How do people feel about AI?

Wave two of a nationally representative survey of UK attitudes to AI designed through a lens of equity and inclusion

Executive summary

AI technologies are proliferating rapidly across society with substantial global investment, yet the discourse and research around public attitudes towards these technologies remain incomplete. The UK is made up of many publics that are differentially impacted by new technologies. Some applications of AI technologies bring up unique concerns and attitudes from those on lower incomes, those from minoritised ethnic backgrounds and those with fewer digital skills.

Current research suffers from two fundamental gaps: 1) it typically treats AI as a single entity rather than examining specific contextual applications, and 2) it does not adequately represent marginalised voices. These blind spots not only hinder responsible design and development but also preclude effective governance and regulation that could address socioeconomic inequalities.

We welcome the recognition in the UK's AI Opportunities Action Plan that 'Government must protect UK citizens from the most significant risks presented by AI and foster public trust in the technology, particularly considering the interests of marginalised groups'; although there is a gap in both specific commitments against this ambition, and in understanding people's perceptions of AI risks and what would foster public trust. The risk of failing to meet this gap was recently stated by the Secretary of State Peter Kyle:

> 'Trust is incredibly important in this whole agenda. We have seen too many times in the past where a fearful public have failed to fully grasp the potential for innovation coming out of the scientific community in this country. We are not going to make that mistake. We understand from the outset that to take the public with us we must inspire confidence.'[1]

To address these broad gaps in understanding, we conducted a nationally representative survey of 3,513 UK residents in November 2024. This survey, which is part of the UKRI-funded Public Voices in AI project,[2] is the second iteration of *How do people feel about AI?,*[3] a national survey of attitudes to AI. Our sample was representative of the UK public across age, sex, income, education and ethnicity, among other demographic factors. To strengthen principles of equity and inclusion in our survey design, we deliberately oversampled groups often underrepresented in survey-based research: those with low digital skills, those on lower incomes, and people of Black or Asian ethnicities.

We asked people about their awareness of, experience with and attitudes towards eight different uses of AI – six of which were repeated from the 2022/23 survey. The two new

[1] 'Artificial Intelligence Opportunities Action Plan - Hansard - UK Parliament' (13 March 2025) <https://hansard.parliament.uk/Commons/2025-01-13/debates/8C036071-5845-443C-B903-57483D552854/ArtificialIntelligenceOpportunitiesActionPlan> accessed 13 March 2025.

[2] 'Public Voices in AI' (ESRC Digital Good Network) <https://digitalgood.net/dg-research/public-voices-in-ai/> accessed 13 March 2025.

[3] Ada Lovelace Institute and Alan Turing Institute, 'How Do People Feel about AI? A Nationally Representative Survey of Public Attitudes to Artificial Intelligence in Britain' (2023) <https://www.adalovelaceinstitute.org/report/public-attitudes-ai/> accessed 6 June 2023.

technologies were applications of AI launched after the data collection for our 2022/23 survey: general-purpose large language models (LLMs) and mental health chatbots. Four applications are already in use: facial recognition in policing, which is well covered in the media, and technologies for assessing eligibility for welfare benefits, cancer risk or loan repayment, which are less visible in public discourse. We also asked about applications of AI that are not yet part of people's everyday experiences, such as driverless cars and robotic care assistants.

For each specific use of AI, we asked people what they believed were the key benefits and concerns, recognising that these perceptions are contextual. People's perceptions vary in three respects: 1) the specific context of each application of AI, 2) the different demographic groups people identify with, and 3) seeing both potential benefit and concern simultaneously. We also asked people about their views around issues such as AI-generated decision making and harms, and how they would like to see these technologies regulated and governed. Key findings relating to public attitudes across these technologies are summarised below.

Key findings

- **Awareness varies significantly across AI applications.** People have a high awareness of technologies commonly featured in media discourse. For example, 93% of the public have heard of driverless cars, 90% have heard of facial recognition in policing and 61% have heard of LLMs. We found that technologies which are increasingly used in public services remain largely invisible to members of the public. Despite their significant potential to impact people, especially the most vulnerable, awareness of AI for assessing eligibility for welfare benefits like Universal Credit (18%), robotic care assistants designed to carry out physical tasks in care settings such as hospitals and nursing homes (24%) and tools assessing how likely a person is to repay a loan, such as a mortgage (24%) is low.
- **General-purpose LLMs such as ChatGPT have acquired rapid levels of awareness and use.** 61% of the public have heard of LLMs and 40% have used them. This is rapid penetration for an AI application that began to receive media coverage only in December 2022. However, openness to the use of these tools is context dependent. For example, 67% of people have used, or are open to using, LLMs for searching for answers and recommendations. This figure drops to 53% for using LLMs to support job applications.
- **Since 2022/23, perceptions of overall benefit for most AI technologies have remained stable, while concern levels have increased.** Where previously benefits outweighed concerns for five of the six technologies (all except driverless cars), in the current wave, benefits outweigh concerns for only three (cancer risk assessment, facial recognition in policing, and assessing loan repayment risk). The rise in concern is particularly notable for the use of AI in assessing welfare eligibility. In 2022/23, 44% of people were concerned by this technology. This has risen to 59% in 2024/25.
- **Different demographic groups have distinct attitudes to applications of AI.** In the case of facial recognition for policing, while 39% of the general population expresses concern about its use in policing, this rises to over half among Black (57%) and Asian (52%) people. Some of their concerns are also more strongly held than the general public: 66% of Black people and 62% of Asian people are concerned by false accusations compared to 54% of the general public. Similarly, people on lower incomes consistently report lower net benefit scores across AI technologies compared

to people on higher incomes (meaning they are more likely to see their concerns around a technology as outweighing their perceived levels of benefit for that technology), even when holding other demographic variables constant. This indicates that income status also influences perceptions of AI technologies.

- **The UK public hold nuanced views on the specific benefits and concerns associated with different uses of AI.** Overall, people most commonly identify speed and efficiency improvements as key AI benefits, while their top concerns centre on overreliance on technology over professional judgement, errors, and lack of transparency in decision making. Within this overall pattern, benefits and concerns vary by AI application. Even for applications with high levels of perceived benefit, people had concerns. For the use of AI in cancer detection, 64% worried about the loss of professional judgement due to overreliance on technology, while for facial recognition, 54% were concerned about fairness due to the risk of false accusations. For the least popular application, driverless cars, 63% saw accessibility as a major benefit. Taken together, these points show how people can simultaneously hold perceptions of benefits and concerns, and how perceptions vary across the specific contexts in which AI technologies are used.
- **The public self-report high exposure to AI-generated harms.** Overall, close to two-thirds of the UK public (67%) have experienced any form of harm a few times, while over a third (39%) have encountered any form of harm many times. The most common harms people report experiencing are false information (61%), financial fraud (58%) and deepfakes (58%).
- **The public expect the government to be equipped in relation to AI safety.** 58% of people believe both an independent regulator and AI companies should be responsible for ensuring AI is used safely, and the majority (over 75%) feel it is 'very important' for the government or independent regulators to have a suite of safety powers, rather than private companies alone having this control. While younger people (18-44) favour company responsibility, those over 55 prefer regulators, reflecting differing levels of trust and expectations based on age. These expectations are important given the risks to safety people already report experiencing, and the pace at which advancements in AI are being made.
- **The public increasingly want laws and regulation in order to be more comfortable with AI technologies.** The majority of the public (72%) indicate that laws and regulations would increase their comfort with AI technologies – an increase from 62% from the 2022/23 survey. This rise in demand for laws and regulation comes at a time when the UK does not have its own set of comprehensive regulations around AI. 65% of people said that procedures for appealing decisions made by AI would make them feel more comfortable with AI, along with 61% who felt getting information on how AI systems made decisions about them would increase their comfort levels. This will be significant in the context of upcoming regulatory decisions, for example, forthcoming UK government changes to the law around automated decision making.

How to read this report

If you are a policymaker or regulator concerned with AI technologies:

- Section 3.2 offers an overview of overall perceptions of benefit for AI technologies and levels of concern. Section 3.3 offers insight into the specific benefits and concerns the public associates with different uses of AI.
- Section 3.4 details public expectations around the governance and regulation of AI. Insights include the mechanisms that would increase people's comfort with the use of AI, as well as concerns around decision making, AI safety, and data sharing and representation. It also includes detail on expectations from different stakeholders, including the government.

If you are a developer or designer building AI-driven technologies, or an organisation or body using them or planning to incorporate them:

- Section 3.1.2 offers insight into personal experiences with LLMs.
- Section 3.3 offers insight into the specific benefits and concerns the public associates with different uses of AI.
- Section 3.4 details public expectations around the governance and regulation of AI. It also includes detail on expectations from different stakeholders, including developers.

If you are a researcher, civil society organisation, public participation practitioner or member of the public interested in technology and society:

- Section 2 details the survey methodology, including important limitations encountered when designing with an intent to improve sample diversity.
- Section 3.1 includes an overview of people's awareness and experience of different AI uses. Section 3.2 offers an overview of overall perceptions of benefit for each technology and levels of concern. We also highlight key demographic differences in these perceptions. Appendix section 5.1 has further details on predictors of net benefit scores for each technology. Public perceptions about specific benefits and concerns are detailed in Section 3.3.
- Section 3.4 details public expectations around the governance and regulation of AI. It also includes detail on expectations from different stakeholders.

1. Introduction

Countries and companies worldwide are investing in rapidly deploying AI technologies, leading to unprecedented advancements in AI capabilities. From DeepSeek R1 to OpenAI's o3, AI models can be used to undertake complex tasks – including reasoning, writing software, generating hyperrealistic images and videos, and engaging in multi-turn open-ended conversations – as well as to contribute to addressing broader challenges such as modernising public services.[4]

When designed responsibly and safely, these technologies have the potential to improve people's lives. However, concerns persist that AI could also exacerbate the socioeconomic inequalities and sense of disempowerment that have had such significant impacts on national political landscapes. Key concerns include job displacement, biases that mean AI tools do not work as intended and risks to safety.

Effectively regulating these trade-offs requires understanding public perspectives, especially as AI becomes increasingly embedded in daily life. The ways individuals across different demographic groups experience and perceive AI provide valuable insights to support its responsible adoption, development and regulation. Without active public involvement, there is a risk of creating an 'AI-cracy', where a small, privileged group controls AI development and governance to the detriment of broader society.[5] If a significant gap is allowed to develop between public expectations around protection from AI impacts and government action, this could risk igniting a public backlash against AI that would significantly limit its potential benefits.

However, there is currently a lack of evidence on how people view AI. Existing studies have two major limitations: they often focus on AI as a single entity or product rather than examining specific applications, and they do not represent or examine marginalised or underrepresented voices. These gaps in understanding hinder the government's monitoring of AI's impact, and consequently its decision making about and effective development of regulation and accountability mechanisms.

To address these gaps, the Ada Lovelace Institute and The Alan Turing Institute conducted a nationally representative survey to assess the UK public's attitudes towards eight AI applications in risk and eligibility assessment, facial recognition, LLMs and mental health chatbots, and robotics. This study, which is part of the UKRI-funded Public Voices in AI project,[6] marks the second iteration of the *How do people feel about AI?* survey.[7] It explores public awareness of AI technologies, concerns, perceived benefits and differences in attitudes across demographic groups. Additionally, to inform policy action, it examines public opinions

[4] Jonathan Bright and others, 'Generative AI Is Already Widespread in the Public Sector' (arXiv, 2 January 2024) <http://arxiv.org/abs/2401.01291> accessed 13 March 2025.
[5] Reema Patel, 'A Framework and Self Assessment Workbook for Including Public Voices in AI (Elgon Social Research and ESRC Digital Good Network) <https://digitalgood.net/dg-research/public-voices-in-ai/> accessed 13 March 2025.
[6] 'Public Voices in AI' (n 2).
[7] Ada Lovelace Institute and Alan Turing Institute (n 3).

and expectations about AI governance and regulation. While previous studies have explored related issues, this survey distinguishes itself through three key features.

1.1. How we define AI

Recognising that AI is a broad and evolving field, and that public perceptions may vary based on specific applications,[8] this study seeks to understand public attitudes towards specifically defined AI use cases. Other research often relies on broad definitions of AI, providing limited application-specific insights,[9,10,11,12] or focuses on singular use cases, such as attitudes towards biometrics in policing and law enforcement.[13] These approaches make it difficult to capture public sentiment comprehensively.

Our survey enables respondents to express both benefits and concerns associated with distinct AI applications. In this wave, we examined public attitudes towards the following AI categories:

- **Risk and eligibility assessments and facial recognition**: Assessing eligibility for welfare benefits, assessing risk of cancer from a scan, assessing risk of repaying a loan, and facial recognition for policing
- **LLMs and mental health chatbots**: General-purpose large language models (LLMs) and mental health chatbots
- **Robotics**: Driverless cars and robotic care assistants.

1.2. Understanding diverse perspectives

This study highlights the perspectives of different demographic groups, especially those marginalised in conversations and research about AI. We focus on:

- people on lower incomes
- digitally excluded people

[8] ibid.
[9] Workday, '2024 Global Study: Closing the AI Trust Gap' (2024) <https://forms.workday.com/en-us/reports/the-ai-trust-gap/form.html> accessed 13 March 2025.
[10] American Psychological Association, '2023 Work in America Survey: Artificial Intelligence, Monitoring Technology, and Psychological Well-Being' (APA, 2023)<https://www.apa.org/pubs/reports/work-in-america/2023-work-america-ai-monitoring> accessed 26 September 2023.
[11] Alec Tyson and Emma Kikuchi, 'Growing Public Concern about the Role of Artificial Intelligence in Daily Life' (Pew Research Center, 28 August 2023) <https://www.pewresearch.org/short-reads/2023/08/28/growing-public-concern-about-the-role-of-artificial-intelligence-in-daily-life/> accessed 13 March 2025.
[12] Ipsos, 'Public Trust in AI: Implications for Policy and Regulation' (2024) <https://www.ipsos.com/sites/default/files/ct/news/documents/2024-09/Ipsos%20Public%20Trust%20in%20AI.pdf> accessed 13 March 2025.
[13] Sam Stockwell and others, 'The Future of Biometric Technology for Policing and Law Enforcement' (Centre for Emerging Technology and Security, March 2024) <https://cetas.turing.ac.uk/publications/future-biometric-technology-policing-and-law-enforcement> accessed 13 March 2025.

- people from minoritised ethnic groups, such as Black, Black British, Asian and Asian British people.

Recognising diverse lived experiences is crucial, as social conditions significantly influence how AI affects individuals and groups of people in different contexts. Privilege and disadvantage shape who can influence and benefit from AI systems. For instance, it is well documented that AI systems can encode a range of biases, including those relating to race, gender and ability.[14] These biases can range from facial recognition technology that classifies White male faces with more accuracy than darker-skinned women,[15] to algorithms that disadvantage women in recruitment.[16] If these inequalities are not addressed in evidence around AI, they risk becoming exacerbated. As AI continues to reshape sectors such as healthcare, law and social welfare, it is vital to include diverse voices in discussions about its future to prevent deepening societal divisions.[17]

1.3. Insights for regulation

Understanding public attitudes towards AI governance is essential for decision-makers shaping policies that support accountability, fairness and transparency and for determining when and how people expect to be protected from any negative impacts of AI. However, few large-scale studies have explored public preferences for AI regulation or the level of explainability expected in AI decision-making processes. Our research aims to fill these gaps by examining:

- mechanisms that increase people's comfort towards the adoption of AI
- concerns about AI-generated decisions and preferences for explainability versus accuracy in AI decision making
- concerns around AI safety and trust in different stakeholders in relation to regulation and governance
- concerns related to data sharing, and representation in decision making.

By exploring these factors, our research aims to inform policymakers, technology developers and regulators on how AI can be developed and governed in ways that reflect societal values and public preferences.

Ultimately, this study contributes to building the evidence base on diverse public attitudes towards distinct AI applications and their regulation.

[14] Meredith Broussard, More than a Glitch: Confronting Race, Gender, and Ability Bias in Tech (MIT Press, 2023).
[15] Joy Buolamwini and Timnit Gebru, 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification', Conference on fairness, accountability and transparency (2018).
[16] Reuters, 'Insight - Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women' (2018) <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G/> accessed 17 March 2025.
[17] Reema Patel (n5)

2. Methodology

This section summarises the study's methodology. For a more detailed account of how we designed our survey and our sampling approach, please refer to our technical report.[18] The survey was approved by the Ethics Committee at The Alan Turing Institute, UK (approval number: 24091908). The survey materials and data will be available as open access resources on publication.

2.1. Sample

The sample was drawn from the National Centre for Social Research Opinion Panel. This is a standing panel of people who have been recruited based on random probability design, meaning they were selected at random to be invited onto the panel. For this survey, a random subsample of 5,650 panel members was invited to take part. Fieldwork ran from 25 October 2024 until 24 November 2024. We achieved a 62% response rate with a final sample of 3,513 participants. The majority of respondents (94%) completed the questionnaire online. Two hundred and twenty-two people were interviewed by telephone either because they do not use the internet or because this was their preference.

We were interested in exploring in more depth the views and experiences of people in the following groups:

- **Individuals on lower incomes, measured as those with an equivalised monthly household income of £1,500 or less.** When asked about their perception of their financial status, the majority of people in this group felt they were either finding things difficult financially or just about getting by.
- **Digitally excluded populations, measured as those with low digital skills.** We used an adapted measure from Lloyds' Consumer Digital Index[19] to understand digital skills. This captures whether participants can perform a range of digital tasks under the broad skills of 'managing information', 'communicating', 'transacting', 'problem solving' and 'creating'. Those who could not do at least one task under each skill were classed as having low levels of digital skills. Many of those interviewed over the phone (65%) fell into this group.
- **People from Black/Black British and Asian/Asian British ethnic backgrounds.**

The focus on these groups was premised on the fact that the UK is made up of many publics that are differentially impacted by new technologies, and that some people and groups are frequently missing from existing research. Some applications of AI technologies bring up unique concerns and attitudes from those living with lower incomes,[20] minoritised ethnic

[18] The technical report can be accessed at: https://github.com/AdaLovelaceInstitute/wave-2---how-do-people-feel-about-AI-

[19] Lloyds Bank, 'UK Consumer Digital Index' (2018) <https://www.lloydsbank.com/assets/media/pdfs/banking_with_us/whats-happening/LB-Consumer-Digital-Index-2018-Report.pdf> accessed 13 March 2025.

[20] Ada Lovelace Institute, 'Access Denied? Socioeconomic Inequalities in Digital Health Services' (2023) <https://www.adalovelaceinstitute.org/wp-content/uploads/2024/07/Ada-Lovelace-Institute-Access-denied.pdf> accessed 3 February 2025.

groups[21] and those with fewer digital skills.[22] Typically, it is difficult to obtain large enough sample sizes for subgroup analysis for these populations in a nationally representative survey alone, a limitation we note in the previous iteration of this survey. We therefore oversampled based on these subgroups, while recognising that people have intersecting identities and often belong to more than one identity group.

Our final sample consists of: 433 (12% of the overall sample) Asian or Asian British participants, 198 (6% of the overall sample) Black or Black British participants, 1,319 (38%) low-income participants and 962 (27%) low digital skills participants. Respondents were all over the age of 18. Unweighted, a total of 1,875 (53%) were female, 1,632 (46%) male, with no sex recorded for six participants. Unlike the previous wave, which looked at participants in England, Scotland and Wales, this survey covered participants across the four nations of the United Kingdom: 2,937 (84%) were from England, 156 (4%) from Northern Ireland, 255 (7%) from Scotland and 165 (5%) from Wales. These proportions are broadly in line with UK population estimates across the four nations.[23]

The data was weighted based on official statistics to match the demographic profile of the UK population and to adjust for sampling probabilities used in the sampling process and non-response to this survey. All figures reported in this study, unless otherwise specified, are adjusted according to this weighting.

2.2. Survey

We told respondents that each of the technologies in this survey uses AI to varying degrees. We provided the following definition of AI:

'Artificial intelligence (AI) is a term that describes the use of computers and digital technology to perform complex tasks commonly thought to require human reasoning and logic. AI systems typically analyse large amounts of data to draw insights or patterns and achieve specific goals. They can sometimes take actions autonomously, that is without human direction. These systems can also be used to generate content like text, images, music or videos.'

The survey first covered general attitudes to new technologies, the ability to carry out certain digital tasks, and access to a smartphone and mobile data. We then asked about awareness of specific uses of AI, experiences with select uses, how beneficial and concerning respondents perceived each use of AI to be, and the key risks and benefits they associated

[21] Ada Lovelace Institute, 'The Citizens' Biometrics Council' (2021) <https://www.adalovelaceinstitute.org/wp-content/uploads/2021/03/Citizens_Biometrics_Council_final_report.pdf> accessed 3 February 2025.
[22] Joseph Rowntree Foundation, 'AI shifts the goalposts of digital inclusion' (JRF, February 2024) <https://www.jrf.org.uk/ai-for-public-good/ai-shifts-the-goalposts-of-digital-inclusion> accessed 13 March 2025.
[23] 'Population Estimates for the UK, England, Wales, Scotland and Northern Ireland' (Office for National Statistics, October 2024) <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/bulletins/annualmidyearpopulationestimates/mid2023#the-population-of-the-uk> accessed 13 March 2025.

with each. We also measured preferences around the governance and regulation of AI technologies, including different elements of decision making, safety, and data sharing and representation.

The specific technologies we asked about were:

- Risk and eligibility assessment technologies and facial recognition:
    - Facial recognition for policing and surveillance
    - Assessing eligibility for welfare benefits
    - Assessing risk of cancer from a scan
    - Assessing risk of repaying a loan
- LLMs and mental health chatbots
    - General-purpose LLMs
    - Mental health chatbots
- Robotics
    - Robotic care assistants
    - Driverless cars

Respondents were asked about all of the above AI use cases. The survey questions can be accessed in our technical report.

2.3. Analysis

We analysed the data between December 2024 and March 2025, using descriptive analyses for all survey variables followed up with two-proportion z-tests for different demographic groups. We then used regression analyses to understand relationships between demographic and attitudinal variables, and perceived benefit of specific technologies (see Appendix 5.1 for further information).

We analysed the data using the statistical programming language R, and used a 95% confidence level to assess statistically significant results. Analysis scripts and the full survey dataset can be accessed on the Ada Lovelace Institute's GitHub site.[24]

In this report, we generalise from a nationally representative sample of the UK population to refer to the 'UK public'. This does not refer to UK nationals, but rather people living in the UK at the time the survey was conducted.

2.4. Limitations

We recognise that our approach has limitations that need to be considered when interpreting our findings. In line with the objectives of the Public Voices in AI programme, and recognising that minoritised groups are often missing from nationally representative survey data, we set out to design a survey that provided information about these groups.

[24] The technical report can be accessed at: https://github.com/AdaLovelaceInstitute/wave-2---how-do-people-feel-about-AI-

We were limited in terms of the sample we were able to access, and this limitation is representative of the broader ecosystem of UK survey providers. We assessed nine separate providers for their ability to deliver a survey that would provide enhanced information about minoritised groups and chose the provider who was best able – through an existing recruitment panel – to access sufficient numbers of people from different backgrounds and offer robust sampling methodologies.

Based on sample availability within the recruitment panel we used, we were able to boost to some extent the number of people from minoritised ethnic backgrounds, on lower incomes and with fewer digital skills. However, we were only able to boost sample sizes for Asian/Asian British and Black/Black British populations sufficiently to enable some subgroup analysis. This limitation was due to both sample availability and budget.

While these sample sizes are greater than those in other surveys, they do not represent the diversity of the UK population. Where they do represent specific minoritised groups, they are broad categories that obscure the diversity within these groupings, for example, Bangladeshi Asian, Pakistani Asian or Chinese Asian experiences. We recognise that using broad ethnicity groupings can risk homogenising the experiences of distinct communities, but we were not able to produce a sufficient sample to accurately represent the attitudes of those people and communities.

Furthermore, while our target was to achieve at least 380 responses for each minoritised ethnic group to enable subgroup analysis, we were only able to achieve this for the Asian/Asian British demographic group. This was due to limitations of the recruitment panel, which had only 353 Black/Black British people registered. Through issuing the survey to the entire population of Black/Black British people, we were able to achieve a sample of 198 Black/Black British respondents. These factors reflect the challenge in surveying minoritised groups, despite intent to design surveys for equity and inclusion and produce data that relates to these groups.

To avoid limiting sample sizes further, we opted to avoid any routing in the survey, meaning all participants saw all of the questions. This meant we had to make trade-offs in terms of survey coverage and we were unable to explore as wide a range of AI applications as we had previously in our 2022/23 study. It also meant we were unable to explore in depth experiences of specific technologies.

Due to limited panel information on key factors impacting digital exclusion, such as access to digital goods, affordability and digital skills, we had to adopt a narrow definition of digital exclusion in our sample. Recruitment was based on a basic measure of internet access, including individuals with no internet access, those who reported using the internet less than once a week, and those who used it weekly but either participated more in phone surveys than online or had not provided an email address.

Our analysis of this group considered a more comprehensive measure of digital skills, as mentioned above. But because the sample composition itself did not fully capture the true extent of digital exclusion, it is likely that individuals with fewer digital skills are underrepresented in our findings, limiting the generalisability of insights into the digitally excluded population.

Although we recognise the importance of considering multiple and intersecting identities, the survey has not engaged deeply with intersectionality as a framework for analysis to understand how the intersection of multiple identities and systems of oppression may impact on experiences of, and attitudes towards, AI technologies. This was not only because of limitations in sampling minoritised groups and reaching targeted sample sizes, but also because our research questions focused primarily at a general population level, supplemented by exploratory subgroup analysis across sociodemographic factors. As such, our findings offer insight into distinct experiences of some communities at a broad level.

To overcome these limitations would require enhanced provision in the UK survey ecosystem, as well as sufficient resource to incentivise participation from a range of minoritised people and groups.

Finally, due to differences in sample composition between this survey and its first iteration, comparative results need to be interpreted with caution. This survey follows a cross-sectional design rather than a longitudinal one, meaning it examines different cross-sections of the public in each wave rather than surveying the same people over time. Moreover, the current survey iteration oversampled based on specific demographics, as detailed above. Weighting has been applied to make the sample representative of the UK public, and where comparisons across survey iterations have been made, they compare a nationally representative Great British public (Wave 1) with a nationally representative UK public (Wave 2) to enable as much comparability as possible. But these comparisons are indicative of trends, and not conclusive evidence of changes over time.

3. Key findings

3.1 Awareness and experience of AI uses

To assess public awareness of and experiences with AI technologies, we asked participants whether they had previously encountered each AI application. Additionally, we asked about their personal experience with general-purpose large language models (LLMs), such as ChatGPT, Gemini, Claude and Llama, among others, and mental health chatbots. Given the rapid rise of generative AI and its increasing prominence in the public domain, we felt it was important to explore direct user experience with these two emerging technologies. Direct experience was not explored for other AI uses, as people would probably find it difficult to say whether they had experienced them or not (e.g., AI used to support decisions around receiving a loan).

3.1.1. Awareness of AI varies substantially depending on its specific application

Overall, awareness of AI technologies varies according to its specific use. For three out of the eight AI uses we asked about, more than 50% of the UK public said they had heard of them before. Figure 1 shows levels of awareness for each of the eight AI uses.

Figure 1. Awareness of uses of AI technology

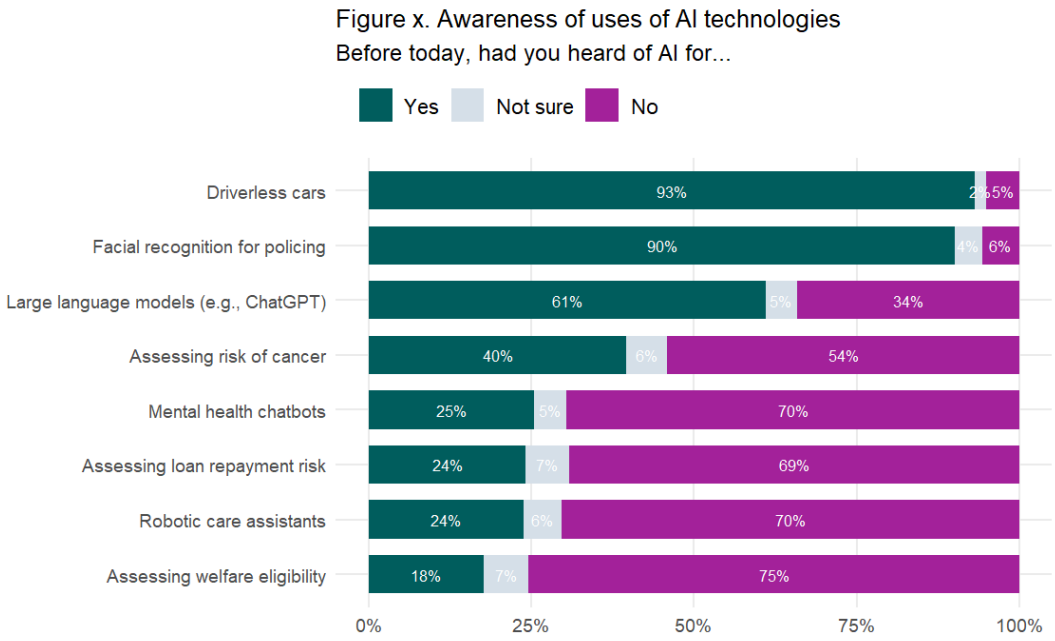

Awareness is highest for driverless cars and the use of facial recognition for policing, with 93% and 90% of the public, respectively, reporting having heard of these technologies before. People are least aware of the use of AI for assessing eligibility for welfare benefits (e.g. Universal Credit), with just 18% having heard of this before. Similarly, people are less aware of other risk and eligibility technologies, such as using AI to assess how likely a person is to repay a loan such as a mortgage, with only 24% aware of them. These results reflect trends similar to those found in our survey in 2022/23, suggesting that public awareness for risk and eligibility technologies has not significantly increased, even as these technologies become integrated into public services[25] and are therefore likely to be impacting large numbers of people.

LLMs and mental health chatbots were a new inclusion in this wave of the survey. Most people (61%) are aware of general-purpose LLMs, an application of AI that has been widely discussed in the media ever since the launch of ChatGPT. This aligns with existing survey-based research, which found that 58% of the UK public have heard of ChatGPT specifically.[26] In contrast, awareness of mental health chatbots is low, with only one in four people (25%) reporting having heard of this application of AI before.

3.1.2. Personal experience with LLMs suggests increasing trends in adoption for everyday tasks

[25] Bright and others (n 4).
[26] Richard Fletcher and Rasmus Kleis Nielsen, 'What Does the Public in Six Countries Think of Generative AI in News?' (Reuters Institute for the Study of Journalism, 2024) <https://reutersinstitute.politics.ox.ac.uk/what-does-public-six-countries-think-generative-ai-news> accessed 13 March 2025.

Personal experience with general-purpose LLMs is mixed. In terms of frequency, people reported using LLMs a few times rather than regularly. The most popular use is searching for answers and recommendations, with a third (33%) of the UK public indicating they have used these technologies at least a few times. This is followed by educational purposes (21%) and everyday tasks such as writing emails (21%). Two-fifths (40%) of the UK public have used LLMs for one or more of the tasks we asked about.

When compared with existing research, these figures suggest an upward trend in the use of LLMs. For example, a 2024 survey across six countries including the UK found that on average 24% of people used generative AI tools for getting information, 9% used it for writing emails and 8% used it for educational purposes.[27] The usage figures in our survey may even be a conservative estimate as it is possible that some people in our sample have used AI without being aware of it, due to existing integration of AI in some search engines. As these tools become more integrated in search engines, we can expect usage to increase.

As opposed to everyday tasks, Figure 2 shows that few people have used general-purpose LLMs for entertainment purposes (14%), supporting job applications (11%) or guidance on issues such as legal disputes or taxation (8%). A considerable proportion of the UK public are also closed off to using LLM-based AI tools for some of the applications we presented. This was most prevalent for supporting job applications, where 39% of people would not want to use LLMs for this.

Those with fewer digital skills and those on lower incomes are slightly more likely to be closed off to the use of LLMs for all the tasks we asked about than those with higher levels of digital skills and those on higher incomes, with this difference being statistically significant.[28] For example, of those not open to using general-purpose LLMs for supporting job applications, 27% do not have basic digital skills and 39% have low incomes (equivalised monthly household income of £1,500 or less). This is in contrast with those that have used LLMs for supporting job applications, or are open to using them for this, where only 16% do not have basic digital skills and 33% have low incomes.

This limited adoption might be related to a range of reasons. First, it might be indicative of an apprehension towards using a general tool for a specialised task, such as getting legal guidance, suggesting personal red lines in terms of in which context AI tools are deemed appropriate. People may feel, for example, that some tasks require human expertise. Second, it may relate to concerns around access and opportunity (for low-income or digitally excluded groups), with some feeling apprehensive about the role of emerging technologies in domains such as legal advice or job applications and/or their ability to use these tools. While our work offers important preliminary insights into public experiences with general-purpose LLMs, we are unable to unpack reasons for limited adoption with our data due to limitations of survey length. Future research should track public experiences with such emerging technologies in more detail.

[27] ibid.

[28] Based on logistic regression analyses predicting whether individuals had used LLMs for each of the tasks provided, or were open to using it, versus those that had not and were not open to using it. Predictors included our boosted demographic groups: Asian participants, Black participants, low-income participants, and those with fewer digital skills.

Figure 2. Experience with large language models

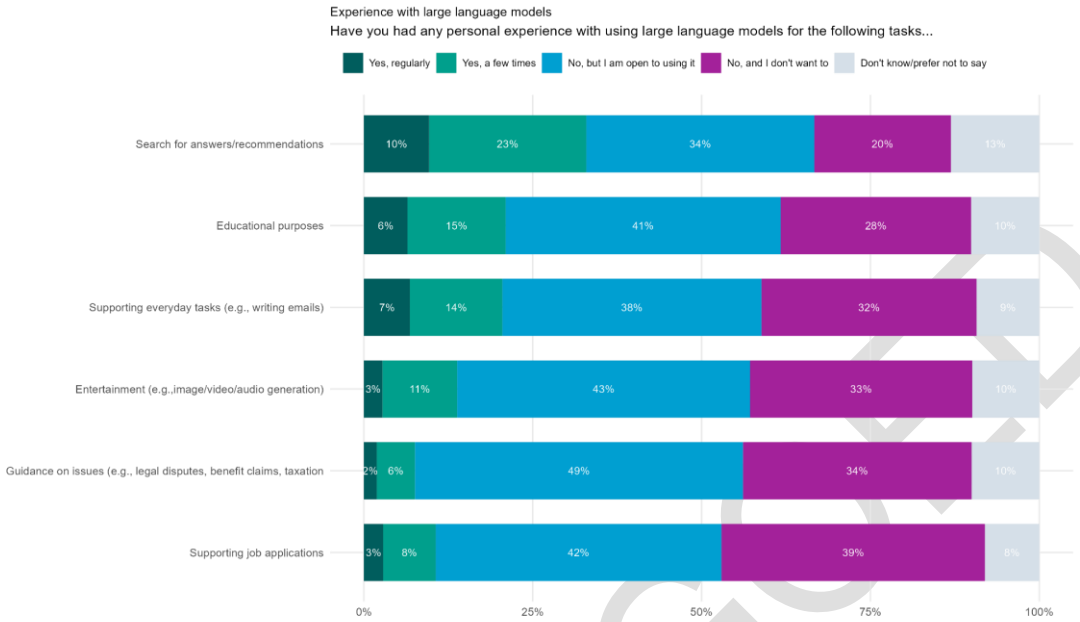

3.1.3. Seven per cent of the public have used a mental health chatbot

We also asked people about their personal experience with mental health chatbots. This was described as a tool that is usually developed by private companies and offered to the public sometimes at a cost, either online or via mobile applications. The chatbots were described as being able to respond to the emotions expressed during an individual's interaction with it to offer mental health support or advice.

Seven per cent of the UK public have at least some personal experience with using a mental health chatbot. Given the relatively niche nature of this AI application, this figure is small but substantial: in real terms it represents approximately two million people in the UK in absolute numbers.[29] As foundation models continue to evolve, mental health chatbots warrant further investigation to understand their potential impact,[30] particularly as these tools are currently available freely online and – as an emerging use case – are not yet subject to specific regulatory oversight.

[29] Based on the Office for National Statistics' mid-year 2023 population estimate of 68.3 million in the UK.

[30] Ada Lovelace Institute, 'Delegation Nation' (2025) <https://www.adalovelaceinstitute.org/policy-briefing/ai-assistants/> accessed 5 February 2025.

3.2. Perceptions of benefit for AI technologies and levels of concern

We asked people to indicate the extent to which they felt each technology in our survey would be beneficial, and separately the extent to which they were concerned by each technology. Overall, we found that the public holds nuanced views about AI, seeing both the benefits and risks associated with different applications.

3.2.1. Perceptions of benefit for AI are high for some applications in diagnostic health and justice

The UK public has high expectations for some AI technologies. In particular, the majority of the public perceive facial recognition in policing (91%) and AI-driven risk assessment for cancer (86%) to be beneficial uses of AI. A majority of the public (63%) also have positive views about LLMs (e.g. ChatGPT).

Expectations of positive impact are lower for other uses of AI, as shown in Figure 3. While optimism around the use of AI in diagnostics for cancer is high, the same is not true for the application of AI in other areas of healthcare, with only 36% of the public perceiving mental health chatbots to be beneficial, and 55% perceiving robotic care assistants to be beneficial. As with our previous survey wave, the only public sector application of AI in our current survey – the use of AI for assessing eligibility for welfare benefits – was also viewed less positively than others, with less than half (47%) of the public perceiving it as beneficial. Similarly, only 45% perceive driverless cars to be beneficial.

Figure 3. The extent to which each AI use is perceived as beneficial

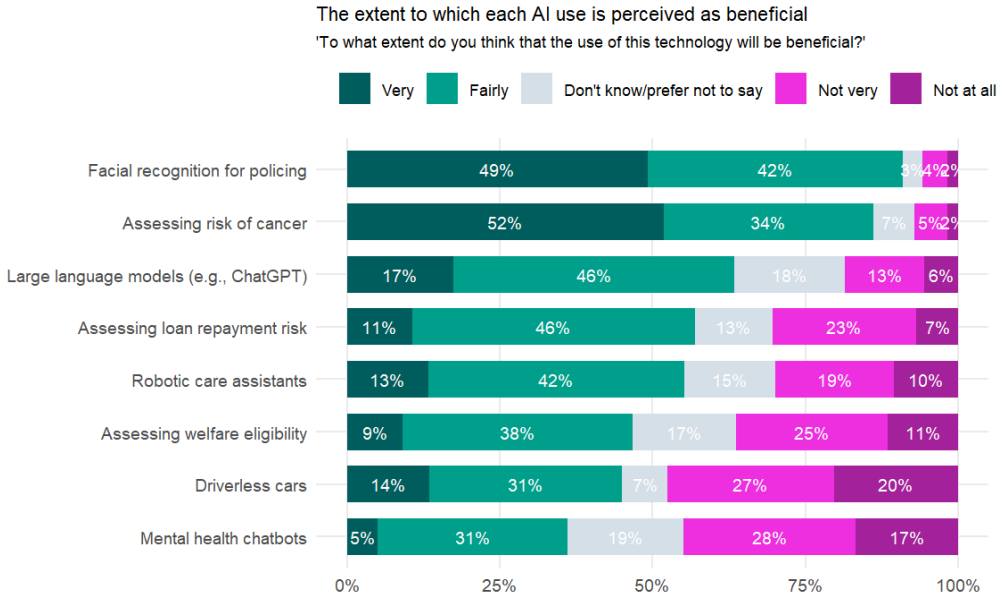

3.2.2. Concerns around AI are substantial, even when expectations of positive impact are high

Overall, the UK public are most concerned about the application of AI in driverless cars (75%), mental health chatbots (63%) and assessing welfare eligibility (59%). Figure 4 shows concern levels for each AI technology. Even for technologies where expected positive impact is high, concern levels are also substantial. For example, nearly two-fifths (39%) of the public are concerned by the use of facial recognition in policing.

Figure 4. The extent to which each AI use is perceived as concerning

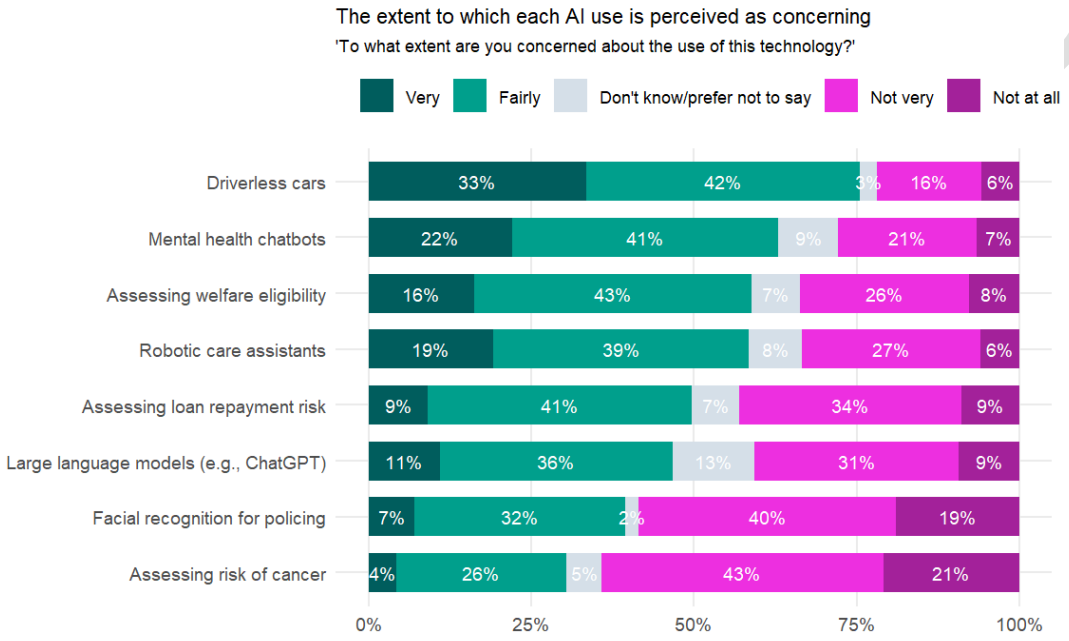

3.2.3. Different demographic groups have distinct attitudes to applications of AI

We observed key demographic differences in the perceived benefits of each AI technology. Black/Black British and Asian/Asian British people are more likely than the national average to view applications of robotics (driverless cars and robotic care assistants), LLMs and mental health chatbots as beneficial. The demographic difference in perceived benefits is most notable for general-purpose LLMs, where 80% of each minoritised ethnic group perceives them as beneficial, compared to 63% of the general population.

People on lower incomes and those with fewer digital skills are less likely than the general public to perceive nearly all of the AI technologies we asked about as beneficial. For instance, only 48% of people on lower incomes felt robotic care assistants could be beneficial compared to 55% of the general population. Among those with fewer digital skills, 41% felt LLMs could be beneficial compared to 63% of the general population.

Figure 5 highlights demographic variances in overall levels of perceived benefit for each technology.

Figure 5. The extent to which each AI use is perceived as beneficial: demographic analysis

We also observed key demographic differences in perceptions of concerns. Over half of all Black (57%) and Asian (52%) people in our sample reported being fairly or very concerned about the use of facial recognition in policing, compared to 39% of the general public. The top three concerns they reported included: 1) the gathering of personal information which could be shared with third parties (for 59% of Black people and 60% of Asian people); 2) causing police to rely too heavily on technology rather than their professional judgment (for 66% of Black people and 58% of Asian people); and 3) the risk of innocent people being wrongly accused if the system makes mistakes (for 62% of Black people and 61% of Asian people). Black and Black British people are also more likely to report concerns around the use of AI to determine eligibility for welfare – 71% find this use of AI very or fairly concerning, compared to 59% of the nationally representative cohort.

For some applications of AI, concern levels are lower among oversampled subgroups. Individuals belonging to low-income groups (42%) and those with low digital skills (38%) are significantly less concerned by general-purpose LLMs compared to the national average (47%). This is in line with their perception of the benefit of these technologies. Those with low digital skills (55%) are also less concerned by mental health chatbots compared to the nationally representative cohort (63%). And Asian and British Asian people (49%) are less concerned about robotic care assistants than the average (58%).

Figure 6 highlights demographic variances in overall levels of concern for each technology.

Figure 6. The extent to which each AI use is perceived as concerning: demographic analysis

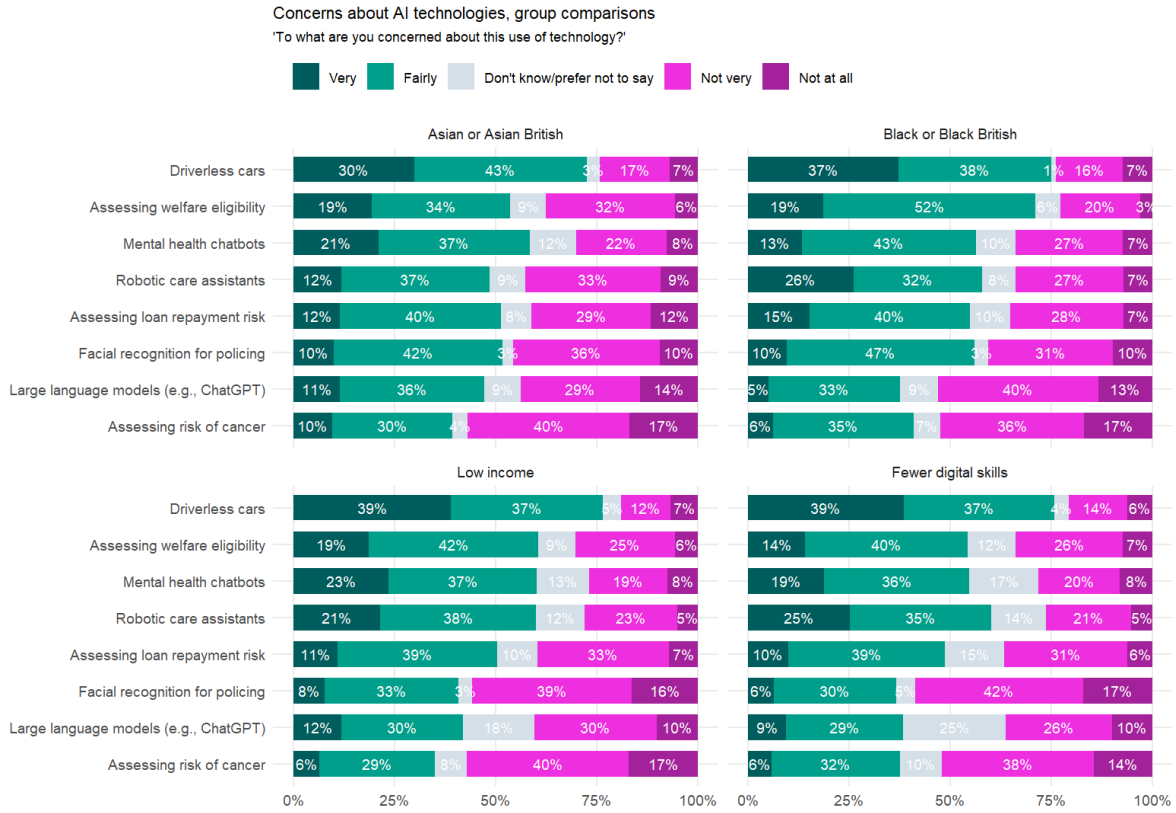

The findings above highlight that attitudes to AI are multifaceted in several aspects. First, AI is not considered to be a single entity, with perceptions varying towards specific applications. For example, the public has positive attitudes towards general-purpose LLMs, with high overall benefit scores and low levels of reported concern, and more negative attitudes towards mental health chatbots. This shows how the context each technology is applied to matters.

Second, people can simultaneously perceive benefits and risks associated with different applications of AI. For each of the applications of AI we asked about, people reported differential levels of both perceived benefit and concern that are not mirror images of each other.

Third, minoritised demographic groups perceive differential levels of benefits and concerns for each AI application, as we explore above. For instance, people on lower incomes have higher levels of concern for many of the applications of AI we asked about compared to the national average. It is therefore important to consider the views of diverse publics when trying to understand public sentiment towards different applications of AI.

3.2.4. While perceptions of beneficial impact have remained stable, overall concern around AI uses has increased since 2022/23

When comparing responses across both waves of our survey, perceptions of benefit have remained relatively stable (except in the case of facial recognition for policing, where perceptions of benefit have increased slightly) while levels of concern have significantly increased across all applications of AI. However, it is important to note that the comparisons

should be read with caution as the sample composition across the two waves is different (refer to the discussion in Section 2.4).

Figures 7 and 8 show a comparison of perceptions of benefit and concern across both surveys for all repeated uses of AI. For example, in 2022/23, 44% of the public were concerned by the use of AI for determining welfare eligibility. This has increased to 59% in 2024/25.

Figure 7. The extent to which each AI use is perceived as beneficial: survey wave comparison

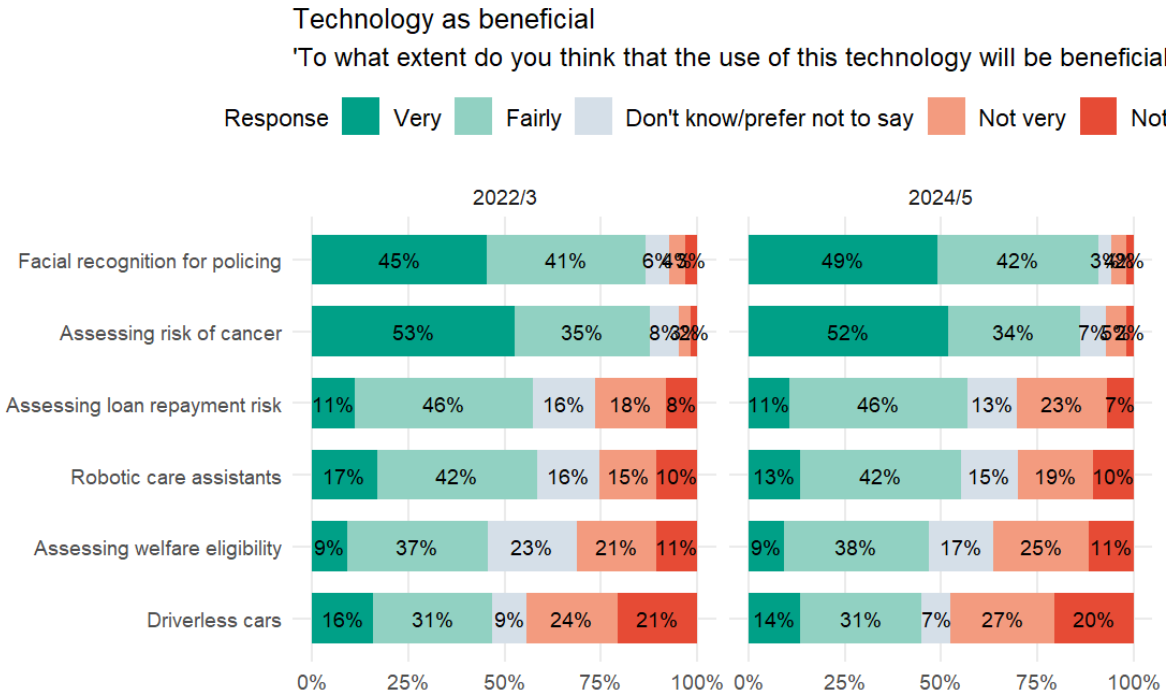

Figure 8. The extent to which each AI use is perceived as concerning: survey wave comparison

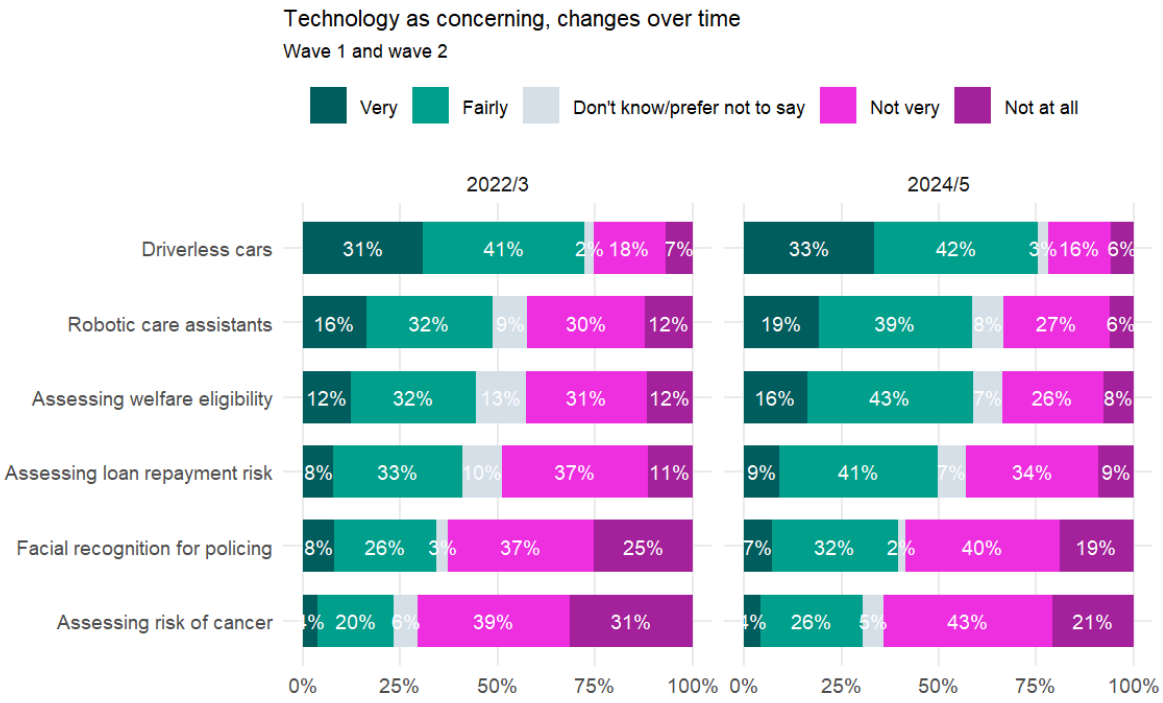

To further understand the relationship between perceptions of benefit and concern around each AI technology, we created a net benefit score for each AI use by subtracting the extent to which each person indicates the AI use was concerning to them from the extent to which they indicate the AI use was beneficial. Positive scores show that perceived benefit outweighs concern, while negative scores show concern outweighs perceived benefit. Scores of zero indicate equal levels of concern and perceived benefit (Figure 9).

We found that for four out of the eight AI uses, perceived benefit outweighs concern: assessing risk of cancer from a scan, facial recognition for policing, LLMs and assessing loan repayment risk. For the remaining four uses, concern outweighs perceived benefit: robotic care assistants, assessing welfare eligibility, mental health chatbots and driverless cars. Looking across both waves of the survey, we find a declining trend in net benefit scores.

For the two new technologies introduced in this wave, general-purpose LLMs like ChatGPT received a positive net benefit score (0.38), while mental health chatbots were viewed negatively (-0.56). As mentioned previously, this might indicate that, while the public sees some value in general-purpose LLMss, there is more hesitation about their application in sensitive areas like mental health support. However, it is important to note that we defined mental health chatbots as those *potentially* powered by LLMs, and people might not be sufficiently aware of this distinction, so this interpretation remains speculative.
The use of AI for assessing risk of cancer and for facial recognition in policing continue to retain high net benefit scores (1.34 and 1.17, respectively), as in the previous survey wave. However, both scores have declined (from 1.58 and 1.23, respectively). AI applications in welfare and driverless cars continue to face scepticism, with net negative scores as in the previous survey wave.

Notably, people are more concerned about robotics than other technologies, and this concern has increased over the last two years. Perceptions of driverless cars have become more negative overall compared to the previous survey wave. And perceptions around robotic care assistants have dipped into negative territory (-0.07), suggesting increased hesitation about their role in caregiving.

The decrease in net benefit scores for risk and eligibility technologies, such as those for welfare benefits and loan repayment, might be indicative of growing concerns about their fairness or effectiveness (explored in the next section).

Figure 9. Net benefit scores

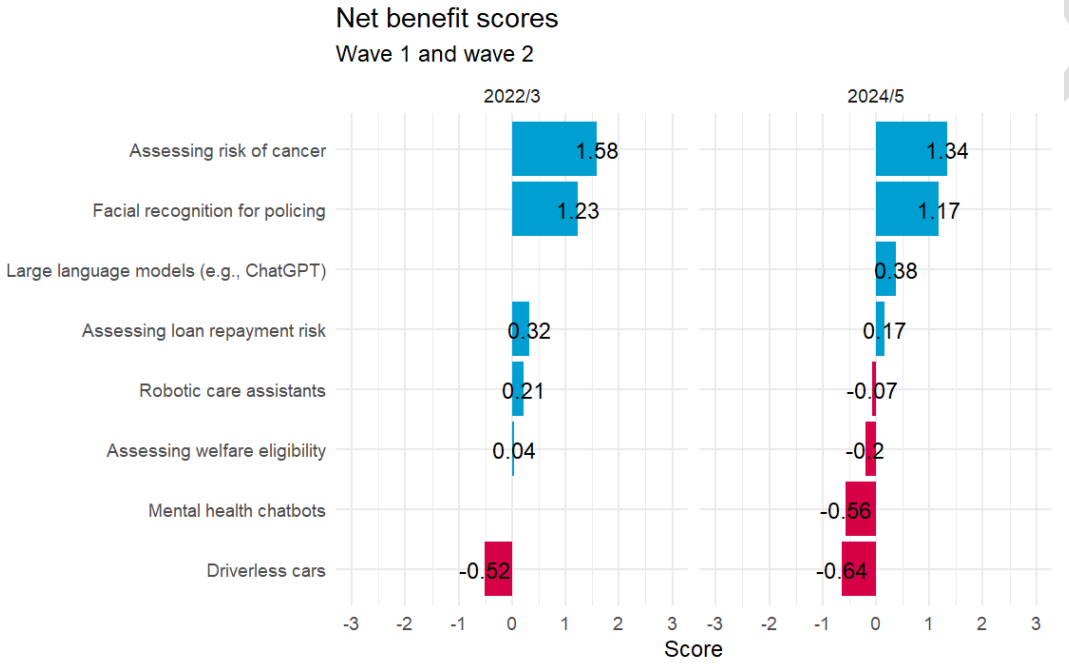

3.2.5. Low-income groups in particular are more likely to feel their concerns outweigh perceived benefits of AI

To understand in more detail the association of specific demographic characteristics on overall attitudes to each application of AI, we conducted a regression analysis examining the extent to which an individual's income, digital skills, awareness of each AI application, age, gender and education level predicts their 'net benefit' score for each technology.

We found that within these factors, income status seems to be driving attitudes. When all other variables are held constant, those on low incomes still have significantly lower net benefit scores than those with higher incomes. This finding suggests that being on a low income may be linked with less acceptance of AI technologies. This could be due to concerns around accessibility, fairness and potential biases in decision making that could impact financial stability – such as through determining eligibility for welfare benefits or loans. It presents a case for understanding in more detail the concerns those on lower incomes have of these technologies and whether and how these technologies can be designed to benefit them.

Appendix 5.1 provides more information about the analyses outlined in this section, including further results showing the effects of demographic and attitudinal differences on perceived net

benefit for each technology.

3.3. Specific benefits and concerns for each AI technology

To further understand some of the underlying reasoning behind the attitudes represented at a general level in the previous sections of this report, we asked respondents to select the specific benefits and concerns they see for each technology, from a multiple choice list. The benefits and concerns included in each list reflected common themes such as efficiency, accuracy, bias and accountability, though each list was specific to each technology. Participants could select as many statements from each list as they felt appropriate, with 'something else', 'none of the above' and 'don't know' options also given for each technology.

Overall, people most commonly identified benefits related to improvements in speed and efficiency of decision making or support, and most commonly express concerns related to overreliance on technologies over professional human judgement, mistakes and a lack of transparency in decision making. Tables 1–6 show the three most commonly chosen benefits and concerns for each technology. A full list of benefits and concerns presented to participants and the percentage of people selecting each can be found in Appendix 5.3. We cluster these by categories of technologies for risk and eligibility assessments and facial recognition, LLMs and mental health chatbots, and robotics.

3.3.1. Risk and eligibility assessments and facial recognition

We asked about the following uses of assessing eligibility and risk using AI: to assess eligibility for welfare benefits, to predict the risk of cancer from a scan, to predict the risk of not repaying a loan, and facial recognition for policing and surveillance.

Table 1: Most commonly selected benefit (Will be an infographic)

Technology	Benefit 1	Benefit 2	Benefit 3
Assessing risk of cancer from a scan	Earlier detection of cancer (85%)	Less human error (64%)	More accurate than doctors (46%)
Assessing loan repayment risk	Faster and easier (58%)	Less likely to discriminate (44%)	Less human error (41%)
Assessing welfare eligibility	Faster and easier (52%)	Save money (43%)	Less human error (39%)
Facial recognition for policing	Faster (89%)	More accurate than professionals (66%)	Save money (51%)

As mentioned above, speed and accuracy commonly feature across perceived benefits for many AI technologies. For example, 85% of the UK public feel earlier detection of cancer is a potential benefit of AI tools that assess risk of cancer from a scan, and 66% of the public feel using facial recognition technologies in policing will be more accurate than police officers at identifying wanted criminals and missing persons. Alongside these, the public feel many AI tools may reduce mistakes made in carrying out the tasks we surveyed. This may be through reducing human error in decision making (e.g. 41% feel loan repayment risk tools will lead to less human error).

Concern around overreliance is most frequently reported for technologies that assess cancer risk (64%), technologies that assess welfare eligibility (60%), technologies that assess loan repayment risk (57%) and technologies that use facial recognition for policing (57%). Across nearly all of these uses of AI, transparency in decision making also features as a commonly reported concern.

Table 2: Most commonly selected concerns

Technology	Concern 1	Concern 2	Concern 3
Assessing risk of cancer from a scan	Overreliance on technology (64%)	Accountability for mistakes (50%)	Transparency in decision making (41%)
Assessing loan repayment risk	Accounting for individual differences (59%)	Overreliance on technology (57%)	Transparency in decision making (54%)
Assessing welfare eligibility	Overreliance on technology (60%)	Accounting for individual differences (60%)	Transparency in decision making (54%)
Facial recognition for policing	Overreliance on technology (57%)	Technology will gather and share personal information (56%)	False accusations (54%)

Table 2: Most commonly selected concerns

3.3.2. Robotics

We asked about the following uses of robotics: driverless cars and robotic care assistants. In the case of driverless cars, nearly two-thirds of the public (63%) feel that improvements to accessibility by making travel by car easier for people who have difficulty driving is a benefit. This highlights positive perceptions, and potentially high expectations, around AI making tasks easier for all of society. For robotic care assistants, approximately half of the public feel faster support (48%) and being less likely than humans to discriminate against some groups of people in society (37%) are key benefits.

Table 3: Most commonly selected benefits for robotics

Technology	Benefit 1	Benefit 2	Benefit 3
Driverless cars	Improve accessibility (63%)	Free up time to do other things (35%)	Result in fewer accidents (34%)
Robotic care assistants	Faster and easier (48%)	Less discriminatory (37%)	More effective than professionals (37%)

People are concerned about a lack of human interaction in AI technologies, the potential overreliance on the technology at the expense of human judgement, and issues of who to hold accountable when the technology makes a mistake. As with benefits, the concerns also vary

depending on where robotics are applied. The public are worried about losing human interaction in applications of AI delivering one-to-one care. For example, 82% of the public are concerned about people missing out on human interactions through the delivery of care via robotic care assistants**.** While the most popular concern for driverless cars relates to unreliability and the accountability when the technology makes a mistake.

Table 4: Most commonly selected concerns for robotics

Technology	Concern 1	Concern 2	Concern 3
Driverless cars	Unreliable (69%)	Accountability for mistakes (66%)	Transparency in decision making (57%)
Robotic care assistants	Loss of human interaction (82%)	Technology being unsafe (59%)	Job cuts (53%)

3.3.3. LLMs and mental health chatbots

We found similar perceptions as those described for other AI uses when we asked about LLMs and mental health chatbots. Improvements in efficiency and faster access to support are the most popularly selected benefits for general-purpose LLMs and mental health chatbots, respectively.

Table 5: Most commonly selected benefits for AI chatbots

Technology	Benefit 1	Benefit 2	Benefit 3
General-purpose LLMs (e.g. ChatGPT)	Efficient by automating tasks (56%)	Tool for learning and skills development (50%)	Enhance creativity (38%)
Mental health chatbots	Faster support (50%)	Improve accessibility (46%)	Feels human, preventing isolation (33%)

As with robotic care assistants, losing human interaction was a popular concern for mental health chatbots. 68% of the public are concerned about mental health chatbots leading to isolation by replacing human interactions. Specific to LLMs, the public are worried about the generation of biased content (50%) or harmful content (47%), as well as attrition in individual problem-solving skills (66%).

Table 6: Most commonly selected concerns for AI chatbots

Technology	Concern 1	Concern 2	Concern 3
LLMs (e.g. ChatGPT)	Reduce problem solving and critical thinking skills (66%)	Biased content (50%)	Generate harmful content (47%)
Mental health	Loss of human	Unclear that you are	Provide misleading

chatbots	interaction (68%)	not interacting with a human (63%)	advice (62%)

3.4. Governance and regulation

This section presents findings relevant to AI regulation and governance. We begin by examining the mechanisms that would make people more comfortable with the use of AI, followed by exploring three broad categories of public concerns with AI: decision making, safety, and data sharing and representation.

3.4.1. Laws and regulation increase most people's comfort with the use of AI

We asked respondents what, if anything, would make them more comfortable with the use of AI. Participants could select multiple options. Figure 10 presents a comparison between the 2022/23 and 2024/25 survey results, illustrating changes in public attitudes towards mechanisms that enhance comfort with AI technologies.

The majority of the public (72%) indicated that laws and regulations would increase their comfort with AI technologies – an increase from 62% in 2022/23. The second most commonly selected mechanism was the ability to appeal AI-generated decisions (65%), highlighting a strong public desire for avenues of redress in AI-driven decision making.

Figure 10. Mechanisms for increasing comfort with AI

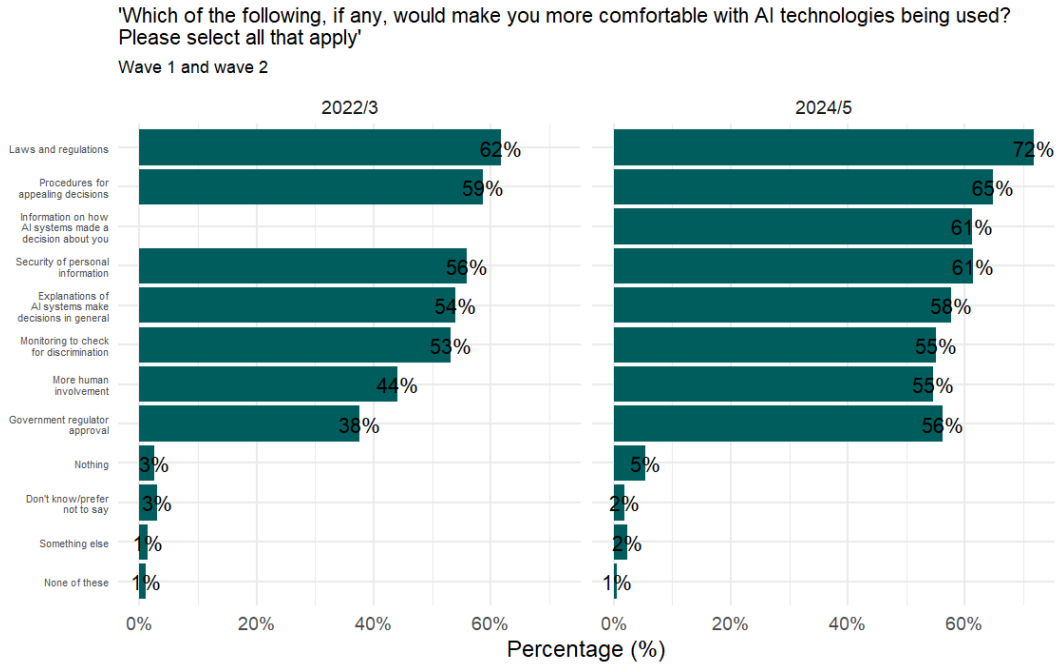

3.4.2. A majority of the public are uncomfortable with AI-based decision making, preferring explainability over accuracy

As highlighted previously, there is some latent discomfort with AI-generated decisions in the absence of an appeal mechanism. To explore this further, we examined people's comfort with AI-generated decisions that impact their lives, and the role that explainability of those decisions would play to allay their discomfort.

Two-thirds of the UK public (63%) are not comfortable with AI systems making decisions that affect their lives. In particular, those with fewer digital skills and lower incomes are slightly more likely than the nationally representative sample to report discomfort with automated decision-making systems, at 69% and 68% respectively, with this difference being statistically significant. Figure 11 shows overall how comfortable people are with AI technologies being used to make decisions that affect their lives.

Figure 11. Comfort with AI decision-making

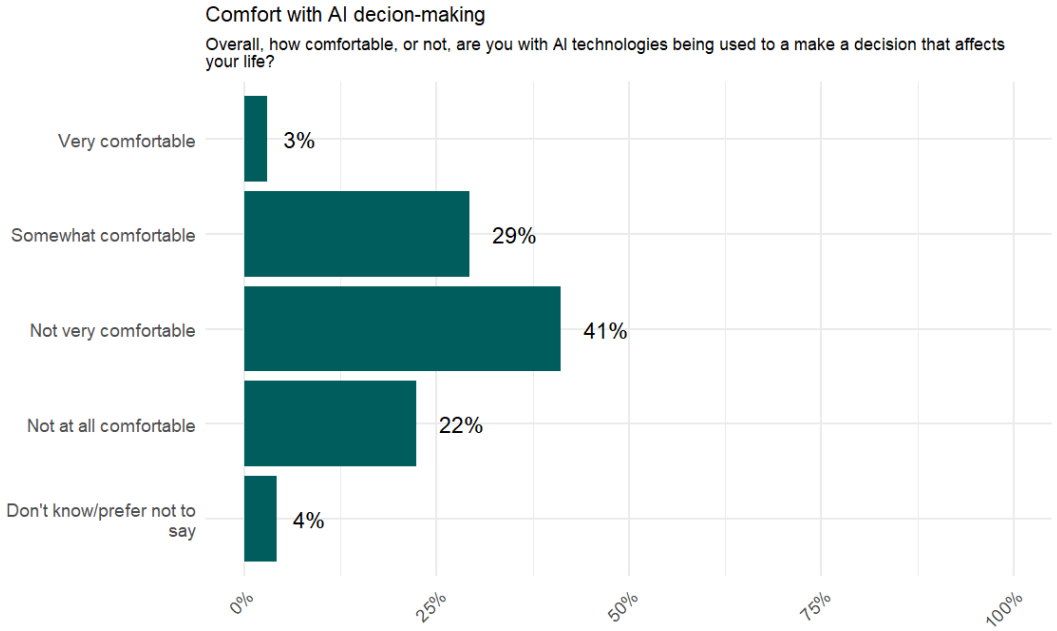

This discomfort sits alongside a preference for explanations to accompany decisions that affect people. To understand how the public make trade-offs between explanations accompanying AI decisions and the accuracy of these decisions, we informed participants that: 'Many AI systems are used with the aim of making decisions faster and more accurately than is possible for a human. However, it may not always be possible to explain to a person how an AI system made a decision.' They were then asked to consider a range of statements around automated decision making and trade-offs between accuracy and explanations to accompany those decisions.

The public have a strong preference for explanations over accuracy. 62% of people think an explanation should always accompany a decision, with 37% feeling humans, not computers, should be the ones making these decisions. Only 8% of people think accuracy is more

important than providing an explanation when it comes to automated decision making by an AI system. These findings are consistent with the previous survey wave, suggesting little change in the last two years in preferences for explanations over greater accuracy in automated decision-making systems. Figure 12 shows the distribution of responses when considering trade-offs between accuracy and explanations.

People's preferences for explainability over accuracy are different across age groups. Older people choose explainability and human involvement over accuracy to a greater extent than younger people. For those aged 18–34, 'sometimes an explanation should be given even if it reduces accuracy' was the most popular response (Figure 12). In contrast, for those aged 55 and above, the most popular response was 'humans should always make the decisions and be able to explain them'. This difference is also consistent with findings from our previous survey wave.

Figure 12. Trade-offs in accuracy and explainability

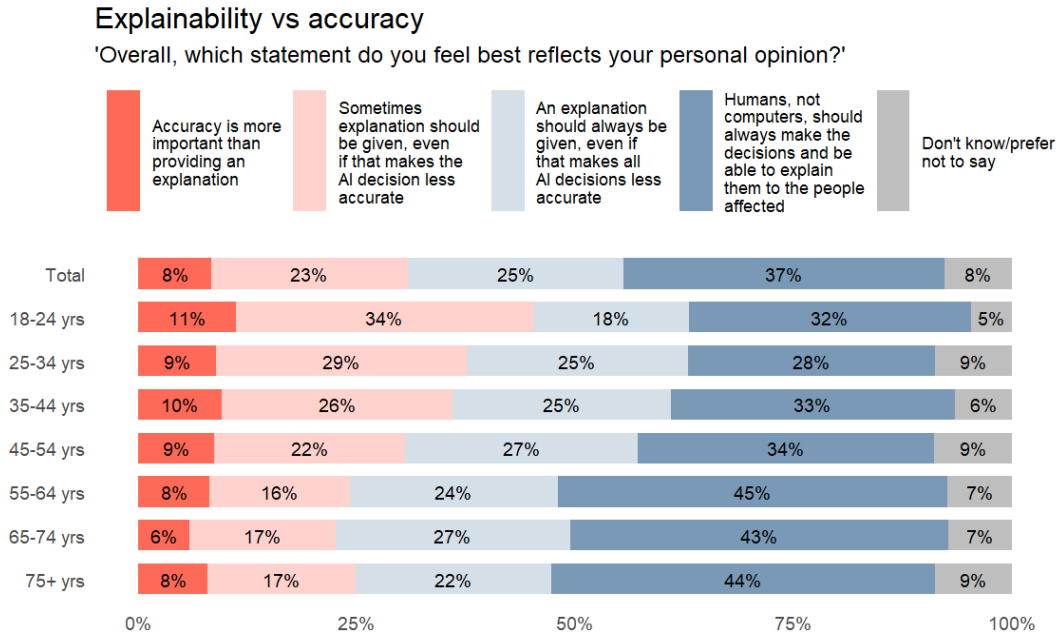

3.4.3. The public have had high exposure to AI-generated harms and strongly support shared public-private (rather than private-only) responsibility for AI safety

We asked the public about their experiences of the following types of online harms that may have been AI-generated: financial fraud or scams, deepfake images or video clips, false information, and content that promotes violence, abuse or hate. Figure 13 shows self-reported personal exposure to these harms.

On average, close to two-thirds of the UK public (67%) have experienced any form of harm a few times, while over a third (39%) have encountered any form of harm many times. Exposure to harm was highest for false information, with 61% of people having experienced this,

followed by financial fraud (58%), deepfakes (58%), and content promoting violence, abuse or hate (39%). However, many individuals are unsure if the harms they encountered online were AI-generated, with at least 20% reporting this for each of the harms we surveyed.

Exposure to these harms is associated with age. Individuals aged 18–24 were more likely than other age groups to report having experienced these harms, with 81% reporting exposure to false information and 85% having encountered deepfakes. In contrast, older age groups reported different patterns of exposure to these harms. Those aged 65–74 reported encountering financial frauds (57%) and false information (53%) more commonly than other AI-generated harms.

Men were more likely to report encountering online harms that were potentially AI-generated than women and this is statistically significant across harms. This aligns with research on public exposure to AI-generated harms such as deepfakes.[31] However, it is important to note that lower self-reported exposure among women could be due to their adaptive behaviours, including limiting their online engagement with others or limiting their own online engagement such as sharing photos or opinions online.[32] Such proactive behaviours, which may be driven by heightened fears about becoming a target of online harms,[33] might successfully reduce exposure to certain types of online harms, including AI-generated harms. However, as described in the limitations section, we do not have sufficient data to carry out detailed analysis.

[31] Tvesha Sippy and others, 'Behind the Deepfake: 8% Create; 90% Concerned. Surveying Public Exposure to and Perceptions of Deepfakes in the UK' (arXiv, 7 July 2024) <http://arxiv.org/abs/2407.05529> accessed 23 September 2024.
[32] Francesca Stevens and others, 'Women Are Less Comfortable Expressing Opinions Online than Men and Report Heightened Fears for Safety: Surveying Gender Differences in Experiences of Online Harms' (arXiv, 27 March 2024) <http://arxiv.org/abs/2403.19037> accessed 13 March 2025.
[33] Francesca Stevens and others, 'Women Are Less Comfortable Expressing Opinions Online than Men and Report Heightened Fears for Safety: Surveying Gender Differences in Experiences of Online Harms' (arXiv, 27 March 2024) <http://arxiv.org/abs/2403.19037> accessed 13 March 2025.

Figure 13. Experience of harms online

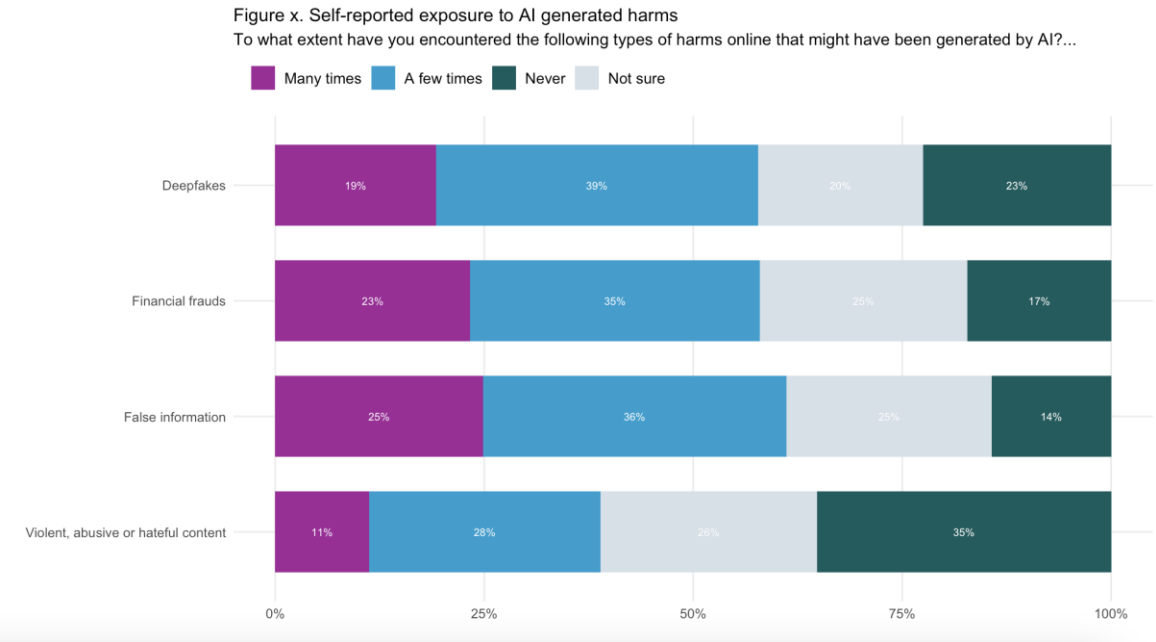

We also asked the public about the extent to which they were concerned about the spread of harmful AI-generated content online. Figure 14 shows self-reported concern about AI-generated harms. An overwhelming majority (94%) of the UK public said that they were either very or somewhat concerned about the spread of such harms online. Given these safety risks and concerns, it is important to understand public expectations around the regulation and governance of AI. We asked respondents who they think should be responsible for the safe use of AI and what specific powers they should have.

Figure 14. Self-reported concern about AI-generated harms

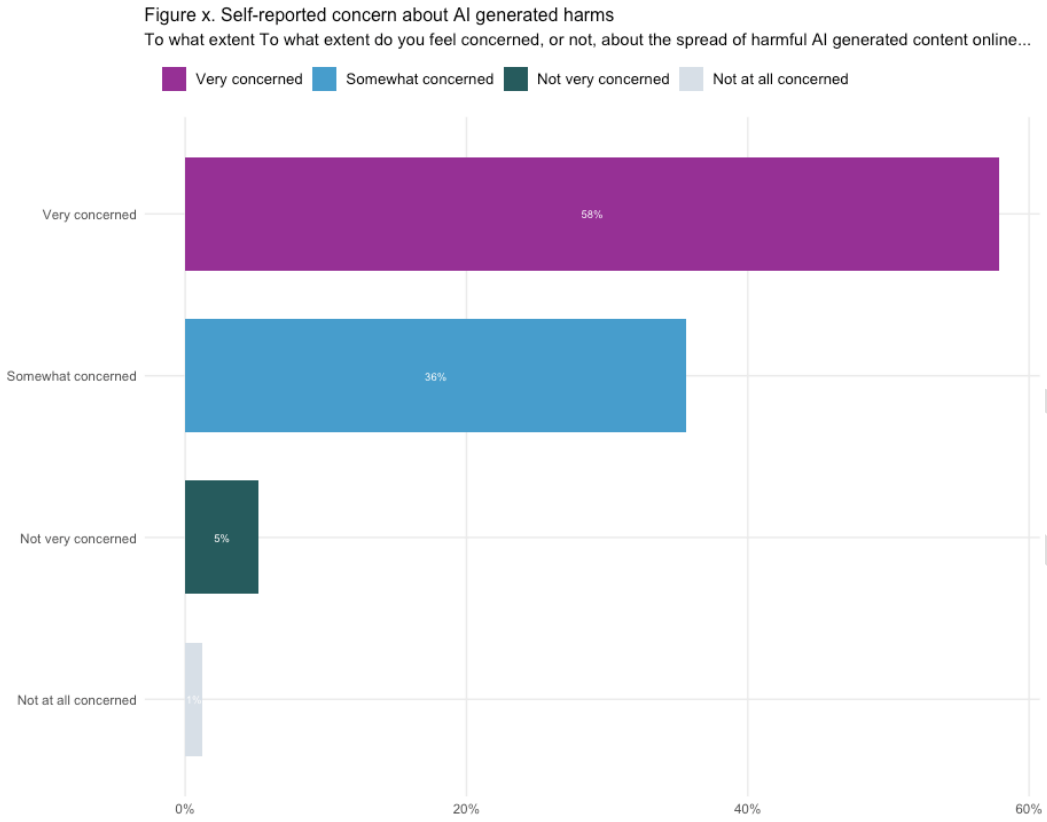

We asked the public about their expectations around the involvement of different stakeholders in AI safety. When asked who they think should be most responsible for ensuring AI is used safely, the majority of the UK public think an independent regulator (58%) and the companies developing AI technologies (58%) should be. Figure 15 shows expectations of responsibility.

Figure 15. Expectation of responsibility

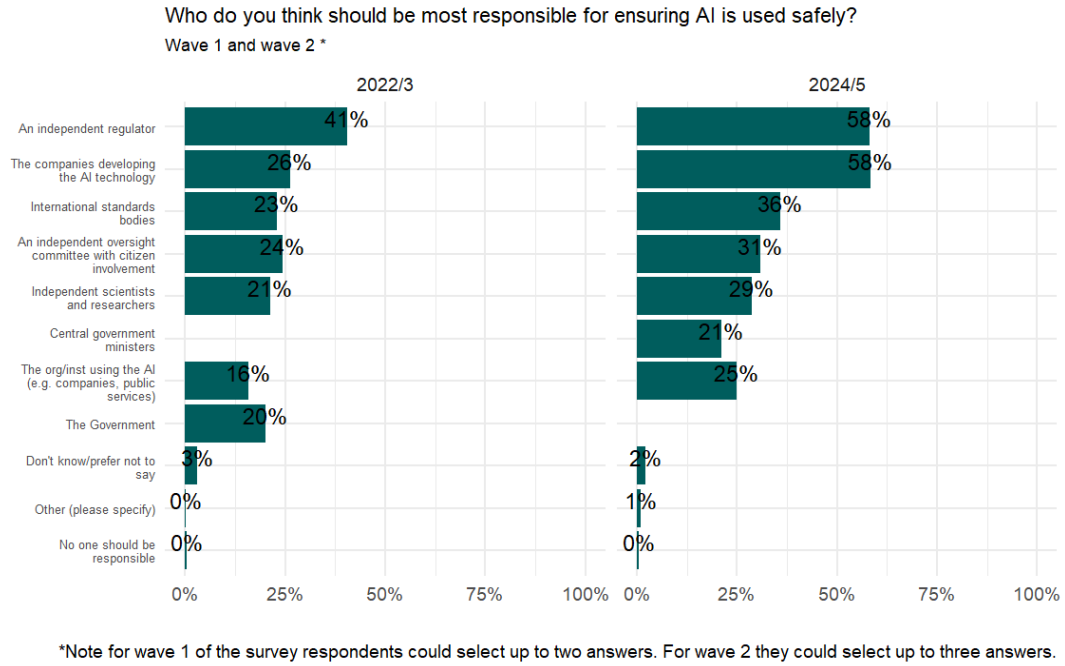

*Note for wave 1 of the survey respondents could select up to two answers. For wave 2 they could select up to three answers.

However, within this overall view, the preference for an independent regulator to be most responsible for ensuring AI is used safely increases with age, while preference for the companies developing AI decreases with age. Those aged 18–44 show a preference for companies over regulators, while those over the age of 55 have a preference for regulators over companies. This pattern of preference was similar to that found in our previous survey, suggesting age may continue to relate to expectations of, and potentially trust in, different organisations and institutions involved in AI development and deployment. Figure 16 shows responses across age groups to this question.

Figure 16. Expectations of responsibility by age group

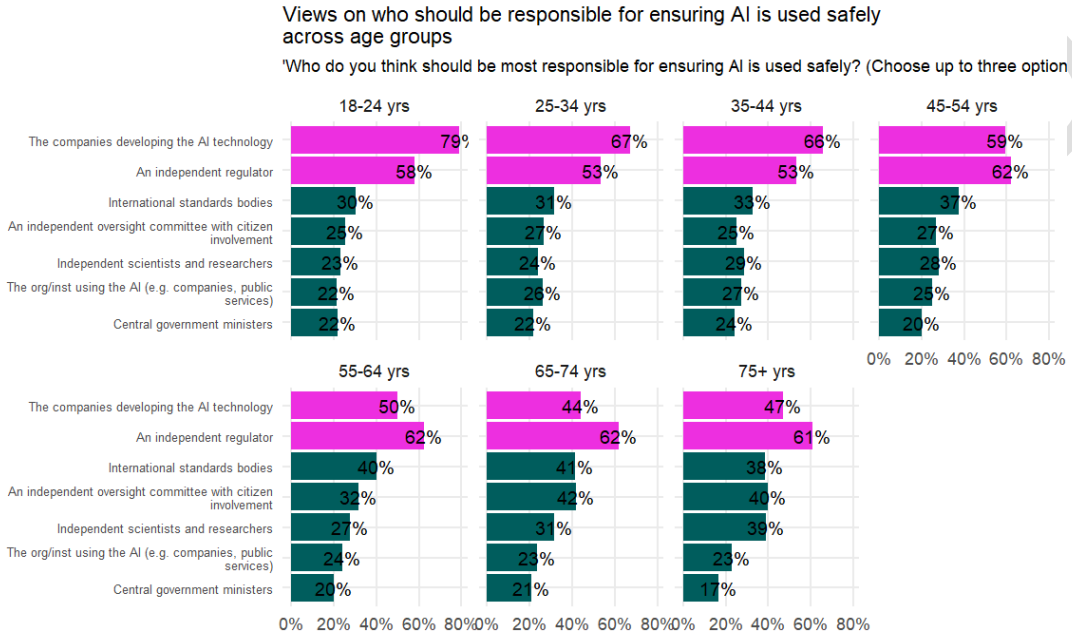

We also noted some regional variation in stakeholder preferences. For example, Northern Ireland shows a stronger preference for an independent oversight committee with citizen involvement (48%) compared to other nations. Further details can be found in Appendix 5.2. However, due to small sample sizes by nation, we cannot investigate regional differences in more depth in this report.

We asked the public how important it was to them that the government or independent regulators have a series of specific powers around the use of AI. These were the power to:

- stop the use of an AI product if it causes harm
- actively monitor the risks posed by AI systems
- develop safety standards on AI use
- access information about the safety of AI systems from developers.

These powers were chosen because they relate to live issues around AI regulation, which may be pertinent for the UK government's potential development of a future AI bill. Currently, there are no legal requirements for AI developers or independent regulators to regularly test or monitor upstream AI foundation models for safety risks, or statutory powers that allow

regulators to restrict the sale of AI products or services outside of narrow sectoral regulation such as that for medical devices.

The public feel strongly about the government or independent regulators, rather than private companies alone, having a suite of powers related to the use of AI. 88% of people believe it is important that the government or regulators – and not just private companies – have the power to stop the use of an AI product if it is deemed to pose a risk of serious harm to the public. Figure 17 shows responses across the powers. The low prevalence of 'don't know / prefer not to say' responses suggests that these preferences are widely held.

Figure 17. Importance of government or regulators having powers around the use of AI

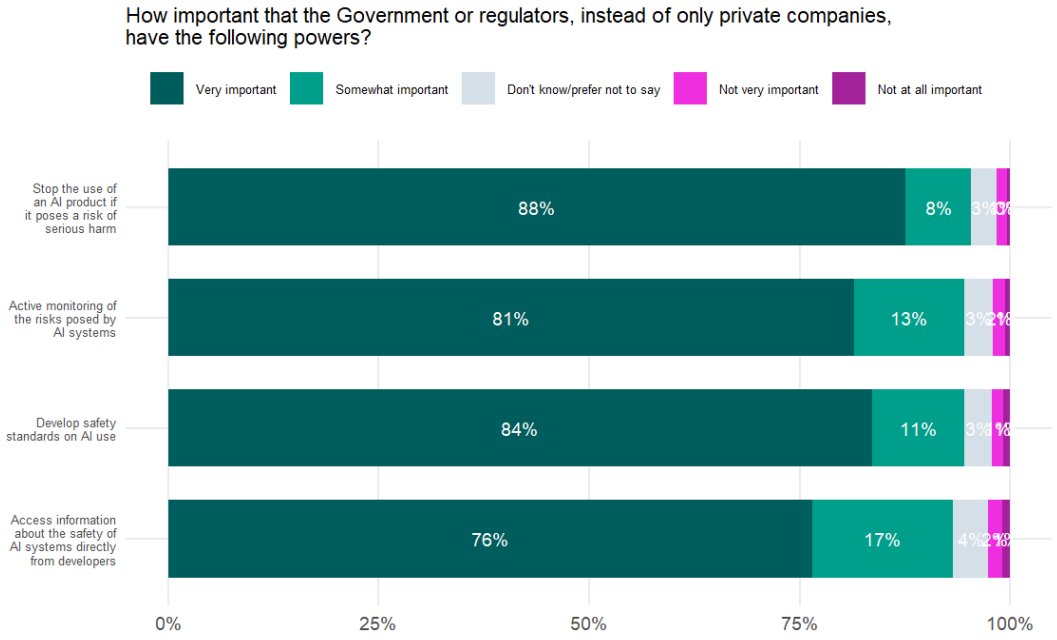

3.4.4. The public is concerned about data privacy, data sharing and lack of representation in decisions about AI

When looking at the range of specific concerns people chose in relation to each AI technology, we found that more people report feeling concerned about the safety and security of their personal information across most uses of AI this year than in our previous survey wave. For example, this year 56% of people are concerned about facial recognition technologies in policing gathering and sharing their personal information with third parties. This figure was 38% in 2022/23. Similarly, 33% of people are concerned about the safety and security of their personal information in relation to AI technologies assessing welfare eligibility, with this figure at only 19% in our previous survey. Table 7 shows the prevalence of personal information concerns for uses of AI in 2022/23 and 2024/25.[34]

Table 7. Concerns around personal information

[34] Concern around the safety and security of personal information was not asked for robotic care assistants and driverless cars and is therefore omitted from the data.

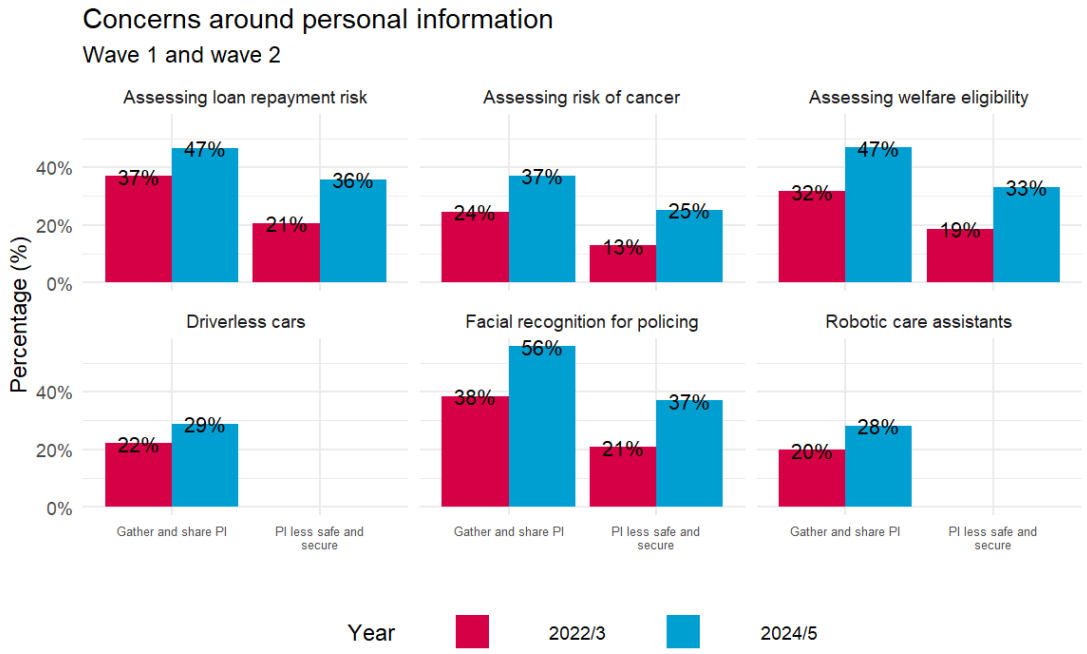

We asked the public about their views on data sharing in the public sector. Most (83%) are concerned by the idea of public sector bodies sharing data about them with private companies to train AI systems. This concern appears to be strongly held with very low proportions of the population feeling they did not know how concerned they were, if at all (3%), as shown in Figure 18. These concerns are important in the UK context where there have been explorations of anonymised NHS data sharing with private companies.[35]

Figure 18. Attitudes to public sector data sharing

[35] Kiran Stacey and Dan Milmo, 'Ministers Mull Allowing Private Firms to Make Profit from NHS Data in AI Push' The Guardian (13 January 2025) <https://www.theguardian.com/society/2025/jan/13/ministers-mull-allowing-private-firms-to-make-profit-from-nhs-data-in-ai-push> accessed 13 March 2025.

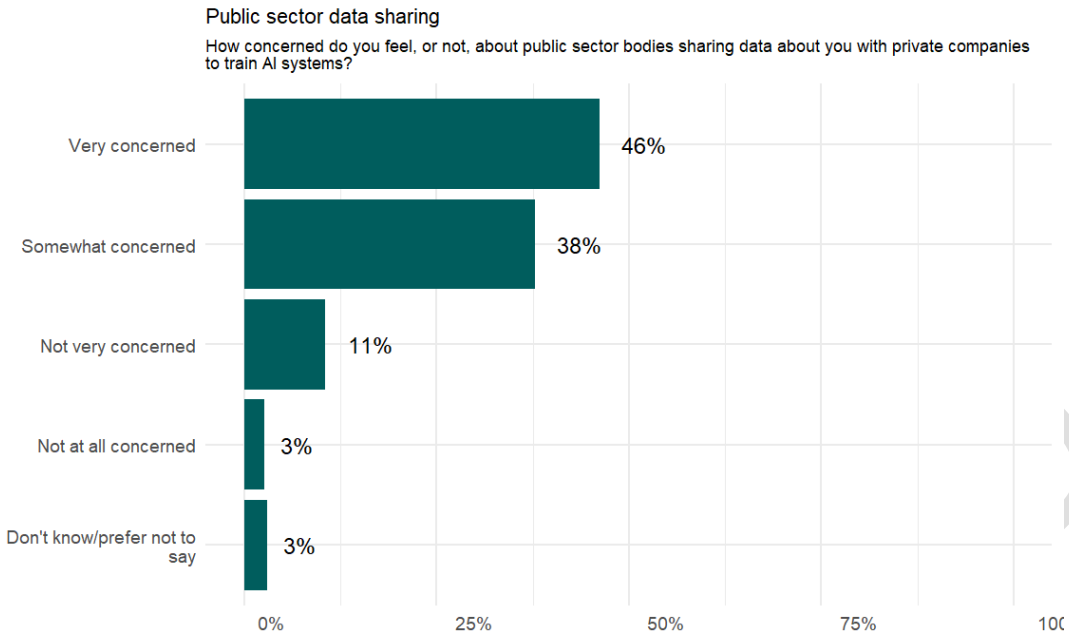

We also asked the public about the extent to which they felt their views and values are represented in current decisions being made about AI and how it affects their lives. Half of the UK public (50%) said that they do not feel represented in this decision making, while just over a quarter (27%) said they do. Not feeling represented increases with age, with 57% of people aged 65–74 feeling unrepresented, compared to 40% of people aged 18–24.

Figure 19 shows the nationally representative distribution of responses to this question alongside breakdowns by age.

Figure 19. Public voice in AI

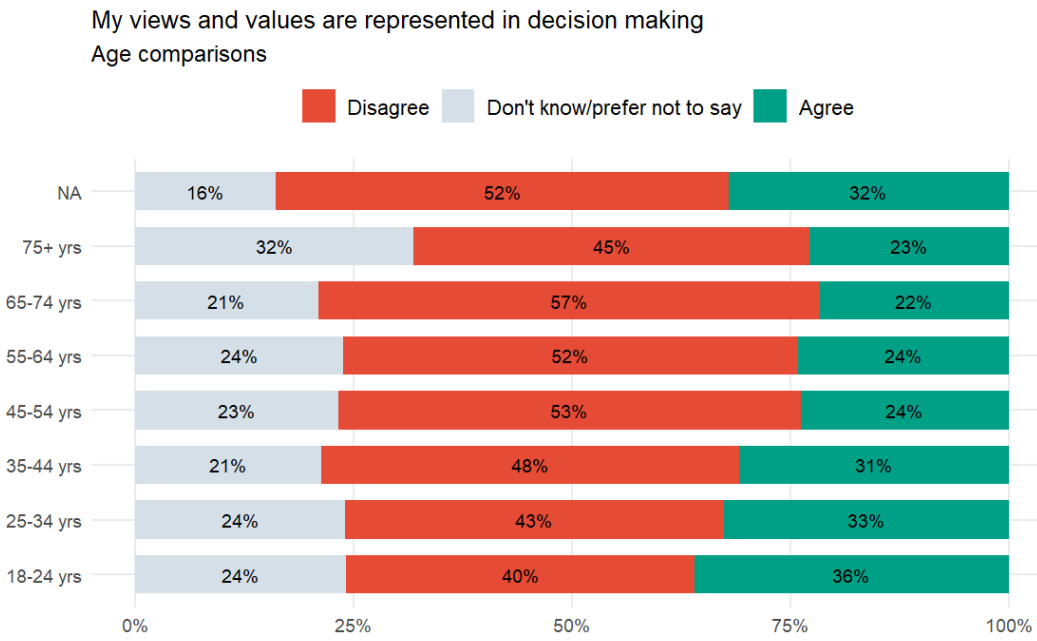

4. Conclusion

This report offers insights into how publics perceive AI technologies, and their expectations around its regulation and governance, with some specific insights about minoritised groups and people. It follows on from our 2022/23 survey, which took place before the emergence of technologies like ChatGPT in public discourse, providing new insights into how attitudes to AI might be changing over time.

The findings reiterate the importance of considering AI technologies in the context in which they are applied. While attitudes towards technologies in health diagnostics – e.g. detecting the risk of cancer from a scan – are largely positive, attitudes towards the use of AI in the delivery of care – e.g. mental health chatbots and robotic care assistants – are largely negative. At the same time, across applications of AI the public can recognise distinct potential benefits and identify areas of concern.

Just as attitudes to AI are multifaceted, so are UK publics. For instance, the survey found that Black/Black British and Asian/Asian British publics are significantly more concerned about facial recognition for policing than the general public. We know from existing evidence that people of colour are more likely to be disproportionately negatively affected by the deployment of these technologies.[36] Decision making around AI should consider the distinct impacts of AI technologies on diverse communities, and seek to embed their views and values – in recognition that different publics offer new insights into whether and how AI is used across different contexts.

The public typically see potential advantages of many applications of AI around improving efficiency and accuracy. However, people are also increasingly concerned about the safety of their personal data, as well as the replacement of humans in decision making. These attitudes point to a need to evidence whether AI systems are meeting public expectations around efficiency and accuracy, and how these systems can address public concerns.

Awareness of AI also fluctuates across applications. Emergent technologies like general-purpose LLMs, as well as less everyday applications such as driverless cars, have relatively high levels of public awareness. At the same time, awareness of more behind-the-scenes applications of AI, such as the use of AI to assess welfare eligibility, is significantly lower. Transparency around the use of AI systems is needed to ensure the public are aware of less visible, but highly impactful, uses of AI – especially those that have the potential to disproportionately affect those who are already marginalised in society.

To ensure that the introduction of AI-enabled systems in public sector services works for diverse publics, policymakers must engage and consult these publics to capture the range of attitudes towards and concerns about AI expressed by different groups. Capturing diverse perspectives may help to identify high-risk use cases, novel concerns or harms, and/or potential governance measures that are needed to garner public trust and support adoption.

[36] Thaddeus L Johnson and others, 'Facial Recognition Systems in Policing and Racial Disparities in Arrests' (2022) 39 Government Information Quarterly 101753.

In the last two years, the public have become more concerned by many applications of AI. At the same time, their preference for laws and regulation has increased. This rise in demand for laws and regulation comes at a time when the UK does not have its own set of comprehensive regulations around AI. The evidence suggests that the public support a multi-stakeholder approach to AI safety, with high expectations of both an independent regulator to ensure AI is used safely and of the companies developing AI technologies.

The UK government has repeatedly delayed consultation on AI legislation to address the potential risks and harms of some of these technologies, which stands in direct contrast to the concerns of the public and their growing desire for regulation. This tension presents a risk of low adoption or even backlash if AI technologies and the protections people are afforded around them do not meet public expectations. Delivering on the commitment in the AI Opportunities Action Plan to 'funding regulators to scale up their AI capabilities, some of which need urgent addressing',[37] will support meeting this expectation, recognising that – in the absence of legislation – regulators will need substantial resources, capabilities and expertise to build consideration of AI into their horizon-scanning, guidance and enforcement.

For AI to be developed and deployed responsibly, the hopes, concerns and experiences of the public need to be accounted for. Decision makers and AI developers need to listen to the voices of the public to ensure AI tools work for people and society, rather than further entrench inequalities in society. For example, in addition to traditional consultation methods (which target industry, academia or policy experts), policymakers should look to include evidence of public views, where it exists, and if appropriate engage diverse publics in public deliberation workshops on policy proposals.

[37] 'Artificial Intelligence Opportunities Action Plan - Hansard - UK Parliament' (13 March 2025) <https://hansard.parliament.uk/Commons/2025-01-13/debates/8C036071-5845-443C-B903-57483D552854/ArtificialIntelligenceOpportunitiesActionPlan> accessed 13 March 2025.

5. Appendix

5.1. Predictors of net benefit scores for each technology

To understand how demographics and attitudinal variables are related to the perceived net benefits of AI, we fitted linear regression models for each individual AI technology using the same set of predictor variables. The dependent variable in each model is 'net benefit', calculated as described above. The independent variables in the models were:

- Age (65 and older compared to younger than 65)
- Sex (male compared to female)
- Education (having a degree compared to not having a degree)
- Awareness of the technology (aware compared to not aware)
- Digital Skills (has digital skills compared to does not have digital skills)
- Low Income
- Black/Black British ethnicity
- Asian/Asian British ethnicity
- Tech pace (self-reported informedness about pace of technology change)
- Tech impact (views about technology making society better or worse)

Figure 20 presents the results for all regressions in a single plot. Each square in the plot represents the expected change in net benefit for a unit increase in the corresponding independent variable on the vertical axis, controlling for all other variables included in the model.

Statistically significant coefficients ($p < 0.05$) are shown in red, while black coefficients denote non-significant coefficients. Coefficient estimates higher than 0 indicate a higher net benefit and conversely coefficients lower than 0 are associated with lower net benefit (or higher concern) on a particular variable.

Figure 20. Predictors of net benefit

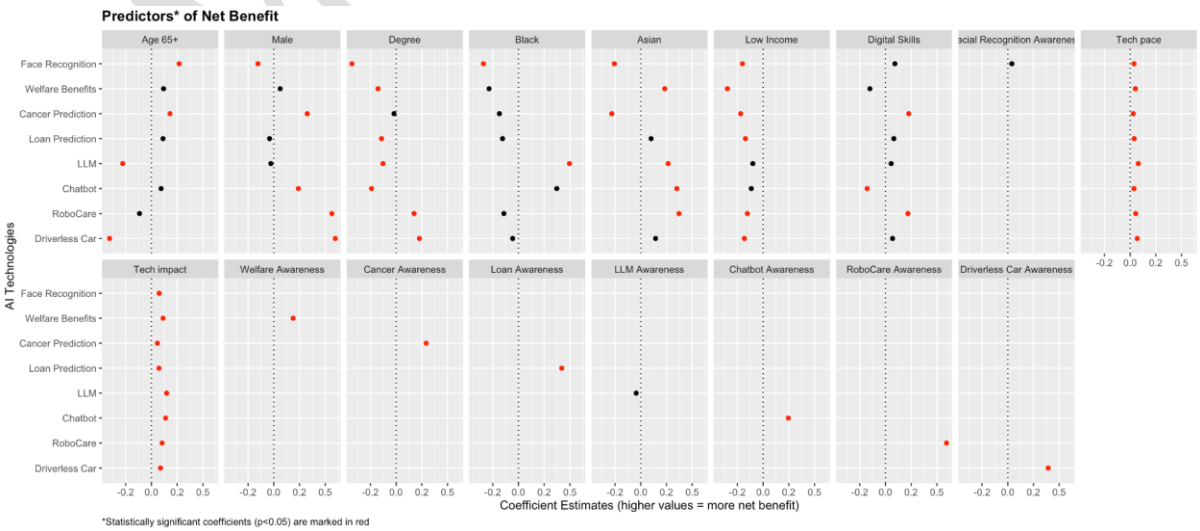

Taking the age variable as an example, respondents aged 65 and over are significantly more likely than those under 65 to believe concerns outweigh benefits for LLMs and driverless cars, indicating greater scepticism towards these emerging technologies. In contrast, male respondents are significantly more likely than women to believe benefits outweigh concerns for nearly all technologies.

Those holding a graduate degree are significantly less likely than those who do not hold such qualifications to believe that benefits outweigh concerns for most AI technologies, except driverless cars and robotic care assistants. This might indicate that greater exposure to risks associated with AI may contribute to a more critical stance on its benefits. As mentioned previously, Black/Black British respondents are more likely than non-Black respondents to believe concerns outweigh benefits for most AI technologies, with the exceptions of LLMs and mental health chatbots. However, this was not statistically significant when controlling for other demographic variables. In contrast, being Asian/British Asian is significantly associated with believing benefits outweigh concerns for most AI applications, except facial recognition and cancer risk prediction.

When all other variables are held constant, those on low income still have significantly lower net benefit scores than those with a higher income for most technologies. This finding suggests experiencing low income may be linked with less acceptance of AI technologies. This could be due to concerns around accessibility, fairness and potential biases in decision making that could impact their lives – such as through determining if they are eligible for welfare benefits or loans. It presents a case for understanding in more detail the concerns those on low income have towards these technologies and whether and how these technologies can be designed to benefit them.

Being aware is strongly associated with believing benefits outweigh concerns across all AI applications, except for its use as facial recognition in policing and general-purpose LLMs. Perceptions regarding the pace and impact of technology on society shows a consistent relationship across technologies, with people who hold more positive views about technology changing society at a good pace, and making society better, being more likely to see net benefits across all eight AI uses.

Figure 20 illustrates how patterns of perceived net benefit vary substantially across demographic groups and attitudinal indicators.

5.2. The four UK nations may have differing preferences around AI governance

We conducted some exploratory analysis into regional differences in attitudes across the four UK nations: England, Northern Ireland, Scotland and Wales. Due to small sample sizes, we did not investigate whether differences were statistically significant across the four regions.

We found that Northern Ireland has a stronger preference for an independent oversight committee with citizen involvement than the other devolved nations. 48% of Northern Irish publics feel an oversight committee with citizen involvement should be responsible for ensuring AI is used safely compared to only 30% of English publics, 32% of Scottish publics and 32% of Welsh publics. This preference for citizen involvement may be linked with greater familiarity with public participation initiatives (e.g. citizens' assemblies[38]) or less trust in other actors. In turn, they are less likely to select an independent regulator to perform this role than the other nations.

[38] 'Home' (*Citizens' Assembly*) <https://citizensassembly.ie/> accessed 13 March 2025.

Scotland is more likely to want to place responsibility on international standards bodies than the other nations, with 42% selecting this option compared to 35% in England, 32% in Northern Ireland and 36% in Wales.

Wales places less responsibility on the companies developing AI technologies (50% choose this option), and more on independent scientists and researchers (37% choose this option) than the other devolved nations.

Figure 21 shows a nation-level breakdown of expectations and preferences around the governance of AI.

Figure 21. Nation-level variation in preferences for AI governance

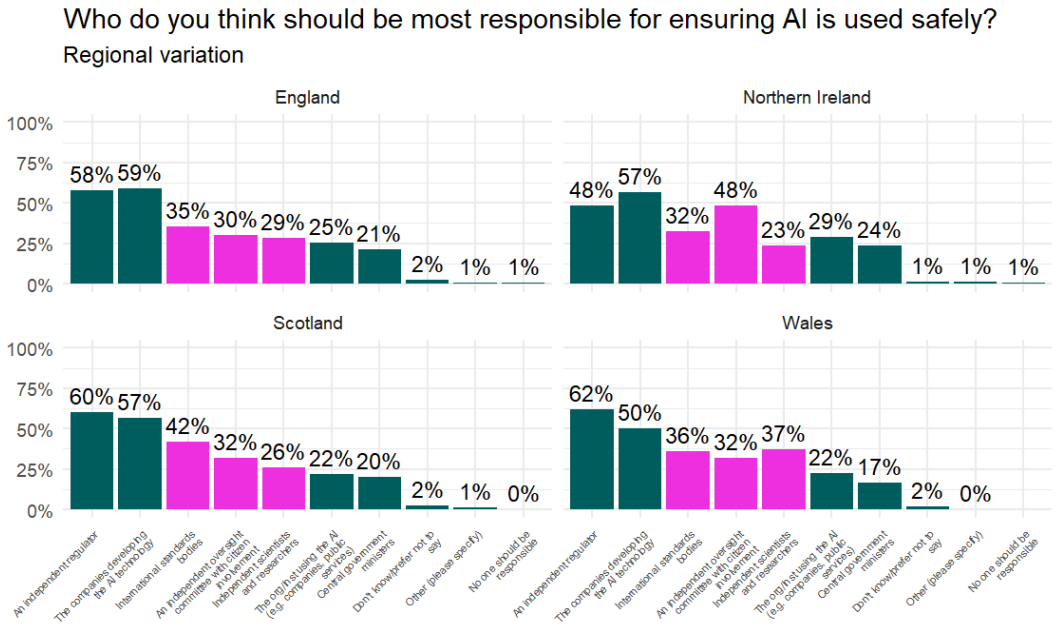

5.3. Specific benefits and concerns for each technology: full list

Table 8: Specific benefits

AI use	Benefit	Per cent (%)
Predicting cancer risk from a scan	Enable earlier detection of cancer, allowing earlier monitoring or treatment	85%
	Be more accurate than a doctor at predicting the risk of developing cancer	46%
	Reduce discrimination in healthcare	32%
	Reduce human error in predicting risk of developing cancer	64%

	Make personal information more safe and secure	9%
	Something else (please specify)	2%
	None of these	3%
	Don't know	6%
	Prefer not to answer	< 1%
Assessing loan repayment risk	Make applying for a loan faster and easier	58%
	Be more accurate than banking professionals at predicting the risk of repaying a loan	30%
	Be less likely than banking professionals to discriminate against some groups of people in society	44%
	Save money usually spent on human resources	36%
	Make personal information safe and secure	9%
	Reduce human error in loan decisions	41%
	Something else (please specify)	2%
	None of these	8%
	Don't know	11%
	Prefer not to answer	< 1%
Assessing welfare eligibility	Be faster than welfare officers at determining eligibility for benefits	52%
	Be more accurate than welfare officers at determining eligibility for welfare benefits	23%
	Be less likely than welfare officers to discriminate against some groups of people in society	39%
	Save money usually spent on human resources	43%
	Make personal information more safe and secure	13%
	Reduce human error in determining eligibility for benefits	39%
	Something else (please specify)	3%
	None of these	10%

	Don't know	12%
	Prefer not to answer	< 1%
Facial recognition for police surveillance	Make it faster and easier to identify wanted criminals and missing persons	89%
	Be less likely than the police to discriminate against some groups of people in society when identifying criminal suspects	66%
	Save money usually spent on human resources	46%
	Make personal information more safe and secure	51%
	Something else (please specify)	23%
	None of these	3%
	Don't know	2%
	Prefer not to answer	3%
Driverless cars	Make travel by car easier	32%
	Free up time to do other things while driving	35%
	Drive with more accuracy than humans	32%
	Be less likely to cause accidents than humans	34%
	Make travel by car easier for some groups (e.g. disabled people or people who have difficulty driving)	63%
	Save some money usually spent on human drivers	25%
	Something else (please specify)	3%
	None of these	19%
	Don't know	6%
	Prefer not to answer	0%
Robotic care assistant	Make caregiving tasks easier and faster	48%
	Be more effective than caregiving professionals at tasks such as lifting patients out of bed	37%
	Be less likely than caregiving professionals to discriminate against some groups of people in society	37%

	Save money usually spent on human resources	36%
	Something else (please specify)	4%
	None of these	14%
	Don't know	12%
	Prefer not to answer	< 1%
Large language models (LLMs)	Serve as a resource for continuous learning and skill development	50%
	Improve efficiency by automating repetitive tasks (e.g. writing emails)	56%
	Enhance creativity by generating ideas	38%
	Save money usually spent on human resources	31%
	Something else (please specify)	3%
	None of these	9%
	Don't know	17%
	Prefer not to answer	< 1%
Mental health chatbots	Serve as a faster way to get mental health support	50%
	Be more accurate than a mental healthcare professional at suggesting treatment options	7%
	Be less likely than mental healthcare professionals to discriminate against certain groups	27%
	Save money usually spent on human resources	28%
	Feel like interacting with a human, helping to prevent feelings of isolation	33%
	Be useful for certain groups of people to use (e.g. those with mobility conditions)	46%
	Something else (please specify)	3%
	None of these	15%
	Don't know	13%
	Prefer not to answer	< 1%

Table 9: Specific concerns

AI use	Concern	Per cent (%)
Predicting cancer risk from a scan	Be unreliable and cause delays to predicting a risk of cancer	26%
	Be less accurate than a doctor at predicting the risk of developing cancer	23%
	Be less effective for some groups of people in society, leading to more discrimination in healthcare	21%
	Make it difficult to understand how decisions about potential health outcomes are reached	41%
	Make it difficult to know who is responsible if a mistake is made	50%
	Gather personal information which could be shared with third parties	37%
	Make personal information less safe and secure	25%
	Cause doctors to rely too heavily on it rather than their professional judgements	64%
	Something else (please specify)	3%
	None of these	7%
	Don't know	6%
	Prefer not to answer	< 1%
Assessing loan repayment risk	Be unreliable and cause delays to assessing loan applications	23%
	Be less accurate than banking professionals at predicting the risk of repaying a loan	25%
	Be more likely than banking professionals to discriminate against some groups of people in society	16%
	Make it difficult to understand how decisions about loan applications are reached	54%
	Make it difficult to know who is responsible if a mistake is made	48%
	Gather personal information which could be shared with third parties	47%
	Make personal information less safe and secure	36%

	Lead to job cuts (for example, for trained banking professionals)	42%
	Cause banking professionals to rely too heavily on the technology rather than their professional judgements	57%
	Be less able than banking professionals to take account of individual circumstances	59%
	Something else (please specify)	3%
	None of these	3%
	Don't know	7%
	Prefer not to answer	0%
Assessing welfare eligibility	Cause delays to allocating welfare benefits	17%
	Be less accurate than welfare officers at determining eligibility for welfare benefits	35%
	Be more likely than welfare officers to discriminate against some groups of people in society	14%
	Make it difficult to understand how decisions about allocating welfare benefits are reached	54%
	Make it difficult to determine who is responsible if there is a mistake	52%
	Gather personal information which could be shared with third parties	47%
	Make personal information less safe and secure	33%
	Lead to job cuts (for example, for trained welfare officers)	45%
	Cause welfare officers to rely too heavily on it rather than their professional judgements	60%
	Be less able than welfare officers to take account of individual circumstances	60%
	Something else (please specify)	4%
	None of these	3%
	Don't know	8%
	Prefer not to answer	0%

Facial recognition for police surveillance	Cause delays in identifying wanted criminals and missing persons	7%
	Be less accurate than the police at identifying wanted criminals and missing persons	13%
	Be more likely than the police to discriminate against some groups of people in society	15%
	Lead to innocent people being wrongly accused if it makes a mistake	54%
	Make it difficult to determine who is responsible if a mistake is made	45%
	Gather personal information which could be shared with third parties	56%
	Make personal information less safe and secure	37%
	Lead to job cuts (for example, for trained police officers and staff)	42%
	Cause the police to rely too heavily on it rather than their professional judgments	57%
	Something else (please specify)	4%
	None of these	7%
	Don't know	4%
	Prefer not to answer	0%
Driverless cars	Not always work, making the cars unreliable	69%
	Make getting to places longer	15%
	Not be as accurate or precise as humans	43%
	Gather personal information which could be shared with third parties	29%
	Be less effective for some groups of people in society than others	26%
	Be difficult to use for some people	45%
	Lead to job cuts (for example, for truck drivers, taxi drivers and delivery drivers)	54%
	Make it difficult to know who is responsible if a mistake is made	66%

	Make it more difficult to understand how the car makes decisions compared to a human driver	57%
	Be more likely to cause accidents than human drivers	42%
	Something else (please specify)	5%
	None of these	3%
	Don't know	3%
	Prefer not to answer	< 1%
Robotic care assistants	Be unreliable and cause delays to urgent caregiving tasks	40%
	Be less effective than caregiving professionals at tasks such as lifting patients out of bed	42%
	Be less effective for some groups of people in society than others, leading to more discrimination	26%
	Be unsafe as it could hurt people	59%
	Make it difficult to know who is responsible for what went wrong if a mistake is made	50%
	Gather personal information which could be shared with third parties	28%
	Lead to job cuts (for example, for trained caregiving professionals)	53%
	Cause patients to miss out on human interaction from human carers	82%
	Something else (please specify)	3%
	None of these	2%
	Don't know	6%
	Prefer not to answer	< 1%
Large language models (LLMs)	Reduce users' own problem-solving skills or critical thinking abilities	66%
	Harm the environment due to high energy consumption	26%
	Be biased because of the data it is trained on	50%
	Be used to generate offensive or harmful content	47%

	Make it difficult to know who is responsible if a mistake is made	46%
	Infringe on copyright because of the data it is trained on	45%
	Lead to personal data being less secure and safe	40%
	Lead to job cuts (for example, due to automating tasks)	42%
	Something else (please specify)	5%
	None of these	3%
	Don't know	12%
	Prefer not to answer	< 1%
Mental health chatbot	Be unreliable and cause delays to getting help	37%
	Be less accurate at suggesting treatment options	49%
	Provide misleading advice, potentially leading to harmful consequences	62%
	Lead to discrimination against certain groups	12%
	Make it difficult to understand how decisions are reached	44%
	Make it difficult to know who is responsible if a mistake is made	46%
	Lead to sensitive personal data being less secure and safe	39%
	Lead to job cuts (for example, for trained mental healthcare professionals)	41%
	Lead to isolation by replacing human to human interactions	68%
	Make it unclear that people are not interacting with a human	63%
	Be relied on too heavily by those using it	57%
	Something else (please specify)	4%
	None of these	2%
	Don't know	8%
	Prefer not to answer	< 1%

5.4. Sample demographics

Table 10: Unweighted sample demographics

Demographic information		Unweighted sample size
Age	1824 yrs	73
	25-34 yrs	421
	35-44 yrs	596
	45-54 yrs	600
	55-64 yrs	635
	65-74 yrs	645
	75+ yrs	509
	NA	34
Digital skills	Has digital skills[39]	2549
	No digital skills	962
	NA	2
Education	Degree level qualification(s)	1635
	No qualifications	401
	Non-degree level qualifications	1450
	Other	14

[39] As per measure specified in: Lloyds Bank, 'UK Consumer Digital Index' (2018) <https://www.lloydsbank.com/assets/media/pdfs/banking_with_us/whats-happening/LB-Consumer-Digital-Index-2018-Report.pdf> accessed 13 March 2025.

	NA	13
Ethnicity	Asian or Asian British	433
	Black or Black British	198
	Mixed or multiple	49
	Other	40
	White British	2515
	White other	221
	NA	57
Sex	Female	1875
	Male	1632
	NA	6
Digital access	Mobile and data	2998
	Mobile, no data	225
	No mobile	284
	NA	6
Household income	Above £1,500 (equivalised) per month	1965
	£1,500 or less (equivalised) per month	1319
	NA	229

Acknowledgements

This report was co-authored by Roshni Modhvadia and Tvesha Sippy, with substantive input from Octavia Field Reid and Helen Margetts.
The research reported here was undertaken as part of Public Voices in AI, a satellite project funded by Responsible AI UK and EPSRC (Grant number: EP/Y009800/1). Support for the Ada Lovelace Institute's work on public survey and deliberative enquiry was provided by BRAID. BRAID is funded by AHRC (Grant number: AH/X007146/1).

Public Voices in AI was a collaboration between: the ESRC Digital Good Network @ the University of Sheffield, Elgon Social Research Limited, Ada Lovelace Institute, The Alan Turing Institute and University College London.