



February 2023

www.research.hsbc.com

By: Mark McDonald, Davey Jose and Camila Sarmiento

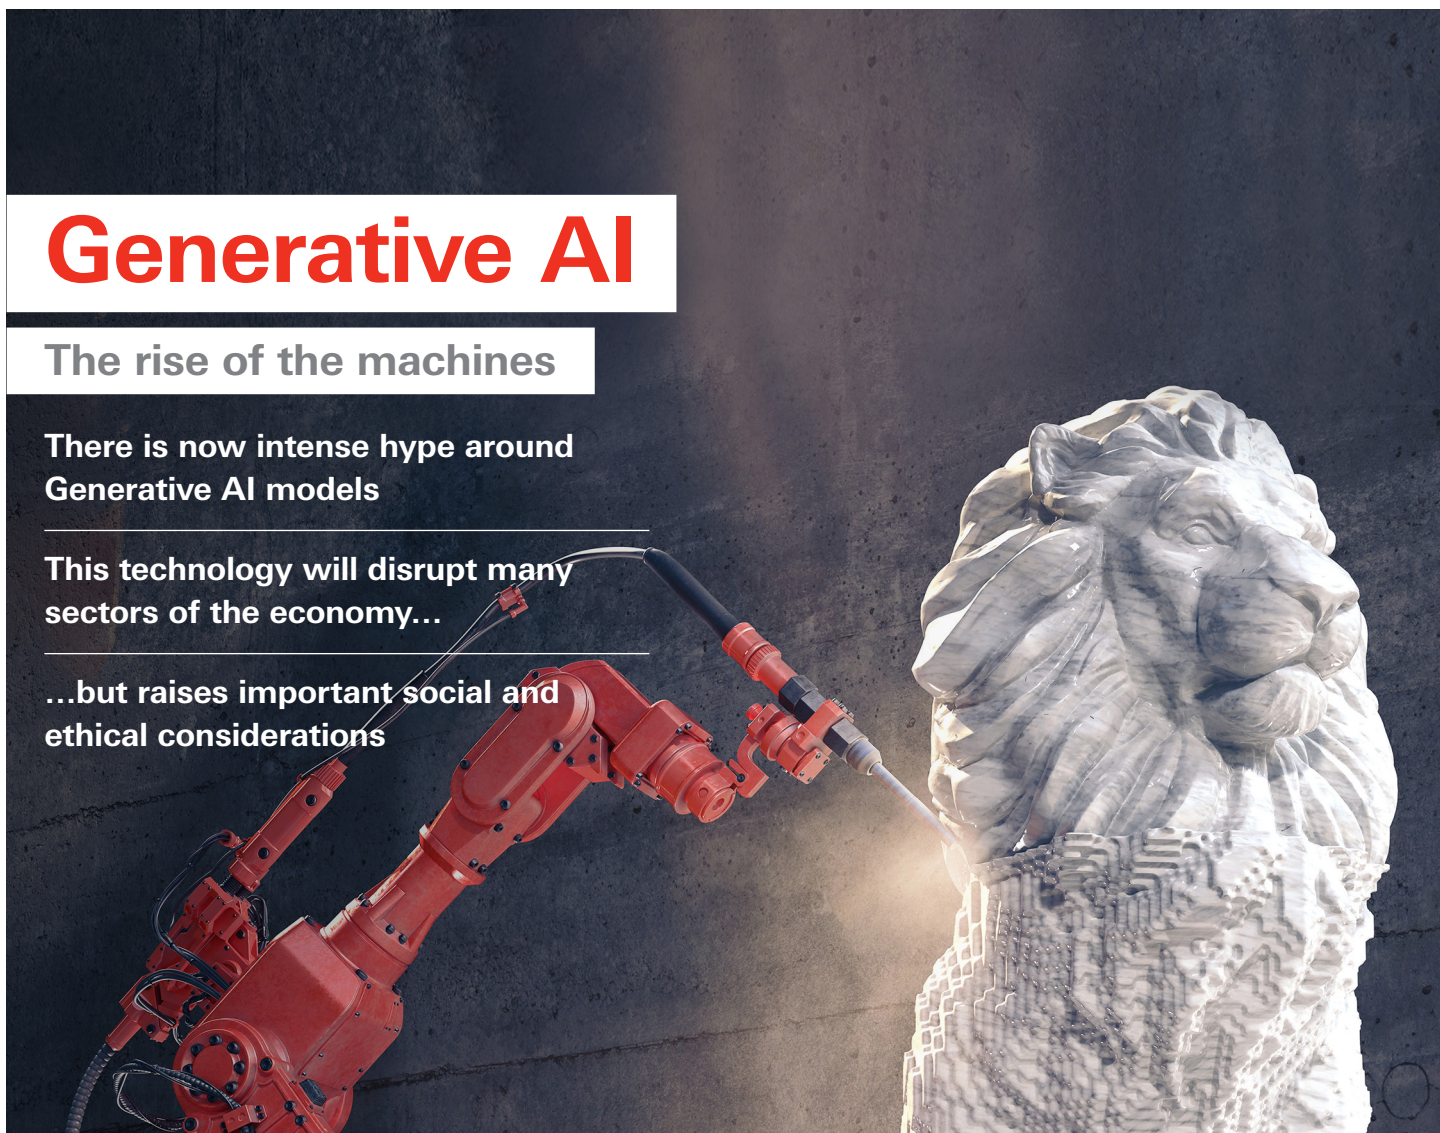
Generative AI

The rise of the machines

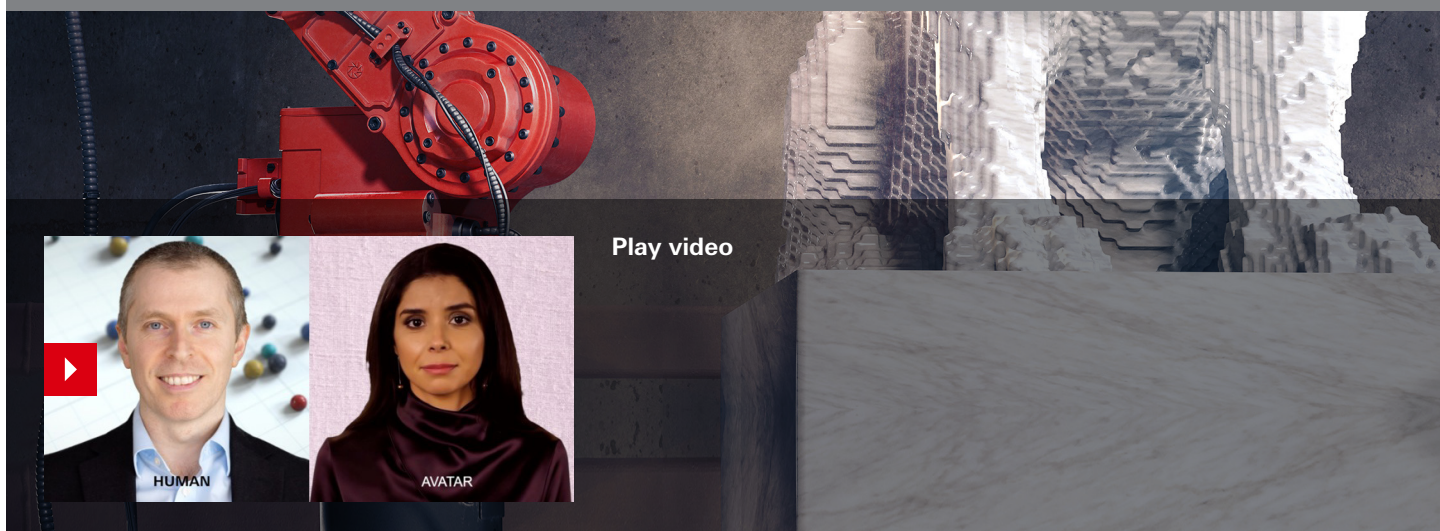
There is now intense hype around
Generative AI models

This technology will disrupt many
sectors of the economy...

...but raises important social and
ethical considerations



This is a Free to View version of a report with the same title published on 23-Feb-23. Please contact your HSBC representative or email AskResearch@hsbc.com for more information.



Play video

Disclosures & Disclaimer: This report must be read with the disclosures and the analyst certifications in the Disclosure appendix, and with the Disclaimer, which forms part of it.

Contents

Executive Summary	2
Generative AI	2
Significant ethical considerations	4
AI-mazing economic opportunities	5
Ethical considerations of Generative AI	6
Social considerations	7
Reliability and accuracy	8
Disinformation	8
Biases	9
Is it coming for my job?	11
Legal implications	13
Regulations	13
Disclosure appendix	15
Disclaimer	16

Executive Summary

Generative AI

Generative AI is one of the hottest topics within machine learning at the moment. There is a significant degree of hype surrounding the field and investment flows are pouring into start-up companies trying to commercialise the technology. But what is it, why is it a hot topic now, and what are the implications for investors in public companies?

What is Generative AI?

Whilst AI (artificial intelligence) has been a hot topic for several years¹, most of the focus until recently has been on predictive machine learning models. Over the last year there has been a surge in interest in, and a staggering improvement in the performance of, so-called **generative models**. These are machine learning models which can produce new content, something which until recently was considered to be the unique purview of humans.

The variety of Generative AI models is huge: the two models which have made the biggest influence on public awareness are DALL-E, which produces images from text, and ChatGPT, which produces text via a chatbot interface. Both of these models were produced by OpenAI, but many of the major tech companies have built/are building their own equivalent generative models. Increasingly, there are high quality open-source generative models which anyone can use and a wide array of start-ups which are commercialising the output of generative models.

This description barely scratches the surface; the types of content which can be generated by Generative AI models is steadily widening. We discuss the most important model categories in the *What can Generative AI do* section in the full report. For those readers who are curious about how it is possible for Machine Learning models, which ultimately need to process numbers, to generate text, image, or audio content, please see the *Geek Squad* also in the full report.

Why is Generative AI a hot topic now?

Generative AI has suddenly captured the popular imagination, with major newspapers writing opinion articles on the technology and social media feeds being inundated with discussions of what is now possible. So why has this happened now?

From a technical perspective, there has been a smooth evolution of what is possible over the past few years. As such, many technologists appear somewhat bemused as to why this technology has suddenly gained such prominence.

Of course, what has really changed in recent months is the ability for the non-technical audience to use this technology. For example, the core technology which powers ChatGPT is a type of machine learning model known as a Large Language Model (LLM). The LLM which ChatGPT uses is not a dramatic improvement on the state of the art technology a year ago, or even two years ago. However, back then these LLMs were only able to be used by people who knew how to code. Whereas now, anyone who knows how to use a website can access ChatGPT.

Can a robot turn a canvas
into a beautiful painting?
... well they can now

¹ There have been several previous waves of interest in AI, but the current wave really began in 2012 following the success of AlexNet in the 2012 ImageNet Challenge.

Generative AI technology is now widely available to the general public

This change may seem small to those who are technically adept themselves; it is, however, transformative. These technologies are now easily accessible by any entrepreneur or content creator so we are seeing a dramatic expansion in the use cases being actively explored. As with many other technologies, it is not the invention of the new technology that causes a significant impact; rather, it is when the technology matures and disperses throughout the population.

What are the implications for investors?

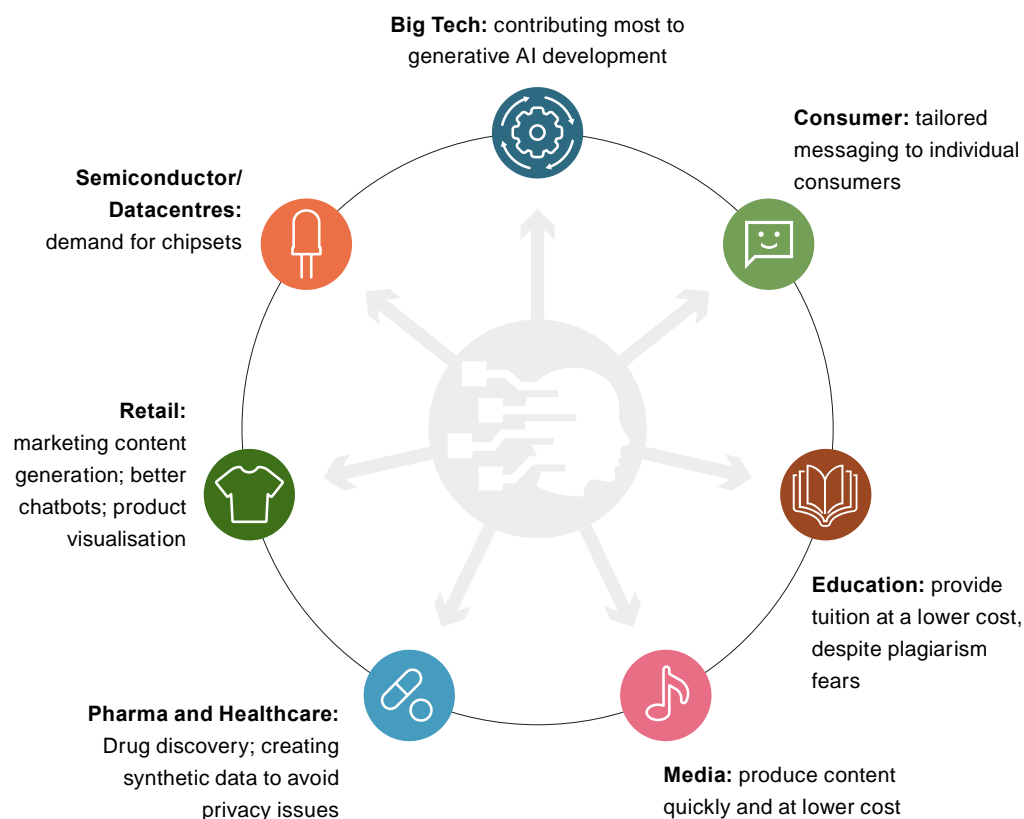
With any nascent technology it is hard to predict the various ways it will ultimately end up being used productively. However, there is the potential for Generative AI to have a significant impact on several sectors of the economy.

Who will be the winners and losers? And is this appropriately priced in?

With all disruptive technologies, there will be winners and losers. For investors, determining the likely winners and losers from disruptive technology is often easier than understanding whether the opportunities and challenges from this technology have been appropriately priced in. In the full report our sector analysts worldwide discuss the ways in which Generative AI is likely to be commercialised in their sectors – see the *Thematic and Sector Implications* section.

- ◆ **Media:** Music generated from AI has already started to appear on streaming platforms. These models will cause the cost of production to collapse for many types of content. This will clearly have a huge impact on any content-generating industry.
- ◆ **Big Tech:** Generative AI is key for Big Tech companies; indeed, these companies are often those contributing the most to progress in the field of Generative AI. The sharp fall in Alphabet's share price in response to the poorly received public demo of Google Bard highlights the potential cost of getting it wrong in this area for Big Tech companies.
- ◆ **Semiconductor/Datacentres:** Generative AI models typically need advanced chipsets; this is particularly true for the training process but increasingly true even when the models are being used in production. This should stimulate demand for those companies who manufacture advanced chipsets.
- ◆ **Pharma and Healthcare:** The impact of Generative AI on the pharma and healthcare sector has been muted thus far. However, these sectors have been keen adopters of other AI techniques and there are several promising areas for the application of Generative AI to pharma and healthcare.
- ◆ **Education:** Tools such as ChatGPT could be a powerful tool for education if used in the right way. As techniques are developed to reduce the likelihood of LLMs giving incorrect information these tools will be able to provide tuition more broadly and at a lower cost. However, at the moment, tools such as ChatGPT are leading to concerns of a significant increase in plagiarism.
- ◆ **Consumer:** There are a number of areas in which Generative AI can be of use to consumer staples companies, all of which stem fundamentally from the ability to tailor messages more tightly to individual consumers.
- ◆ **Retail:** The retail sector has seen widespread adoption of Predictive AI techniques. Generative AI will have applications in retail, but these are likely to be less disruptive than Predictive AI techniques have been.

The impact of Generative AI on sectors



Source: HSBC

Significant ethical considerations

Like all technological developments, there is the possibility of Generative AI techniques being misused. There are several key ethical issues associated with generative natural language processing (NLP) that are worth considering. Some of the most important include:

1. **Bias:** Generative AI models can perpetuate and even amplify existing biases in the data they are trained on. For example, a model trained on a biased dataset of news articles might generate text that reflects those biases. This could perpetuate harmful stereotypes and discrimination.
2. **Privacy:** Generative AI models can be used to generate realistic synthetic data, raising concerns about the protection of individuals' privacy.
3. **Misuse:** Generative AI models can be used to create deepfakes or other manipulated content, which can be used to spread misinformation or cause harm.
4. **Impact on Employment:** Generative AI models can be used to automate the creation of content, raising concerns about the impact on employment in industries such as media and entertainment.
5. **Lack of Transparency:** Generative models can be complex and hard to interpret, making it difficult to understand how they arrived at a certain output, which can lead to lack of trust in the model and its decisions.
6. **Inclusion:** Generative models might not be accessible to everyone. This could lead to issues of social and economic inequality, as individuals who are not able to access these tools might be at a disadvantage.

AI-mazing economic opportunities

There is the potential for a dramatic increase in productivity from Generative AI. Large tracts of the economy rely in some way on content generation. Even for industries that one might not immediately think of being influenced by this technology there are likely to be considerable resources devoted to producing marketing content. Generative AI will massively reduce the cost of producing such content.

Did you spot our guest author for the previous section?

For example, the **section you have just read titled *Significant ethical considerations* was actually written by ChatGPT** – the recently released AI model from OpenAI. This model also was responsible for the somewhat ‘cheesy’ title of this subsection.

Good content

The first thing to notice here is that the content generated is actually very good: the six ethical issues associated with Generative AI which the model describes are all valid considerations. Furthermore, the explanation of each issue which ChatGPT discussed is a good summary of the issue. However, there are some important ethical issues that ChatGPT did not identify. We discuss the ethical issues of Generative AI in more detail (and the discussion is written by a human expert in this case!) in the *Social Considerations* section in the full report. This section will also focus on the social and governance issues that plague Generative AI.

Appropriate content

The second point is that the tone of the output is appropriate. If we had not explicitly highlighted the section that was generated by AI it is unlikely that readers would have noticed.

Same content, multiple formats

Another important consideration is the ease with which existing content can be adapted for new formats. We asked ChatGPT to instead produce an extended version of this information in the form of a twitter thread; asking for it to use language that is likely to lead to high engagement.

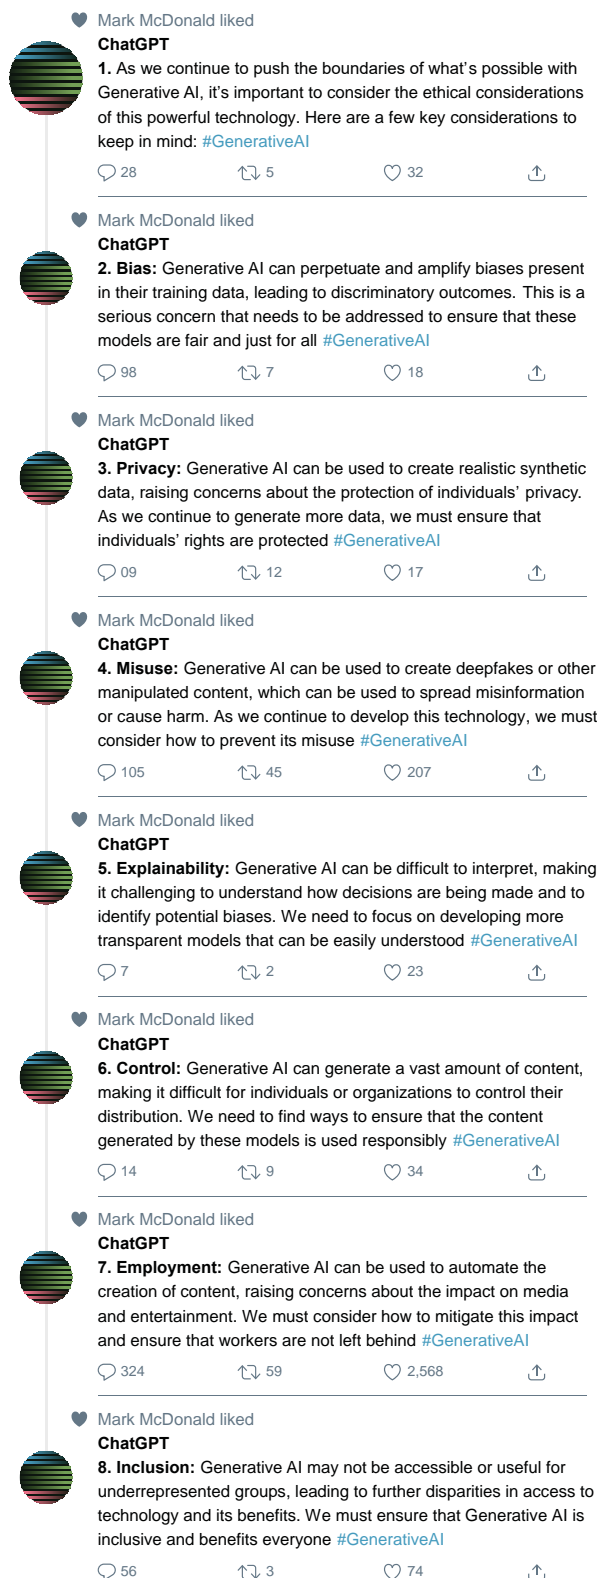
Finally, we asked ChatGPT to put this information into the form of a transcript for a TikTok video. The model happily obliged, even peppering the transcript with directions for what should be on screen during different sections of the video such as “*cut to a shot of a person scrolling through a social media feed*”.

Already a useful tool

Clearly the technology is not yet perfect. It is, however, already at the point where it can be an extremely useful tool for anyone generating content. Still not convinced? Large sections of the [video](#) accompanying this piece were entirely produced by AI. Please check out the video to see what’s now possible with very low cost.

Ethical considerations of Generative AI

We asked ChatGPT to generate a Twitter thread ...



... and to script a TikTok video post

[Opening shot of a computer screen with the text "Generative AI: The Ethical Implications"]
Voiceover: As we continue to push the boundaries of what's possible with Generative AI, it's important to consider the ethical implications of this powerful technology.

[Cut to a shot of a person scrolling through a social media feed]
Voiceover: Bias. Generative AI can perpetuate and amplify biases present in their training data, leading to discriminatory outcomes.

[Cut to a shot of a person typing on a computer]
Voiceover: Privacy. Generative AI can be used to create realistic synthetic data, raising concerns about the protection of individuals' privacy.

[Cut to a shot of a deepfake video]
Voiceover: Misuse. Generative AI can be used to create deepfakes or other manipulated content, which can be used to spread misinformation or cause harm.

[Cut to a shot of a person looking confused while trying to understand a computer screen]
Voiceover: Explainability. Generative AI can be difficult to interpret, making it challenging to understand how decisions are being made and to identify potential biases.

[Cut to a shot of a person scrolling through social media on their phone]
Voiceover: Control. Generative AI can generate a vast amount of content, making it difficult for individuals or organizations to control their distribution.

[Cut to a shot of a person sitting in front of a camera with a worried expression]
Voiceover: Employment. Generative AI can be used to automate the creation of content, raising concerns about the impact on employment in industries such as media and entertainment.

[Cut to a shot of a person from an underrepresented group using a computer]
Voiceover: Inclusion. Generative AI may not be accessible or useful for underrepresented groups, leading to further disparities in access to technology and its benefits.



Source: HSBC, Getty Images

Social considerations

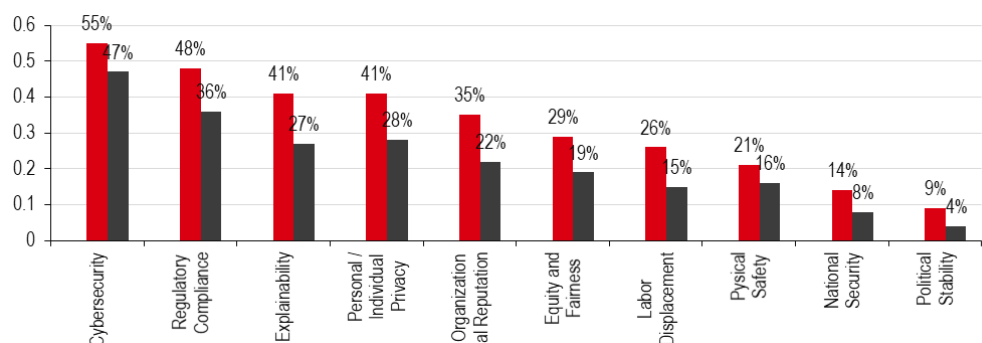
- ◆ Generative AI has many potential applications for business, but social impacts need to be carefully assessed
- ◆ Content generated can raise concerns around biases, ethics, and fairness, and expose corporates to various risks
- ◆ Generative AI can also be influential for job markets, education systems and regulatory priorities

As Generative AI undergoes rapid development, its impacts on broader society are rapidly evolving. The technology can bring many advantages including increased efficiency, improved quality of content generated and even cost benefits. Generative AI can help businesses make better decisions, be more creative and improve their operations. However, despite the value Generative AI can bring, as with many new technologies, social concerns are apparent. Social biases are often reflected in models and thus can perpetuate these biases further, raising concerns around ethics and fairness. Additionally, unemployment, misinformation, and privacy risk often surround the AI space and, to no surprise, are of particular relevance to Generative AI as it is rapidly integrated across regions, sectors and in everyday life. As the industry matures, these risks can dissipate, but they should be closely watched, cautiously considered, and fundamentally understood by investors.

The social impacts of Generative AI should be closely watched by investors and corporates

As investors pour money into the Generative AI sector, we think it is pertinent to assess and address the risks around the potential for corporate and societal impacts. Indeed, while many corporates identify social risks associated with AI, they are less likely to be already taking action to mitigate these risks (Chart A1). To help investors navigate this complex and fast-changing landscape, we look at several of these risks in more detail below.

A1. AI risks corporate consider relevant vs are taking steps to mitigate



Source: Artificial Intelligence Report 2022, HSBC

Reliability and accuracy

Generative AI is not fool proof

In producing content – images, longer text, emails, social media content, voice recordings, program code, or structured data – Generative AI applications use machine learning to make predictions based on previous data, which can reveal problems with accuracy and reliability.

Instances of errors generated by platforms like OpenAI's ChatGPT, are not uncommon. For example, news sites report that CNET editors fed the bot article outlines and when the draft was reviewed and edited, errors were found in more than half of the articles published by the editor-bot teams: errors ranged from minor mistakes to issues that required "substantial correction".

Even when it comes to objective facts, Generative AI can also make mistakes. As OpenAI says on its website "*ChatGPT sometimes writes plausible-sounding but incorrect or nonsensical answers*". This was also demonstrated during the launch of Google's new AI chatbot, Bard. The system made a factual error saying that James Webb Space Telescope snapped "the very first image of a planet outside our solar system," even though the milestone happened 17-years before the telescope launched.

Generative AI may spread misinformation...

While these errors show imperfections in AI-generating systems and raise questions over its real world applications, it also raises questions around **the spreading of misinformation**. If you are unknowledgeable about the launch of the James Webb Space Telescope, you are less likely to catch the inaccurate facts. Indeed, this is further explored in a Northwestern University study that used ChatGPT to generate artificial medical research abstracts to test whether scientists can spot them. While the scientists identified the majority of generated abstracts, they incorrectly identified 32% of generated abstracts and 14% of genuine abstracts.² This shows that while a less knowledgeable group might not catch flaws and discrepancies in Generative AI, the technology also has the ability to fool experts.

In a society with extensive communication platforms, information is traveling fast and it is becoming increasingly difficult to verify information online, especially with entrance of Generative AI. While Generative AI technology can have many corporate benefits, the lack of reliability and accuracy can pose risks to corporate integrity and reputation, and impact the bottom line. Indeed, after Google's Bard showed the inaccurate answer on the telescope, parent company Alphabet experienced an 8% drop in its share price that day.



ChatGPT sometimes writes plausible-sounding but incorrect or nonsensical answers

OpenAI website

Disinformation

Generative AI also can facilitate the spreading of false or misleading ideas

In addition to the imperfections and factual errors made by Generative AI, the technology also raises the risk that it may be used to spread disinformation: the deliberate intention of spreading false or misleading ideas. Indeed, ChatGPT will produce work around false narratives when asked. For example, researchers at NewsGuard asked ChatGPT to write responses based on false or misleading ideas, and the bot complied about 80% of the time³.

² Abstracts written by ChaptGPT fool scientists, Nature, 2023

³ The Next Great Misinformation Superspreader, NewsGuard, 2023

Because Generative AI can pose risks of increased sharing of disinformation, they can disrupt the credibility of news and information in general. For as long as we have known, photographs, videos and audio have been trusted primary sources of information. Increasing the sharing of false and misleading information can create mistrust in people for what they see and hear online if they cannot tell what is real or not.

Platforms are stepping up efforts to mitigate these risks

AI generating platforms are trying to mitigate this. For example, OpenAI policies prohibit the use of its technology to build products that target “*Illegal or harmful industries, misuse of personal data, promoting dishonesty, deceiving or manipulating users or trying to influence politics,*” and use their Moderation API tool to warn or block certain types of unsafe content, but this is not perfect. Additionally, their new AI classifier tool can be used to help determine text written by human vs artificial intelligence to identify misinformation campaigns. While these innovations are needed to build trust in the model, they often lag behind technological advances. We think innovations and regulatory policies will develop to combat these risks eventually, nonetheless corporates will need to enhance their due diligence and strategies around verifying online content.

Deepfakes

Generative AI poses risks to cybersecurity

Deepfakes use Generative AI to make artificial videos of fake events by combining images to make new footage that depict events or statements that didn't actually happen. While deepfakes do have positive applications in many industries, particularly for modelling human behaviour, there are also many possibilities for abuse. This AI-generated content can pose serious ethical implications, in our view, since it can be very difficult, and in some cases impossible, to distinguish between real and fake information. **Deepfakes could be used for identity theft, extortion, reputational damage, ridicule, intimidation and harassment.**⁴

As Generative AI for deepfakes continues to advance, these societal risks will become more important to address. While there have been innovations to help identify deepfakes, they are not yet keeping pace with the technology itself. Some efforts include researchers at University of California, Berkeley creating software for deepfake detection, and OpenAI working on a way to watermark text. Indeed, given the social risks with deepfakes, more controls are likely to be investigated and required in the future.

Illicit activities

While AI can be used as a tool to benefit cybersecurity, for example in detecting anomalies or predicting cyber threats, Generative AI can be misused in a way that harms cybersecurity. For example, Generative AI can simulate very realistic phishing emails and phone calls, which can make it simpler for individuals to be tricked into revealing sensitive information or becoming vulnerable to malware.

Biases

Generated AI models emulate biases from past human content they are trained on

Concerns about biases in artificial intelligence are not new. Indeed, social biases including those around sex, race and disabilities, are often absorbed by AI systems from their training data. Since they are trained on past human content they tend to emulate biased language they were exposed to. A report by Stanford University exemplifies this, where an AI model training on BookCorpus (a dataset of books from e-book websites), often produced toxic text (defined as rude, disrespectful or unreasonable comments), most likely as a result of the significant number of romance novels containing explicit content. Models training on internet text that had filtered out toxic content, produced significantly less toxic content (see Chart A2). Indeed, it is important to note that in defining what is toxic and not, the study can contain biases itself. Nonetheless, the fact that that toxic information was omitted when the content was filtered, shows how important the data informing the bot is to its behaviour. The study additionally finds larger models are more capable of reflecting biases in their training data⁵. This means that although language models continue to grow and develop, so does their severity to biases.

⁴ Australian Government, eSafety Commissioner

⁵ AI Index Report 2022, Stanford University Human-Centered Artificial Intelligence

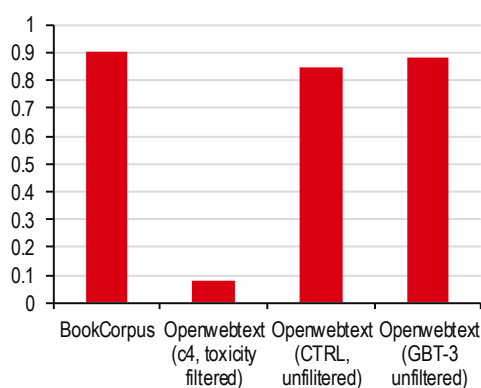
There are widespread examples of Generative AI replicating social biases. For example, when Generative AI models produce outputs which reflect men in more senior positions, or pictures of women as teachers vs men as professors. These biases can help perpetuate inequalities when these AI models are not inclusive or useful to underrepresented groups.

Additionally, the engineers and developers who create AI systems have frequently received blame in their role with creating these biases. For example, several studies show ChatGPT exhibiting left-leaning biases when fed political questions. A possibility for the source of these biases can be from an overrepresentation of data from sources like academic literature that are often left-leaning. However, another possibility is that biases of the programmers are reflected in the model itself. Tech employees overwhelmingly support Democrats and tend to identify as liberal⁶. Tech employees overwhelmingly support Democrats and tend to identify as liberal⁶.

Can a more diverse group of programmers help eliminate biases?

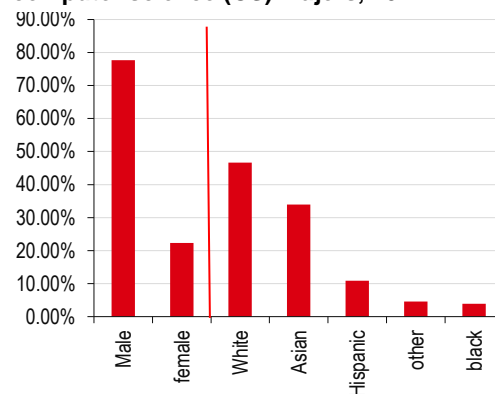
This raises other questions around the role of diversity in Artificial Intelligence. For example, given that tech and the AI space is overwhelmingly favoured by males over females (Chart A3), gender biases may be unintentionally perpetuated⁷. As we discussed in our *Diversity Challenges* report (2 December 2021), research shows that when teams have one or more members who represent the gender, ethnicity, culture, generation or sexual orientation of the team's target end user, the entire team is at least 144% more likely to say that it understands that end user⁸. In our view, diversity can be a useful tool within the Generative AI space in aligning stakeholder expectations with the model and by doing so, also eliminating biases.

A2. Toxicity in language models by language set



Source: Gehman et al., 2021; Rae et al., 2021; Welbl et al., 2021.; 2022 AI Index Report HSBC

A3. US diversity in bachelor degree computer science (CS) majors, 2021



Source: CRA Taulbee Survey 2022, HSBC
 *Racial diversity does not include non-US residents

People are trying, but more effort is needed

Efforts are being made to combat biases in artificial intelligence, but they have not proven successful. Companies like OpenAI are working on filtering out hate speech, and more and more research is being conducted to understand how AI biases develop. Indeed, NeurIPS, one of the largest AI conferences, collected 4 times more research papers on fairness and bias in 2021 than it had in 2017⁹. Additionally, there has been significant research investment in developing ethics diagnostic metrics and benchmarks to understand system biases. However, there is little publically available information. There is also concern that AI-related degree programmes fail to equip designers with appropriate knowledge of ethics.

⁶ Predispositions and the Political Behavior of American Economic Elites: Evidence from Technology Entrepreneurs, Stanford Graduate School of Business, 2017

⁷ CRA Taulbee Survey 2022, Computing Research Association (CRA)

⁸ S Hewlet et al, Diversity's Positive Impact on Innovation and Outcomes, INSEAD, The Global Talent Competitiveness Indx 2018

⁹ AI Index Report 2022, Stanford University Human-Centered Artificial Intelligence

Generative AI biases present major challenges for corporates that are looking to integrate this technology within corporate strategy. Indeed, the last thing corporates are looking for is for their automation tool to cause problems such as reputational damage or a legal issue. However, despite these risks, efforts to address ethical concerns associated with using AI are limited. According to a McKinsey survey, while 29% and 41% of respondents recognise “equity and fairness” and “explainability” as risks when adopting AI, only 19% and 27% are taking steps to mitigate those risks (Chart A1). We think this will continue to grow as the risks and their repercussions grow and materialise.

Is it coming for my job?

James Pomeroy
Economist
HSBC Bank plc

We argued in *Robots and Jobs* (1 November 2021) that the impact of automation on the labour market is likely to be substantial in the decade to come. Pandemic-related shocks to worker availability and wages are likely to mean more businesses are incentivised to think about how technological developments can help them run more effectively or with fewer staff. And so for Generative AI, the worries immediately turn to whether jobs are on the line. Will previously human performed processes be displaced by AI generation?

Historically, new technologies have accelerated the transition of roles in the labour market

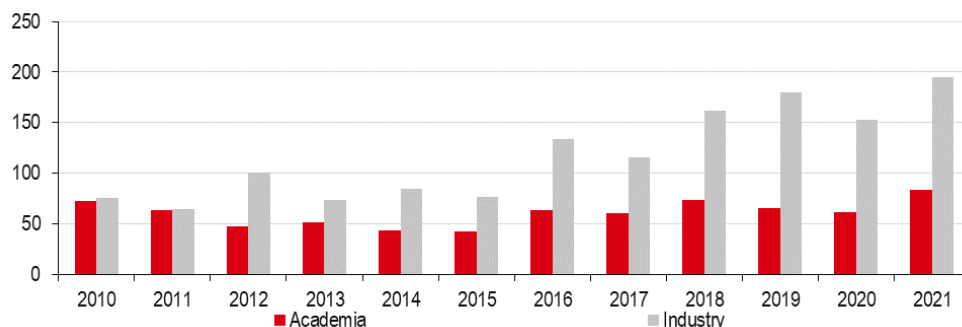
For all the concerns about jobs being lost, history tells us that this phenomenon has typically been overblown, where new technologies have simply been accelerating the transition of roles in the labour market. Before the days of computers there used to be people whose job was “calculator”. The advent of cheap computing power means that those specific jobs don’t exist anymore, but this instead acted as a platform for other jobs, such as in analysis, to be created. In financial markets a wave of electronic trading has replaced human jobs as market makers in many asset classes. However, trading activity has increased notably since then and it is likely that many institutions have far more people employed in the trading process than before, but now it’s quants, database administrators and engineers rather than traders.

and labour markets could experience shifts in skillsets

Indeed, there are many roles that may be vulnerable to Generative AI – particularly those that have a larger number of process components rather than creative or problem solving ones, such as customer service and administrative roles. Rapid developments in AI clearly pose risks to those sorts of roles if problems can be solved much more quickly. There will still be a role for people sharing information, generating product and even thinking about how to get solutions out of AI-supported processes, but they may require different skill sets than before. For example, the workforce for a customer chat service could be replaced by several computer science experts making sure the bot is properly functioning. Indeed, while AI PhD graduates had more often gone into academia 10 years ago, they now overwhelmingly choose to work in industry (Chart A4). In our view this shows a shift in skills now in demand in the labour market.

The impact on the labour market of all of these changes in technological adoption will depend on a number of things – most notably the speed at which new jobs appear, whether training allows people to move between roles and how quickly AI processes both improve and are embedded into existing work processes.

A4. Employment of AI PhDs in academia and industry in North America over time



Source: CRA Taulbee Survey 2022, 2022 AI Index Report HSBC, HSBC

Nonetheless, artificial intelligence can also benefit the sphere of knowledge work without making roles redundant. We argued in [Will AI take your job?](#) (18 January 2023) that it's more likely that AI will take away parts of jobs, especially in the first instance. Almost all knowledge workers – be it lawyers, teachers or even research analysts have part of their role that requires some 'grunt work' which can now easily (and more quickly) be done via AI generation as this report showcases.

Therefore, we expect the outcomes of a greater role of AI in work to be twofold: firstly, a lift to productivity from the speed at which outputs are able to be created. However, due to the nature of picking up productivity gains in knowledge work, this may not show up clearly in official output per hour statistics. Instead, it may be visible in the second impact: more leisure time. The number of hours worked each week in most developed economies stopped falling in the aftermath of the GFC and with these breakthroughs, we could see working weeks fall notably, lifting leisure time and spend.

Education

Generative AI leaves many questions around the social implication for educational systems. Indeed, Generative AI has the potential to drastically impact educational systems and the introductions of technology to classrooms has elicited both excitement and distress. Generative AI has received significant backlash regarding concerns about the negative impact on student learning. For example, New York City Department of Education blocked ChatGPT on all devices and networks, due to fears around encouraging quick answers, rather than critical thinking and problem solving skills, and fear over the safety and accuracy of content created. Concerns and confusion around what cheating and plagiarism means for ChatGPT are not clear – see the *Education* section in the full report.

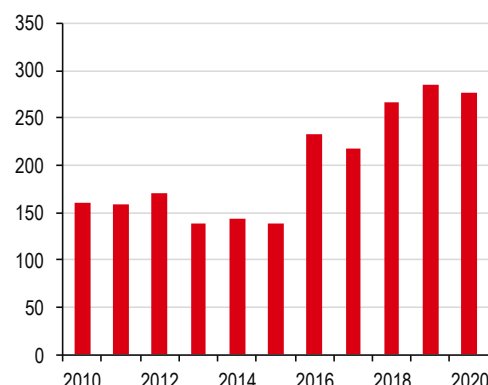
Despite criticism by education institutions, Generative AI has also been embraced by some. For example, teaching students how to work with Generative AI to understand biases, misinformation, and building critical thinking skills around how to make Generative AI text better.

While it's too early to tell the impact that Generative AI can have on the educational system, education and labour markets are intertwined. A change in demand for new skills will often be followed by educational systems changing in light of student learning preferences. Indeed, there has been significant growth in AI interest within higher education (Chart A5). Of computer science (CS) specialisations, AI and Machine Learning (AI/ML) are the most popular among PhD students; 1 in every 5 CS PhD students specialise in AI/ML (Chart A6). Between 2010 and 2020, the number of CS PhD graduates with AI/ML specialties grew by 72%¹⁰.

Generative AI has caused both excitement and distress in classrooms

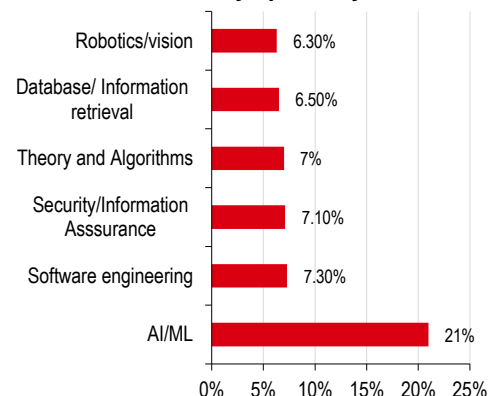
¹⁰ Artificial Index Report 2022, Stamford University Human-Centered Artificial Intelligence

A5. Growth in AI CS PhDs in the US



Source: CRA Taulbee Survey 2022, 2022 AI Index Report HSBC, HSBC

A6. New CS PhDs by speciality, 2021



Source: CRA Taulbee Survey

Legal implications

Generative AI raises many questions around **originality and property rights**. While Generative AI products are derivatives of existing text or images, the text or images created are not exactly like any previous content. Because businesses could soon be using Generative AI for the majority of written and image based content, Generative AI likely will open up a new frontier around content ownership and intellectual property protections.

Indeed, they are already being seen. Getty Images filed a lawsuit against Stability AI, an AI art generator, for infringement on its intellectual property rights. As Generative AI becomes more mainstream so will these concerns.

Privacy issues also exist for users of Generative AI. These can surface for both the AI maker in gathering information for training the bot, but also for data that a user enters into the model. For example, data entered into the bot could be used by the bot and AI maker for a variety of purposes, including to improve the AI service. Because Generative AI can be used by many services, including those which handle confidential information, such as a law practice or a financial service provider, using an application like this in business could raise privacy risks.

Regulations

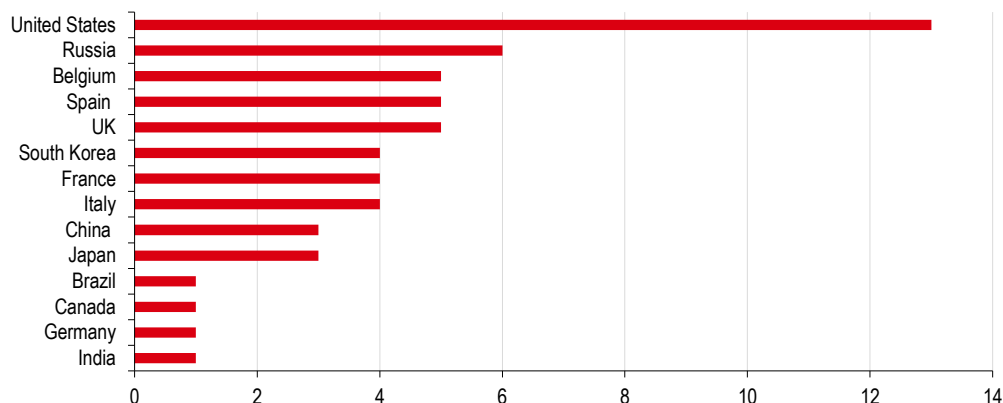
In efforts to address ethical and social concerns surrounding AI, intergovernmental, national and regional organisations have worked towards developing strategies and policies around AI governance. As the AI umbrella expands to include new applications for Generative AI – and new challenges, we think regulatory involvement will be unavoidable.

Indeed, the number of AI-related bills has increased significantly since 2016. Of the 25 countries analysed by Stanford University, 18 new laws containing mentions of AI were passed in 2021. The United States leads AI legislation across the group passing 13 AI-related bills in five years, followed by Russia, Belgium, Spain and the UK¹¹ (Chart A7). Many of the more recent bills are focused on social concerns. For example, in Belgium an Ethics Committee was created to submit advice on artificial intelligence tools used for digitalisation activities. In the US, the Identifying Outputs of Generative Adversarial Networks Act (IOGAN Act) directed the National Science Foundation to support research dedicated at identifying the outputs such as deepfakes and other comparable technologies.

¹¹ Artificial Index Report 2022, Stanford University Human-Centered Artificial Intelligence

These regulations show measures in place to regulate the industry. We believe responsible technology practices require effective and agile governance – both within an organisation and across the regulatory and public policy landscape.

A7. Number of AI related bills passed into law by country, from 2016 to 2021



Source: 2022 AI Index report, HSBC

Conclusions

Generative AI has many potential applications for business, and they are no doubt exciting. However, beyond its immediate uses, there are some serious social concerns that are worth understanding. While Generative AI is in the early stages, industry stakeholders are more and more acknowledging the implications that Generative AI can have on society, including labour markets, education, legal systems, regulation and communication platforms. We think corporates and investors should be proactive and diligent in monitoring their exposure to the industry, and identifying and mitigating social risk including those around fairness and biases, explainability, privacy, reliability and accuracy.

We think risks posed by Generative AI could ease as the market matures, regulatory initiatives increase, and knowledge of the technology rises.

This is a Free to View version of a report with the same title published on 23-Feb-23. Please contact your HSBC representative or email AskResearch@hsbc.com for more information.

HSBC Global Research clients can access the full note which also contains thematic and sector implications, covering the technology, media, education, healthcare & pharma, retail, consumer staples, leisure and banking sectors, along with an explanation on how Generative AI models work.

Disclosure appendix

The following analyst(s), who is(are) primarily responsible for this document, certifies(y) that the opinion(s), views or forecasts expressed herein accurately reflect their personal view(s) and that no part of their compensation was, is or will be directly or indirectly related to the specific recommendation(s) or views contained in this research report: Mark McDonald, Davey Jose, Camila Sarmiento, James Pomeroy, Nicolas Cote-Colisson, Charlene Liu, Ritchie Sun, CFA, Christina Chen, CFA, Frank Lee, Ali Naqvi, Joseph Thomas, Charles Pullan, Shubhangi Gupta, Ph.D., Yifeng Liu, Sezgi Oezener, CFA, Rajesh Kumar, Jeremy Fialko, CFA, Ankur Agarwal, CFA and Robin Down

This document has been issued by the Research Department of HSBC.

HSBC and its affiliates will from time to time sell to and buy from customers the securities/instruments, both equity and debt (including derivatives) of companies covered in HSBC Research on a principal or agency basis or act as a market maker or liquidity provider in the securities/instruments mentioned in this report.

Analysts, economists, and strategists are paid in part by reference to the profitability of HSBC which includes investment banking, sales & trading, and principal trading revenues.

Whether, or in what time frame, an update of this analysis will be published is not determined in advance.

For disclosures in respect of any company mentioned in this report, please see the most recently published report on that company available at www.hsbcnet.com/research.

Additional disclosures

- 1 This report is dated as at 23 February 2023.
- 2 All market data included in this report are dated as at close 21 February 2023, unless a different date and/or a specific time of day is indicated in the report.
- 3 HSBC has procedures in place to identify and manage any potential conflicts of interest that arise in connection with its Research business. HSBC's analysts and its other staff who are involved in the preparation and dissemination of Research operate and have a management reporting line independent of HSBC's Investment Banking business. Information Barrier procedures are in place between the Investment Banking, Principal Trading, and Research businesses to ensure that any confidential and/or price sensitive information is handled in an appropriate manner.
- 4 You are not permitted to use, for reference, any data in this document for the purpose of (i) determining the interest payable, or other sums due, under loan agreements or under other financial contracts or instruments, (ii) determining the price at which a financial instrument may be bought or sold or traded or redeemed, or the value of a financial instrument, and/or (iii) measuring the performance of a financial instrument or of an investment fund.

Disclaimer

Issuer of report
HSBC Bank plc

This document has been issued by HSBC Bank plc, which has based this document on information obtained from sources it believes to be reliable but which it has not independently verified. Neither HSBC Bank plc nor any member of its group companies ("HSBC") make any guarantee, representation or warranty nor accept any responsibility or liability as to the accuracy or completeness of this document and is not responsible for errors of transmission of factual or analytical data, nor is HSBC liable for damages arising out of any person's reliance on this information. The information and opinions contained within the report are based upon publicly available information at the time of publication, represent the present judgment of HSBC and are subject to change without notice.

This document is not and should not be construed as an offer to sell or solicitation of an offer to purchase or subscribe for any investment or other investment products mentioned in it and/or to participate in any trading strategy. It does not constitute a prospectus or other offering document. Information in this document is general and should not be construed as personal advice, given it has been prepared without taking account of the objectives, financial situation or needs of any particular investor. Accordingly, investors should, before acting on it, consider the appropriateness of the information, having regard to their objectives, financial situation and needs. If necessary, seek professional investment and tax advice.

The decision and responsibility on whether or not to purchase, subscribe or sell (as applicable) must be taken by the investor. In no event will any member of the HSBC group be liable to the recipient for any direct or indirect or any other damages of any kind arising from or in connection with reliance on any information and materials herein.

Past performance is not necessarily a guide to future performance. The value of any investment or income may go down as well as up and you may not get back the full amount invested. Where an investment is denominated in a currency other than the local currency of the recipient of the research report, changes in the exchange rates may have an adverse effect on the value, price or income of that investment. In case of investments for which there is no recognised market it may be difficult for investors to sell their investments or to obtain reliable information about its value or the extent of the risk to which it is exposed. Some of the statements contained in this document may be considered forward looking statements which provide current expectations or forecasts of future events. Such forward looking statements are not guarantees of future performance or events and involve risks and uncertainties. Actual results may differ materially from those described in such forward-looking statements as a result of various factors.

This document is for information purposes only and may not be redistributed or passed on, directly or indirectly, to any other person, in whole or in part, for any purpose. The distribution of this document in other jurisdictions may be restricted by law, and persons into whose possession this document comes should inform themselves about, and observe, any such restrictions. By accepting this report, you agree to be bound by the foregoing instructions. If this report is received by a customer of an affiliate of HSBC, its provision to the recipient is subject to the terms of business in place between the recipient and such affiliate. The document is intended to be distributed in its entirety. Unless governing law permits otherwise, you must contact a HSBC Group member in your home jurisdiction if you wish to use HSBC Group services in effecting a transaction in any investment mentioned in this document.

Certain investment products mentioned in this document may not be eligible for sale in some states or countries, and they may not be suitable for all types of investors. Investors should consult with their HSBC representative regarding the suitability of the investment products mentioned in this document.

HSBC and/or its officers, directors and employees may have positions in any securities in companies mentioned in this document. HSBC may act as market maker or may have assumed an underwriting commitment in the securities of companies discussed in this document (or in related investments), may sell or buy securities and may also perform or seek to perform investment banking or underwriting services for or relating to those companies and may also be represented on the supervisory board or any other committee of those companies.

From time to time research analysts conduct site visits of covered issuers. HSBC policies prohibit research analysts from accepting payment or reimbursement for travel expenses from the issuer for such visits.

HSBC Bank plc is registered in England No 14259, is authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority and is a member of the London Stock Exchange. (070905)

© Copyright 2023, HSBC Bank plc, ALL RIGHTS RESERVED. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the issuing entity name. MCI (P) 017/01/2023, MCI (P) 027/10/2022

[1208576]