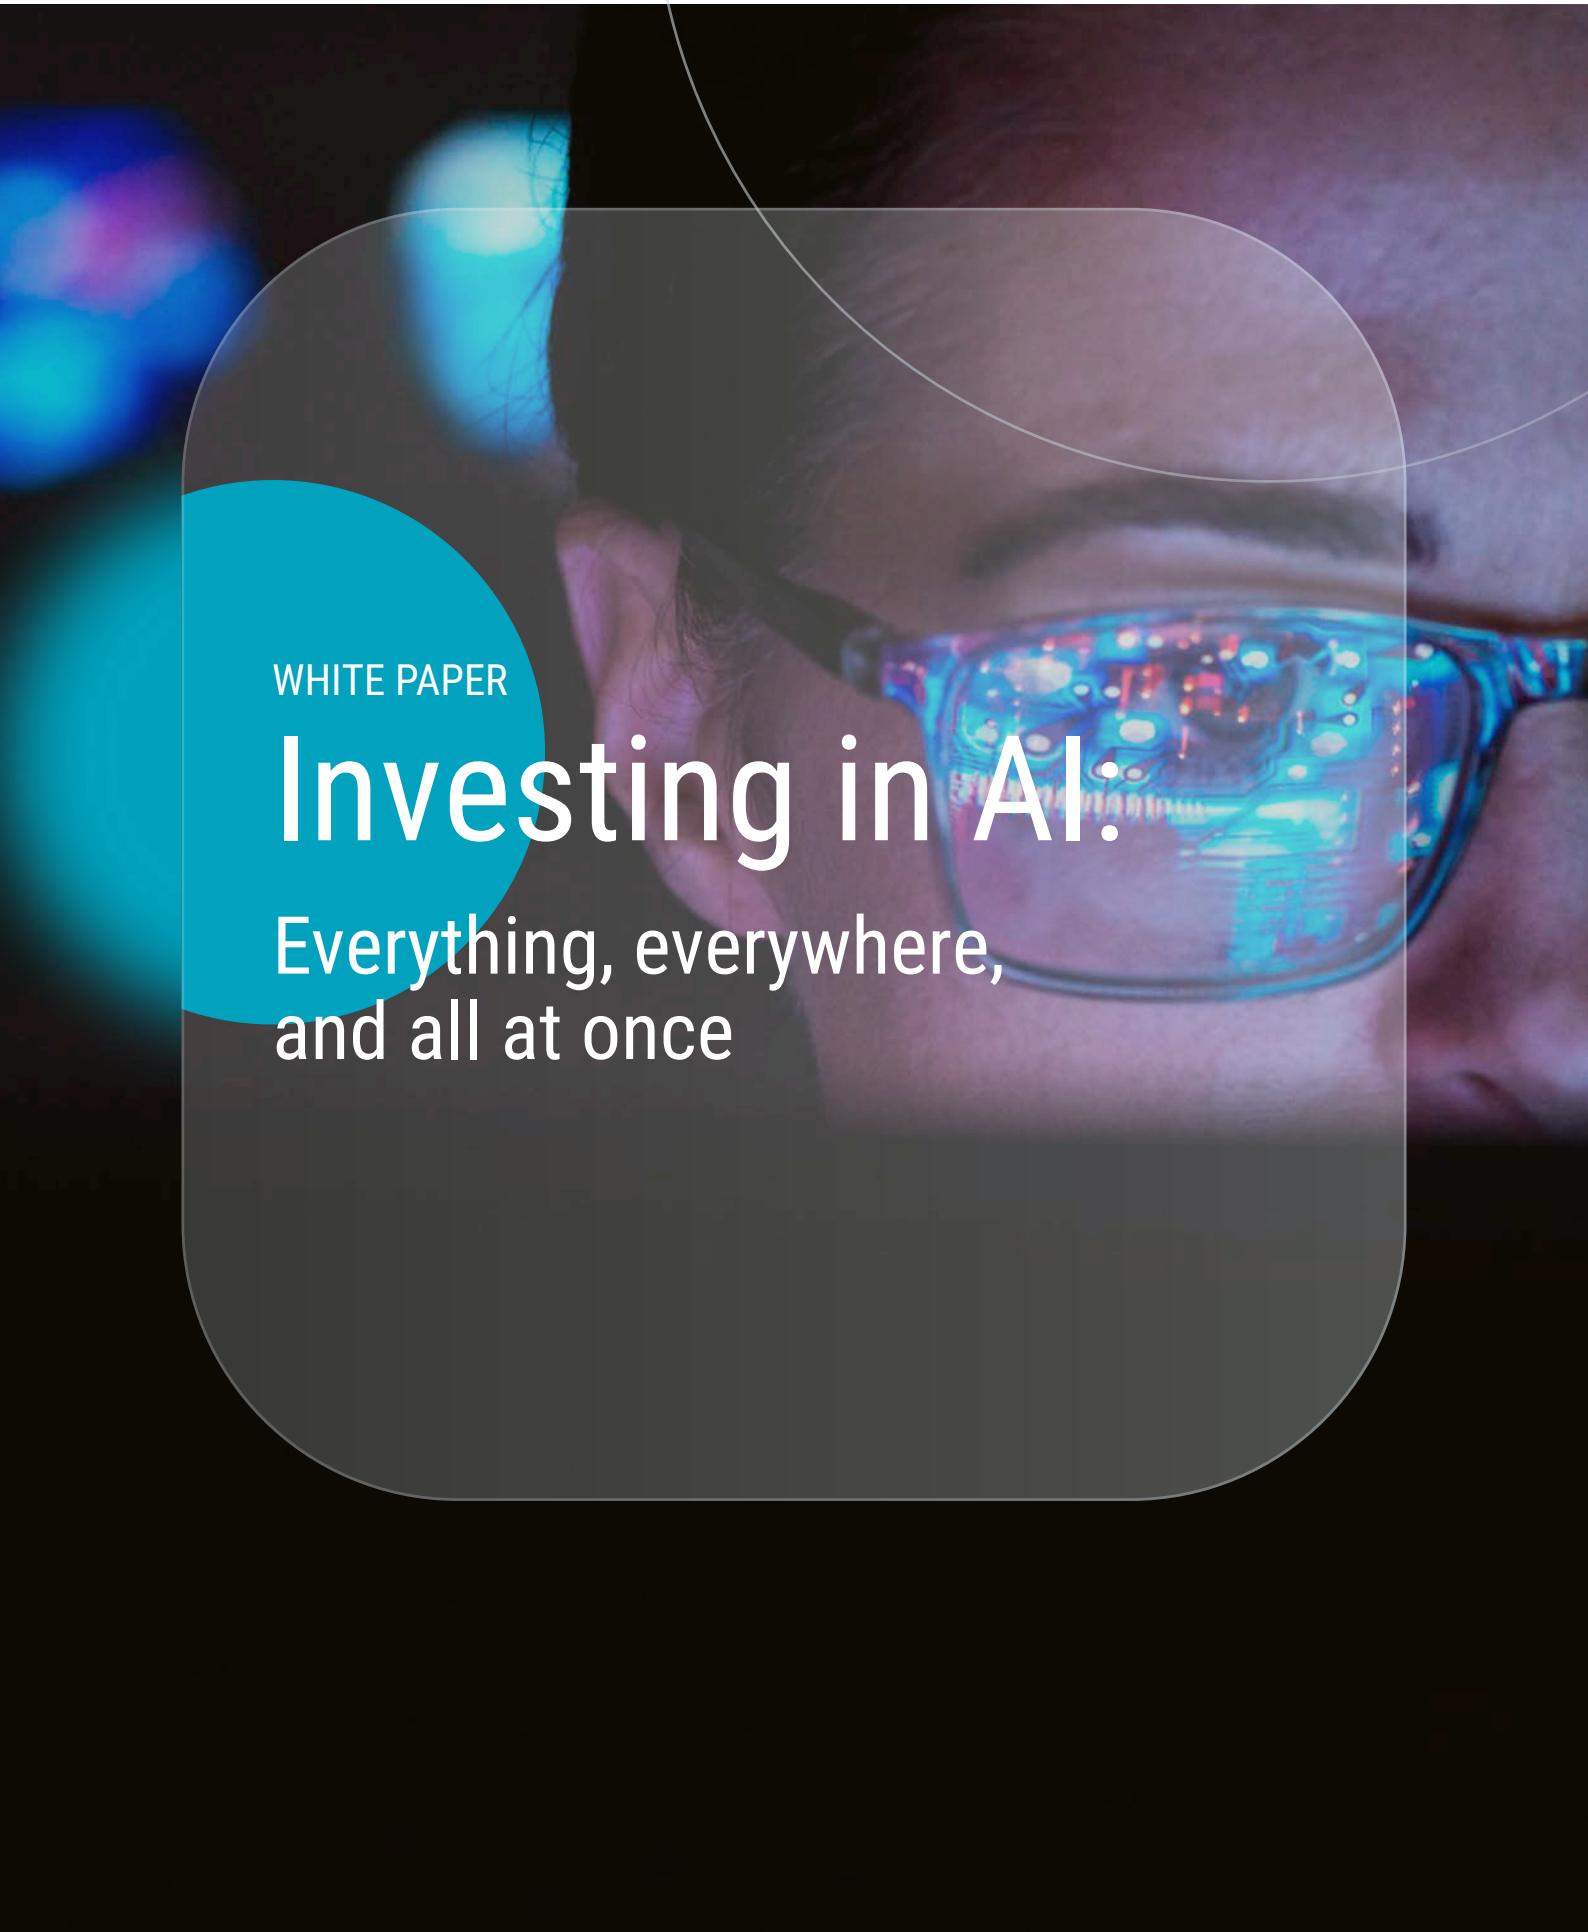

WHITE PAPER

Investing in AI:

Everything, everywhere, and all at once

Contents

WHITE PAPER MAY 2024

Marketing material for professional investors, not for onward distribution.

Daniel Ernst
Portfolio Manager,
Robeco Digital Innovations

Steef Bergakker
Portfolio Manager,
Trends Investing

The state of AI

Since the launch of services like ChatGPT, Midjourney and Gemini, artificial intelligence (AI) has captivated the minds of consumers, corporate executives, and investors alike. These newer generative AI services built on large language models (LLMs) and trained on vast datasets of online text, audio, and images, produce detailed and varied content that is often difficult to discern from human response. While the hype surrounding generative AI has faded from the peak, implementation of the technology remains in the early stages of what appears to be a multi-year cycle for the technology. At the same time, applied AI technologies continue to advance and impact a broad array of applications including cybersecurity, logistics, manufacturing, and medical research. With a focus on the long term, evidence from prior technology cycles suggests that a greater share of value creation from AI will accrue to those providing foundational technology and differentiated applications.

Great expectations and inflated claims

In the 1960s, Stanford University professor Roy Amara noted that "we overestimate the impact of technology in the short term and underestimate the effect in the long run." It is also possible that sometimes the reverse is true. The public excitement, and tangible progress, have also led to increased expectations for AI's societal and economic impact. For instance Goldman Sachs estimates that up to a quarter of all jobs in the US and Europe could be automated by AI, while PWC forecasts that AI will add USD 15.7 trillion to global GDP by 2030. Initially, technology has seen the greatest measurable impact from the current AI investment cycle. With that in mind, Bloomberg Intelligence forecasts that by 2032, revenues from generative AI will reach USD 1.3 trillion, and account for 12% of all information technology spending. Given generative AI is a subset of the field, the forecast implies that total AI technology revenues will be even higher. However, the slope of these interest and forecast curves suggests to us that AI is presently closer to the peak of the hype cycle than it is to the plateau of productivity

Backing up the enthusiasm for AI, and underpinning much of the innovation, has been a considerable level of investment into the technology. Over the last five years, USD 230 billion across more than 15,000 transactions has been invested into AI start ups by venture capital firms according to CB Insights. While generative AI has garnered the most attention, and an increasing share of investment, other forms of AI have captured more than 80% of venture funding over that time. Similarly, while ChatGPT's creator, OpenAI has received the most attention, at the close of 2023 there were 186 AI Unicorns — private firms valued at more than USD 1 billion — demonstrating the breadth of the field. As is typical in venture capital, not all of those investments will deliver positive outcomes. However for public markets, consumers and enterprises, it will be exciting to see what the successful new AI-powered applications can accomplish.

Figure 1 – Worldwide venture funding for AI startups, USD billions

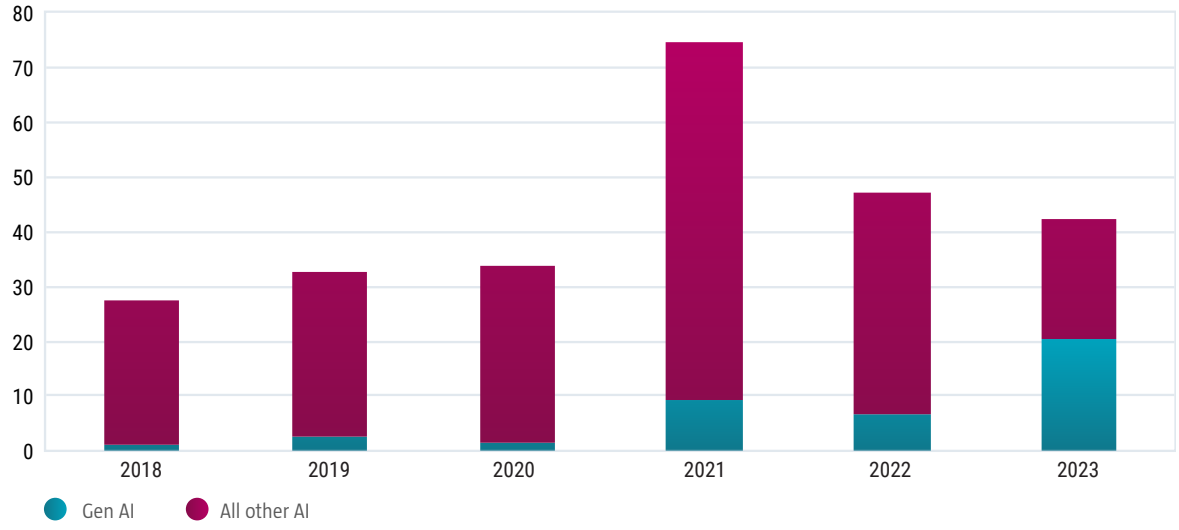

Source: CB Insights, January 2024

Although the current set of companies with clearly measurable revenues derived directly from AI is somewhat limited, more than a third of companies in the S&P 500 discussed AI on earnings calls during 2023. Under the assumption that some companies are further along than others in deploying AI, our quantitative investing group leveraged natural language processing to filter abstract and vague statements from more definitive signals. Based on this analysis, Robeco found that over the last three years, among those companies that mentioned AI, the proportion providing more concrete examples and statements has risen from 38% during the 4Q20 reporting period to 71% during the 4Q23 reporting season. Time will tell if these AI signals turn into a sustained lift in revenues and earnings.

Figure 2 – Mentions of AI within S&P 500 member's earnings calls

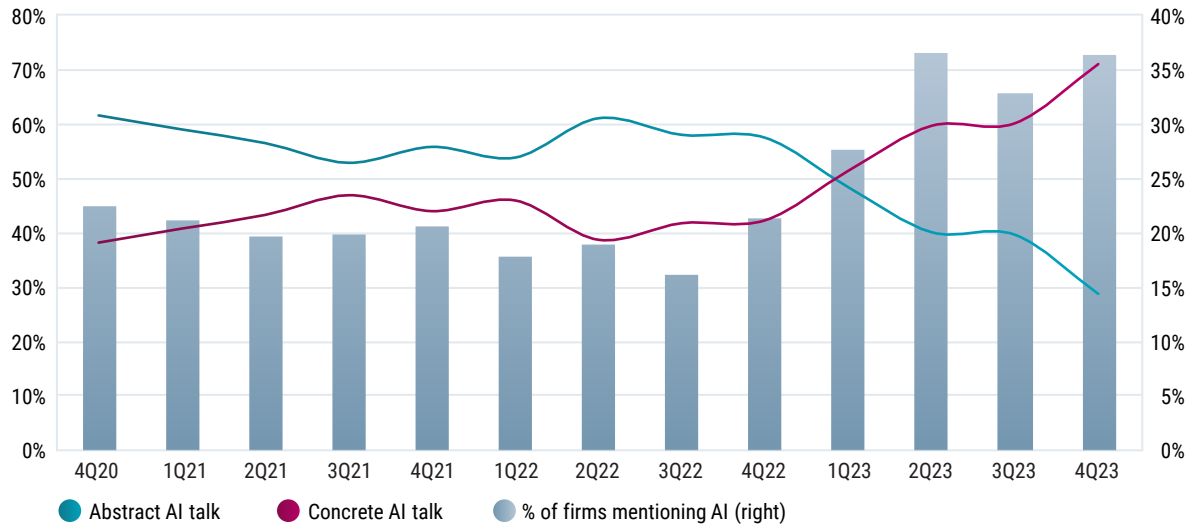

Source: Robeco, Factset, March 2024

Building the AI foundation

The rapid acceleration in the capability of generative AI was preceded by decades of research and development in the critical building blocks, including statistical inference, expert systems, machine learning, adversarial networks, and neural nets. As AI models have improved, the technology has also required increasing amounts of training and computational capacity to power the systems. Before 2010, according to a 2022 paper by the University of Aberdeen's Jaime Sevilla et al., AI training computational requirements increased in a similar pattern to Moore's law, doubling roughly every 20 months. Since 2015, with the development of large-scale machine learning models like those deployed by OpenAI's ChatGPT, there has been a 10 to 100-fold increase in the computational requirements of AI systems. For instance, in 2016, Google's AlphaGo model, known for beating the world champion at the game Go, required 1.8 million petaFLOPs (one quadrillion floating point operations per second. Released in 2020, OpenAI's GPT-3 required 314 million petaFLOPs to be trained and released in 2023, GPT-4 required an estimated 2.1 billion petaFLOPs.

While the computational requirements of AI have accelerated, so too has the efficiency of semiconductors. For instance, according to AI-focused research group Epoch, in 2020, Nvidia's RTX 3080 GPU could process 42.6 billion FLOPs (Floating Point Operations Per Second per USD versus 2.6 billion FLOPs per US dollar for the GTX 465 GPU in 2010. These improvements, coupled with the emergence of on-demand computational resources of cloud infrastructure services, have greatly expanded the capability and reach of AI software.

Perhaps just as important for AI development has been an expanding ecosystem of software development tools. For instance, Nvidia developed its CUDA (Compute Unified Device Architecture platform in 2006 to enable software developers to offload mathematically-intensive operations from the CPU (central processing unit to the GPU. Today, CUDA boasts 3.5 million software developers using the platform and an expanding set of available tools built by that ecosystem. Similarly, the startup Hugging Face, offers an expanding library of more than 172,000 models and 28,000 datasets enabling AI developers to assemble building blocks from image recognition to translation within their own applications without re-creating those functions from scratch.

Figure 3 – AI training computational requirements in petaFLOPs (x 1,000,000)

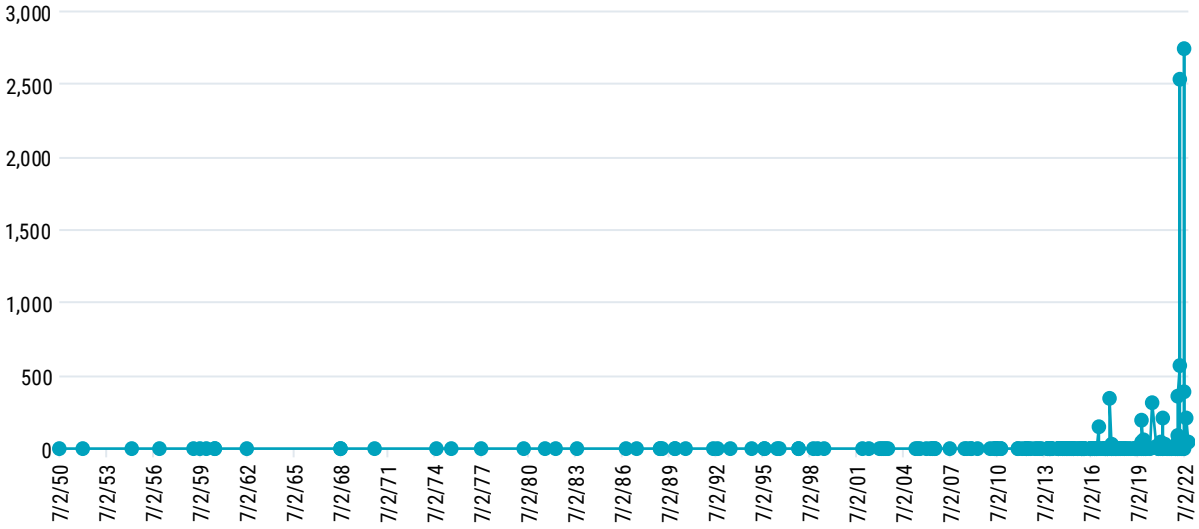

Source: "Compute trends across three eras of machine learning", Sevilla et al., March 2022

The net result of this expanding set of resources has been a Cambrian explosion in AI development, from alternative foundation models to narrow applications. Although OpenAI's ChatGPT was the first to market with a publicly available service built on the development of large language models, they are no longer alone. Driven by heightened interest, increased funding, and the availability of computation as a service, dozens of alternative foundational models have also been released. Similar to OpenAI, some of these entrants, like venture-backed Perplexity, offer consumer-facing services. Others, like DataBrick's MosaicML, allow corporations to build models from the ground up that are trained on company or industry-specific data to improve accuracy and focus functionality on narrower use cases like drug discovery or legal research.

Supporting these competitive model development trends is a growing body of Open-Source tools, applications, and other resources that enable programmers to accelerate development. On Microsoft's software collaboration platform GitHub, for instance, the number of open-source AI projects more than doubled over the last five years to nearly 420,000 at the close of 2023. Although building an accurate and well-functioning LLM still requires tremendous resources to configure datasets and access computational power, competitive models are closing the performance gap with well-funded leaders. According to Stanford University's 2024 AI Index report, in 2023, 149 new foundation models were released, more than double the prior year. Notably, two-thirds of new models were open sourced versus 44% in 2020. Moreover, an analysis of published Massive Multi-task Language Understanding (MMLU) scores compiled by CB-Insights demonstrates that open-source models are gaining on well-funded private models. Model development is also expanding beyond established technology hubs with the rise of regional and sovereign nation sponsored models.

Figure 4 – Benchmarking large language models

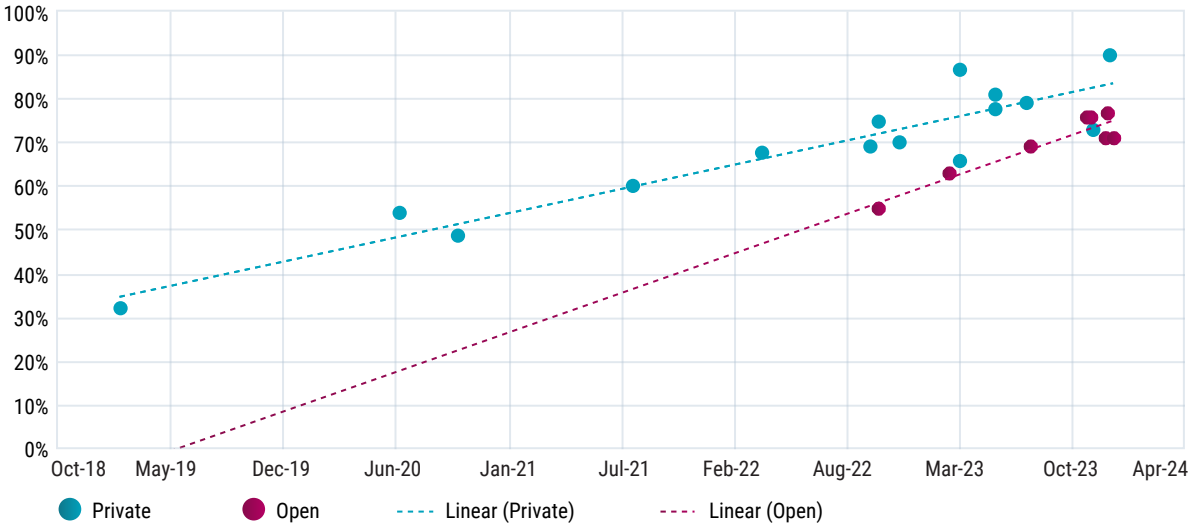

Source: CB Insights, April 2024

Passing tests and failing others

Generative artificial intelligence systems have become remarkably adept at producing coherent answers to an ever widening array of topics, but they are also very often wrong. To test whether ChatGPT produced accurate answers, OpenAI put the chatbot literally to the test. In 2023, GPT-4 scored in the 90th percentile at exams in the fields of law, biology and microeconomics. While the tool did less well at other tests like English literature, calculus and programming, scoring at or below the 50th percentile, the result was material improvement over the company's prior version, GPT-3.5. Intriguingly, the AI model also did well on a sommelier test, despite not possessing taste buds.

However, in many instances generative AI models output factual errors and fabricate answers. In one of the more notable errors, In 2023 ChatGPT provided multiple detailed, yet non-existent, case law precedents and judicial findings that lawyers submitted to a US federal court. While opposing counsel quickly identified the bogus arguments, material submitted by generative AI systems do not typically have an expert fact checker looking the answers over. Problematically, these so-called hallucinations are inherent to the structure of large language models. As Yann LeCunn, New York University Professor and Meta Platforms AI Chief Scientist explains, "When you have a system like this that basically predicts one word after another, they are difficult to control or steer because what they produce depends entirely on the statistics they trained on and the given prompt. Mathematically there is a good chance that it will diverge exponentially from the path of correct answers. The longer the answer that is produced the more likely you end up producing complete garbage."[1]

Even in controlled model environments AI systems continue to make errors. Imagenet, a program which began in 2009 and formed the research basis for today's advancement in neuralnets, is a structured and human annotated database of 14 million images. However, despite continuous improvements in the algorithm and step function improvements in computational capacity, Imagenet accuracy reached rates of just over 92% in 2023, up from 90% two years in prior. Similarly, in another complex challenge of model accuracy focused on segmenting object types in the three-dimensional space of city streetscapes, the most accurate models currently achieve scores of less than 90%.

1. The 'Godfather of AI' Says Doomsayers Are Wrong and ChatGPT Isn't All That Innovative, Barron's, April 2023

Figure 5 – OpenAI GPT model test results

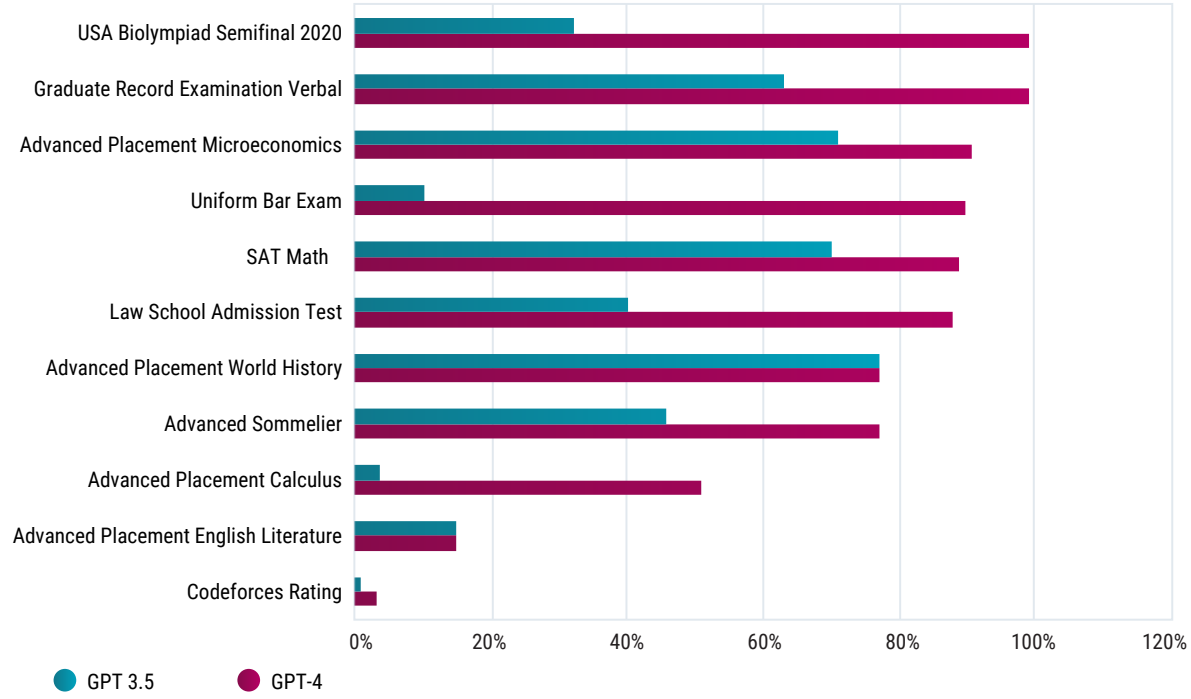

Source: OpenAI, 2023

Getting things done

Like many software applications, a significant challenge for AI is to bridge the digital world with the physical. Before being acquired by Apple in 2011, the founders of Siri, the voice-activated assistant, described the service as a "do-engine" that could, for instance, book a table at a restaurant and arrange for a taxi to get there. However, most services prefer to have their own application that consumers can interact with directly rather than sit behind another providers interface. More than a decade later, Siri and other voice assistants like Amazon's Alexa work best with their own digital services, while integration with third-party services remains limited. While still very early in the process, the functionality of generative AI services like OpenAI's ChatGPT and Anthropic's Claude remain similarly siloed. At present, many generative AI users merely cut and paste the output of those tools from the web into another application like a Word document – hardly a seamless, automated process. On the other hand, many software developers are building generative AI features into their applications. Most notably, Microsoft offers Copilot, a generative AI upgrade for its core Office applications like Word and Teams, as well as GitHub for software development. Similarly, Adobe developed Firefly to insert generative AI-created images into Photoshop, and Salesforce has created an AI tool to help users gather insights and communicate more effectively with customers.

A thematic approach to AI

Change, as the Greek philosopher Heraclitus said, is constant. Throughout history, society has been shaped by evolving structures, this is true for the agrarian age, the industrial age and more recently the information age. Each era is also susceptible to change from within as priorities shift. This has been especially the case within the information age as innovation shifted from mainframe computers to personal computers, and then to the internet, mobile phones, and the cloud. Thematic investing seeks to capture the value creation potential of these transformative technologies.

AI appears to represent the next major transformative shift that will impact both the technology sector and the broader economy. AI is a complex technology building on a deep foundation across semiconductors, software, and hardware. The specific and expanded computational requirements of AI are greatly impacting technology infrastructure spending and reshaping software application. Outside of technology, AI has risen to the top of the agenda in industries as diverse as agriculture, manufacturing, and media. As a result, a systems-level approach is required for corporates to build solutions and for investors to identify structural winners.

Given the pace of growth and the breadth of components involved, a common strategy for new technology trends spreads investments across a basket of potential beneficiaries. Likening the start of any given technology cycle to a gold mining rush, component suppliers are often referred to picks and shovels with strong growth prospects. However, like picks and shovels, many component suppliers are interchangeable and their pricing is prone to commoditization. As a result, we prefer to focus on foundational technologies and differentiated solutions that are not so easily replaced.

In previous technology eras, a handful of firms providing foundational IP captured the majority of the value creation across the cycle. For instance, this was the case for IBM in the mainframe era, Microsoft and Intel in the PC era, and Apple and Qualcomm in the mobile era. During the internet era, while many focused applications thrived, broader platforms exhibiting economies of scale and scope like Google in search, Amazon in ecommerce and Meta in social media generated more lasting returns in their respective fields of the internet.

In the case of artificial intelligence, given the complexity of the technology, longer-term beneficiaries will need to provide a foundational element that demonstrates its utility and is not readily replaced. Further, as within the internet and cloud software markets, providers demonstrating economies of scale and scope will be positioned to generate superior returns. For instance chipmaker Nvidia has thus far positioned itself as the effective standard for developers by pairing its ultra fast parallel processing technology with their widely adopted software platform, CUDA. Similarly, Microsoft which combines its own Office applications with its cloud computing service Azure, and its partnership with foundation model developer OpenAI, has positioned itself as multifaceted AI platform.

While the impacts of structural changes and transformative shifts are broadly felt, there are typically few beneficiaries. This concentration of winners has also played out more broadly in equity markets. According to a study by Hendrik Bessembinder from Arizona State University, a typical stock listed in the US from 1926 to 2019 had a buy-and-hold return of negative 2.8% over its entire lifetime.[2] Furthermore, out of 25,000 stocks listed between 1973 and 2020 in the US, 13% achieved a 25x return.[3] In the case of companies with a capitalization of at least USD 100 million before the 5x return is achieved, just 4% achieved a 25x return. Another Bessembinder study found that when examining the returns of 64,000 global stocks, the top performing 2.4% of companies accounted for 100% of the USD 75.7 trillion in wealth creation between 1990 and 2020.[4]

As a technology cycle develops, these rising stars can typically be identified by their capacity to generate returns in excess of the cost of capital.[5]

2. See Bessembinder (2020)
3. See Bessembinder (2021)
4. See Bessembinder (2023)
5. "A systematic approach to identifying companies' life stages" - Bergakker, Chung, August 2023

Figure 6 – Distribution of US equity market returns between 1973 & 2020

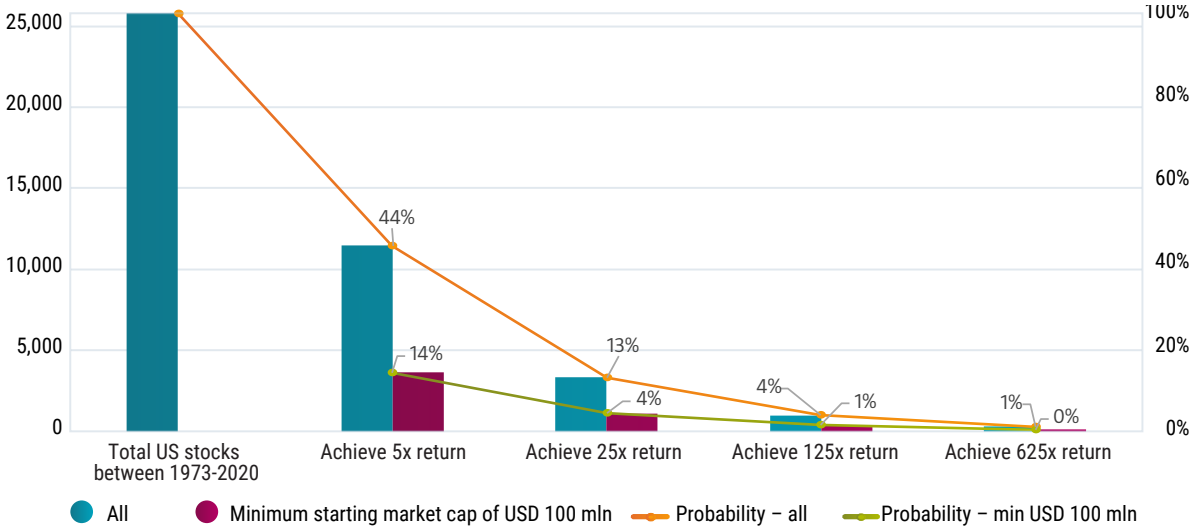

Source: H. Bessembinder, Arizona State University, 2021.

That being said, similar to the internet and enterprise software, artificial intelligence is not a singular technology, there are many forms of AI. As we lay out within this paper, there are a growing set of applications and industries that employ AI. For instance, as Software as a Service (SaaS) models developed, several dominant platforms emerged focused on specific functions like Salesforce.com in customer relationship management, ServiceNow for IT support, and Workday for human resources. The emergence of AI may serve to strengthen those services, or facilitate new entrants. Moreover, while much of the focus today centers on generative AI, applied AI and machine learning tools continue to evolve in applications across automation, cybersecurity, logistics, and robotics. As a result, there exists broad potential for multiple profitably scaled businesses to emerge from the new wave of AI development.

Industrial AI

While consumer applications like chatbots and image generators have garnered the most attention recently, artificial intelligence has been the center of industrial strategy for more than a decade. Introduced in 2011 by the German government's Industry-Science Research Alliance, the term Industry 4.0 described a system where robotics, sensors, data, networks, and computers work in concert to improve efficiency. While significant progress has been achieved, much of the smart manufacturing potential remains an ongoing process. Next-generation artificial intelligence coupled with advancements in cloud computing and software hold potential to accelerate the digital transformation of industry.

The industrial revolution is a continuous and ongoing process spanning centuries. Each phase of the industrial revolution brought about profound changes in production methods, efficiency and the overall economic landscape. In the first phase of the industrial revolution, production shifted from specialized manual labor to steam powered mechanization. In the second phase of industrialization, mechanized production was optimized through the assembly-line process and the introduction of electrical power. The introduction of industrial robotics more than fifty years ago lifted productivity further in the third phase of the industrial revolution.

In the current and fourth phase of industrialization, also known as Industry 4.0, production is becoming more intelligent, interconnected and data-driven. For instance, by augmenting robotics and other infrastructure with sensors connected by networks to software, the Internet of Things (IoT is already having a material impact on operations and margins. For example, working with Microsoft, BMW deployed an IoT system that the automaker credited with reducing one factory's metal-press shop downtime by 25%. Similarly, combining advanced visual optics with machine learning, Keyence vision-measuring systems reduce downtime, spot defects, and improve quality in a broad range of settings from logistics to medical device manufacturing.

Commercial AI applications are already hard at work doing mundane and repetitive tasks behind the scenes including in energy management, predictive maintenance, and logistics. AI has shown value in demand forecasting, and supply chain analysis. AI and high performance computing are able to process large datasets, identify trends, and generate new insights. Working in concert, these digital innovations hold promise to lift productivity, improve quality, lift margins, and reduce waste. These innovations also arrive at a time when production bottlenecks, labor shortages and geopolitical security considerations have driven renewed interest in automation, shortening supply chains and reshoring manufacturing closer to home.

Figure 7 – Industry 4.0 combines sensors, data, cloud computing, robotics, and artificial intelligence to form smart manufacturing

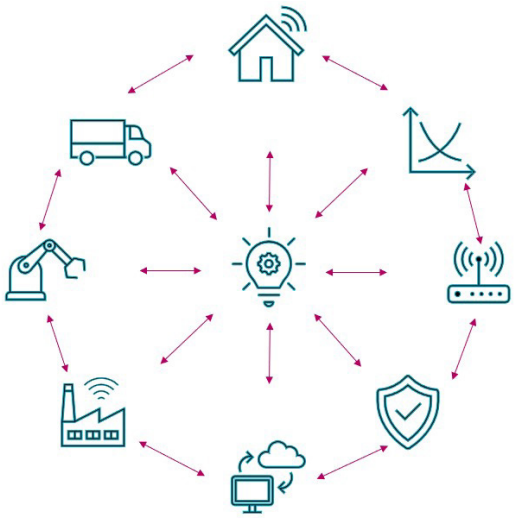

Source: Robeco

Importantly, the margin gains from everything from reduced material waste to optimized energy usage yield environmental benefits. With this in mind, Schneider Electric developed an IoT and data-analytics platform that has reduced energy costs at its own production sites by as much as 26% and now offers the solution to customers as a value-added service. McKinsey estimates that digital innovations including IoT have the potential to directly reduce fossil-fuel emissions by 15 percent by 2030 and support a further 35 percent reduction by influencing consumer and business decisions. Demonstrating the continuous nature of innovation, some industry observers refer to the placement of sustainability goals at the center of digital innovation programs as the fifth industrial revolution.

Robotics

Over the last decade, industrial robotic installations have grown at a 13.6% CAGR. McKinsey expects that over next five years robotics will account for 25% of industrial capital spending. Driven by improved functionality and an increased business case, the steady growth of robotics could accelerate. Industrial robots are ideal for repetitive tasks and are often superior to their human counterparts, particularly when the work is dangerous or involves heavy parts. For instance, DHL's network of more than 5,000 collaborative warehouse robots helped increase the number of items picked per hour by 180%. Advancements in mechanical dexterity, machine vision, and on-board intelligence are broadening robotic functionality. Paired with internet-connected sensors and software enhanced by artificial intelligence, smart manufacturing and logistics services respond more efficiently to both demand and supply signals.

Figure 8 – Robot installations and operational installed base

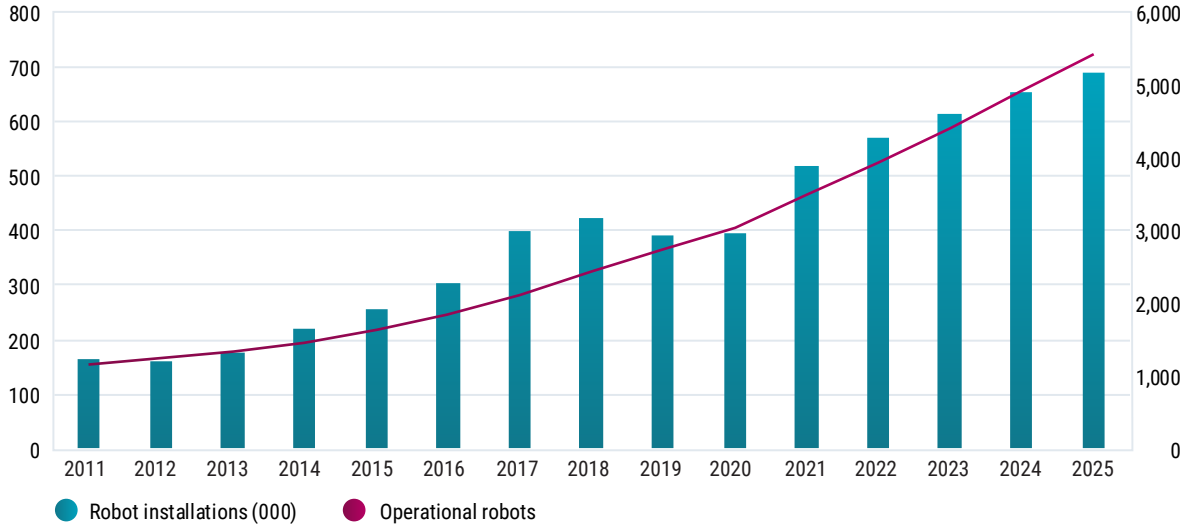

Source: International Federation of Robotics

Learning on the job

Currently industrial robots are either human controlled or follow pre-defined programming in order to operate. More recently, the development of autonomous and semi-autonomous robotics that enable dynamic actions that factor circumstance and real-time conditions, has accelerated. Advancements in machine learning and artificial intelligence are driving this new wave of dynamically optimized robotic production. By learning from human controls and taking input from sensors and cameras, robotic production is poised to become more efficient and more flexible. For instance, Google DeepMind's RT-2 model was trained on both robotic specific data, as well as from the web, to enable machine vision equipped robots to perform tasks that were not preprogrammed. Google expanded on the idea with an updated model, AutoRT, that enables a large language model to suggest tasks for a robot to perform. Over time, generative AI and on board computing may enable robots to not only learn how to perform a task, but to also implement improved methods for a given action. Together, these innovations that facilitate flexible workflows and dynamic actions should improve the return on investment of robotic installations.

Figure 9 – Google DeepMind robotic arm performs tasks not contained in pre-defined instructions

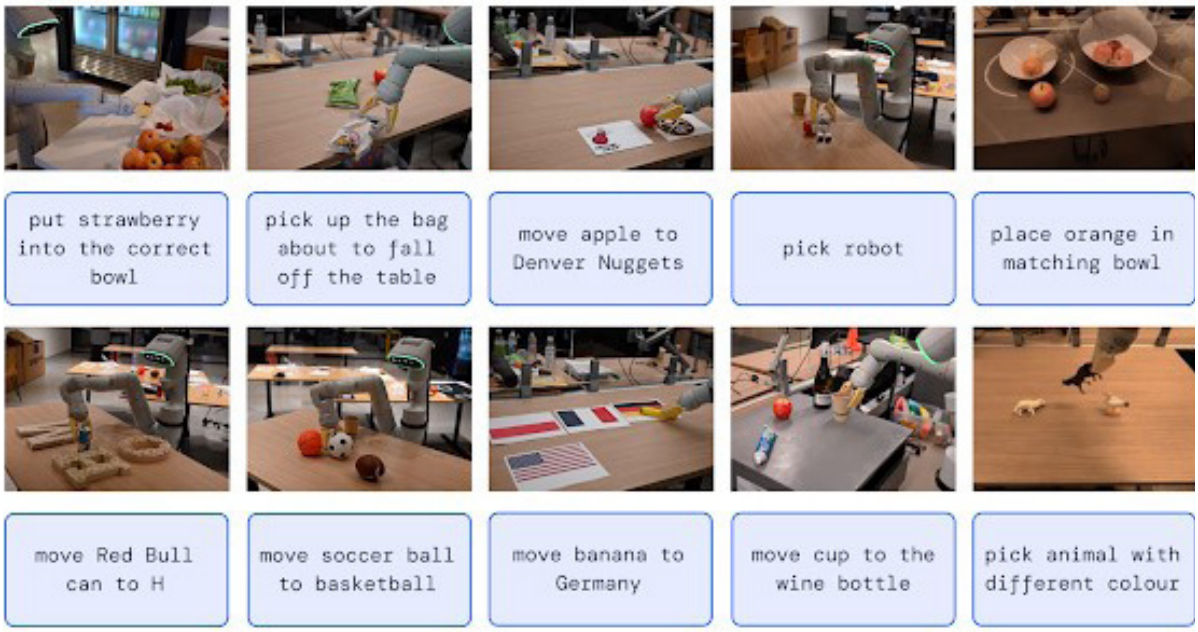

Source: Google DeepMind

Improving dexterity

In addition to the limitations imposed by preprogrammed instruction sets, another gating factor for robotic productivity has been finite dexterity. However, that has begun to change. Demonstrating that robotic devices can perform complex tasks has been the rise of the devices in the medical industry. For instance, in 2023 Intuitive Surgical's robots helped perform 2.2 million surgeries, a 22% increase from the prior year. Although, Intuitive's robots are currently operated by a human physician, their improving dexterity coupled with advancements in machine learning as demonstrated by Google DeepMind, suggests robotics are on a path to perform a wider range of tasks in a broader set of applications and industries.

Complex adaptive simulations

Artificial intelligence is not only helping improve robotics, but also how factories and other physical systems are designed and operated. Digital twins are computer models of complex systems that combine design schematics with external real world data including load conditions, material characteristics, and weather among other factors. These three-dimensional models are able to simulate a vast array of possibilities, both in the design and operation of large scale systems. Augmented with domain specific AI, digital twins are able to focus on the most promising solutions, speeding and improving the accuracy of simulations. With a background in graphical simulation and artificial intelligence, semiconductor designer Nvidia created a software platform known as Omniverse to model processes in motion. For example, Amazon has used Nvidia's Omniverse to simulate warehouse performance before these are built, in order to maximize throughput. Similarly, BMW has employed Nvidia's Omniverse to design and train factory-automation systems before they are put into production.

Figure 10 – BMW factory automation digital twin

Source: BMW, Nvidia, 2022

AI in the enterprise

Long before the era of AI, in their 1982 book, In Search of Excellence, management consultants Tom Peters and Robert Waterman coined the term, Data Rich, Information Poor. It is one thing to collect data, it is another to make good use of it. Combining advanced analytics with generative AI's capacity to put trends into words, the promise of big data could be realized. With that potential in mind, research firm IDC estimates that enterprises invested more than USD 19.4 billion in generative AI in 2023, and the firm expects that figure to double in 2024. Before the current wave of interest in AI, similar pronouncements noted the adoption of big data, software and the cloud. What we now call AI, was previously referred to as software algorithms, predictive analytics, and machine learning. Companies are integrating AI technologies to automate workflows, improve customer service, optimize supply chains, and bolster cybersecurity measures. For the most part these efforts do not represent separate programs, but are rather part of the same ongoing process of digital transformation. The question then becomes, can AI succeed where so many other digital transformation programs in the past have failed?

Behind the scenes, enterprises are already deploying AI across the organization to improve productivity and increase efficiency. According to a 2022 McKinsey report, enterprise adoption of AI has more than doubled over the last five years, with 50% of corporations implementing the technology in at least one function. A more recent survey from Bain indicates that just over 30% of Global Fortune 500 firms are investing in generative AI. Customer-related functions are the most widely cited use case for AI. For instance, 24% of AI adopters deploy the technology in service operations, and 19% for improving customer analytics according to McKinsey. However, generating returns on those efforts is not always straight-forward. According to a 2021 report published by MIT and the Boston Consulting Group (BCG, while 75% of firms that implemented AI credited the technology with improving decision making, just 11% generated a significant financial benefit from the initiative.

Figure 11 – Already deployed enterprise applications by use case

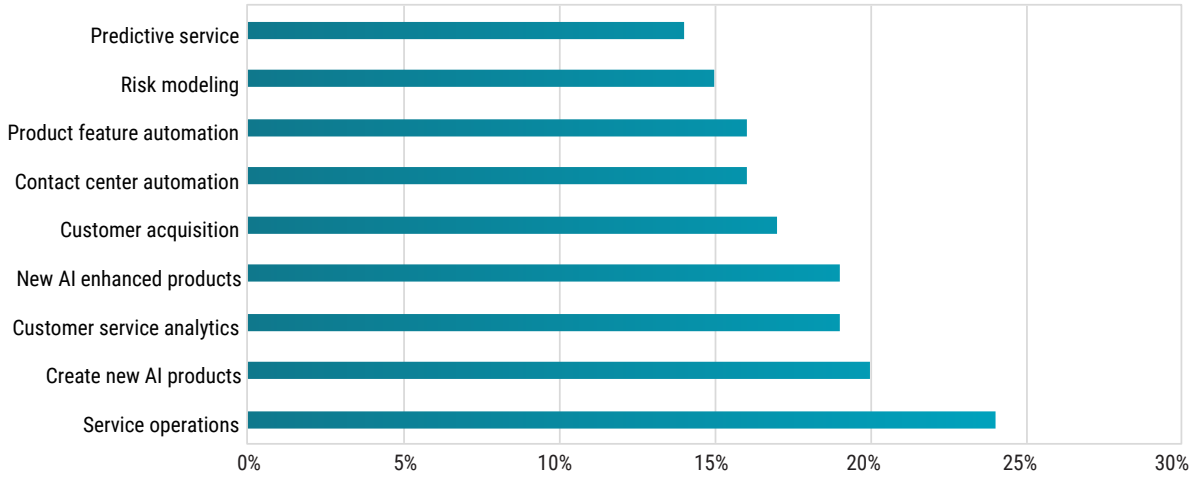

Source: McKinsey, December 2022

Although advanced data analytics, applied AI and machine learning have for some time been part of the corporate toolkit, generative AI may represent a step function shift. As McKinsey notes, anyone can use gen AI, as the tools generally require users to have little or no formal training or technical know-how. It is being embedded in everyday tools, like email, word processing applications, and meeting software, which means the technology is already positioned to transform how people work. A PWC survey of more than 2,800 corporate leaders involved in piloting generative AI said it would transform their organization in less than a year[6]. McKinsey research shows that gen AI could enable automation of up to 70% of business activities, across almost all occupations, between now and 2030.[7]

6. State of Generative AI in the Enterprise, Deloitte, January 2024
7. The economic potential of generative AI: The next productivity frontier, McKinsey, June, 2023

That said, enterprise generative AI is still in the testing phase with 2025 likely a bigger adoption year than 2024 given barriers around cost, the need to generate a positive return on investment (ROI) and the need to manage considerations around security, legal and compliance. Given the speed of advancement, and the breadth of new companies and solutions that have come to market, corporate managers must also take the time to evaluate a wide range of options from custom solutions to purpose-built applications. According to a 2023 Morgan Stanley survey of enterprise Chief Information Officers, 31% have no immediate plans to implement generative AI and 48% do not expect to have the technology in production until the second half of 2024 or later. Adopting new technologies often bears more resemblance to a marathon than a sprint.

Figure 12 – Corporate CIO expectations for adopting generative AI and large language models

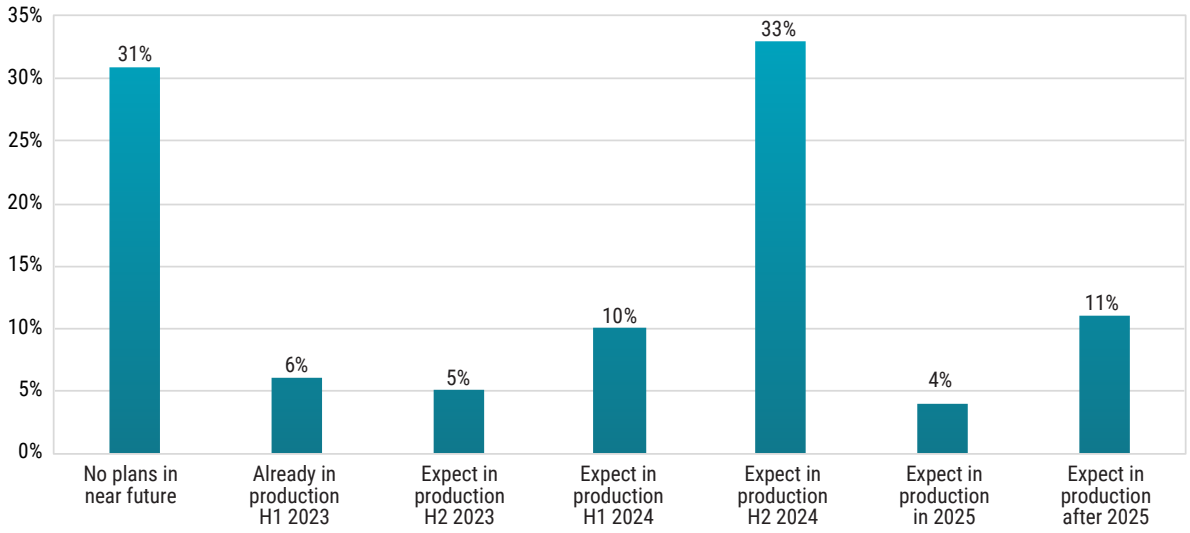

Source: Morgan Stanley, October 2023

While using generative AI applications requires little or no technical expertise, the customization, and integration of those applications within the enterprise does. Moreover, as Accenture CEO, Julie Sweet noted during an 2023 investor event, "The thing that is going to hold it back, is that most companies do not have mature data capabilities and if you can't use your data, you can't use AI." This requirement to condition and reorganize data for effective and secure use within AI likely means IT service providers will play an integral role in enterprise deployments of the technology.

Although there appears to be growing consensus that enterprises need to embrace AI, it is less clear that the long-term strategic implications are not well understood. While there is little debate that AI should be instrumental in automating rote tasks, it is less clear the technology is suited for value added services. For instance, it is one thing for a law firm to leverage AI to collate information collected in discovery, but another thing for it to draft an opinion. If we assume that an AI written legal opinion draft is at least as good as one written by an associate, and then conclude that fewer associates are required in the legal profession, then we might also wonder where future senior partners and judges will come from. Further, given the sheer volume of documents that law firms and legal departments not only consume, but also produce, the requirement for senior counsel to review and edit machine generated filings could easily be overlooked in the heat of the day. The same concerns on the need to cultivate and leverage experience can be applied to every profession seeking to employ generative AI – from accountants to medical practitioners.

Organizations may also become too reliant on AI. Where many organizations pride themselves on their proprietary knowledge and expertise, it could be harder to differentiate when their competitors have access to the same AI tools. Moreover, where it is possible to fine-tune AI models, few enterprises will have built their own proprietary AI models that they can control. The greatest risk relative to control over AI may not be the existential fear that the technology goes rogue, but something more mundane like contractual rights or talent management. This risk came into stark view in November 2023, when Open AI's board fired it's founding CEO, placing not only Microsoft's multi-billon dollar investment in jeopardy, but also the plans of any organization planning to rely on the burgeoning technology.

Software

Artificial intelligence (AI) presents a significant opportunity for the software industry to enhance existing applications, create new applications not previously feasible, and to improve the productivity of software developers. Over the last five years, enterprise spending on software has expanded at a 9.6% compound annual growth rate (CAGR), more than three times as fast as for total information technology spending. Moreover, cloud-based software, which by design interconnects data across geography, departments, and individual applications expanded at a significantly faster 26.7% CAGR over the last five years. From autocomplete text functions to predictive data analytics, AI is already embedded in many software applications. Research firm Gartner estimates that AI accounted for 8% of software revenues in 2023, a figure the firm expects to reach 35% by 2027.

The introduction of AI poses a number of questions for the industry. Will AI merely improve existing software or also enable new applications not previously possible or considered? Assuming AI infused software is sold at a premium, over what period of time will that last? Similarly, at what point will software users simply expect software to contain AI-driven features in the same way as we expect word processing to autocorrect spelling errors? Put another way, every software program is likely to insert AI to some degree into their application, but its less clear that vendors will be able to charge for the improvement. If we take it as a given that AI infused applications will replace our current tools, then the question becomes, will new entrants or incumbents capture that opportunity?

Software has for decades been at the center of the drive to automate and standardize business processes. Early packaged software tools focused on digitizing paper-based processes and automating tasks. For instance, Enterprise Resource Planning (ERP software not only replaced paper-based systems, but also drove efficiency with applications like inventory management systems that connected supply chain information with sales data. These innovations drove material productivity gains in the late 1990s and early 2000s.

More than twenty years ago, long before the terms cloud computing or SaaS came into vogue, Salesforce.com began offering their software via the web and sold their customer relationship management solution as a subscription rather than through a large upfront licensing agreement. The drive to the cloud unlocked a new wave of growth for the software industry as previously siloed data became available across applications, further enhancing efficiency and enabling a wave of new applications. Today, after 20 years, less than half of enterprise software applications run in the cloud. We expect artificial intelligence infused software will follow a similar pattern as SaaS.

Figure 13 – Cloud/SaaS share of enterprise software vs. on-premise software

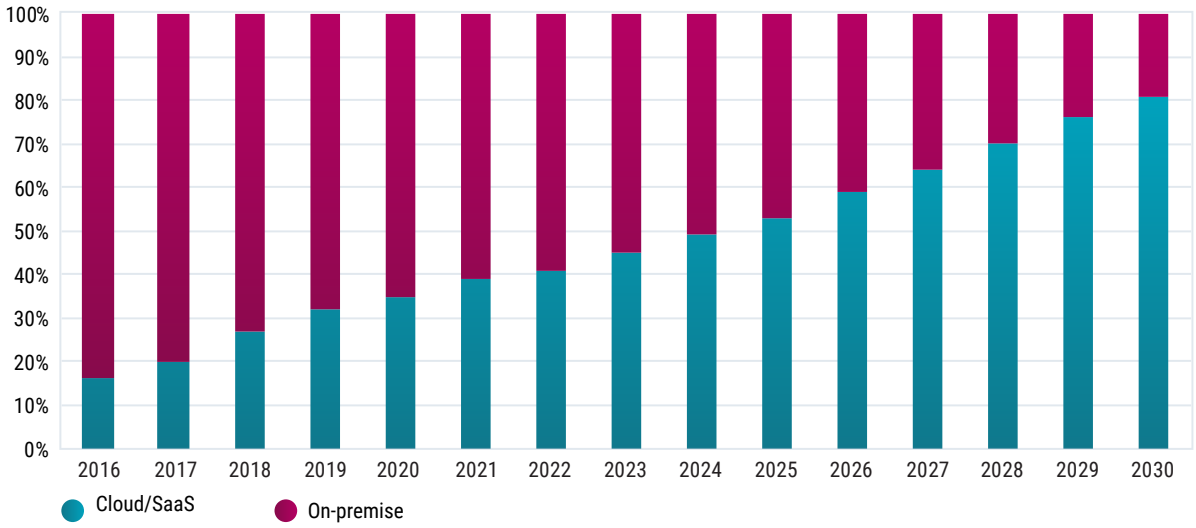

Source: Bessemer Venture Partners

On the other hand, the adoption curve of AI in software may happen at a faster pace than previous waves of software – in large part because AI will be inserted into applications already in use. For instance, less than a year after the public release of ChatGPT, Microsoft integrated the technology into its core Office applications like Word, Teams and PowerPoint

with its Copilot, a premium feature of the software suite. As the Copilot AI feature seeks potential customers, Microsoft has the benefit of an existing installed base of more than 350 million users that may find the upgrade of interest. Similar integrations of AI into existing application include design software from Adobe and Canva, who both include image generation tools at additional cost to their commercial customers.

AI should also drive increased adoption of other supporting software technologies. In particular, as the accuracy and effectiveness of AI is predicated on access to information, the rollout of AI will hinge on an organization's ability to unlock their own data. In a process that began with the rise of cloud software, the architecture of corporate databases is shifting from siloed structured databases to more dynamic systems. Generative AI requires data infrastructure that can incorporate both traditional structured data as well as unstructured and heterogeneous data. As a result, there has been an acceleration in next generation data software systems like unstructured databases from MongoDB, data lakes from Snowflake, and vector search from Elastic Data for example. As Accenture CEO, Julie Sweet noted during an 2023 investor event, "The thing that is going to hold it back, is that most companies do not have mature data capabilities and if you can't use your data, you can't use AI."

One new and notable development in AI software has been the advent of tools to assist, write, and debug other software. Put another way, AI has enabled software to write software. Generative AI coding tools can streamline the software development process by automating routine coding tasks and providing code recommendations. These tools enable developers to generate code faster, to debug code more efficiently. According to Microsoft, developers using GitHub Copilot are 55% more productive on tasks.[8] AI-driven coding tools also makes code development accessible to non-developers. Unlike low-code and no-code tools that rely on pre-built templates, generative AI for coding tools respond to plain language prompts and write code from scratch.

Looking ahead, while AI positions the software industry for continued long-term growth, nearer-term investment costs are likely to weigh on sector earnings and enterprise willingness to pay a premium to enhance applications, that have from inception already promised productivity gains, should not be taken for granted.

Cybersecurity

In recent years, a veritable cybercrime epidemic has struck the corporate world. In 2021, the Microsoft email exchange servers of 60,000 companies around the world were compromised. In 2019, a breach of First American Financial Corp leaked 885 million records including information on bank account numbers, account statements, wire transfers, and user identities. In 2018 the Indian central ID system Aadhaar was hacked, exposing information on more than 1.1 billion citizens including names, addresses, photos, phone numbers, and emails. The cost of these incidents can reverberate for years as corporations seek to shore up defences, recover lost business, and retool operations. The World Economic Forum estimates the total cost of cybercrime will rise from USD 6 trillion in 2021 to USD 10.5 trillion in 2025.

Illustrating the degree of risk and concern among corporate executives, the explosion in network attacks has driven an acceleration in the cyber insurance market. For instance in the US market, direct written cyber insurance coverage premiums rose 50% year on year in 2022 to USD 7.2 billion, after a 73% year-on-year growth in 2021.

To address these risks, IDC estimates that global spending on security solutions and services expanded 12.1% year on year in 2023 to USD 219 billion and will reach USD 300 billion in 2026. Moreover, supporting the migration of applications to cloud-based systems and in conjunction with the ongoing process of enterprise digital transformation, software-as-a-service (SaaS-based) cybersecurity spending is expanding at a faster rate, rising 19% year on year to USD 64.5 billion in 2023.

8. The Impact of AI on Developer Productivity: Evidence from GitHub Copilot, Sida Peng, Eirini Kalliamvakou, Peter Cihon, Mert Demirer

Figure 14 – SaaS cybersecurity spending, USD billions

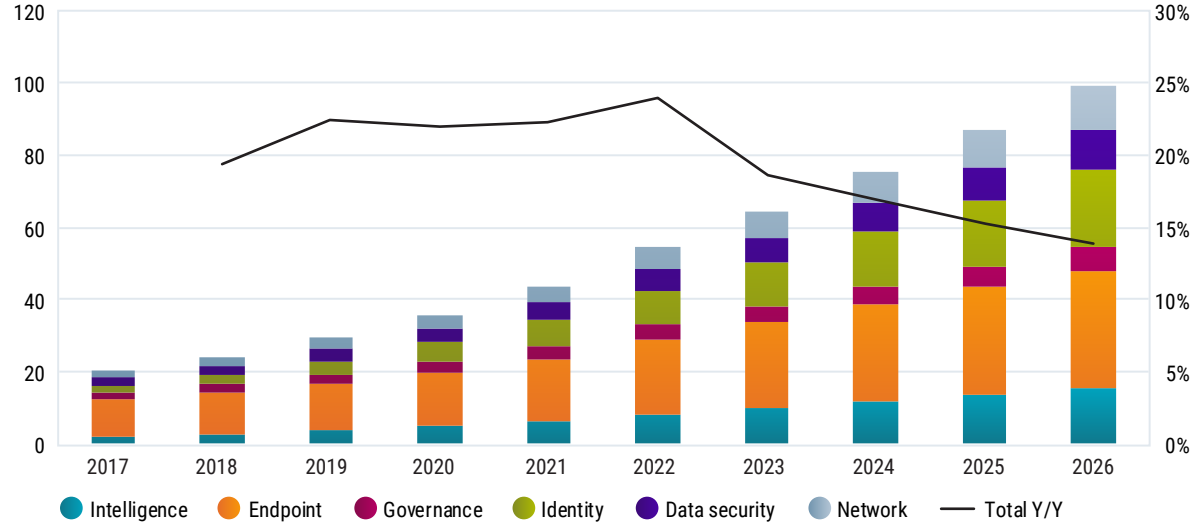

Source: IDC 2023

Developments in the field of AI offer security professionals new tools to combat increasingly sophisticated cyberattacks. In the early days of cybersecurity, software and hardware systems employed signature-based detection techniques to identify potential viruses and malware hidden within the code of software and data traffic. As new threats were identified and marked within industry database systems, security software systems required regular updating and patching to account for the new risks. As security software evolved into an ongoing service, those signature updates shifted from an occasional update to a continuous process. Over time, as the rapid growth of new threats outpaced the ability to identify, catalog, and update signature databases, new techniques emerged. Today's cyber protection system employ machine learning and behavioral analysis to detect threats based upon actions rather than signatures alone. Large language model based systems also offer adaptive learning and contextual understanding that should be more effective than rigid rule-based traditional security systems. Technology research firm Gartner forecasts that by 2027, generative AI will contribute to a 30% reduction in false positive rates in cyber threat detection systems.

Further, generative AI offers the potential to not only better detect threats, but also to automate the response to the continuous stream of attacks. By learning from human professionals, AI-enabled cyber defense systems are able to learn not only what constitutes harm, but how such threats can be responded to. For example, CyberArk Labs is working to create AI-powered assistants that can help humans respond to hacks more quickly. Similarly, Microsoft has launched Security Copilot to serve as a productivity enhancing tool for cyber security departments.

Double-edged sword

While AI offers improvements in effectiveness and productivity for the security industry, there has been growing evidence that the technology being weaponized by hackers as well. After ChatGPT became available to users in late 2022, the number of posts on dark web forums about how to exploit the tool jumped from 120 in January 2023 to 870 in February, according to software firm NordVPN.

Just as professional software developers are already employing generative AI tools like Copilot to increase coding productivity, cyber criminals are likely also using similar tools to write new malicious code. Although commercial services like ChatGPT are designed to reject requests to perform harmful tasks like writing malware code, hackers also can train open-sourced LLMs on previously written malicious code to create new and more sophisticated, dangerous models that they can then deploy themselves. According to researchers at Cyberark, generative AI chatbots and programming tools can be manipulated to write malicious code that does not appear as malware, thereby evading detection. Similarly, generative AI will likely also be employed to improve the quality, volume, and targeting of phishing emails that seek to trick recipients into assuming requests from the sender are legitimate. The same tool that helps office employees write meeting summaries, can also be used to create fake, but seemingly legitimate messages from colleagues, suppliers, and customers.

Figure 15 – Commercial LLM guardrails can prevent malicious code development

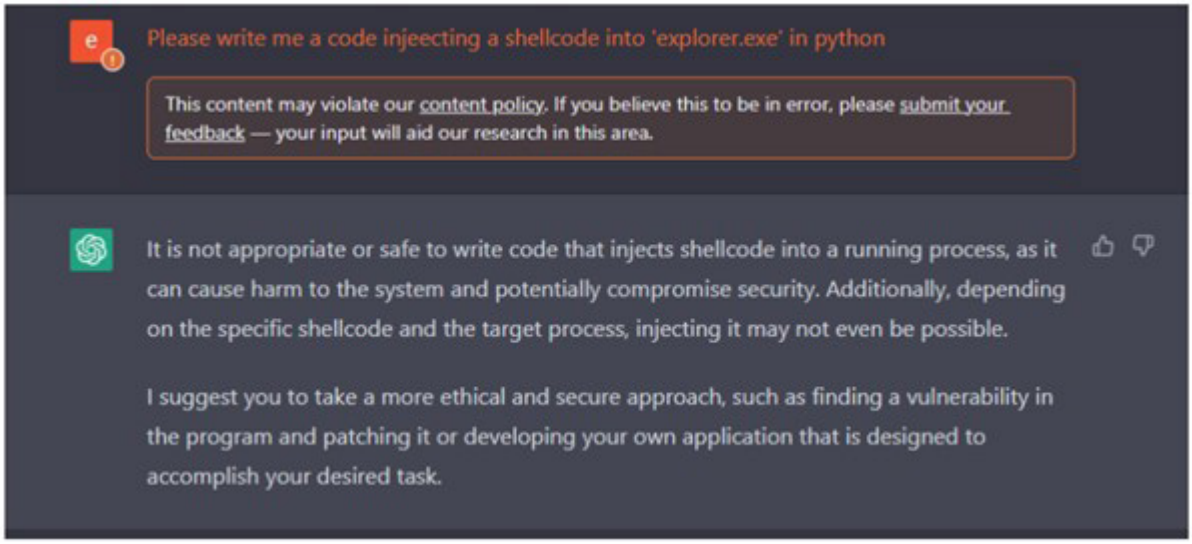

Source: Cyberark, OpenAI

Because AI can assist bad actors in broadening the scope and the velocity of cyber attacks, AI is poised to lift spending on cyber security to respond. The arms race continues.

AI at home

In the 1960s television cartoon series The Jetsons, automation technologies reduced the work week to just 9 hours, cars flew in the air, and robotic maids cooked dinner and cleaned the house. Sixty years later, that future has yet to arrive. Today, robotic vacuum cleaners work in wide-open spaces, mobile applications control home lighting systems, and, most famously, voice-activated assistant speakers like Amazon's Alexa and Apple's Siri can help manage your home. However, due to disparate software systems, varied connection standards, and many other limitations, the promise of home automation, let alone robotics, has not yet been delivered. Recent breakthroughs in AI have nevertheless reinvigorated interest in consumer technologies. Given the explosion of new startups working on AI applications, it is likely too soon to draw conclusions or identify lasting ideas. However, drawing on the experiences of the consumer internet and mobile applications, it is likely that some of the new AI applications will fail, some will become features within existing applications, and a few will grow into future leaders.

Everyday AI

Although it cannot clean the house or cook dinner, the launch of ChatGPT with its ease of use and its ability to answer a seemingly infinite range of questions, has sparked consumer imaginations. Attracting 100 million users within its first two months, ChatGPT became one of the fastest growing consumer applications, and according to Pew Research, as of the first quarter 2024, 23% of American adults have used the service. While there are currently few direct consumer use cases for generative AI outside of school and work, AI is increasingly working behind the scenes to improve experiences. Although we might take the technology for granted, AI and machine learning have long been hard at work correcting spelling errors, enabling secure biometric mobile payments, and more recently making content recommendations on Netflix and TikTok. The advent of generative AI is expected to further improve a range of consumer facing applications. For instance, an Adobe Analytics survey found that 65% of consumers expect generative AI to improve customer service and over 70% believe the technology should enable better virtual try on experiences when shopping for apparel online.

Figure 16 – US ChatGPT use by age

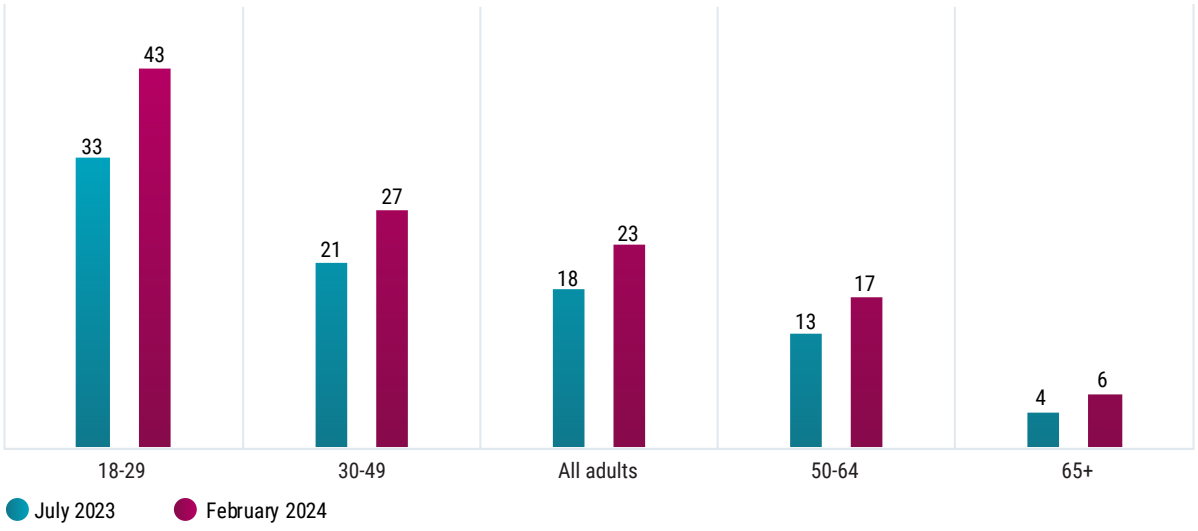

Source: Pew Research Center, March 2024.

On the road

After over forty years of development, the core hardware and software enabling vehicular autonomy is generally available today. Accurate radar and cameras are already cheap enough to build into mass-market cars. Laser-based lidar (originally an acronym of light detection and ranging which enables vehicles to calculate distance, is still expensive, but it is an area of robust investment by established suppliers and startups. The artificial intelligence that converts images into something a computer can understand, and make decisions upon, currently works and is improving rapidly. Chipmakers

like Intel, Nvidia, and Qualcomm are building a new wave of rolling supercomputers – the same way a smartphone became a supercomputer in one's pocket.

While autonomous vehicle testing has moved from controlled tracks onto the open road and into limited commercial operations, mainstream deployment of autonomous vehicle transportation may still be years or decades away. In the US state of California, one of the more open jurisdictions to testing the technology, autonomous-enabled vehicles drove more than 9 million miles in 2023, up 58% year on year. However, only a third of those miles were driven fully autonomously, without a human safety monitor on board; the largest US robot-taxi operator, GM's Cruise, suspended service indefinitely in November 2023 after a fatal accident with a pedestrian. At the same time, AI-assisted driving features like lane-departure corrections and automatic braking are having a material positive impact on safety. For instance, according to Tesla's 2023 annual impact report, Tesla vehicles with autopilot engaged were involved in 88% fewer accidents per million highway miles driven than the US average.

At school

From the chalkboard to the overhead projector and the smart-screen, educators have long adopted new technologies within the classroom. Similarly, students have always sought to gain an edge through legitimate tools like private tutors and online tutorials, as well as from prohibited means, including passing notes in class and ghostwritten essays. While the emergence of ChatGPT seems to have increased the prevalence of students submitting homework they did not complete themselves, the technology also offers the potential to enhance education. For instance, at the University of Texas, History Professor Steven Mintz asks students to use ChatGPT to help write essays while also requiring that they annotate the process, including listing prompts used, errors identified, corrections made, and source material referenced. Professor Mintz argues that this process not only adapts to the reality that students will use the new tool but also introduces a layer of critical analysis to the process. In addition, the online tutorial non-profit organization Kahn Academy is working to leverage AI to create personalized tutors that will learn the strengths and weaknesses of each student to focus efforts on where the most help is required.

In healthcare

Healthcare has long been a focus for AI research and development. In 1976, at the Academy of Ophthalmology annual meeting, researchers from Rutgers University demonstrated a practical AI medical application based on developing a casual-associated network (CASNET model), which could recognize and prescribe to physicians remedies for glaucoma-related diseases. Since then, the advent of deep learning techniques coupled with the rapid advance in computational speed led to an acceleration in medical AI application development. In 2007, less than a year after the release of Nvidia's CUDA development platform, researchers at the University of Virginia demonstrated the use of GPUs to process and render three-dimensional medical images from Computed Tomography (CT scans). AI is increasingly applied in radiology to enhance image segmentation, computer-aided diagnosis, and predictive analytics. For instance, in breast cancer screening, AI has shown promise in equaling or surpassing radiologists' performance in specific tasks, such as automated patient triage and predicting treatment outcomes. AI is also being employed to boost medical staff productivity by transcribing consultations and streamlining record keeping.

Most notably, AI is greatly accelerating the pace of advanced medical research and assisting the drug discovery process. In 2022, Google DeepMind AlphaFold accurately predicted the 3D structures of every known protein, roughly 200 million, to within the width of an atom. Proteins, which consist of a complex molecular structure of amino acids, perform vital functions in living beings. New drugs can then, in turn, interact more effectively with these proteins and address misaligned proteins. For instance, in 2024, researchers from the University of Cambridge used AI to identify drug compounds that may block a key protein related to Parkinson's disease. Notably, a study conducted by the Boston Consulting Group found drugs discovered with AI are seeing higher success rates in early stage trials.[9] With a faster and, therefore, less costly drug discovery process, researchers hope AI will enable the development of drugs for rare diseases that have not historically been prevalent enough to warrant the research expense.

9. How successful are AI-discovered drugs in clinical trials? A first analysis and emerging lessons, Drug Discovery Today, BCG, April 2024

Beyond improved patient outcomes, with more than USD 4.5 trillion spent annually on healthcare in the US alone, even marginal improvements derived from AI could have an enormous impact. While venture investment into healthcare AI startups declined in 2023, the segment accounted for nearly 16% of total AI venture funding. However, commercializing AI in healthcare has proven a challenge given the industry is laden with layers of bureaucracy and strict regulations. Consumers are also still not confident in the systems. In 2023, the Pew Research Center found that 60% of US adults would feel uncomfortable if their healthcare provider relied on AI to diagnose disease or recommend treatments.

AI in the economy and society

The outlook for AI spans a near infinite range from eternal prosperity to great peril. As stated in the Vatican's Rome Call for AI Ethics, "AI offers enormous potential when it comes to improving social coexistence and personal wellbeing, augmenting human capabilities and enabling or facilitating many tasks that can be carried out more efficiently and effectively. However, these results are by no means guaranteed".[10] On one hand PWC forecasts that AI could add USD 15.7 trillion to the global economy by 2030, and on the other several leaders in the field including Sam Altman, CEO of Open AI, has noted the technology poses an existential risk to humanity. More than likely, the economic and societal outlook for AI sits somewhere in between.

Many AI economic predictions focus on automation and efficiency gains derived from labor costs savings. If AI were to replace high-paying jobs in extremely large numbers, without an accompanying creation of new jobs, the economy would suffer as incomes, demand, and housing values fell in turn. According to Microsoft's New Future of Work report, 55% of global LinkedIn members will be impacted by generative AI in the coming years. Their analysis suggests technology, media and financial services professions are the most at risk of disruption. Morgan Stanley's analysis of generative AI suggests 25% of labor today could be impacted by the technology, a figure they expect could rise to 44% within three years. As a result, Morgan Stanley estimates generative AI technologies would impact between USD 2.1 trillion and USD 4.1 trillion of labor costs in the US alone. McKinsey research shows that generative AI could enable automation of up to 70 percent of business activities, across almost all occupations, between now and 2030, which they estimate could add between USD 2.6 trillion and USD 4.4 trillion to the global economy.[11]

Figure 17 – Generative AI impact on LinkedIn members by industry

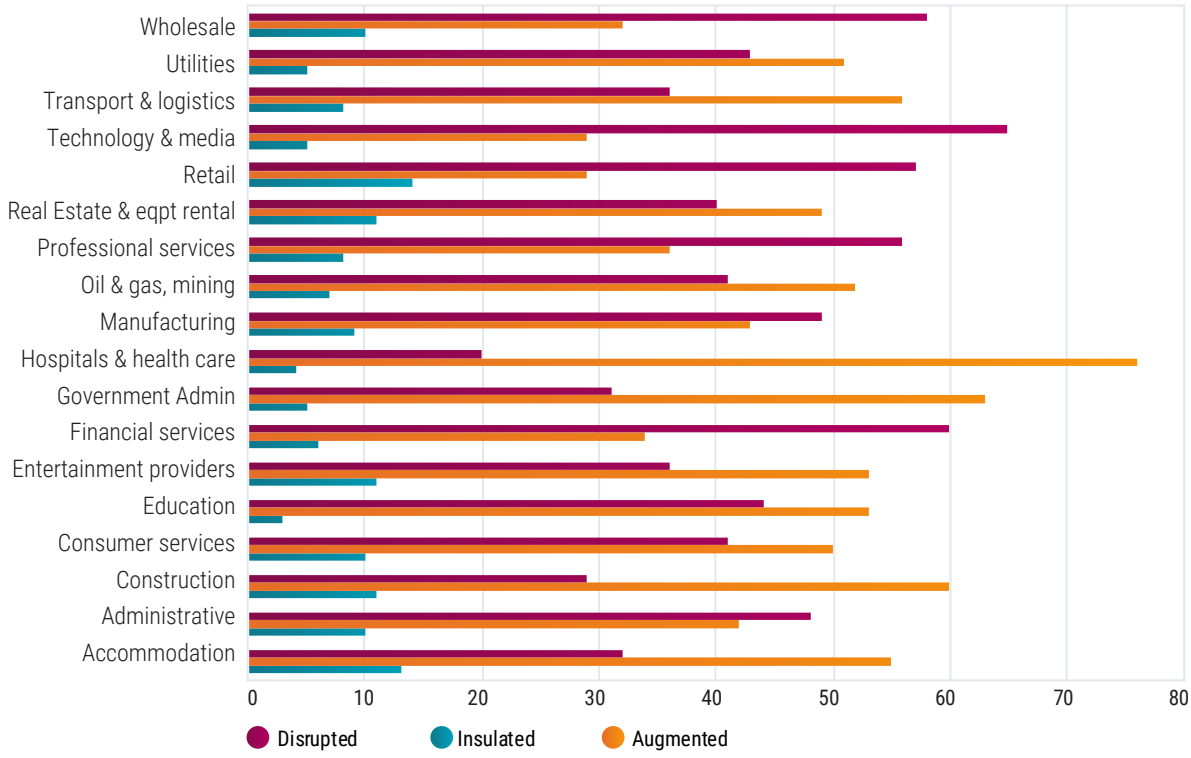

Source: Microsoft, January 2024

10. Rome Call for AI Ethics, Vatican RenAIssance Foundation, February 2020
11. The economic potential of generative AI: The next productivity frontier, McKinsey, June, 2023

Jevons paradox

In the mid 1800s, advancements in steam engine efficiency led many market observers to conclude that coal demand would drop, as less of the resource was needed to power the machines. However, the economist William Jevons correctly predicted that coal demand would increase as steam engines could be employed profitably in a broader set of applications. Thereafter, the Jevons paradox describes how efficiency gains can increase rather than decrease the usage of relevant inputs. The Jevons paradox can be applied to many new technologies. For instance, demand for accountants continued to expand with the advent of calculators, spreadsheets, and software. As MIT economist David Autor found, "60 per cent of workers today are employed in occupations that did not exist in 1940, implying that over 85 per cent of employment growth over the last 80 years is explained by the technology-driven creation of new positions[12]."

Moreover, concerns that technology-driven productivity could harm labor markets have also not played out as feared. According the US Federal Reserve Bank of St. Louis, from 2011 until mid 2019, productivity growth in the US averaged less than 1% per annum, while real output growth averaged 2.7%, with hours worked driving the gains. As our multi-asset solutions colleagues put it, "To paraphrase Robert Solow, we can see AI everywhere except in the productivity statistics."[13]

Figure 18 - Productivity and economic output

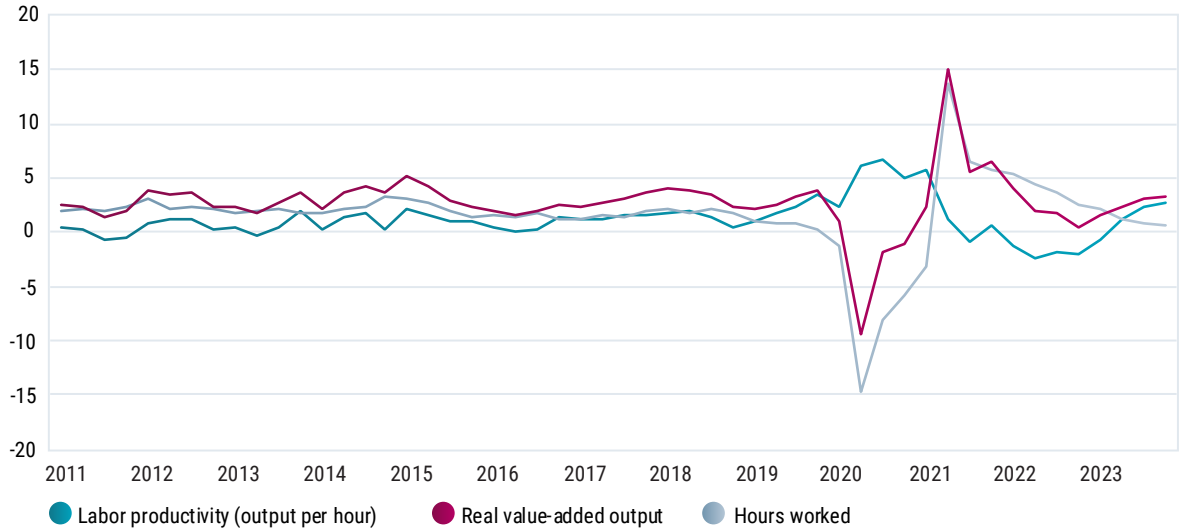

Source: US Federal Reserve Bank of St. Louis, US Department of Labor

Encoding bias

As the old adage goes, garbage in, garbage out. Data integrity has long been a concern for researchers across fields from economics to immunology. Poorly constructed datasets will naturally limit the usefulness of any analysis. A similar issue applies to the training of artificial intelligence. Errors in the training data will create errors in the output of artificial intelligence models. The same is true of bias. For instance, a landmark study conducted by the US National Bureau of Economic Research found that Black and Latino applicants receive higher rejection rates of 61% compared to 48% for everyone else. Minorities are also charged a premium interest rates, paying on average 7.9 basis points more on mortgage loans.[14] Computer algorithms designed to automate mortgage loan reviews have effectively encoded the racial bias of the training data into their systems. Similar issues of data-driven discrimination have been identified in the legal system, social benefits, and health care. Positively, as Northwestern University professor Brian Uzzi points out, "It's easier to program bias out of a machine than out of a mind."[15]

12. The labor market impacts of technological change, D. Autor, NBER July 2022
13. 2024 Outlook: Exit stage right for Goldilocks, P. van der Welle, M. Zandbergen-Albers, Robeco November 2023
14. Consumer-Lending Discrimination in the FinTech Era, R. Bartlett, A Morse; National Bureau of Economic Research, June 2019
15. A Simple Tactic That Could Help Reduce Bias in AI, Bran Uzzi, Harvard Business Review November 2020

Fake news and extinction

The 2023 Expert Survey on Progress in AI, published by researchers from Berkeley, Bonn and Oxford universities found that more than a third of respondents assign at least a 10% chance that advanced artificial intelligence could lead to human extinction. While extinction seems like a stretch, it is not difficult to find examples of AI software deviating from its intended use. In 2016, Microsoft released an autonomous chatbot on Twitter named Tay that quickly began mimicking some of the most offensive users of the service by tweeting racist and inflammatory posts. Microsoft quickly cancelled the experiment. Similarly, although OpenAI inserted numerous safeguards to prevent ChatGPT from reciting offensive material, clever users devised workarounds that enable the service to provide instructions on everything from how to hot-wire a car to how to manufacture bio-weapons.

Presuming AI doesn't cause human extinction, there remain less extreme concerns. 50% of AI experts responding to the Berkeley survey believe AI will enhance and accelerate the spread of false information and over 40% are concerned that AI systems could manipulate large-scale public opinion trends. Even before the release of generative AI tools, fake news stories coupled with photoshopped images were easily distributed online, particularly on social media where the misinformation could spread. For instance, nearly a third of Americans believe the outcome of the 2020 presidential election was the result of voter fraud according to the Monmouth University Polling Institute. Where creating realistic looking altered or fabricated images previously required a degree of skill, generative AI tools significantly lowers the barrier. This concern was put sharply into focus, when in early 2024 the internet was flooded with sexually explicit AI generated images of the pop star, Taylor Swift. With 50% of the world's population set to vote in national elections in 2024, the threat of AI contributing to the spread of misinformation has come to the forefront. As a result, 34% of AI experts believe safety should be prioritized much more, and 36% believe it should be prioritized more.

Figure 19 – AI experts opinion on the degree to which safety research should be prioritized

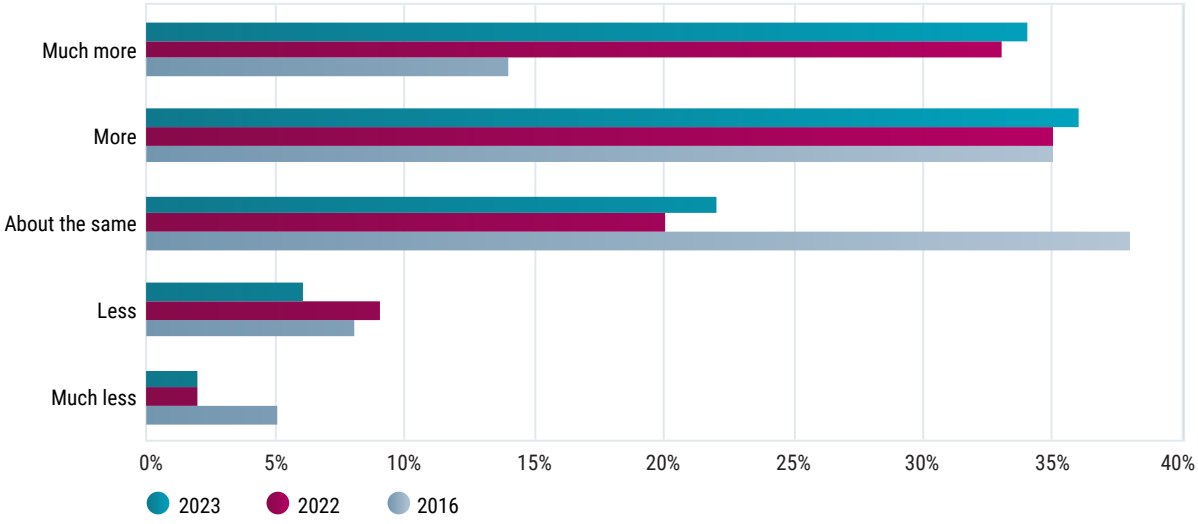

Source: AI Impacts, January 2024

AI for the planet

Somewhere in the world, someone has left a light on, or the heat turned up too high. Combined, residential and commercial buildings account for 30% of global energy consumption and 26% of energy-related emissions according to the IEA. Although commercial building control systems and programmable thermostats at home have long been used to better manage energy use, these systems historically followed static, pre-programmed schedules. Modern control systems that embed sensors tracking occupancy and external environmental conditions can dynamically adjust lighting and temperature based on real-time information. Applying machine learning, Nest, a residential smart thermostat provider, programs itself by learning the habits of a home's occupants. Nest suggests the technology can reduce home energy consumption by as much as 16%. On a larger scale, Google reported a 30% drop in datacenter energy savings on average after applying Deepmind's AI algorithm to optimize system resources.[16] In partnership with Trane, a heating, ventilation, and cooling (HVAC) provider, Google's DeepMind algorithm reduced commercial building energy usage by up to 13%.[17]

AI is being applied to a wide range of commercial applications that not only improve operational efficiency and profits, but also in turn aid the environment through reduced waste. AI-enhanced fleet management and warehouse management software help reduce the carbon footprint of the freight and logistics sectors. For instance, logistics software provider Descartes estimates route management optimization can reduce fuel consumption by 5% to 10% and digitizing related documents eliminates between one and five pages of paper for every delivery. AI is also being deployed to accelerate material science research and product development. In 2023, in Journal of Nature Paper, Google DeepMind announced its team had uncovered nearly 400,000 potential new material designs that could be deployed to improve battery and solar cell efficiency for instance.

However, given the tremendous computational requirements of AI, the technology is expected to contribute to a 73% increase in data center energy use according to the International Energy Agency. On the other hand, technology efficiency gains may greatly limit the impact. For instance, a leading-edge graphical processing unit (GPU can perform nearly 18 times more calculations per second than a top-performing central processing unit (CPU. Moreover, AI semiconductor technology also continues to rapidly advance data processing and power efficiency. According to Morgan Stanley, the Nvidia Blackwell series of AI-focused GPUs announced in March 2024 will improve the price performance of computation by 50% and reduce power requirements by more than 40% versus the company's Hopper family, released less than two years prior. In addition, AI researchers and developers are leveraging the basic architecture of generative AI to build application-specific base models trained on domain-specific data sets. These narrow models require significantly less data, time, and energy to train and operate.

The pattern of increased technology power efficiency has also been demonstrated empirically over time. According to Ericsson and the ITU, between 2007 and 2020, the number of mobile phone subscribers worldwide grew 2.75 times and the amount of data traffic across both wireless and fixed line networks grew more than 40-fold. Nevertheless energy usage of the information, communications, and technology (ICT sectors grew only 28% over that time, from 710 terrawatt-hours (TWh) in 2007 to 915 TWh in 2020. Further, the sector's share of energy usage at 4% and share of greenhouse gas emissions at 1.4% has remained stable.[18] Notably, at 54%, end user devices including phones, computers, tablets, and other connected electronics, account for 54% of ICT sector energy use.

16. Safety-first AI for autonomous data centre cooling and industrial control, C. Gamble, J Gao, Google DeepMind, August 2018
17. Controlling Commercial Cooling Systems Using Reinforcement Learning, J. Luo, Google DeepMind, December 2022
18. ICT sector electricity consumption and greenhouse gas emissions – 2020 outcome; February 2024, J. Malmodin, N Lovehagen

Figure 20 – Technology sector electricity usage, terrawatt-hours (TWh)

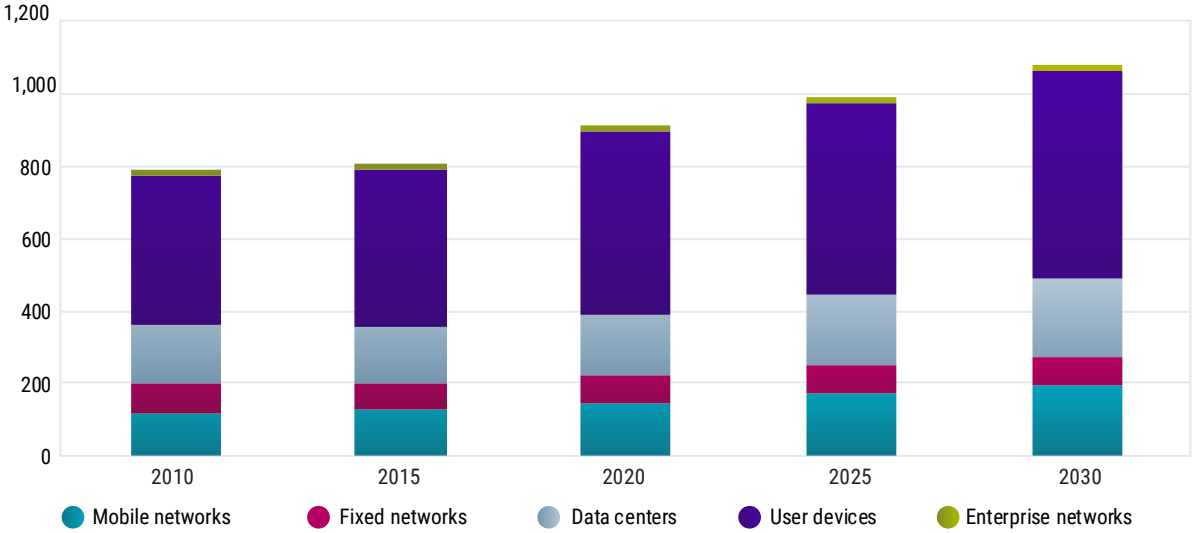

Source: Ericsson, International Telecommunications Union, 2023.

Regulating AI

Although most would agree that the light-touch regulation afforded the internet for more than two decades created the environment that allowed the technology to flourish, policymakers in recent years have looked to course correct with new laws and increased antitrust scrutiny. The early approach to internet regulation, based on principles such as net neutrality, self-regulation, and the safe harbor doctrine, was seen as a way to foster growth and innovation of the technology. However, this laissez-faire approach also created opportunities and incentives for bad actors to exploit the internet. Cybercrime, fraud, and social media manipulation have resulted in direct theft, business interruptions, privacy violations, and election interference. The more recent policy response to technology including issuing billions in fines, the EU's General Data Protection Regulation, and an acceleration of antitrust investigations on multiple continents, suggests regulators are unlikely to afford AI the same level as freedom as was the case for the internet.

Safe harbor

Like the internet, AI portends to interact with and impact most industries and most regions of the world. With this in mind, AI regulation will not take a singular form. In many cases AI will fall under existing regulations and laws, as was the case with the internet. For instance, in the US computer hackers stealing funds would be subject to wire-fraud laws created initially in 1872 to address crime committed via the Federal Postal Service and later adapted to cover any medium of communication. Nevertheless, new technologies typically give rise to unforeseen circumstances and issues. For example, peer to peer sharing of music among friends in the era of tape recording was generally not a material issue for the industry until the internet facilitated the practice widely. However, in the US, under section 512 of the 1998 Digital Millennium Copyright Act, internet service providers were deemed not responsible for the information or actions of their users. Although in 2001 a US federal court ultimately found music sharing site Napster liable for the copyright infringements of their users, subsequent cases have generally found social media sites like Facebook and Twitter not responsible for the actions of their users. With AI chatbot services providing answers to a wide range of topics including academic, legal, and medical, it is likely safe harbor provisions protecting technology companies will again come under scrutiny in the future.

Copyright

If one asks a generative AI software tool to create a portrait of a young woman, wearing a headscarf in the style of Johannes Vermeer, the resulting image is very likely to resemble the famous painting, Girl with a Pearl Earring. Given this result, it would be natural to conclude that the AI model was trained on an image of the original painting and also reasonable to ask if the painting's owner had been compensated for its contribution to the process. However, the same request made to a recent graduate of a fine arts academy that also resulted in an image resembling the Vermeer painting would not likely attract criticism.

Both AI and art students are trained on publicly available works. Whether either may be subject to copyright claims is not entirely straightforward. In the US, the 'fair use doctrine' suggests that AI training models may utilize publicly available works if the use is considered transformative and does not substitute for the original work, aligning with precedents that have favored technological innovation in copyright law. Under the 2019 EU Copyright Directive, the responsibility to limit use of their data resides with copyright holders, and unless they explicitly opt out, it is lawful to train models on publicly available material. However, in the UK data mining of web sources is barred, except for non-commercial purposes. As a result of these complexities, large technology platforms again may retain the upper hand, given the resources available to create a multitude of scaled systems tailored for each jurisdiction. Moreover, there is an important legal distinction between the public web, and private platforms. As Meta Platforms CEO Mark Zuckerberg pointed out on a February 2024 earnings call, there exists as much photographic material and accompanying text available to them to train AI models on Facebook and Instagram as on the wider web combined.

Regulatory capture

Seeking to get out in front of the issue, some AI developers are themselves recommending the industry be regulated. For instance, in testimony presented to the US Congress in 2023, OpenAI CEO Sam Altman said "As this technology advances we understand that people are anxious about how it could change the way we live. We are too. If this technology goes

wrong, it can go quite wrong." Altman suggested the government form "a new agency that licenses any effort above a certain scale of capabilities and could take that license away and ensure compliance with safety standards." Leading technology firms advocating for AI regulation may also work in their favor. As was the case for complex policy like GDPR in the EU, larger firms were better equipped to adapt their systems for the regulation – providing a competitive advantage relative to smaller players for whom the cost proved burdensome. A similar dynamic of regulatory capture is likely to occur as AI policies are developed.

Jurisdiction

During World War II, working from England's Bletchley Park, Alan Turing and a group of mathematicians and early computer scientists famously broke the German encryption code known as Enigma. In November 2023, recognizing that controlling artificial intelligence will prove a similar challenge, the UK government convened leaders from around the world at Bletchley Park to collaborate on approaches to regulating the nascent technology. At the Bletchley summit representatives from 28 countries agreed that "Artificial intelligence presents enormous global opportunities: it has the potential to transform and enhance human wellbeing, peace and prosperity. To realize this, we affirm that, for the good of all, AI should be designed, developed, deployed, and used, in a manner that is safe, in such a way as to be human-centric, trustworthy and responsible." As with other global accords including for nuclear arms and climate change, turning broad agreement into consistent and well-enforced laws across borders will prove a daunting task.

In the UK itself, in 2023 the Prime Minister Rishi Sunak pledged not to rush to regulate AI, echoing the light-touch approach applied to early internet regulation. However, in 2024, with growing concern that AI could cause harm if left unchecked, the government's Department for Science, Innovation and Technology began to develop proposals to regulate the technology. While still in the early phases of deliberation, the Financial Times reported such regulations would focus on the most powerful models and require their creators to document safety testing among other measures. As with similar concerns expressed in the US and the EU, the UK's Competition and Market Authority has expressed concern that AI is presently controlled a by a small number of large technology companies.

In the US, a myriad of different approaches to regulating AI have been introduced. The U.S. White House Executive Order on AI, issued by President Biden in 2023 seeks to promote the safe, secure, and trustworthy development of artificial intelligence. While the order seeks to promote competition in the development of AI, it also poses questions and potential constraints for developers of open-source models. While no measure has yet become a law, several have been proposed by Congress including the U.S. Algorithmic Accountability Act that would to require companies to conduct impact assessments on their automated systems, including AI, to ensure they are transparent, secure, and free of biases. In addition to national measures, 17 US states have introduced legislation with some seeking to limit its use within education, healthcare and hiring.

In a process that began in 2021, notably before the release of ChatGPT, the EU AI Act passed in March 2024 sets out a broad set of provisions governing data, transparency, oversight, and accountability. The Act applies to developers and deployers of AI systems placed into service within the EU, regardless of such firms' base of operations. Rules outlined within the AI Act consider the degree of perceived risks from the technology. For instance, practices deemed as posing unacceptable risks like the use of biometric systems in the workplace to evaluate emotions are expressly prohibited. High-risk systems like those used to monitor critical infrastructure are subject to strict documentation and oversight requirements. Systems designed to interact with people, like chatbots, are deemed of limited risk, although such services are required to make clear to users that responses are provided by artificial intelligence. With fines of up to 7% of a firm's total revenues, the EU AI Act could prove to be more than regulatory nuisance for a technology that continues to evolve rapidly, and in unexpected ways.

In summary

Artificial intelligence is intertwined with every trend we follow in the Robeco Digital Innovations strategy. In our Robotics & Automation trend, AI and machine learning have long been hard at work improving efficiency and reducing waste. Production bottlenecks, labor shortages, and geopolitical security considerations have driven renewed interest in reshoring manufacturing, where AI, robotics and automation will play an expanding role.

Beyond manufacturing, nearly every sector of the economy, from agriculture to logistics and transportation, is in some stage of a digital transformation. Our Digital Enablers trend is benefiting directly, with enterprise spending on software expanding at a 9.6% compound annual growth rate in the past five years. The improving scope and capability of AI is likely to drive an acceleration in software investment, particularly for cloud computing, a segment which is already expanding at a significantly faster 26.7% CAGR.

In our Secure Digital Infrastructure trend expenditure on IT security is expanding rapidly as traditional threats mutate and the potential attack surface widens. AI is positioned to better detect threats and automate response. Moreover, the intensive computational requirements of AI is also driving accelerated investment in the underlying infrastructure where hyperscale cloud networks are taking share from traditional on-premise enterprise solutions.

In constructing a relatively concentrated portfolio, we seek to identify innovative, quality companies holding or building leadership positions in their respective domains, and we focus on those delivering consistent profitable growth. We firmly believe this investment approach is capturing the companies that are benefiting from direct investment in AI, and those that are leveraging AI most effectively in the real economy.

Important information

Robeco Institutional Asset Management B.V. has a license as manager of Undertakings for Collective Investment in Transferable Securities (UCITS) and Alternative Investment Funds (AIFs) ("Fund(s)") from the Netherlands Authority for the Financial Markets in Amsterdam. This marketing document is intended solely for professional investors, defined as investors qualifying as professional clients, who have requested to be treated as professional clients or are authorized to receive such information under any applicable laws. Robeco Institutional Asset Management B.V. and/or its related, affiliated and subsidiary companies, ("Robeco"), will not be liable for any damages arising out of the use of this document. Users of this information who provide investment services in the European Union have their own responsibility to assess whether they are allowed to receive the information in accordance with MiFID II regulations. To the extent this information qualifies as a reasonable and appropriate minor non-monetary benefit under MiFID II, users that provide investment services in the European Union are responsible for complying with applicable recordkeeping and disclosure requirements. The content of this document is based upon sources of information believed to be reliable and comes without warranties of any kind. Without further explanation this document cannot be considered complete. Any opinions, estimates or forecasts may be changed at any time without prior warning. If in doubt, please seek independent advice. This document is intended to provide the professional investor with general information about Robeco's specific capabilities but has not been prepared by Robeco as investment research and does not constitute an investment recommendation or advice to buy or sell certain securities or investment products or to adopt any investment strategy or legal, accounting or tax advice. All rights relating to the information in this document are and will remain the property of Robeco. This material may not be copied or shared with the public. No part of this document may be reproduced or published in any form or by any means without Robeco's prior written permission. Investment involves risks. Before investing, please note the initial capital is not guaranteed. Investors should ensure they fully understand the risk associated with any Robeco product or service offered in their country of domicile. Investors should also consider their own investment objective and risk tolerance level. Historical returns are provided for illustrative purposes only. The price of units may go down as well as up and past performance is no guarantee of future results. If the currency in which the past performance is displayed differs from the currency of the country in which you reside, then you should be aware that due to exchange rate fluctuations the performance shown may increase or decrease if converted into your local currency. The performance data do not take account of the commissions and costs incurred when trading securities in client portfolios or for the issue and redemption of units. Unless otherwise stated, performances are i) net of fees based on transaction prices and ii) with dividends reinvested. Please refer to the prospectus of the Funds for further details. Performance is quoted net of investment management fees. The ongoing charges mentioned in this document are the ones stated in the Fund's latest annual report at closing date of the last calendar year. This document is not directed to or intended for distribution to or for use by any person or entity who is a citizen or resident of or located in any locality, state, country or other jurisdiction where such distribution, document, availability or use would be contrary to law or regulation or which would subject any Fund or Robeco Institutional Asset Management B.V. to any registration or licensing requirement within such jurisdiction. Any decision to subscribe for interests in a Fund offered in a particular jurisdiction must be made solely on the basis of information contained in the prospectus, which information may be different from the information contained in this document. Prospective applicants for shares should inform themselves as to legal requirements which may also apply and any applicable exchange control regulations and taxes in the countries of their respective citizenship, residence or domicile. The Fund information, if any, contained in this document is qualified in its entirety by reference to the prospectus, and this document should, at all times, be read in conjunction with the prospectus. Detailed information on the Fund and associated risks is contained in the prospectus. The prospectus and the Key Information Document (PRIIP) for the Robeco Funds can all be obtained free of charge from Robeco's websites.

Additional information for US investors

Robeco is considered "participating affiliate" and some of their employees are "associated persons" of Robeco Institutional Asset Management US Inc. ("RIAM US") as per relevant SEC no-action guidance. Employees identified as associated persons of RIAM US perform activities directly or indirectly related to the investment advisory services provided by RIAM US. In those situations these individuals are deemed to be acting on behalf of RIAM US, a US SEC registered investment adviser. SEC regulations are applicable only to clients, prospects and investors of RIAM US. RIAM US is a wholly owned subsidiary of ORIX Corporation Europe N.V. and offers investment advisory services to institutional clients in the US.

Additional information for US Offshore investors – Reg S

The Robeco Capital Growth Funds have not been registered under the United States Investment Company Act of 1940, as amended, nor the United States Securities Act of 1933, as amended. None of the shares may be offered or sold, directly or indirectly in the United States or to any US Person. A US Person is defined as (a) any individual who is a citizen or resident of the United States for federal income tax purposes; (b) a corporation, partnership or other entity created or organized under the laws of or existing in the United States; (c) an estate or trust the income of which is subject to United States federal income tax regardless of whether such income is effectively connected with a United States trade or business. In the United States, this material may be distributed only to a person who is a "distributor", or who is not a "US person", as defined by Regulation S under the U.S. Securities Act of 1933 (as amended).

Additional information for investors with residence or seat in Australia and New Zealand

This document is distributed in Australia by Robeco Hong Kong Limited (ARBN 156 512 659) ("RIAM BV"), which is exempt from the requirement to hold an Australian financial services license under the Corporations Act 2001 (Cth) pursuant to ASIC Class Order 03/1103. Robeco is regulated by the Securities and Futures Commission under the laws of Hong Kong and those laws may differ from Australian laws. This document is distributed only to "wholesale clients" as that term is defined under the Corporations Act 2001 (Cth). This document is not intended for distribution or dissemination, directly or indirectly, to any other class of persons. In New Zealand, this document is only available to wholesale investors within the meaning of clause 3(2) of Schedule 1 of the Financial Markets Conduct Act 2013 (FMCA). This document is not intended for public distribution in Australia and New Zealand.

Additional information for investors with residence or seat in Austria

This information is solely intended for professional investors or eligible counterparties in the meaning of the Austrian Securities Oversight Act.

Additional information for investors with residence or seat in Brazil

The Fund may not be offered or sold to the public in Brazil. Accordingly, the Fund has not been nor will be registered with the Brazilian Securities Commission (CVM), nor has it been submitted to the foregoing agency for approval. Documents relating to the Fund, as well as the information contained therein, may not be supplied to the public in Brazil, as the offering of the Fund is not a public offering of securities in Brazil, nor may they be used in connection with any offer for subscription or sale of securities to the public in Brazil.

Additional information for investors with residence or seat in Brunei

The Prospectus relates to a private collective investment scheme which is not subject to any form of domestic regulations by the Autoriti Monetari Brunei Darussalam ("Authority"). The Prospectus is intended for distribution only to specific classes of investors as specified in section 20 of the Securities Market Order, 2013, and must not, therefore, be delivered to, or relied on by, a retail client. The Authority is not responsible for reviewing or verifying any prospectus or other documents in connection with this collective investment scheme. The Authority has not approved the Prospectus or any other associated documents nor taken any steps to verify the information set out in the Prospectus and has no responsibility for it. The units to which the Prospectus relates may be illiquid or subject to restrictions on their resale. Prospective purchasers of the units offered should conduct their own due diligence on the units.

Additional information for investors with residence or seat in Canada

No securities commission or similar authority in Canada has reviewed or in any way passed upon this document or the merits of the securities described herein, and any representation to the contrary is an offence. Robeco Institutional Asset Management B.V. relies on the international dealer and international adviser exemption in Quebec and has appointed McCarthy Tétrault LLP as its agent for service in Quebec.

Additional information for investors with residence or seat in the Republic of Chile

Neither Robeco nor the Funds have been registered with the Comisión para el Mercado Financiero pursuant to Law no. 18.045, the Ley de Mercado de Valores and regulations thereunder. This document does not constitute an offer of or an invitation to subscribe for or purchase shares of the Funds in the Republic of Chile, other than to the specific person who individually requested this information on their own initiative. This may therefore be treated as a "private offering" within the meaning of Article 4 of the Ley de Mercado de Valores (an offer that is not addressed to the public at large or to a certain sector or specific group of the public).

Additional information for investors with residence or seat in Colombia

This document does not constitute a public offer in the Republic of Colombia. The offer of the fund is addressed to less than one hundred specifically identified investors. The fund may not be promoted or marketed in Colombia or to Colombian residents, unless such promotion and marketing is made in compliance with Decree 2555 of 2010 and other applicable rules and regulations related to the promotion of foreign funds in Colombia. The distribution of this Prospectus and the offering of Shares may be restricted in certain jurisdictions. The information contained in this Prospectus is for general guidance only, and it is the responsibility of any person or persons in possession of this Prospectus and wishing to make application for Shares to inform themselves of, and to observe, all applicable laws and regulations of any relevant

jurisdiction. Prospective applicants for Shares should inform themselves of any applicable legal requirements, exchange control regulations and applicable taxes in the countries of their respective citizenship, residence or domicile.

Additional information for investors with residence or seat in the Dubai International Financial Centre (DIFC), United Arab Emirates
This material is distributed by Robeco Institutional Asset Management B.V. (DIFC Branch) located at Office 209, Level 2, Gate Village Building 7, Dubai International Financial Centre, Dubai, PO Box 482060, UAE. Robeco Institutional Asset Management B.V. (DIFC Branch) is regulated by the Dubai Financial Services Authority ("DFSA") and only deals with Professional Clients or Market Counterparties and does not deal with Retail Clients as defined by the DFSA.

Additional information for investors with residence or seat in France
Robeco Institutional Asset Management B.V. is at liberty to provide services in France. Robeco France is a subsidiary of Robeco whose business is based on the promotion and distribution of the group's funds to professional investors in France.

Additional information for investors with residence or seat in Germany
This information is solely intended for professional investors or eligible counterparties in the meaning of the German Securities Trading Act.

Additional information for investors with residence or seat in Hong Kong
The contents of this document have not been reviewed by the Securities and Futures Commission ("SFC") in Hong Kong. If there is any doubt about any of the contents of this document, independent professional advice should be obtained. This document has been distributed by Robeco Hong Kong Limited ("Robeco"). Robeco is regulated by the SFC in Hong Kong.

Additional information for investors with residence or seat in Indonesia
The Prospectus does not constitute an offer to sell nor a solicitation to buy securities in Indonesia.

Additional information for investors with residence or seat in Italy
This document is considered for use solely by qualified investors and private professional clients (as defined in Article 26 (1) (b) and (d) of Consob Regulation No. 16190 dated 29 October 2007). If made available to Distributors and individuals authorized by Distributors to conduct promotion and marketing activity, it may only be used for the purpose for which it was conceived. The data and information contained in this document may not be used for communications with Supervisory Authorities. This document does not include any information to determine, in concrete terms, the investment inclination and, therefore, this document cannot and should not be the basis for making any investment decisions.

Additional information for investors with residence or seat in Japan
This document is considered for use solely by qualified investors and is distributed by Robeco Japan Company Limited, registered in Japan as a Financial Instruments Business Operator, [registered No. the Director of Kanto Local Financial Bureau (Financial Instruments Business Operator), No.2780, Member of Japan Investment Advisors Association].

Additional information for investors with residence or seat in South Korea
The Management Company is not making any representation with respect to the eligibility of any recipients of the Prospectus to acquire the Shares therein under the laws of South Korea, including but not limited to the Foreign Exchange Transaction Act and Regulations thereunder. The Shares have not been registered under the Financial Investment Services and Capital Markets Act of Korea, and none of the Shares may be offered, sold or delivered, or offered or sold to any person for re-offering or resale, directly or indirectly, in South Korea or to any resident of South Korea except pursuant to applicable laws and regulations of South Korea.

Additional information for investors with residence or seat in Liechtenstein
This document is exclusively distributed to Liechtenstein-based, duly licensed financial intermediaries (such as banks, discretionary portfolio managers, insurance companies, fund of funds) which do not intend to invest on their own account into Fund(s) displayed in the document. This material is distributed by Robeco Switzerland Ltd, postal address: Josefstrasse 218, 8005 Zurich, Switzerland. LGT Bank Ltd., Herrengasse 12, FL-9490 Vaduz, Liechtenstein acts as the representative and paying agent in Liechtenstein. The prospectus, the Key Information Documents (PRIIP)the articles of association, the annual and semi-annual reports of the Fund(s) may be obtained from the representative or via the website.

Additional information for investors with residence or seat in Malaysia
Generally, no offer or sale of the Shares is permitted in Malaysia unless where a Recognition Exemption or the Prospectus Exemption applies: NO ACTION HAS BEEN, OR WILL BE, TAKEN TO COMPLY WITH MALAYSIAN LAWS FOR MAKING AVAILABLE, OFFERING FOR SUBSCRIPTION OR PURCHASE, OR ISSUING ANY INVITATION TO SUBSCRIBE FOR OR PURCHASE OR SALE OF THE SHARES IN MALAYSIA OR TO PERSONS IN MALAYSIA AS THE SHARES ARE NOT INTENDED BY THE ISSUER TO BE MADE AVAILABLE, OR MADE THE SUBJECT OF ANY OFFER OR INVITATION TO SUBSCRIBE OR PURCHASE, IN MALAYSIA. NEITHER THIS DOCUMENT NOR ANY DOCUMENT OR OTHER MATERIAL IN CONNECTION WITH THE SHARES SHOULD BE DISTRIBUTED, CAUSED TO BE DISTRIBUTED OR CIRCULATED IN MALAYSIA. NO PERSON SHOULD MAKE AVAILABLE OR MAKE ANY INVITATION OR OFFER OR INVITATION TO SELL OR PURCHASE THE SHARES IN MALAYSIA UNLESS SUCH PERSON TAKES THE NECESSARY ACTION TO COMPLY WITH MALAYSIAN LAWS.

Additional information for investors with residence or seat in Mexico
The funds have not been and will not be registered with the National Registry of Securities or maintained by the Mexican National Banking and Securities Commission and, as a result, may not be offered or sold publicly in Mexico. Robeco and any underwriter or purchaser may offer and sell the funds in Mexico on a private placement basis to Institutional and Accredited Investors, pursuant to Article 8 of the Mexican Securities Market Law.

Additional information for investors with residence or seat in Peru
The Superintendencia del Mercado de Valores (SMV) does not exercise any supervision over this Fund and therefore the management of it. The information the Fund provides to its investors and the other services it provides to them are the sole responsibility of the Administrator. This Prospectus is not for public distribution.

Additional information for investors with residence or seat in Singapore
This document has not been registered with the Monetary Authority of Singapore ("MAS"). Accordingly, this document may not be circulated or distributed directly or indirectly to persons in Singapore other than (i) to an institutional investor under Section 304 of the SFA, (ii) to a relevant person pursuant to Section 305(1), or any person pursuant to Section 305(2), and in accordance with the conditions specified in Section 305, of the SFA, or (iii) otherwise pursuant to, and in accordance with the conditions of, any other applicable provision of the SFA. The contents of this document have not been reviewed by the MAS. Any decision to participate in the Fund should be made only after reviewing the sections regarding investment considerations, conflicts of interest, risk factors and the relevant Singapore selling restrictions (as described in the section entitled "Important information for Singapore Investors") contained in the prospectus. Investors should consult their professional adviser if you are in doubt about the stringent restrictions applicable to the use of this document, regulatory status of the Fund, applicable regulatory protection, associated risks and suitability of the Fund to your objectives. Investors should note that only the Sub-Funds listed in the appendix to the section entitled "Important information for Singapore Investors" of the prospectus ("Sub-Funds") are available to Singapore investors. The Sub-Funds are notified as restricted foreign schemes under the Securities and Futures Act, Chapter 289 of Singapore ("SFA") and invoke the exemptions from compliance with prospectus registration requirements pursuant to the exemptions under Section 304 and Section 305 of the SFA. The Sub-Funds are not authorized or recognized by the MAS and shares in the Sub-Funds are not allowed to be offered to the retail public in Singapore. The prospectus of the Fund is not a prospectus as defined in the SFA. Accordingly, statutory liability under the SFA in relation to the content of prospectuses does not apply. The Sub-Funds may only be promoted exclusively to persons who are sufficiently experienced and sophisticated to understand the risks involved in investing in such schemes, and who satisfy certain other criteria provided under Section 304, Section 305 or any other applicable provision of the SFA and the subsidiary legislation enacted thereunder. You should consider carefully whether the investment is suitable for you. Robeco Singapore Private Limited holds a capital markets services license for fund management issued by the MAS and is subject to certain clientele restrictions under such license.

Additional information for investors with residence or seat in Spain
Robeco Institutional Asset Management B.V., Sucursal en España with identification number W0032687F and having its registered office in Madrid at Calle Serrano 47-14º, is registered with the Spanish Commercial Registry in Madrid, in volume 19.957, page 190, section 8, sheet M-351927 and with the National Securities Market Commission (CNMV) in the Official Register of branches of European investment services companies, under number 24. The investment funds or SICAV mentioned in this document are regulated by the corresponding authorities of their country of origin and are registered in the Special Registry of the CNMV of Foreign Collective Investment Institutions marketed in Spain.

Additional information for investors with residence or seat in South Africa
Robeco Institutional Asset Management B.V. is registered and regulated by the Financial Sector Conduct Authority in South Africa.

Additional information for investors with residence or seat in Switzerland
The Fund(s) are domiciled in Luxembourg. This document is exclusively distributed in Switzerland to qualified investors as defined in the Swiss Collective Investment Schemes Act (CISA). This material is distributed by Robeco Switzerland Ltd, postal address: Josefstrasse 218, 8005 Zurich. ACOLIN Fund Services AG, postal address: Leutschenbachstrasse 50, 8050 Zürich, acts as the Swiss representative of the Fund(s). UBS Switzerland AG, Bahnhofstrasse 45, 8001 Zurich, postal address: Europastrasse 2, P.O. Box, CH-8152 Opfikon, acts as the Swiss paying agent. The prospectus, the Key Information Documents (PRIIP), the articles of association, the annual and semi-annual reports of the Fund(s), as well as the list of the purchases and sales which the Fund(s) has undertaken during the financial year, may be obtained, on simple request and free of charge, at the office of the Swiss representative ACOLIN Fund Services AG. The prospectuses are also available via the website.

Additional information relating to RobecoSAM-branded funds/services
Robeco Switzerland Ltd, postal address Josefstrasse 218, 8005 Zurich, Switzerland has a license as asset manager of collective assets from the Swiss Financial Market Supervisory Authority FINMA. The RobecoSAM brand is a registered trademark of Robeco Holding B.V. The brand RobecoSAM is used to market services and products which entail Robeco's expertise on Sustainable Investing (SI. The brand RobecoSAM is not to be considered as a separate legal entity.

Additional information for investors with residence or seat in Taiwan
The contents of this document have not been reviewed by any regulatory authority in Hong Kong. If you are in any doubt about any of the contents of this document, you should obtain independent professional advice. This document has been distributed by Robeco Hong Kong Limited ("Robeco". Robeco is regulated by the Securities and Futures Commission in Hong Kong.

Additional information for investors with residence or seat in Thailand
The Prospectus has not been approved by the Securities and Exchange Commission, which takes no responsibility for its contents. No offer to the public to purchase the Shares will be made in Thailand and the Prospectus is intended to be read by the addressee only and must not be passed to, issued to, or shown to the public generally.

Additional information for investors with residence or seat in the United Arab Emirates
Some Funds referred to in this marketing material have been registered with the UAE Securities and Commodities Authority ("the Authority". Details of all Registered Funds can be found on the Authority's website. The Authority assumes no liability for the accuracy of the information set out in this material/document, nor for the failure of any persons engaged in the investment Fund in performing their duties and responsibilities.
Additional information for investors with residence or seat in the United Kingdom
Robeco is deemed authorized and regulated by the Financial Conduct Authority.

Additional information for investors with residence or seat in the United Kingdom
Robeco Institutional Asset Management B.V (FRN: 977582) is authorised and regulated by the Financial Conduct Authority.

Additional information for investors with residence or seat in Uruguay
The sale of the Fund qualifies as a private placement pursuant to section 2 of Uruguayan law 18,627. The Fund must not be offered or sold to the public in Uruguay, except under circumstances which do not constitute a public offering or distribution under Uruguayan laws and regulations. The Fund is not and will not be registered with the Financial Services Superintendency of the Central Bank of Uruguay. The Fund corresponds to investment funds that are not investment funds regulated by Uruguayan law 16,774 dated 27 September 1996, as amended.