



Global Counsel

Regulating Generative AI

OCTOBER 2023

Table of contents

05	Introduction
06	Executive summary
12	Research approach
15	Context
20	Familiarity
24	Attitudes
34	Use cases
38	Regulation
49	In focus: Healthcare
53	In focus: Financial services
57	Looking ahead to the next policymaking cycle

Introduction

Over the decade in which Global Counsel has analysed global technology policy, Artificial Intelligence (AI) has persisted as the regulatory challenge with the greatest resonance both to the technology sector, but also almost every other sector of the economy.

The legal, ethical and economic questions raised by AI are as pertinent for the companies developing AI models and systems, as they are for the deployers of AI in financial services, manufacturing and in the public sector.

In 2019, we hosted our flagship “Politics of AI” conference which looked at the global debate around cross-cutting AI governance principles. These were ultimately consolidated into the OECD’s AI Principles, which have acted as the benchmark for global AI governance ever since. In the years following, where there was divergence, this was typically over how to apply these principles rather than whether to apply them at all.

The launch of ChatGPT in November 2022 shattered this consensus. It is the closest that the technology sector has come to an “iPhone Moment” since the smartphone transformation of the late 2000s. However, the impact of ChatGPT goes well beyond the smartphone experience, in that it has not only disrupted the commercial, investor and technological landscapes, but also immediately transformed policy, regulatory and legislative debates.

This change coincides with political upheaval in Europe and North America in 2024, with elections in the US, the EU and (likely) the UK. These act naturally both as a break on current policy initiatives but also the starting point of a new political and policy cycle in which debates over generative AI will be played out.

The objective of this report is to understand the likely direction of travel for regulation during the next political cycle from 2024 till 2028/29. It is not intended as a contribution to live regulatory debates over initiatives such as the EU’s AI Act but looks ahead to the next US administration, European Commission and British government. The report specifically looks at generative AI, with clear implications for foundation models, large language models and frontier AI, but does not assess narrower AI models and systems. This is both for reasons of brevity and analytical consistency, but more importantly because it is generative AI and the reaction to it which will drive policymaking in the coming years.

The report combines original research with insights and perspectives from Global Counsel’s policy specialists in technology, financial services and healthcare, as well as teams on the ground in Washington DC, Brussels, Berlin and London.

The debate around the impact and regulation of AI will continue to move at pace. We hope this report gives our clients and wider network the grounding to engage in that debate.

CONAN D'ARCY

Senior Practice Director, Technology, Media and Telecoms



Executive summary

KEY FINDINGS

1

Public awareness of generative AI was high but deep understanding was largely limited to regulators and research institutions.

There were notable demographic differences - older people and women typically described themselves as less familiar - and confusion about how it differs from other forms of AI and other technologies. This pattern was the same across the US, Germany and the UK. Opinion formers varied in their understanding, with legislators typically having lower understanding than regulators. There were also basic definitional disagreements, such as how generative AI, LLMs and foundation models are distinct.

2

There was little evidence of serious anxiety about generative AI amongst the public while, in contrast, opinion formers voiced deep concerns.

The public expressed a wide range of concerns about generative AI, but these were rarely front-of-mind. In contrast, opinion formers were quick to identify a range of risks and the urgent need to address them. Some were shared between the two groups, such as joblessness, misinformation and data protection, while others, most notably the potential for bias in generative AI models and the environmental impact of generative AI, were far more pronounced with opinion formers. This is a reminder that policy elites may be primed to see their function as being ultimately to define the acceptable parameters for generative AI's evolution.

3


Use cases fundamentally shaped how generative AI was perceived.

Abstract benefits did not resonate strongly with either the public or opinion formers. While there was a consensus in favour of 'permit but regulate' across all use cases, attitudes varied significantly between different applications of generative AI in healthcare, financial services and other settings. The extent to which consumers and opinion formers are comfortable with different use cases and, by extension, the companies deploying them, was determined by a number of factors. These include the extent to which use cases are already subject to regulation, particularly sectoral frameworks in healthcare and financial services. This helps identify the areas in which the pressure for new regulation may first emerge.


4

Most opinion formers argued that policymakers should prioritise applying existing frameworks and addressing near term risks.


While opinion formers identified a handful of areas for legislative reform, such as copyright frameworks in the EU, the majority felt that the priority should be applying existing legal frameworks, such as the GDPR. Likewise, most were cautious about prioritising existential risks, with some sceptical about industry motives in this regard, and arguing that near term risks should be prioritised. Opinion formers pointed to geographical differences: the US - primarily driven by fear of Chinese competition - is prioritising competitiveness, the EU has a greater focus on upholding European values in AI development, most obviously through the AI Act, and the UK is currently positioned somewhere between the two. This suggests that regulatory approaches may be similarly divergent.



28% of German adults claim to have used ChatGPT



52% of UK adults think that humans are more likely to be biased than generative AI



75% of US adults would prefer some form of regulation of social media accounts being run entirely by AI



Implications for businesses

Businesses face a major challenge of educating stakeholders, most obviously legislators, about generative AI.


There is a demand for more information from politicians and government officials and a risk that - in the absence of such education - legislative frameworks will not align sufficiently with the technologies they are seeking to regulate. For example, the strong association of both the public and opinion formers with ChatGPT creates a risk that a 'chatbot regulation' emerges that fails to appreciate the diversity of generative AI.

Generative AI companies should avoid conflating a strong consumer brand with their political brand.

This is particularly important in light of more deeply-held concerns amongst opinion formers than the public, as well as a strong scepticism about self-regulation. Businesses will be expected to not only show that they are compliant with existing frameworks, such as data protection or consumer laws, but also that they are engaged on the development of new regulatory agendas. In Europe at least, consumer approval is unlikely to trump the instinct to control technology deployment through regulation.

There is, therefore, an onus on businesses to demonstrate the relevance of generative AI use cases to salient policy challenges faced by governments.

This includes a clear 'contribution' to economic growth and positive societal outcomes. For companies operating in already highly regulated sectors, like healthcare and financial services, the expectation is that regulatory enforcement of existing sectoral rules is more likely than new legislation. For companies operating outside of these sectors, there will be a high threshold for reassuring policymakers in the absence of a regulatory 'halo effect' and a need to communicate safety processes.





Businesses should anticipate a patchwork of regulatory initiatives according to policy area and geography, including:



Election interference

There is significant concern ahead of next year's elections and this will intensify as they draw closer. Should a significant controversy arise, this could prompt calls for regulatory interventions. Electoral events are also likely to place a spotlight on any perceived political bias in generative AI models.



Copyright

Copyright will be a prominent issue, both in legal test cases and potential legislative reform, particularly in the EU.



Licensing systems

Proposals for licensing systems have gained most traction when linked to the debate about existential risks with few indications to date that governments will move to a licensing model to regulate near term risks.



Open & closed source

The open source debate is unlikely to prompt governments or regulators to 'pick a winner' between open and closed models, though some governments may follow the French example of investing in the open source industry.



Data protection

Data protection has already emerged as a live regulatory issue for generative AI. Legislative reform is possible, particularly if the EU proposes a revision of the General Data Protection Regulation (GDPR), but, in the first instance, further enforcement cases in Europe against generative AI companies are likely.



Restrictions on Chinese AI

There is a strong prospect of further US restrictions targeting Chinese companies, building on the Biden Administration's semi-conductor export controls. As with 5G and semi-conductor policy, the US is likely to pressure Western allies to mirror its policies towards Chinese tech.




International regulator

A global regulator or comprehensive international governance system appears unlikely, though it is plausible that a testing body loosely based on the Intergovernmental Panel on Climate Change (IPCC) or the European Organization for Nuclear Research (CERN) could be established in the coming years.

Research approach

To explore familiarity and attitudes towards generative AI, Global Counsel conducted a programme of primary research, spanning three regions (UK, US and EU), and qualitative, quantitative and deliberative methodologies. All fieldwork was conducted between July and October 2023.



Public



CITIZENS' JURY

Deliberative workshop to explore familiarity with generative AI and understand how views change as participants are exposed to new information.

A three-hour face-to-face session held in the UK (London).

12 members of the public, recruited to reflect a spread of ages, gender, ethnicity, socio-economic group, area of residence (inc. urban, suburban, rural), engagement with digital technology and voting history.



SURVEY

Quantitative survey in the US, Germany (as largest EU member state) and UK to understand incidences of different views among the public and to enable robust comparison between and within countries.

Online survey up to 15 minutes.

Nationally representative sample of n=1,000 in each of Germany, US and UK (with a total sample of n=3000). Quotas applied for gender, age, region and other demographic variables.

Opinion Formers



INTERVIEWS

Qualitative interviews in Washington, DC, Brussels, Berlin and London to explore the views of experts in relation to generative AI, as well as likely regulatory implications.

In-depth interviews, lasting 30-45 minutes, conducted either face-to-face or via Zoom.

48 interviews completed with politicians, government and EU officials, regulators, academics, journalists and business leaders with relevant experience in the technology, healthcare and financial services sectors across the UK, US and EU.

In the research with the public (and after initial, unprompted views had been sought), the following definition of generative AI was used:



Generative AI is a type of artificial intelligence technology capable of generating highly realistic text, images, video or other content in response to prompts made by users. Notable examples of generative AI systems include ChatGPT, Bard and Midjourney. Whereas traditional AI systems are primarily used to analyse data and make predictions, generative AI goes a step further by creating entirely new content (based on its training data)."

In the interests of brevity, charts in this report occasionally feature abbreviated versions of the questions and answer options shown to respondents in the survey. Full data tables are available on request. Note that percentage figures based on public survey data may not always add to 100% due to rounding.

Context

THE PRE-2022 REGULATORY LANDSCAPE

Although AI had been subject to regulatory scrutiny for several years, tangible efforts were largely confined to supranational 'soft' regulation - primarily guidance, principles and voluntary standards¹.

OECD AI PRINCIPLES

The OECD's AI Principles encompass concepts of transparency, security and fairness. The Principles, were adopted by the OECD's 36 member states in May 2019.

WEF GUIDELINES FOR AI PROCUREMENT

The WEF's Guidelines, published in September 2019, set out a series of recommendations for governments when deploying AI systems in the public sector.

UNESCO RECOMMENDATION ON THE ETHICS OF AI

UNESCO adopted the Recommendation on the Ethics of AI in November 2021, setting a normative framework to guide the development of national-


Similarly at a national and regional levels, AI regulation prior to 2022 was predominantly principles-based. The EU was a notable exception, since it had published its AI White Paper in February 2020², which laid out policy options for a common EU regulatory framework, with a focus on high-risk applications.

Building upon nationwide and supranational regulatory frameworks were instances of more targeted sectoral regulation. For example, in financial services, under MIFID II, which came into effect in 2018, investment firms have been required to test algorithmic trading models and seek authorisation prior to deployment. Similarly, in healthcare, as early as 2019 the US Food and Drug Administration (FDA) published proposals to adapt existing guidance to allow for AI with adaptive machine learning capabilities³. Building on this with an action plan in 2021, these plans allowed the FDA to approve a colon cancer diagnostic with the ability to improve diagnostic accuracy over time.



AUSTRALIAN AI ETHICS FRAMEWORK

Australia's 2019 AI Ethics Framework established eight principles around wellbeing, human-centred values, fairness, privacy and security, reliability and safety, transparency and explainability, contestability, and accountability.



JAPANESE SOCIAL PRINCIPLES OF HUMAN-CENTRIC AI

The Japanese Social Principles of Human-Centric AI were published in 2019, proposing seven guiding principles demanding AI development be human-centric, and consider education, privacy, security, competition, transparency, and innovation.⁴



AI RISK MANAGEMENT FRAMEWORK

The National Institute of Standards and Technology (NIST) published its framework in January 2023, offering organisations a voluntary resource to manage risks and promote trustworthiness.



AI WHITE PAPER

The AI White Paper, published in March 2023, establishes non-statutory principles on safety, transparency, fairness, accountability and contestability. The government has indicated it could introduce a statutory duty for regulators to have regard to the principles.



AI AND DATA ACT

Originally tabled in June 2022, when implemented the Act would ensure that high-impact AI systems meet standards on safety and human rights, focusing on prevention of harm to individuals and mitigation of systemic bias.



AI ACT

Proposed in 2021 and subject to ongoing legislative processes at the time of writing, the AI Act looks to classify and regulate AI models based on three categories of risk of harm. New amendments have been tabled which could specifically regulate generative AI.



INTERIM MEASURES: MANAGEMENT OF GENERATIVE AI SERVICES

The Interim Measures entered into force in August 2023 and require developers of AI models accessible to the public to adhere to content moderation and labelling obligations, user safety requirements and strict reporting duties.⁷

CHATGPT AND THE GENERATIVE AI EXPLOSION

Prior to November 2022, the public’s first-hand experience of AI was largely limited to chatbots, often in customer services. Where AI was utilised for chatbots (as opposed to pre-programmed bots), research indicated frustration amongst consumers unable to reach human respondents⁵. Elsewhere, publicly accessible experimental chatbots saw notable failures. Microsoft’s Tay AI chatbot, released in March 2016, was suspended just hours following its launch after users ‘trained’ it to make offensive statements. Similarly in 2021 Scatter Lab’s ‘Luda Lee’ AI chatbot became embroiled in controversy because of its use of offensive language.

The perception of the capabilities of generative AI, and AI more broadly, changed drastically with the launch of OpenAI’s ChatGPT in November 2022. Within just two months of its launch, it was estimated that ChatGPT had reached the 100m user milestone, becoming the fastest platform to reach this landmark⁶.

Following ChatGPT’s launch, other companies were quick to accelerate deployment of their own generative AI models which had yet to be released publicly. In early 2023, Google and Meta launched their

Bard and LLaMA services respectively, while Microsoft integrated its ChatGPT-powered Bing Chat into its search functions.


The explosion in the accessibility of generative AI models inevitably attracted the attention of policymakers and regulators. What has particularly focused the minds of government and regulatory authorities has been that, unlike with other recent rounds of innovation, notable figures in the tech sector have called for regulation from the outset. One group, including Elon Musk and Steve Wozniak, signed a petition in March 2023 calling for a six-month pause on AI development. Similar concerns were raised by Geoffrey Hinton, a pioneer of AI research at Google, who stressed that governments needed to put “a lot of thought into how to stop it [advanced AI systems] going rogue.”

THE POST-2022 REGULATORY LANDSCAPE

Amidst growing recognition of the need for regulation, authorities worldwide have redoubled efforts. Where regulatory efforts had been initiated prior to 2022 but not yet operationalised, these have taken on a renewed urgency and, in many cases, underwent a process of strengthening. For instance, Article 28b was added by the European Parliament to its draft

of the EU’s AI Act in June 2023, introducing new provisions targeted at foundation and generative AI models.

However, despite a growing recognition of possible dangers being shared by authorities internationally, national regulatory initiatives are yet to converge and there remains the prospect of a fragmented regulatory outlook for businesses to navigate.





AI REGULATORY TIMELINE

- INTERNATIONAL REGULATORY DEVELOPMENT
- NATIONAL OR REGIONAL REGULATORY DEVELOPMENT
- INDUSTRY DEVELOPMENT

2019

- February: Japanese Social Principles of Human-Centric AI published
- May: OECD AI Principles published
- September: WEF Guidelines for AI Procurement White Paper published
- November: Australian AI Ethics Framework published

2020

- February: European Commission White Paper on AI published

2021

- April: EU AI Act proposed
- November: UNESCO Recommendation on the Ethics of AI

2022

- June: Canadian AI and Data Act published
- October: US Blueprint for an AI Bill of Rights published
- November: ChatGPT publicly launched

2023

- January: US NIST AI Risk Management Framework published
- March: UK AI White Paper published
- March: Google Bard launched
- March: Tech leaders sign AI pause petition
- May: G7 Hiroshima Leaders' Communiqué
- July: Chinese Interim Measures for the Management of Generative Artificial Intelligence Services enacted


Familiarity

SUMMARY

- Public awareness of generative AI was high for a novel technology, with usage of different tools (particularly ChatGPT) apparently widespread.
- Nevertheless, there were notable demographic differences - older people and women typically felt less familiar - and potential for confusion about what generative AI is and how it differs from other forms of AI and other technologies.
- Opinion formers varied in their understanding and there were basic definitional disagreements, such as how generative AI, LLMs and foundation models are distinct.

FIG. 1: FAMILIARITY WITH GENERATIVE AI (UK)

% OF UK ADULTS CLAIMING TO BE FAMILIAR






- NEVER HEARD OF IT
- HEARD OF IT BUT DON'T KNOW ANYTHING ABOUT IT
- KNOW A LITTLE ABOUT IT
- KNOW A FAIR AMOUNT ABOUT IT
- KNOW A LOT ABOUT IT

Public awareness of generative AI was, prima facie, high for a novel technology: around 9 in 10 in the UK, Germany and the US claimed to have heard of it. Awareness of generative AI was higher than for other technologies, including the “metaverse”, cryptocurrency and non-fungible tokens (NFTs). Likewise, self-reported usage of generative AI was strikingly high given that ChatGPT launched only nine months before the fieldwork was conducted.

Nevertheless, there were notable demographic differences, with older adults feeling significantly less familiar than younger adults and women feeling less familiar than men. Indeed, many - particular older adults - were yet to experience generative AI tools. While 57% of 18-24s in Germany claimed to have used ChatGPT before, this was only 14% of those aged 55+. There were also some clear geographic differences, with self-reported usage notably lower in the UK than in Germany and the US. More broadly, generative AI appeared to be strongly associated with ChatGPT with self-reported usage of other existing tools, including Dall-E, Google Bard and Bing Chat, much less widespread.

FIG. 2: USAGE OF GENERATIVE AI SYSTEMS

% OF ADULTS IN EACH MARKET CLAIMING TO HAVE USED EACH TECHNOLOGY AT LEAST ONCE

			
ChatGPT (by OpenAI)	23%	28%	27%
Bing Chat	9%	10%	20%
Google Bard	7%	12%	18%
Dall-E (by OpenAI)	7%	9%	12%
Microsoft CoPilot	4%	9%	13%
LLaMa 2 (by Meta)	4%	5%	10%
Claude (by Anthropic)	4%	6%	9%
Stable Diffusion	4%	5%	9%
Wysa	4%	5%	8%
LuminaSynth ProgeniAI	3%	4%	9%

There also appeared to be notable potential for public confusion and misunderstanding. Participants in the Citizens’ Jury often conflated generative AI with other forms of AI, and there was also a tendency to confuse generative AI with more established technologies (e.g. traditional search engines or, as demonstrated in Fig. 4, ‘robot’). Several participants expressed surprise on experiencing a live demonstration of ChatGPT, despite having claimed to have some understanding of generative AI. That 9% of German adults claimed to have used Microsoft CoPilot (not publicly available at the time of fieldwork) and a similar proportion of US adults claimed to have used LuminaSynth Progenai (an entirely fictitious tool) provides further evidence of the scope for misunderstanding.

My only experience with it... is internet searching with Bing. Which to me works well, is very fast and seems quite accurate, and gives you a lot of places to search.”

UK PUBLIC

What comes to mind is video games... that you can prompt and decide who you are in your game... There’s also this chatbot I’ve been hearing about that can write your dissertation for you.”

UK PUBLIC

Opinion former interviews also revealed marked differences in the levels of understanding. Within this group, parliamentarians were on one end of the spectrum - with limited knowledge - and technical AI expert bodies on the other. With departmental officials and regulators somewhere in between, there is a major focus on recruiting more technical experts and upskilling on AI.

Even developers and the research community themselves do not understand how decisions are made – it’s an emerging field of research. Policymakers are even further from understanding what is going on.”

INDUSTRY STAKEHOLDER

Generative AI is not artificial general intelligence, or AGI. AGI is an AI system that is smarter than humans. We’re a lot further from AGI than most people think.”

GOVERNMENT AGENCY

On definitions, ambiguity persists over how generative AI, LLMs and foundation models interrelate and are distinct; and how to define ‘frontier AI models’. This has been evident in the final stages of the negotiations over the EU’s AI Act where MEPs have looked to introduce new provisions which distinguish between foundation models, general purpose AI systems and generative AI. Indeed, there was a concern from opinion formers in Brussels that the shift in focus following the launch of ChatGPT risked diverting the EU’s AI Act negotiations away from what they perceived to be more significant risks around non-generative forms of AI, such as those processing biometric data for facial recognition.

Likewise, there was no consensus on defining frontier AI with contrasting definitions given by industry, regulators and government stakeholders. Agreeing a common definition is one the objectives for the AI Safety Summit in London.

Implications for businesses

The rapid success of ChatGPT means that the starting point for the debate is more developed than has been the case for other nascent technologies. However, given misunderstandings about the technology and definitions, businesses should not assume an advanced level of understanding, especially amongst consumers.

The potential for misunderstanding, combined with a sense from some policymakers of the need to act sooner rather than later, means that businesses need to educate policymakers urgently or risk new policy and legislative agenda which do not accurately reflect the technology nor have the desired policy outcomes.

The strong association among both the public and opinion formers between generative AI and ChatGPT creates a risk that a reductive 'chatbot regulation' emerges, that fails to appreciate the diversity of other forms of the technology.


Attitudes

SUMMARY

- The public was largely reserving judgement about generative AI and yet to form strong views, with little evidence of 'moral panic'.
- Whilst recognising a wide range of both benefits and risks after deliberation and debate, these were not well-established and rarely front-of-mind for the public. Indeed, the public did not seem to share well established concerns about generative AI, such as bias.
- By contrast, opinion formers (and policymakers in particular) had significant concerns about generative AI, suggesting a disconnect with public views. Their concerns included mass joblessness, reinforcing bias and discrimination, cybersecurity, liability regimes and environmental impacts.

FIG. 3: FAVOURABILITY TOWARD GENERATIVE AI

% OF ADULTS IN EACH MARKET FEELING FAVOURABLE OR UNFAVOURABLE



With public understanding of generative AI still emerging, it is unsurprising that many were yet to form very strong views on the topic. In the UK, for example, only 14% of the public felt either very favourable or very unfavourable, with the remaining population having had more moderate attitudes or yet to form a clear view.

The German public was, on balance, the most favourably disposed towards generative AI, slightly ahead of the US, with the UK the least enthused. As reflected by the top-of-mind associations in Fig. 4, many have mixed feelings about generative AI: seen as both 'scary' and 'dangerous', but also 'helpful' and 'amazing'.

FIG. 4: TOP-OF-MIND ASSOCIATIONS WITH GENERATIVE AI (US)

ONE WORD THAT COMES TO MIND ABOUT GENERATIVE AI (25 MOST COMMON WORDS FROM US RESPONDENTS)




FIG. 5: FAVOURABILITY TOWARD GENERATIVE AI

% OF ADULTS IN EACH MARKET FEELING FAVOURABLE OR UNFAVOURABLE

	AGE		GENDER		VOTING HISTORY			
	YOUNGER ADULTS (18-34)	OLDER ADULTS (55+)	MEN	WOMEN	LAB ('19)	LIB DEM ('19)	CON ('19)	
	+19%	-7%	+8%	-7%	+6%	+14%	-1%	
	+24%	-6%	+19%	+2%	BIDEN ('20)		TRUMP ('20)	
					+24%		+1%	
	+31%	+7%	+27%	+9%	SPD ('21)	CDU ('21)	AFD ('21)	GRUNE ('21)
					+28%	+39%	+2%	+10%

When prompted, most members of the public recognised both potential opportunities and risks created by generative AI: only 17% of German adults, for example, did not consider there to be any benefits, while only 9% did not have any concerns. As outlined in the next chapter, attitudes to the technology also varied significantly depending on the specific use case.

As well as differences between countries, there were some notable differences within them. Groups typically more familiar with generative AI also tended to feel more favourable, with younger adults more positive than older adults and men more positive than women. There were also some political differences, though these did not conform neatly to a traditional left/right split. In the UK, there was only a minor difference in attitudes between Labour and Conservative voters.

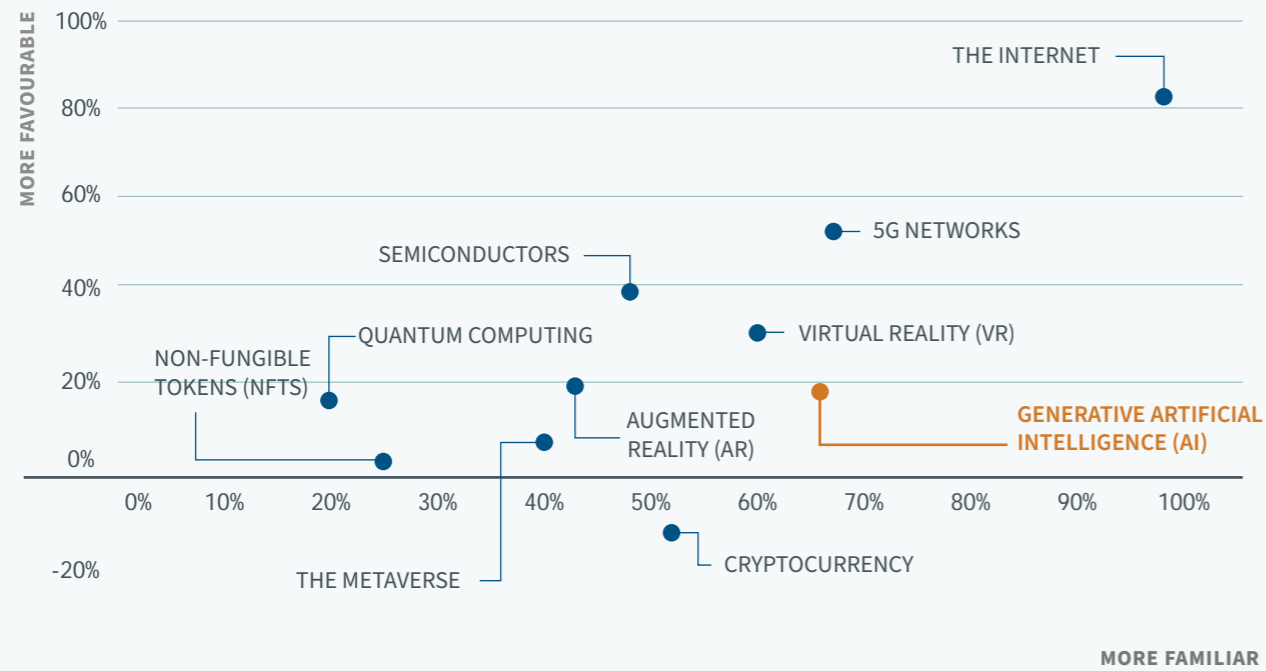
In Germany and the US, there appears to be a correlation between support for traditional, 'mainstream' parties and greater favourability towards generative AI. In contrast, supporters of 'populist' parties are more negative towards generative AI. This sentiment is likely indicative of a broader scepticism within these voting segments and the rhetoric of their political leaders towards large businesses.

When comparing generative AI to other technologies, it had a reasonably middling reputation. The UK, US and German public typically felt more favourably about a range of technologies including 5G, virtual reality and even quantum computing than generative AI - but felt more positive about generative AI than they did about the metaverse, NFTs and cryptocurrency. The geographic differences in attitude toward generative AI are largely consistent with broader trends, with UK adults typically feeling less positive about selected technologies than their US and German counterparts. This is consistent with Global Counsel's research on attitudes towards the metaverse⁸ and other third party sources⁹, which find a sceptical attitude towards digital technology and large businesses in the UK.

FIG. 6: FAVOURABILITY AND FAMILIARITY TOWARD DIFFERENT TECHNOLOGIES (DE)

% OF GERMAN ADULTS FEELING FAVOURABLE MINUS % FEELING UNFAVOURABLE /

% OF GERMAN ADULTS CLAIMING TO KNOW AT LEAST A LITTLE



The evidence suggests that there is scope for public opinion on generative AI to shift as familiarity and usage increase. The general trend is for favourability to increase in line with familiarity. It is not, however, necessarily the case that attitudes toward generative AI will become warmer as the public learns more. Cryptocurrencies, for example, have a poor reputation (especially outside the US) despite being relatively familiar to consumers. Indeed, while around a third of respondents in each market changed their views about generative AI between the start and end of the survey, the proportion that became more favourable was roughly the same as the proportion that became less favourable.



I'm actually going to use the word 'excited.' The reason I'm using the term...is that, in the right hands, I think it could be incredibly good for the world. In the wrong hands, it could be disastrous."

UK PUBLIC

In contrast, the attitude of most opinion formers was less balanced and strikingly more negative than the public, with a number of policymakers arguing that it was "their job" to be pessimistic about the technology given the scale of potential risks. Opinion formers repeatedly brought up the experience of social media as an example of why they were wary about generative AI, reinforced by a flurry of public interventions by prominent industry stakeholders warning about the existential risk of frontier AI.



The feeling that we should pause on the development of AI was in part because people were so worried that we missed the boat with regulating social media."

THINK TANK

BENEFITS

The vast majority of the public in Germany (83%), the US (75%) and the UK (73%) recognise at least one benefit to generative AI. On prompting, the most resonant benefits are: saving consumers' time by automating certain tasks; increasing productivity and efficiency of businesses and public services; making certain products and services more accessible; and (particularly for German adults) resolving customer support queries.

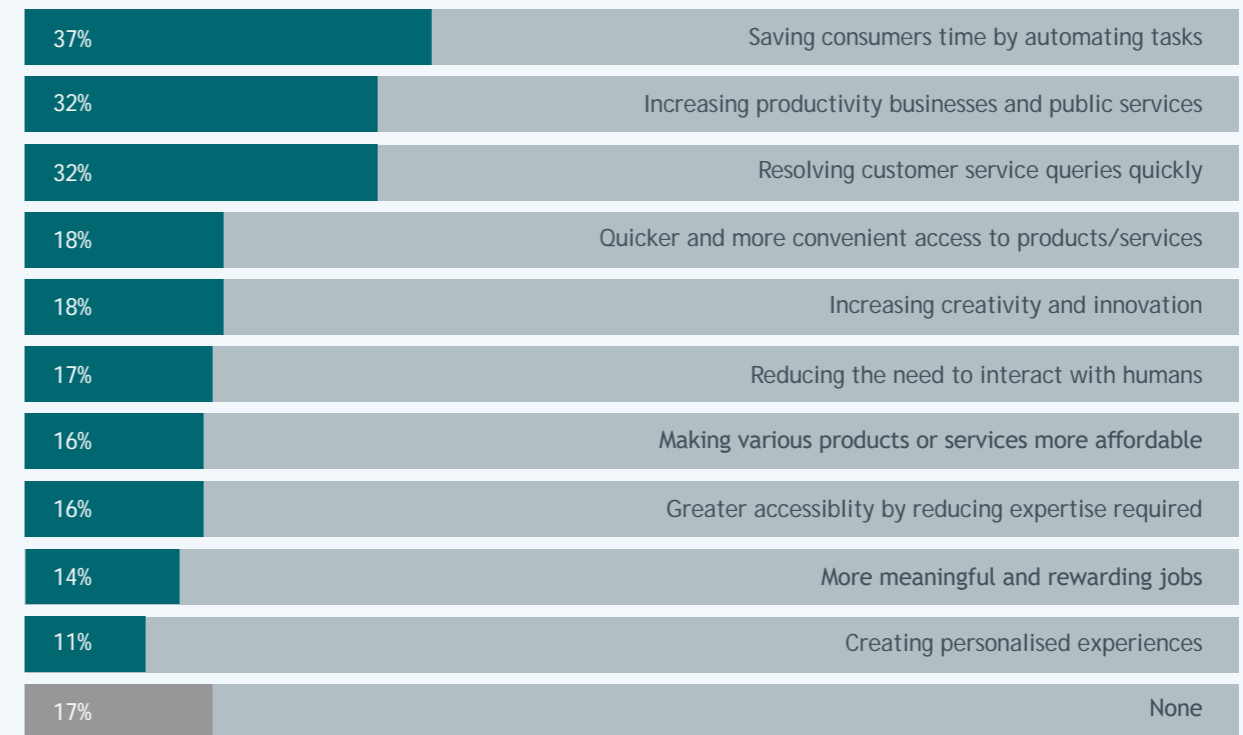


The main benefit is saving time – it is speedy. Time is money...Most of us use the National Health Service (NHS), which is funded by public money and it's not infinite."

UK PUBLIC

FIG. 7: MOST IMPORTANT BENEFITS OF GENERATIVE AI (DE)

% OF GERMAN ADULTS SELECTING EACH BENEFIT AS ONE OF THREE MOST IMPORTANT



Nevertheless, the potential benefits of generative AI are not always front-of-mind. When asked to suggest what the potential benefits might be, participants in the Citizens' Jury struggled to think of many spontaneously beyond saving consumers time, especially if they had not used tools like ChatGPT before. There was recognition of other benefits, but typically only after prompting and more in-depth demonstrations of the technology.



Generative AI is an incredibly helpful source of information – it can offer a personalised experience, service, and is a great source of knowledge."

UK PUBLIC

Opinion formers similarly needed prompting on benefits and were quick to revert back to concerns about generative AI unless explicitly asked not to. Like the public, saving time and increasing efficiency were regularly cited as key advantages to deploying generative AI, with some arguing that this could free up workers for more rewarding tasks. One unique benefit that did not feature in the Citizens' Jury was the potential of generative AI for democratisation of information. Interviewees asserted that generative AI would bring quick and easy access to knowledge to a wider variety of people, although with a caveat that this information is not yet always verified and accurate.

"We could have a world that is free from drudgery if this can be done in a responsible way."

GOVERNMENT AGENCY

"Ultimately, I don't think it's down to the public to decide whether these technologies are safe for them. It's our job to safeguard the public."

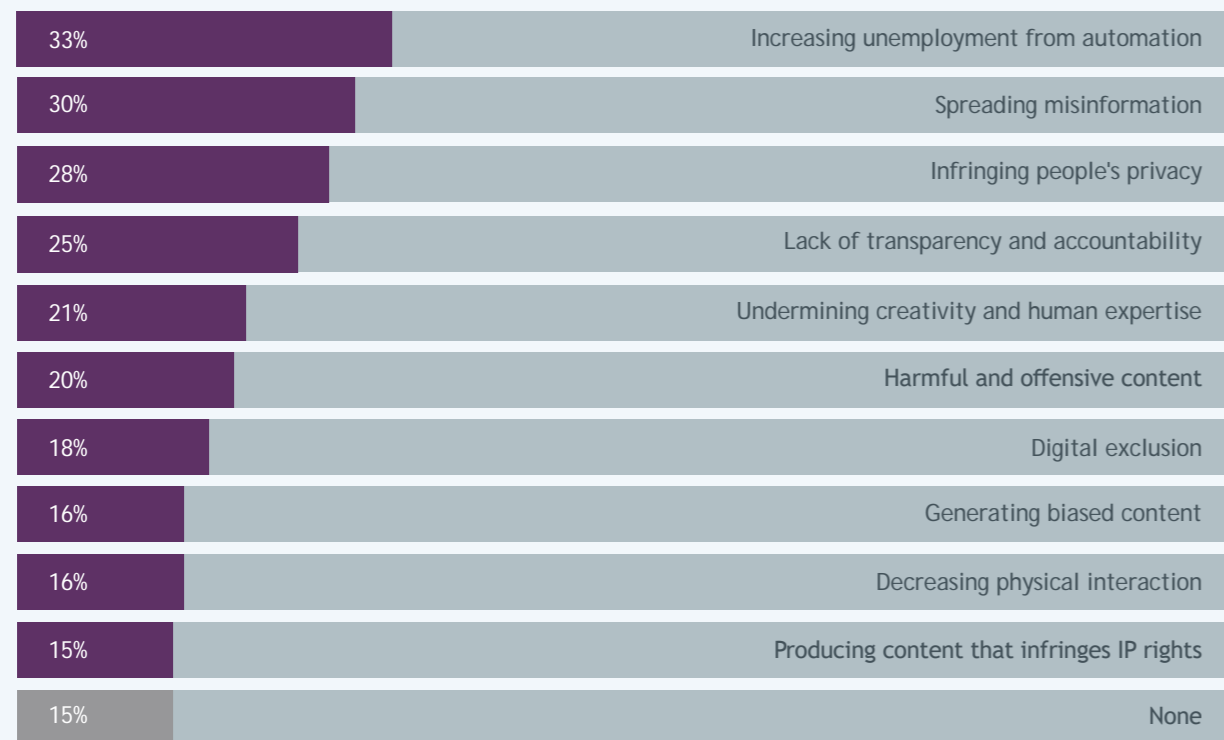
CLINICAL ACADEMIC

CONCERNS

Despite recognition of potential benefits, public anxiety about generative AI was evident with around 9 in 10 members of the public in Germany, the US and the UK having at least one concern about the technology. The most top-of-mind and widely-held concerns were potential joblessness, as well concerns reflecting well-established debates about technology (and digital technology in particular) and are not necessarily unique to generative AI: the spread of misinformation and infringement of privacy. This demonstrates the extent to which the public will view generative AI through the prism of well-established technologies they use on a day to day basis, such as social media.

FIG. 8: MOST CONCERNING RISKS AROUND GENERATIVE AI (US)

% OF US ADULTS SELECTING EACH RISK AS ONE OF THREE MOST CONCERNING




By contrast, more novel risks unique to generative AI around undermining intellectual property (IP) or reinforcing biased content were generally less concerning to consumers. In contrast to the longstanding policy consensus on AI and the concerns raised by opinion formers, the public considered humans to be more likely to be biased when providing information than generative AI

technologies. More broadly, there appeared to be a relatively high level of public trust in the accuracy and partiality of generative AI technology, possibly as a result of limited understanding of how it works and the underlying data it is reliant on, but also suggesting a significant gulf in opinion between the public and policymakers.

FIG. 9: PERCEPTIONS OF BIAS IN GENERATIVE AI TECHNOLOGY (UK)

% OF UK ADULTS CLAIMING EACH STATEMENT COMES CLOSER TO THEIR VIEW



There were a wide range of different concerns, each prioritised by different audiences. For example, older adults (55+) were more likely to be concerned about the spread of misinformation, while younger adults (18-34) were more likely to have concerns about increasing unemployment, possibly reflecting greater personal interest in the job market. German adults were significantly more likely to identify the exclusion of older people and the digitally excluded as a concern than their US and UK counterparts - which may be due to prominence of English language in the first wave of generative AI systems.

Opinion formers shared public concerns on misinformation and risk to the labour market. On the former, there was particular concern about the use of generative AI during the upcoming elections in the US, UK and EU in 2024, as well as the potential of generative AI creating illegal content, ranging from child pornography, for example, to phishing content. However, one regulator working on online safety commented that there was limited immediate evidence of harm to date. On the issue of cybersecurity, multiple interviewees raised fears about compromised security, both in terms of online cyberattacks but also generative AI creating code to detect vulnerabilities in critical national infrastructure, such as telecoms.

"My main concern is the ability to make lies look more and more plausible. I've noticed since Trump came into office it is like you can't believe what people say."

UK PUBLIC

“
Photoshop was a dangerous tool when it first came out too, being used to make all models skinnier than they already were. Generative AI will go through the same phases...”

REGULATOR

“
There’s a difference between showing it’s possible to reach a harmful output and it happening at scale. We’re seeing more of the first camp at the moment.”

REGULATOR

On the labour market, opinion formers reflected well-trodden debates about generative AI and AI more broadly leading to mass joblessness, aligning with concerns reflected in the public survey. There was a fear that generative AI might lead to lower paid, less stable and less fulfilling jobs, though this was partly qualified by those, as noted above, who argued that generative AI tools could improve the quality of work for many. There was no consensus on which level of workers would be hit hardest by the introduction of generative AI into work. Although much of the media emphasis to date has been on generative AI replacing lower skilled jobs, there is also emerging evidence to suggest that there will be significant transformation elsewhere in the hierarchy¹⁰. Some opinion formers argued that there would likely be a ‘hollowing out’ of medium skilled workers like accountants, for example, who would not as easily be retrained as the lower skilled workers.

“
The higher up the workforce chain you go, the more complex the replaced tasks are, and the more margin for unpredictable behaviour and potential harms.”

RESEARCH INSTITUTION

Several stakeholders argued that the political implications will be significant, even if the spectre of mass joblessness is not realised and is uneven across different industries and regions. Some opinion formers argued that a policy response will be required to accommodate those who have been affected, with significant medium to long term fiscal implications, particularly in welfare and benefits systems.

“
We’ll have wholesale industries that will be permanently scarred by this – it will require tax and benefits systems reforms.”

THINK TANK

A sense that generative AI might prompt greater inequality was not restricted to the labour market. Some raised alarm about the potential for two tiers of service provision to arise, such as in the healthcare sector. They argued this could see people on lower incomes interacting with generative AI systems, while high income households continue to access human-led services. In a recent example, Singaporean courts are piloting offering of generative AI services for certain litigants who are not otherwise legally represented¹¹.

Opinion formers were also aligned with the public on their concern about transparency and lack of accountability, and whether developers were being sufficiently clear with companies deploying their technology about the risks and associated liability. Health stakeholders were clear that liability for safety would be the primary factor driving policy in light of a lack of clarity about whether accountability should be placed on a medical provider or organisation, a health practitioner, the technology developer, or someone else entirely.

A final concern raised by opinion formers but not prominent with the public was the environmental impact of generative AI. Stakeholders pointed to the energy resources required for training and operating AI models, which they argued are “much, much higher than prior technologies”, creating significant demands on data centres and cloud computing capacity, a concern seemingly supported by recent academic studies¹². With uptake of generative AI and advanced AI systems set to grow in the coming years, this concern is likely to increase and build on existing debates in the EU, where the environmental impact of blockchain technologies was a major point of debate in the EU’s regulation of crypto-assets.

TRUST IN INSTITUTIONS




In line with the perceived significance and impact of generative AI technology among the public, the governance of generative AI and who will do doing the governing were considered critical. Participants in the Citizens’ Jury talked of the importance of who is ‘in control’ of the technology, as well as expressing concern about it ‘falling into the wrong hands’.

The survey suggests that trust in different institutions in relation to generative AI is relatively limited. No institutions are trusted at least moderately by a clear majority of the population across the three markets, with many trusted only by a small proportion of the public.

“
We need an independent body that works with the government – a body set up who know what they’re talking about.”

UK PUBLIC

FIG. 10: TRUST IN DIFFERENT INSTITUTIONS IN RELATION TO GENERATIVE AI
% OF ADULTS IN EACH MARKET TRUSTING EACH INSTITUTION AT LEAST MODERATELY

			
Friends and family	43%	65%	49%
A relevant charity, consumer group or campaigner	39%	41%	34%
Universities, academics and other researchers	33%	51%	38%
An independent regulator	29%	39%	26%
A relevant government department	27%	34%	30%
Public bodies using generative AI technology	22%	43%	30%
Businesses using generative AI technology	20%	32%	30%
Businesses developing generative AI technology	20%	38%	34%
The government	16%	27%	21%
Other consumers using generative AI technology	15%	29%	27%

Trust in institutions was generally highest in Germany, while it was lowest in the UK, broadly reflecting overall favourability toward generative AI and suggesting a clear link between trust in institutions and trust in the technology. While the UK and German public typically trusted regulators more than businesses developing or using generative AI technology, implying a lack of confidence in industry self-regulation, the reverse was true in the US where businesses were more widely trusted than regulators. Friends and family being the most trusted sources of information in all three markets demonstrates how important word-of-mouth is likely to be in shaping attitudes to generative AI.

Although some felt that self-regulation had an important role to play, the vast majority of opinion formers treated businesses' voluntary commitments with a high degree of caution. Several UK stakeholders were sceptical about the impact of initiatives like the

White House Commitments, while others in the US argued that the absence of regulation meant that the profit motive would be given precedence over safety controls. Opinion formers in Brussels in particular raised concerns that industry-led initiatives could prove to be attempts to 'close the ecosystem' by setting standards that smaller competitors could not match.



Companies' voluntary commitments are... a distraction technique...open letters from industry executives are well-intentioned, but also an opportunity to show off how amazing they are, and suggest they've built an almost God-like intelligence".

GOVERNMENT OFFICIAL

Use cases

Implications for businesses

Abstract, generalised benefits struggle to resonate in the absence of specific use cases. Businesses should lead with tangible applications of generative AI in order to build support and inform audiences about the possibilities of the technology.

Given the higher levels of concerns about generative AI amongst opinion formers than the public, businesses need to be careful not to over-interpret the impact of a strong consumer brand with political and regulatory audiences.

There is a strong degree of scepticism about self-regulation, drawing on the experience of previous waves of tech regulation. The onus is therefore on businesses to show that they are not only compliant with existing frameworks, such as data protection or consumer laws, but also that they are engaged on the development of new regulatory agendas.

Given that policymakers' experiences of prior technologies, such as social media, are influencing perceptions of generative AI, businesses will need to demonstrate how they are different from predecessor technologies – most obviously on safety by highlighting novel techniques, such as the role of post training and practices such as red teaming.

SUMMARY

- Use cases fundamentally shape how generative AI is perceived by the public and opinion formers. Attitudes to different applications of generative AI in healthcare, financial services and other settings varied significantly.
- The extent to which consumers and opinion formers are comfortable with different use cases and, by extension, the companies deploying them, appeared to be determined by a number of factors. This includes the extent to which use cases are already subject to regulation, particularly sectoral frameworks in healthcare and financial services.
- Use cases of greatest concern to the public included AI-run social media accounts, reflecting wider public concern about social media.

Across all use cases surveyed, there was a consensus that some form of regulation of generative AI applications is necessary; of the fifteen use cases tested, a large majority in each market (at least 2 in 3) favoured some form of restriction, ranging from minor restrictions to outright bans. However, those calling for outright prohibitions were in the minority, with no more than 1 in 3 preferring to ban any use case completely.



"I'm really excited about the potential of generative AI in R&D... but use cases become controversial quite quickly if anything is patient facing."

POLITICAL ADVISER

Beneath this basic consensus lay a number of important nuances and differences. Those who were more familiar with and favourable towards generative AI in general (particularly younger adults and men) also tended to favour less significant restrictions. German citizens were less likely than their British and American counterparts to favour heavy restrictions.

There was also some variation in attitudes between different use cases, with the balance of public opinion being in favour of more significant restrictions for some use cases compared to others. For example, whereas 29% of US adults were in favour of banning entirely the use of generative AI to run social media accounts, only 10% supported a ban for use in fraud detection and prevention. The extent to which the

public supported restrictions on different use cases was influenced by a complex mixture of different factors and this was also reflected in responses from opinion formers.

HOW NOVEL IT WAS PERCEIVED TO BE

Use cases that appeared similar to existing tools used by consumers tended to provide reassurance, possibly in part due to an inability to distinguish between generative AI and existing technologies. For example, the use of generative AI to provide patients with automatic diagnoses of health conditions was considered by some to be similar to patients using search engines ("Dr Google") and, as a result, less concerning. Similarly, little distinction was drawn between existing chatbots and the use of generative AI to provide highly responsive, tailored customer support services.

Opinion former interviews revealed a similar trend, with one policymaker offering high levels of support for using ChatGPT "to write letters in another language and in a particular style or document", in a way that replicates the existing functions of online translation services.

THE EXTENT OF HUMAN CONTROL

A critical component of confidence was whether or not there was some degree of human oversight. Use cases where the technology was making ultimate decisions or providing a service, such as making investment decisions, without human intervention were more concerning than use cases where the technology was providing "merely" guidance or information, such

as financial advice or health information. In those circumstances, there was concern about people with vulnerabilities being exploited or negatively affected without an additional layer of expert human oversight, such as doctors.



"My biggest piece of advice for businesses in this space is that you have to focus on using your tools to solve real world issues that people are experiencing right now."

POLITICAL ADVISER

A number of opinion formers stressed that generative AI should be seen as a complementary tool for humans, rather than an autonomous end to end process. One example given was designing complex engineering artifacts - engineers often have to produce many design alternatives and compare them to the chosen model, which is a long and arduous process that could be significantly streamlined by generative AI.

WHETHER IT ADDRESSED (OR EXACERBATED) A SALIENT ISSUE

Use cases that addressed a prominent problem - for example, applications that help with workforce challenges and, therefore, pressures on health systems - were more likely to be supported by the public. Opinion formers also spoke strongly in favour of the potential for generative AI to alleviate certain regulatory concerns, for example, creating synthetic data to train online safety technologies like age estimation tools.



"There has to be some kind of restriction - social media is full of kids."

UK PUBLIC

By contrast, use cases that fed into existing concerns - for example, interacting with children and other vulnerable users on social media - were more likely to be met with instinctive opposition. Opinion formers raised concerns specifically about misinformation online and speculated that generative AI could make these problems more acute. Use cases that were not obviously related to any particularly salient or relevant issues, such as developing an advertising campaign

for brands, tended to be more likely to be met with indifference and, therefore, tolerance.



"A lot of healthcare staff take up a lot of their time with admin tasks relating to their patient care - it would free up time for them to do other things."

UK PUBLIC

HOW SIGNIFICANT ANY IMPACT WAS LIKELY TO BE

Use cases which were assumed to have a significant impact were more likely to lead to demands for greater restrictions than those whose impact was assumed to be minor. 73% of UK adults believe that generative AI should only be allowed to make decisions without human oversight where the impact of the outcome is likely to be minor.

Opinion formers expressed similar caution about generative AI deployment without human oversight. For example, several were excited about the application of generative AI for certain tasks in health - for example, in spurring on R&D or relieving the strain of administrative tasks for practitioners - but were universally wary of applying AI to decision making about people's future treatment. Similarly, several opinion formers made the case for using generative AI to accelerate government processes and improve public services, but identified limits to this, such as determining entitlements to certain welfare benefits.



"Who is this for? It is for the banks to make more money out of you."


UK PUBLIC

WHO IS PERCEIVED TO BENEFIT

Whether or not there was a clear benefit to consumers also appeared to be an important driver of attitudes to a particular use case. In some cases, such as providing financial or legal advice, wider accessibility and greater affordability resonated with the public. Similarly, the benefits to patients of automating administrative tasks currently conducted by healthcare staff were intuitive. In other cases, though, the ultimate consumer benefit was less clear. For example,

FIG. 11: PREFERRED REGULATION OF USE CASES (UK)

% OF UK ADULTS INDICATING PREFERRED LEVEL OF REGULATION



calculating highly-tailored credit scores based on wide-ranging data sources was an application that was not perceived to benefit consumers - and was instead regarded as primarily for the benefit of banks as a means of making money out of consumers.



When we think about who loses out, it will be pretty much all minority communities”

ACADEMIC

For opinion formers, the question of “who benefits” was a central preoccupation, with wide-ranging fears about generative AI leading to an increase in inequality. Although one opinion former in the US suggested that “when the tide rises, all ships rise”, the vast majority of other interviewees vehemently disagreed with this, suggesting that it was far more likely that generative AI would exacerbate inequalities within society. For example, opinion formers pointed out that it is already the case that marginalised and less affluent groups are less likely to receive healthcare. This could, in turn, mean they are under-represented in the health data that ultimately will underpin generative AI tools, potentially reinforcing existing inequalities in the health sector.



Everyone should be able to benefit from this development in technology - not just those with huge resources.”

TRADE UNION

LEVELS OF EXISTING REGULATION

A number of opinion formers said that they felt most optimistic about the use of generative AI in highly regulated sectors because extensive regulation means they are “ahead of the game on controls and trust” relative to other fields. In healthcare, for example, while frameworks will need to be adapted, opinion formers were unanimous that existing regulators that hold existing powers to regulate AI as a “medical device”, together with data protection laws, provided solid foundations.



Companies in financial services, insurance, human resources...will be the winners, because they will already have existing infrastructure to comply with the emerging regulation in this area...”

INDUSTRY STAKEHOLDER

Regulation

SUMMARY

- There was robust public support for regulation of generative AI, though there was limited understanding of the likely policy levers, underlining the gulf between the public and opinion former debate.
- A number of critical tensions were apparent. For example, while opinion formers broadly felt that the focus should be on existing legal frameworks, they identified a handful of areas likely to face reform, such as EU copyright frameworks.
- Many opinion formers felt there was a false dichotomy between tackling immediate risks from AI and generative AI, and focusing on existential future risks from frontier AI, and that both could be approached in tandem.
- Opinion formers pointed to tensions between prioritising competitiveness and ensuring AI adheres to local norms. The US – due to fear of Chinese competition – is prioritising competitiveness while the EU has a greater focus on upholding European values in AI development.

Implications for businesses

There is an onus on businesses developing and deploying generative AI systems to demonstrate the relevance of generative AI use cases to salient policy challenges of governments and regulators, and a clear “contribution” to economic growth and/or positive societal outcomes.

For companies operating outside of already highly regulated sectors, such as healthcare and financial services, there will be a higher threshold for reassuring policymakers and a greater need to communicate safety and risk mitigation policies and processes.

There was widespread public demand for oversight and regulation of generative AI - reflecting its perceived significance, the pace of development, and the risks considered particularly concerning. Only a minority of the public in the US, UK and Germany assumed that there was already significant regulation of generative AI in place. While there was a clear preference for some form of restrictions on different applications of generative AI, the public did not have clear expectations of the precise shape or nature of regulation.

“It’s very early days and it’s a little bit Wild West at the moment... Maybe it’s just happening too quickly for regulation to keep up with.”

UK PUBLIC

On prompting, however, there was significant support for specific regulatory initiatives. This was particularly widespread for technology companies being required to obtain consent prior to using personal data, acquiring a licence before their systems could be sold or marketed, paying compensation for using creators’ IP, and displaying a watermark on any AI-generated content. The UK public was particularly supportive of these regulatory initiatives, in line with greater concern about generative AI more broadly and other large technology companies¹³.

In opinion former interviews, five key policy tensions emerged that will shape how policy, regulation and legislation play out following elections in the US, EU and UK in 2024.

ARE EXISTING REGULATORY FRAMEWORKS SUFFICIENT OR ARE NEW LAWS NEEDED?




In Europe, regulators will look to address many of the immediate challenges posed by generative AI through greater enforcement of existing regulation, most notably the GDPR. Indeed, some opinion formers argued that data protection authorities in Europe need increased resources rather than new powers in order to apply the GDPR to emerging technologies like generative AI.

“We don’t need more powers but we do need more people and more visible enforcement.”

REGULATOR

FIG. 12: NET SUPPORT FOR SELECTED POLICY AND REGULATORY INITIATIVES

% OF ADULTS IN EACH MARKET SUPPORTING EACH INITIATIVE MINUS THOSE OPPOSING IT

			
Requirement to obtain consent before using personal data	+66%	+60%	+49%
Requiring watermarks for generative AI content	+62%	+57%	+48%
A licensing system	+64%	+46%	+49%
Payments to creators for using their IP	+57%	+52%	+44%
Creating a new AI safety regulator	+57%	+35%	+40%
Banning AI models providing information on illegal activity	+44%	+34%	+30%
Obligation to make all models “open source”	+34%	+31%	+27%
Introducing a new tax on the use of generative AI models	+39%	0%	+15%

DATA PROTECTION AND GENERATIVE AI

The EU and UK versions of the GDPR contain a range of provisions designed with AI in mind, such as the right for citizens to request a review of a solely automated decisions, as well as a range of broader provisions which are relevant to AI deployment, such as transparency, user consent, and the right to information. However, it is currently unclear how certain principles under the GDPR will apply to generative AI. For example, do the principles of transparency apply to training datasets used for generative AI models? And how do principles, such as the right to be forgotten, apply to generative AI applications?

As a result, data protection has been one of the earliest policy areas where deployment of generative AI has been tested. The most notable case was in Italy where its data protection authority, Garante per la protezione dei dati personali, temporarily paused OpenAI’s processing of Italian users, and the European Data Protection Board established a taskforce in response.

“Regulators already have what they need to protect people with existing legislation. For example, if copyright law was just enforced, then most challenges currently being faced would be reduced.”

POLITICAL ADVISER

A similar dynamic is apparent in online safety frameworks, which the UK and EU are currently implementing. In the UK’s Online Safety Bill, generative AI content will be considered within scope and subject to the regulatory enforcement powers of the UK online safety regulator, Ofcom, if it is posted on a regulated service or where there is a user-embedded AI tool within such a service. Ofcom will, like data protection authorities, have to prioritise the allocation of its resources and the extent to which it focuses on generative AI content is likely to be determined by extent to which clear instances of harms emerge.

On the question of resourcing, European opinion formers expressed high levels of concern about whether existing regulators are adequately resourced to monitor, scrutinise and undertake enforcement. There was, however, little support for the idea of creating new AI regulators at either the European or national level. One regulator described the lack of funding available as “terrifying”. An opinion former in Brussels

said that the AI Act would “live or die” according to whether it could actually be enforced or not. A similar debate was apparent in the UK where the Equality and Human Rights Commission was regularly namechecked as lacking sufficient resource to undertake its responsibilities under the AI White Paper. There was not a consensus on how to remedy this, particularly over whether industry should increasingly fund regulators through a levy.

In the US, the lack of comprehensive federal legislation means there are significant gaps for regulating generative AI, though there are a number of frameworks that could address certain aspects. At the state level, a number of states such as California, Texas and Colorado have passed or are in the processing of passing privacy and content moderation laws. At the federal level, there is voluntary guidance such as the National Institute of Science and Technology (NIST) AI Framework and the AI Bill of Rights supported by a patchwork of existing rules and regulatory powers for federal agencies. The Federal Trade Commission (FTC), Department of Justice (DOJ), Equal Employment Opportunity Commission (EEOC), and the Consumer Financial Protection Bureau (CFPB) issued a joint statement in April 2023 that they intend to apply their existing legal authorities to the use of automated systems¹⁴.



A new dedicated AI agency should not eat up all the energy in the room right now. I think the current 20 existing departments and agencies already doing work on AI should be given more resources.”

ACADEMIC

In this context, several US opinion formers argued that the debate over a new dedicated AI agency should not distract from the need to utilise existing regulatory powers. They argued that the capabilities and resources of existing federal agencies should be strengthened to ensure they can effectively address and evaluate AI specific use cases that fall within their jurisdiction. However, partisan divides over the appointments to and funding of federal agencies meant that several opinion formers were sceptical that bodies like the FTC could be sufficiently resourced to undertake ambitious regulation of generative AI. Other opinion formers were supportive of federal level AI regulation in the absence of a dedicated regulator, building on the work of Senate Majority Leader Chuck Schumer (D-NY).

SHOULD POLICYMAKERS PRIORITISE IMMEDIATE HARMS OR EXISTENTIAL RISKS?

There is a clear divide between those opinion formers arguing for prioritising immediate harms, such as worker disempowerment, bias and disinformation, and other stakeholders who instead argue that the focus should be on existential risks that might arise in the future, such as weapon design or AI reaching a point where its capabilities exceed those of humans.



One of the things I am most concerned about is the ability of generative AI to amplify disinformation and to influence elections. Targeted messaging as part of political campaigning is something we just don't even fully understand yet.”

POLITICAL ADVISER

ELECTION CAMPAIGNS AND GENERATIVE AI

The prospect of major elections in the US, EU and UK has focused minds on the potential of generative AI tools to produce misleading content. This concern draws on the legacy of the Cambridge Analytica controversy, as well as high profile deepfake incidents, such as an audio clip falsely depicting UK opposition leader Keir Starmer during the UK Labour Party’s annual conference in October 2023.

Věra Jourová, European Commission Vice-President for Values and Transparency, warned in September 2023 of “the high potential of such realistic AI products for creating and disseminating disinformation. The risks are particularly high in the context of elections.” In the first instance, Jourová and other policymakers have encouraged online platforms to take voluntary measures to address this challenge. These could draw on certain countermeasures which already exist, including automated detection tools like GPTZero and DetectGPT.

However, there are efforts in parallel to introduce regulatory requirements. For example, the European Parliament’s negotiating position on the AI Act includes measures to treat AI systems used for electoral purposes as “high-risk” and therefore subject to a high compliance threshold. Should a major controversy arise during the 2024 elections, this would likely generate significant impetus behind further regulation.

In the UK, for example, a series of stakeholders have led the debate around existential risk, including the Centre for the Governance of AI, the Future of Humanity Institute at Oxford University, the Centre for the Study of Existential Risk and the Centre for Long Term Resilience. On the other side of the debate, many parliamentarians and sectoral regulators are pushing for the UK government to be more focused on near term risks. The agenda for the UK AI Safety Summit suggests an increasing shift towards the former, with the focus on malign actors being aided by AI capabilities in biological or cyber-attacks and critical system interference and advanced systems that need to be aligned with values and intentions.

The view that policymakers will be forced to choose between immediate challenges or existential risks was posed as a false dichotomy by policymakers in Brussels. Opinion formers argued that generative AI developers were focusing on frontier risks in order to divert attention from other near-term restrictions, such as those within the EU’s AI Act, a view some industry stakeholders disputed.



It’s like climate change. If you make an issue too abstract and too big, people won’t know how to go about tackling it. What we need to focus on is regulation of the issues here and now.”

REGULATOR



The open letters by industry executives are written by privileged people who don’t need to worry whether they will qualify for a bank loan because of deployed AI, and have the time to worry about more existential risks.”

GOVERNMENT OFFICIAL



You can’t say everyone in the world (or everyone above a certain threshold or parameter) is going to have to be licensed by me. That’s a fool’s errand.”

FORMER REGULATOR

At the centre of this debate is the question of introducing licensing regimes, though there is a lack of clarity over which AI systems would need to be licensed. A number of opinion formers responded positively towards this concept, arguing that this model has been successful in other sectors. Some larger companies like Microsoft have also publicly called for such for some form of licensing system. Stakeholders broadly agreed that this idea was best suited to advanced foundation models and existential risks rather than near term generative AI systems, with some suggesting that processing power could be an appropriate threshold (see box below).

PROCESSING POWER AND GENERATIVE AI

Semiconductors and processing power are fundamental to building AI capabilities, as they determine the speed and efficiency of training the foundation models that underly generative AI. Advanced chips enable the handling of vast datasets, accelerating research and development exponentially, and therefore a nation with superior chip technology and processing infrastructure can see more rapid AI advancements and applications.

Processing power is, therefore, a heavily contested issue in the context of AI regulations and export controls. UK officials have indicated appetite for a potential threshold of computing power, measured in FLOPs, above which systems would face regulation and export controls, and there remains significant political interest in monitoring and placing restrictions on chip exports. Mustafa Suleyman, one of the co-founders of DeepMind, similarly argued that Washington should limit the sales of chips to uphold global AI usage standards.

However, the fact that similar risks manifest themselves in a range of AI technologies and not just generative AI (or advanced foundation models), raises questions of consistency. For example, biased outcomes can be a feature of risk prediction models as much as generative AI and foundation model outputs. Many opinion formers warned about the competition implications of a licensing system, often pointing to financial services to argue that it can have lock-in effects for larger players and create a major barrier to entry for competitors. One former regulator argued that licensing systems work best where there is a scarcity of resources, such as with spectrum, and that with the proliferation of open source generative AI models, such scarcity does not exist.

A licensing system implies a focus on risk at the level of generative AI systems and/or the companies which operate them. Alternative proposals include regulating at the infrastructure level (the following section considers how the US is targeting this layer in its approach to Chinese technology providers) or, as is the approach of the EU's AI Act, on use cases and outcomes. The argument in favour of the latter is that the wide diversity of applications that can arise from a single system would inevitably create issues in terms of deciding which to prioritise for scrutiny, and difficulty for the developers themselves who cannot carry out an impact assessment for every plausible use case. A further element of complexity comes from how generative AI systems can be altered to be used for a higher risk purpose than originally foreseen.



A risk-based approach works far better than focusing on particular products, because that assumes that the intended use will be the only use. With generative AI, that is not always the case.

GOVERNMENT OFFICIAL

HOW SHOULD COUNTRIES BALANCE INTERNATIONAL COMPETITIVENESS AND PROTECTING CORE VALUES?

As with all nascent technologies, there is a tension for governments between, on the one hand, nurturing and championing their domestic technology sector and ensuring that the productivity gains from AI are fully exploited and, on the other hand, upholding values, such as consumer safety, data protection and protection of copyright. Countries' different evaluations of how to balance these two factors has led to varying approaches around the world so far.

For example, in the US, the fear of losing the "AI race", particularly to China, has taken precedence in generative AI policy debates and this was reflected in discussions with US opinion formers. While some warned that this "gives countries the excuse to do whatever they want to 'win'", hence undermining the responsible use and development of AI, most opinion formers agreed that maintaining the US' technological lead was a major priority, especially given that any "slowing down" for safety reasons would not be mirrored by China.

In implementing this policy approach, the Biden administration has introduced restrictions at the infrastructure layer. These include export controls targeting chips and chipmaking tools, an outbound investment program to screen transactions to China, and rules limiting expansion in China by companies receiving subsidies under the 2022 CHIPS and Science Act. There is further speculation that the US administration will tighten restrictions around the use of US hyperscaler cloud providers by Chinese companies using advanced semi-conductors. To date, the policy interventions have targeted the supply chain for Chinese companies but, in the medium term, there is an open question of whether the "trusted vendor" concept in 5G policy, which saw Chinese vendors excluded from several Western markets, could be applied to AI.

While the EU is increasingly seeing critical technologies through the prism of economic security, opinion formers argued that the appetite for replicating the US approach was relatively low. This is both because of the diverse and nuanced positions towards China within Europe, and a desire from certain member states, such as France, to ensure the EU has a policy towards China that is autonomous of the US. This has been demonstrated by the opposition by a number of EU member states to proposals for an EU outbound investment screening mechanism, as well as the modest nature of the EU's "critical technologies" strategy.¹⁵



Despite what is said in outward policy speeches, the UK is not realistically going to be competing with the US or China when it comes to creating new AI technology. We need to accept that.

POLITICAL ADVISER

The outlook in the EU and UK is also shaped by the fact that Europe is likely to remain an importer of advanced AI systems in the near to medium term. In light of this, both UK and German opinion formers raised doubts about whether European companies can be competitive, at the infrastructure level at least, with China and the US, and, indeed, whether it is desirable to invest public funds in an attempt to do or, instead, focus on the application layer. EU opinion formers raised concerns about the dangers of building an AI economy highly dependent on non-European companies that could withdraw access if they so chose, are not necessarily aligned with "EU values" and whose training data is predominantly in English and from American sources.


This has two main consequences for the outlook for EU policy on generative AI under the next European Commission. There is likely to be a major focus on developing "sovereign capabilities" in foundation models, drawing on local content and languages, with EU and national funding schemes diverted towards developing local capabilities. The Chips Act and Green Deal Industrial Plan under the current Commission provide templates for how these could be approached, particularly the use of greater flexibility in state aid regimes.

Furthermore, there is also greater appetite for regulatory intervention in order to ensure that generative AI develops in accordance with "EU values". This is already being seen in the passage of the EU Act, but is likely to be complemented, for example, by enforcement of the GDPR by national data protection authorities and the European Data Protection Board. There is a notable prospect, given the strength of licence holders in EU policymaking, that the EU's Copyright Directive could be reformed to enhance protections for licence holders with regards to generative AI (see break out box), a point which was raised repeatedly by Brussels opinion formers.



I think the Copyright Directive is quite outdated. When it was being adopted, we were thinking about the use of images for creating literal memes. Nobody was thinking about AI.

POLITICAL ADVISER



COPYRIGHT AND GENERATIVE AI

Generative AI's ability to create original content raises questions for traditional copyright frameworks, as it blurs the lines between human-created and machine-generated works. The multitude of different inputs into training data also make significantly less clear the link between an input of copyrighted material and the subsequent outputs of a generative AI system.

The debate is the latest in a succession of legal and policy challenges concerning intellectual property since the development of the internet, from early cases, such as *A&M Records, Inc. v. Napster, Inc.*,¹⁶ through to the legislative protections awarded in the EU's Copyright Directive.¹⁷ The debate has primarily focused around the creative industries, particularly in light of the strikes in 2023 by Hollywood screen writers, but other sectors such as pharmaceuticals are considering potential patent protections for drug development.

The debate hinges on a policy choice between prioritising the development of a local generative AI sector or sovereign capabilities to preserve local language and culture, which would imply greater text and data mining exemptions, or championing the commercial interests of local creative sectors, implying more limited exemptions. In Brussels, there is an ongoing debate regarding whether the use of copyrighted works to train AI models is considered copyright infringement or if this falls under the scope of text and data mining exceptions in the Copyright Directive. In the UK, the Intellectual Property Office is poised to release a code of practice on this matter in the near future.

SHOULD AN INTERNATIONAL REGULATORY BODY BE CREATED?

There have been extensive calls for international coordination to regulate generative AI and frontier AI, both from the private sector, as well as governments, with some suggesting that an international regulatory body should be established.

Opinion formers were broadly sceptical about the prospects of an international regulator, particularly suggestions that the International Atomic Energy Agency should serve as a model. They argued that while an international organisation would be useful, it did not seem practically feasible for several reasons.



I'm personally not sure about the practicality or desirability of an international regulator. I see more advantage in a scientific and technical auditing organisation that can set minimum baselines that are implemented on a national level."

GOVERNMENT OFFICIAL

Opinion formers in the US argued that there was no incentive in the US to give an international organisation any weight. Indeed, apart from the OECD BEPS process, the Western world has largely failed to coordinate on global tech regulation in the past decade, marked by divergences between the absence of regulation in the US and the interventionist, and often extra-territorial, approach of the EU. In AI policy specifically, this trend seems set to continue with the EU's AI Act due to be approved by co-legislators in the first half of 2024 and no immediate prospect of US technology regulation. A major sticking point in the coming years will be how to reconcile the EU's AI Act with the G7 and other international processes. Opinion formers also raised the complexities of the composition of any new international organisation, most notably how to approach the membership of China, a challenge demonstrated by the controversy over whether the UK should invite China to its AI Safety Summit in November 2023. Ultimately the UK government did invite the Chinese government, who subsequently accepted.



I am all for the coordination at the global level but from a legal point of view none of this is legally enforceable. On the other hand, an EU regulation is law. It's applicable."

POLITICAL ADVISER



Given we're expecting generative AI to be a bedrock of economic growth...will there be much appetite for sharing capabilities with other nations? Probably not. So what does that make the international development landscape look like?"

THINK TANK

That said, an international testing centre acting more akin to CERN than the IAEA, will almost certainly be on the agenda for the Summit in London in November 2023, and international standards-making bodies will continue to provide benchmarks for AI companies. Brussels-based opinion formers also broadly supported this idea, suggesting, as Commission President von der Leyen has recently, that an "IPCC for AI" would be a useful new entity and one which could, in the longer term, galvanise international regulatory coordination on generative AI.

DO GOVERNMENTS NEED TO DECIDE BETWEEN OPEN AND CLOSED SOURCE MODELS?

A growing debate in AI policy is the divide between open source and closed source (or proprietary) AI. Proponents of the former camp assert that AI tools should be transparent and widely accessible so as to avoid concentration of market power in just a few companies. They argue that this would also drive safety as models and systems can be tested to ensure their robustness. Advocates of proprietary models counter that it is reckless to open up access at this point to incredibly powerful tools when they could then easily fall into the hands of malicious actors who can fine-tune them for purposes other than which they were built. They argue that this also opens up unresolvable questions about accountability and liability: who is ultimately responsible if open source technology is misused?



The democratization that Facebook enabled brought with it violations of personal privacy, crushing of competition, challenges to truth and trust all because of the distributed nature of the network. Open source is the same kind of situation except infinitely more powerful."

ACADEMIC

THE IAEA MODEL AND GENERATIVE AI

UN Secretary-General Antonio Guterres, UK Prime Minister Rishi Sunak, and OpenAI CEO Sam Altman are just some of the stakeholders that have endorsed the proposal to establish an international AI watchdog similar to the International Atomic Energy Agency (IAEA).

However, there is also growing criticism to suggest that the IAEA model may not be a neat fit for AI. The Bulletin of Atomic Scientists has argued that while both nuclear technology and AI have potential global impacts, AI's specific threats and modalities differ significantly from the clear dangers of nuclear weapons. For effective AI governance, a comprehensive understanding of its existential risks is required, rather than an ad-hoc approach similar to the nuclear regulation history. Moreover, it took approximately four years from the proposal to the actual establishment of the IAEA, and another four years before the first agreement on policies; given the rapid development of AI, a more agile approach is widely preferred.



The open source model only works if you have ways to hold the company liable and accountable for those harms... For systems where we only understand a small part of their capabilities, maybe it's best not to release this to the public."

ACADEMIC

Some countries are specifically supporting open source models. For example, France has chosen to champion open source AI systems, with investments of €40 million announced in June 2023 for an open "digital commons". Henri Verdier, the French ambassador for digital affairs, has been clear that France sees open source as a way to address current disadvantages that European countries face at the generative AI infrastructure and systems levels, and allow French companies to compete better at the application level. The French government is also understood to have intervened in the AI Act negotiations to argue that the way in which the rules are applied to open source models should be re-assessed.

While some opinion formers were highly critical of the open source model, the broad thrust of opinion was that the dichotomy between open and closed is unjustifiably stark. Experts were quick to point out in interviews that "although they call it open source, it's really just API based". Indeed, the direction of travel seems to be facilitating (and even encouraging) a competitive coexistence of the two in tandem. For example, the UK Competition and Markets Authority (CMA) published a report in September 2023 arguing that a mix of open and closed source foundation models would be a positive outcome and allow the foundation model market to trend towards positive competitive outcomes.¹⁸ However, this could evolve in the future if open source and closed source capabilities start to diverge and proprietary models become significantly more powerful and command more market power.



I'm less bothered by whether the technology is proprietary or open source. What is more important is the practical result: is it going to be harmful?"

REGULATOR

Implications for businesses

Businesses should anticipate a range of policy interventions on generative AI following next year's elections, though these are likely to evolve at different paces:



Copyright

Copyright will be a prominent issue, both in legal test cases and potential legislative reform, particularly in the EU.



Data protection

Data protection has already emerged as a live regulatory issue for generative AI. Legislative reform is possible, particularly if the EU proposes a revision of the General Data Protection Regulation (GDPR), but, in the first instance, further enforcement cases in Europe against generative AI companies are likely.



Election interference

There is significant concern ahead of next year's elections and this will intensify as they draw closer. Should a significant controversy arise, this could prompt calls for regulatory interventions. Electoral events are also likely to place a spotlight on any perceived 'political' bias in generative AI models.



Licencing systems

Proposals for licencing systems have gained most traction when linked to the debate about existential risks with few indications to date that governments will move to a licensing model to regulate near term risks.



Open & closed source

The open source debate is unlikely to prompt governments or regulators to 'pick a winner' between open and closed models, though some governments may follow the French example of investing in these open source industry.




International regulator

A global regulator or comprehensive international governance system appear unlikely, though it is plausible that a testing body loosely based on the Intergovernmental Panel on Climate Change (IPCC) or the European Organization for Nuclear Research (CERN) could be established in the coming years.



Restrictions on Chinese AI

There is a strong prospect of further US restrictions targeting Chinese companies, building on the Biden Administration's semi-conductor export controls. As with 5G and semi-conductor policy, the US is likely to pressure western allies to mirror its policies towards Chinese tech.



In focus: Healthcare

SUMMARY


- Healthcare stakeholders shared many of the same concerns as in other sectors, with privacy and data protection prominent. However, there was a greater intensity of concern in light of the severity of the consequences in the event of error and anxiety about liability frameworks.
- Opinion formers agreed on the need for stronger regulation but recognised the role of cross-cutting frameworks, such as the AI Act, rather than new health-specific rules, for achieving this aim.
- A specific challenge is on change control in light of the fact that generative AI models can be updated post-approval, resulting in altered performance, which runs counter to the traditional regulatory model of approval at the point of market access.

AI has been deployed across health, care and life sciences for decades, with areas such as medical imaging, pathology, drug development, mental health support and fraud detection all seeing deployment of AI-based systems and algorithms. Relative to this, generative AI is largely undeveloped, reflecting in part a lack of very large health data sets required for development. In spite of this, interest in generative AI is rapidly expanding, with significant innovation emerging across diagnostic support, document analysis and patient screening.

Core opportunities and risks as identified earlier in the report were mirrored in healthcare. For instance, potential health use cases for generative AI such as summarising clinical notes, automating insurance claims, and development of public health campaigns all reflect core potential benefits identified in other sectors. Primarily, these include efficiency, productivity and supporting creativity.

FIG. 13: PREFERRED REGULATION OF DIFFERENT USE CASES IN HEALTHCARE (US)

% OF US ADULTS INDICATING PREFERRED LEVEL OF REGULATION



However, opinion formers' concerns about the potential harms of generative AI in healthcare were particularly pronounced relative to other sectors. This reflects a naturally low risk-tolerance in healthcare. Most notably, this included what was described as the "insidious" potential for generative AI to worsen health inequalities for disenfranchised groups by amplifying biases in underpinning medical data. One opinion former argued that generative AI had the potential to "regurgitate medical tropes" for specific

demographics. Other concerns included the potential for generative AI to credibly present or even hallucinate false, biased or misleading health information. For example, recommending incorrect dosage or instructions for medication.

Opinion formers were broadly supportive of human oversight of generative AI as mitigation to such harms. However, they noted that both patients and clinicians are poorly equipped to navigate risks of generative AI.

“
Proposals for use cases which would take humans out of the loop fundamentally misunderstand what the human does... A lot of it just misses the fact that it's about a human interacting and reading between the lines.”

PUBLIC HEALTHCARE BODY

Opinion formers across the EU, UK and US all agreed that further regulation - building or adding to existing regulation - for generative AI in healthcare is needed. Health regulators in all three markets have in place medical device regulations for clinical uses of AI such as diagnostics, alongside health data protection legislation. Alone, these are not viewed as sufficient and a holistic extension of these or wider privacy and data protection, liability and online safety laws could play a key role in regulating wider health uses. Even recognising this, across all markets, there was a clear preference that existing regulators be strengthened and co-ordinated, rather than introducing new regulators.

“
I don't think we're the place yet to regulate these things because if you regulate it, what are your regulatory tools? We need to develop those regulatory tools first.”

PUBLIC HEALTHCARE BODY

“
Centralised regulation is not going to work, because there are so many stakeholders and the issues are so complex that everyone has to step up, including government, health systems and groups such as the American Civil Liberties Union.”

PUBLIC HEALTHCARE BODY

It was recognised that prospects for this differ widely across markets, reflecting attitudes to the balance of risk and innovation. In the US, while the FDA has been proactively working on AI for many years, specific concerns were raised that their remit alone was not sufficient. As an

example, the US Health Insurance Portability and Accountability Act (HIPAA) protects only a limited subset of personal data that could be relevant to healthcare, with the FTC responsible for wider data protection. With more cross-cutting privacy and data protection regulations proving politically contested, opinion formers noted that the US was unlikely to favour more holistic regulatory approaches to generative AI in health.

In contrast, opinion formers argued that the EU was more likely to successfully mitigate health risks, including through the expansion and implementation of existing regulation, albeit as a less attractive market for health innovators. In line with other sectors, an opinion former identified that existing frameworks would likely be utilised. This includes a combination of health specific and cross-cutting regulations including EU Medical Device Regulations, the AI Act and the GDPR. The UK was considered likely to take a middle ground.

“
The EU spirit is more to protect the public rather than to unleash all of the economic potential. And that's because healthcare has been identified as a high risk use of AI.”

INDUSTRY STAKEHOLDER

“
In the UK and EU we regulate first and innovate second – in the US, a lot more of the tech development will happen as they regulate second.”

HEALTHCARE ADMINISTRATOR

While opinion formers identified similar regulatory challenges as elsewhere, a specific concern was change control to maintain quality and safety when an AI technology or its data inputs are updated. Traditional models of regulating for safety rely on approval at the point of market access. Unlike traditional products like medicines, generative AI can be updated post-approval, resulting in altered performance. Even relative to other AI products or digital health apps, generative AI poses a particular challenge, as regulators and developers are unable to access or adapt the underpinning model governing outputs. Generative AI heightens questions around who should be responsible for monitoring and change control, and how this would be delivered.

Practical questions raised by challenges such as change control are compounded by the sheer volume of generative AI products that could be utilised along a care pathway. One opinion former identified that healthcare organisations could become more responsible for regulation and change control as data custodians of outcome data, with others supporting a distributed set of responsibilities, given the complexity of the landscape.

“
Are we saying that for that one organisation, a regulator needs to verify and manage change control for twenty devices? Or can an organisation do all of them at the same time based on the same data?”

CLINICAL ACADEMIC

Finally, the question of responsibility was prominent in healthcare as a high-risk, often litigated sector. Without proactive regulation, liability for patient safety is likely to be a powerful driver for how users and developers approach generative AI in healthcare. While the US has historically been regarded as litigious, this factor was noted across all three markets. One opinion former noted the


key role that new EU “collective redress” legislation, which allows for the role that class action lawsuits could play. In contrast, in the US, at present liability primarily falls to physicians where they use decision-support tools. Any policy reforms in this area will be key in shaping how readily generative AI is taken up in healthcare settings.

“
In the US, all the liability is transferred to the clinician who's getting the decision support, and generally, the decision support companies have been immune. That's something that I think is outmoded and needs to change”

CLINICAL ACADEMIC

“
Where is the liability? Do you bet on the AI or not? For instance, taking someone off medication – is that a good idea if you don't know exactly how that decision was made?”

HEALTHCARE ADMINISTRATOR



In focus: Financial services

SUMMARY


- Financial services stakeholders argued a number of concerns about generative AI were more pronounced in finance, such as those concerning bias and discrimination, where the impact of an erroneous AI decision could have significant consequences for individuals.
- Opinion formers also reported a standalone set of concerns related to financial services. These included the potential for generative AI to facilitate economic crime, deepen inequality issues associated with access to finance, and the potential for market manipulation.
- The industry has an advanced regulatory starting point in Europe with legislation like the Consumer Credit Directive already referring to the AI Act in the EU, and the Consumer Duty in the UK. By contrast, stakeholders expected litigation rather than legislation to set precedents for generative AI in financial services.

The consumer-facing side of financial services has significant potential for exploring generative AI use cases and there is growing excitement in industry about the technology being used for a range of

applications. These include improving customer service, unlocking financial advice for those that cannot afford it and fraud detection and prevention.

FIG 14: PREFERRED REGULATION OF USE CASES IN FINANCIAL SERVICES (DE)

% OF GERMAN ADULTS INDICATING PREFERRED LEVEL OF REGULATION



OPINION FORMER CONCERNS

Complexities inevitably arise in the nexus between generative AI and one of the most highly regulated sectors in the economy. In wholesale markets, the application and use of generative AI, whilst presenting clear benefits, raises a whole set of different issues. Opinion formers were clear that there will need to be a cautious regulatory approach to avoid arbitrage, sensitive data protection issues, market manipulation and many other unintended consequences.

Some of the issues raised mirror broader debates, but there was a greater intensity of concern in the financial services sector. These include data protection, where the highly sensitive nature of personal information prompted caution from opinion formers about the deployment of generative AI with customer accounts. There were also significant

fears about bias and discrimination and questions about how this would impact potential financial services products. Likewise, liability was particularly sensitive for financial services opinion formers, as was the need to regulate “critical third parties” like hyperscale cloud services, in light of the potential consequences for consumer protection and financial stability, where there is less scope to accept unpredictability and a black box effect.

Concerns about the concentration of power in a few select AI companies were common across different sectors, but in the financial services sector specifically, some felt that this could have profound consequences - potentially even triggering a financial crisis in the future. Chair of the US Securities and Exchange Commission, Gary Gensler, has spoken publicly about this, arguing that “most of our regulation is about individual institutions, individual banks, individual money market funds

[and] individual brokers... [But] this is about a horizontal matter where many institutions might be relying on the same underlying base model”.¹⁹

There was also a category of issues where concerns raised by opinion formers were specific to financial services. These include economic crime, which has been increasing over recent years. This is likely to be compounded by the use of deepfakes and trendspotting by criminals, which could be used to defraud consumers and businesses or, with a more general application, used to potentially manipulate markets.

Financial inclusion was also a significant issue for opinion formers. They raised concerns about the potential for generative AI to further exacerbate existing financial inclusion challenges and to further entrench social inequalities if used to restrict access to financial products for marginalised groups. Advisory tools could also have the opposite effect by making access to financial advice significantly cheaper and easier, though existing regulatory restrictions may make this difficult to apply in practice.

A final concern raised in interviews was market manipulation, which has two facets. The first is generative AI misinterpreting market trends, the other is nefarious actors utilising generative AI tools to manipulate markets to their own advantage. While this is a subject of constant vigilance, regulators may struggle to understand market manipulation until that manipulation has already been committed.

REGULATORY APPROACHES

The starting point for the debate on AI regulation in financial services differs between jurisdictions. The EU is in a unique position where many of the horizontal issues related to generative AI are being addressed in the AI Act and sectoral regulation is already beginning to slot in. The recently negotiated Consumer Credit Directive (CCD) already refers to the AI Act by putting the use of AI in consumer credit assessments as a “high risk”. Despite this, it seems likely that there will be appetite for more regulation; what is less clear is when this will happen and through what vehicles. Elsewhere, in the UK, the implementation of the Consumer Duty and the Senior Managers regime both allow the Financial Conduct Authority to regulate generative AI deployment in financial services without further rules, although further guidance will be needed on how exactly these powers are to be used.

In contrast, the US has not yet embarked upon legislating for AI and in financial services, there is no clear consensus on whether new legislation is needed, or if existing regulation and protections should be updated to reflect the risks posed by AI. It is therefore likely that major issues will be addressed through the state and federal court systems in the medium term. At an international level, the Financial Stability Board and Bank for International Settlements both faced criticism from opinion formers for being slow to respond to emerging regulatory issues associated with new technologies.

“The momentum we saw in the US previously has stalled. Before we saw bipartisan interest and push – now we are seeing different camps in Congress and not necessarily along partisan lines.”

THINK TANK

Across all interviews, it was clear that while there are existing frameworks in place that can be used to regulate generative AI in financial services, particularly in the EU and UK, more will be needed - but opinion formers advocated a cautious, iterative approach rather than seeing need for a more comprehensive reform of existing financial services rules. On enforcement, it is likely that both the EU and UK will prioritise consumer protection and fraud issues ahead of other policy issues.


“It is inevitable more [regulation] will be needed, but it would be good to start with what we already have.”


INDUSTRY STAKEHOLDER

The implementation of open banking has been an international effort in recent years with Australia, India, the EU and the UK all putting significant resources into developing open data requirements. While never the intended aim of open banking or finance, it facilitates substantial amounts of data standardisation, enabling the collation and comparison of consistent data sets across products and across providers. Generative AI’s unprecedented ability to bring together disparate datasets could potentially lead to advantageous outcomes in terms of personal finance advice, bringing out insights that would hitherto have been difficult to access, particularly should open finance initiatives in both the UK and EU progress further.

“Could open finance be a pre-requisite for generative AI in financial services to work well?”

INDUSTRY STAKEHOLDER





Looking ahead to the next policymaking cycle



Returning to the objective of this report, our clear conclusion is that businesses developing and deploying generative AI should anticipate political scrutiny, regulatory enforcement and, in certain jurisdictions, legislation after the elections in 2024.

However, the political dynamics that generative AI policy will play into are not static and will continue to evolve, shaping whether and how regulation and legislation emerge. Differences in political systems will have an important bearing on where regulation and legislation happen. The structured nature of policymaking in Brussels with clear, five-year policy cycles supported by a cross-party, consensual system and strong technocratic institutions with clear legislative competences, increases the likelihood that the EU will remain the first mover in tech and AI regulation under the next Commission. Barring a major landslide for one party, the partisan, adversarial nature of US politics in particular limits the likelihood of federal generative AI reforms, regardless of who wins the White House.

The debate will also be influenced by external events that shape political perception of generative AI and its potential harms. One example is whether generative AI companies become beset with reputational crises in the same way that social media companies have. The Snowden leaks and the Cambridge Analytica scandal propelled the EU's "techlash" while the tragic case of Molly Russell was a significant factor in the UK's approach to online harms, which ultimately became the Online Safety Bill. Likewise, legal cases could prove influential. Test cases such as *Getty vs. Stability AI* or the fallout from legal rulings on Section 230 could expose gaps in existing legislation and create momentum behind reform.

Commercial dynamics will also be important, particularly on regulatory enforcement. As we saw with the launch of ChatGPT, the extent to which regulators focus on generative AI is directly linked to its commercial success and the number of consumers using generative AI applications. The development of a "killer app" (or apps) in the coming years could have such an effect, as could the impact of technological change on influential incumbent industries and interest groups, whether through alleged IP infringements or joblessness, which may mobilise as a result.

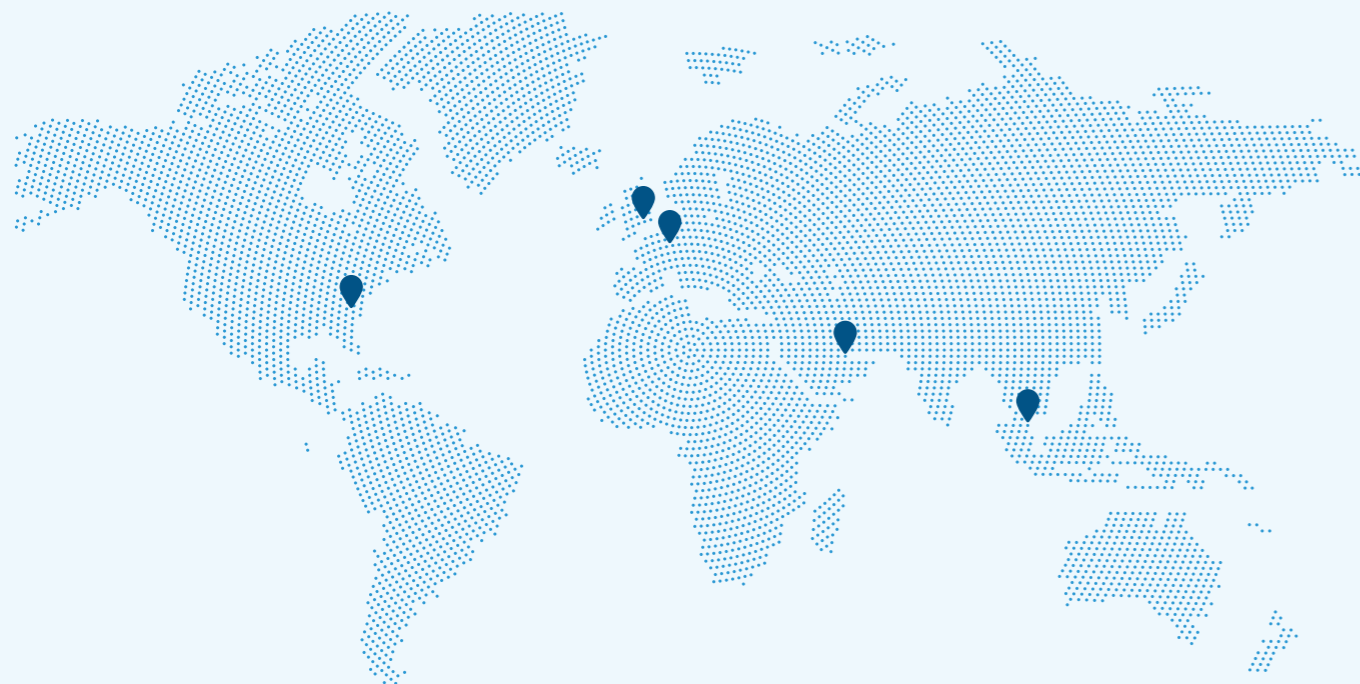
At Global Counsel, we will continue to track and analyse these debates with this report acting as the first in a series of insights anticipating the impact of the 2024 elections on global technology policy.

About Global Counsel

Global Counsel is a strategic advisory business.

We help companies and investors across a wide range of sectors anticipate the ways in which politics, regulation and public policymaking create both risk and opportunity – and to develop and implement strategies to meet these challenges. Our team has experience in politics and policymaking in national governments and international institutions backed with deep regional and local knowledge.

Our offices in Brussels, London, Singapore, Washington DC and Doha are supported by a global network of policymakers, businesses and analysts.



Endnotes

1. See [OECD \(2019\)](#); See [WEF \(2019\)](#); See [UNESCO \(2021\)](#)
2. See [European Commission \(2020\)](#)
3. See [FDA \(2019\)](#)
4. See [Cabinet Secretariat \(2019\)](#); See [Department of Industry, Science and Resources \(2019\)](#)
5. See [Verint \(2022\)](#)
6. See UBS via [Reuters \(2023\)](#)
7. See [NIST \(2023\)](#); See [DSIT \(2023\)](#); See [Canadian Parliament \(2022\)](#); See [European Commission \(2021\)](#); See [PWC \(2023\)](#)
8. See [Global Counsel \(2022\)](#)
9. See [CDEI \(2022\)](#)
10. See [Science \(2023\)](#)
11. See [Straits Times \(2023\)](#)
12. See [BBC \(2023\)](#)
13. See [Global Counsel \(2022\)](#)
14. See [FTC \(2023\)](#)
15. See [European Commission \(2023\)](#)
16. See [US Copyright Office \(2001\)](#)
17. See [Official Journal of the EU \(2019\)](#)
18. See [CMA \(2023\)](#)
19. See [Financial Times \(2023\)](#)

Global Counsel Ltd

E: info@global-counsel.com
www.global-counsel.com

LEAD AUTHORS

Conan D’Arcy,
Senior Practice Director
✉ c.darcy@global-counsel.com

Megan Stagman,
Director
✉ m.stagman@global-counsel.com

Raphael Malek,
Head of Research and Insight
✉ r.malek@global-counsel.com

David Song,
Associate Director, Financial Services
✉ d.song@global-counsel.com

Florence Chalker,
Senior Associate, Health and Life Sciences
✉ f.chalker@global-counsel.com

CONTRIBUTORS

Bart Myners
Constantine Arvanitis
Emma Morris
Jack Keevill
Jeffrey Afianmagbon
Mark Loughridge
Matilda Milne
Natasha Dixon
Mohamed Hamze
Theresa Duemchen
Ugonma Nwankwo



Global Counsel

© GLOBAL COUNSEL 2023

Although Global Counsel makes every attempt to obtain information from sources that we believe to be reliable, we do not guarantee its accuracy, completeness or fairness. Unless we have good reason not to do so, Global Counsel has assumed without independent verification, the accuracy of all information available from official public sources. No representation, warranty or undertaking, express or implied, is or will be given by Global Counsel or its members, employees and/or agents as to or in relation to the accuracy, completeness or reliability of the information contained herein (or otherwise provided by Global Counsel) or as to the reasonableness of any assumption contained herein. Forecasts contained herein (or otherwise provided by Global Counsel) are provisional and subject to change. Nothing contained herein (or otherwise provided by Global Counsel) is, or shall be relied upon as, a promise or representation as to the past or future. Any case studies and examples herein (or otherwise provided by Global Counsel) are intended for illustrative purposes only. This information discusses general industry or sector trends, general market activity and other broad economic, market or political conditions. It is not research or investment advice. This document has been prepared solely for informational purposes and is not to be construed as a solicitation, invitation or an offer by Global Counsel or any of its members, employees or agents to buy or sell any securities or related financial instruments. No investment, divestment or other financial decisions or actions should be based on the information contained herein (or otherwise provided by Global Counsel). Global Counsel is not liable for any action undertaken on the basis of the information contained herein. No part of this material may be reproduced without Global Counsel's consent.