



BIS Bulletin

No 86

Survey evidence on gen AI and households: job prospects amid trust concerns

Iñaki Aldasoro, Olivier Armantier, Sebastian Doerr, Leonardo Gambacorta and Tommaso Oliviero

23 April 2024

BIS Bulletins are written by staff members of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed here are those of their authors and not necessarily the views of the BIS or the Federal Reserve Bank of New York. The authors are grateful to Leo Goldman, Gen Li and Giulio Cornelli for excellent analysis and research assistance, Claudio Borio and Priscilla Koo Wilkens for comments, and to Louisa Wagner for administrative support.

The editor of the BIS Bulletin series is Hyun Song Shin.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2024. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN: 2708-0420 (online)

ISBN: 978-92-9259-755-9 (online)

Survey evidence on gen AI and households: job prospects amid trust concerns

Key takeaways

- *A representative survey shows that almost half of US households use generative artificial intelligence (gen AI) tools. The use of and knowledge about gen AI are significantly lower among women, the elderly and households with lower income or educational attainment.*
- *Respondents expect gen AI to bring more opportunities than risks for job prospects, especially among men and younger, more educated and higher-income households. Nonetheless, all groups trust gen AI less than humans, especially in the provision of financial and medical services.*
- *Survey participants express concern over the risks of data breaches and data abuse and overwhelmingly support the regulation of AI. Consistent with previous surveys, respondents trust government agencies and financial institutions more than big techs to safeguard their data.*

Generative artificial intelligence (gen AI) is poised to affect everyday lives profoundly. Millions are already exploring gen AI to create text, music and videos, and a growing number of firms in all sectors are integrating gen AI tools into their business operations.¹ Despite the remarkable speed of adoption, little is known about consumers' uses of and attitudes towards gen AI. Furthermore, the level of trust households place on decisions made by gen AI and whether this trust depends on the context or provider remain unclear.

This Bulletin sheds light on these questions with novel data from the Survey of Consumer Expectations (SCE). The SCE, a representative high-quality survey of US household heads, is widely used to measure households' expectations about inflation, the labour market and their finances. By leveraging a special module on gen AI recently added to the SCE, this Bulletin finds that almost half of all US households use gen AI and that usage and knowledge are significantly higher among men, younger individuals and households with higher income or educational attainment. These groups are also more optimistic that AI will bring more opportunities than risks for their job prospects. However, the vast majority of respondents trust gen AI less than humans to provide services, especially if provided by big techs, in part reflecting users' privacy concerns. Households also overwhelmingly favour regulation. These results can inform the debate on how gen AI might affect economic inequality as well as on the need for adequate privacy and data regulation.

The Survey of Consumer Expectations

The SCE is a high-quality monthly, internet-based survey produced by the Federal Reserve Bank of New York. Launched in 2013, it has been used extensively to help researchers and policymakers understand how expectations are formed and how they affect consumer behaviour. The SCE uses a 12-month rotating nationally representative panel of approximately 1,300 US household heads. New respondents are drawn each month to match demographic targets from the American Community Survey, and they remain on the panel for up to 12 months before rotating out. The main objective of the survey is to collect

¹ See for example Briggs and Kodnani (2024).

expectations for a wide range of economic outcomes (eg inflation, income, spending, household finance, employment and housing). The survey reports detailed demographic information, including the respondents' gender, age, income and education.


We use a recent addition to the SCE survey. The February 2024 survey included an additional ad hoc module to understand who is using gen AI, how households assess the impact of gen AI on their job prospects and what they see as the major issues regarding trust and data privacy. This module asked detailed questions on respondents' use of gen AI, the opportunities and risks they see, concerns regarding trust and privacy, as well as the need for regulation. In what follows, we use this information to investigate whether these aspects vary across households based on their age, gender, education and income.²

Who uses generative AI?

Almost half of the respondents indicate that they have used gen AI tools, such as ChatGPT, Google Bard or DALL-E, at least once in the past 12 months (Graph 1.A). A quarter of respondents report using these tools at least once a month, while 14% use them at least once a week.

AI demographic divide? Households' use and knowledge of gen AI

Graph 1



Panel A reports the share of respondents and 95% confidence intervals to the question "How often have you used artificial intelligence tools (such as ChatGPT, Google Bard, DALL-E, ...) in the past 12 months?"; a dummy variable has been constructed taking a value of zero if the response is "Never" and a value of one if the response is "Less than once a month", "Once a month", "Once a week" or "More than once a week". The graph reports the share of respondents for which the dummy takes on the value of one. Panel B reports the average score and 95% confidence intervals to the question "How much do you know about artificial intelligence tools (such as ChatGPT, Google Bard, DALL-E, ...)?", with scores from 1 (lowest) to 7 (highest knowledge). Panel C reports the average score to the question "Over the next 12 months, how likely are you to use an artificial intelligence tool in the following contexts? For each of them, please report the likelihood on a scale from 1 (very unlikely that you will use such tools) to 7 (very likely)."

Source: Federal Reserve Bank of New York, *Survey of Consumers Expectations*; authors' calculations.

Use differs markedly across demographic groups. For one, men are substantially more likely than women to have used gen AI in the past 12 months (50% vs 37%), a pattern that echoes broader findings on a gender gap in the use of (financial) technology (Chen et al (2023)). A large difference exists between households with a college degree or higher and those without a college degree (52% vs 32%), as well as between younger cohorts and respondents aged 60 and over (49% vs 26%). The pronounced difference between younger and older respondents mirrors the "digital divide" found in other contexts, a divide that could stem from the elderly's limited perceived benefits of new technology (Doerr et al (2022), Armantier

² The sample used in this Bulletin contains 893 respondents, of which: 51% are women, 22% are 60 years or older, 59% have at least a college degree, and 55% report an annual household income of at least \$75,000.

et al (2024)). Finally, higher-income respondents, ie households whose family income is above \$75,000, report more frequent use than lower-income respondents (50% vs 36%). Individuals with higher education and income likely have greater access and exposure to gen AI tools, potentially due to job requirements and comfort with technology, leading to more frequent use.

Greater use of gen AI tools goes hand in hand with greater self-assessed knowledge about gen AI. When asked about how much they know about gen AI tools on a scale from 1 (I know nothing at all about artificial intelligence tools) to 7 (I know a lot), households report an average value of three (Graph 1.B). Only 10% of respondents reported a value of five or higher. Similar to use, self-reported knowledge is higher among men, those under 60 and households with higher education or income.

Going forward, households expect to use gen AI tools more in relatively low-stake activities. Regarding use in the next 12 months (Graph 1.C), households' responses suggest they are more likely to use it for education and training purposes, for job tasks, as well as for leisure activities (for example, writing, drawing or creating videos). In contrast, respondents indicate that they are less likely to use gen AI tools to obtain financial and medical advice, areas that require a relatively higher degree of trust in experts, as discussed below.


Gen AI and job prospects

No area has attracted more attention than the potential effects of AI on jobs. Unlike previous automation waves that predominantly affected occupations requiring manual labour, gen AI is expected to have the largest impact on knowledge workers. Against this background, there is an ongoing debate on whether gen AI will rebuild the middle class or lead to large-scale job displacement and impoverishment. Recent evidence suggests that, at least in specific tasks, gen AI has the potential to raise worker

Highway to automation or stairway to job security? Gen AI and job prospects

In per cent

Graph 2



Panel A reports the average responses to the following questions: 1 "What do you think are the chances that artificial intelligence will increase your productivity at work?", 2 "What do you think are the chances that artificial intelligence will help you find new job opportunities?", 3 "What do you think are the chances that you will lose your current job because of artificial intelligence tools?" and 4 "What do you think are the chances that your salary in your current job will decrease because of artificial intelligence tools?" Respondents could indicate their assessment on a scale of 0 to 100%. Panel B reports average probabilities and 95% confidence intervals by household groups to questions 1 (red dot) and 2 (blue cross). Panel C reports average probabilities and 95% confidence intervals by household groups to questions 3 (red dot) and 4 (blue cross).

Source: Federal Reserve Bank of New York, *Survey of Consumers Expectations*; authors' calculations.

productivity.³ What is not well understood is whether gen AI represents an opportunity or a risk for job prospects. Such expectations could be highly consequential for the macroeconomic impact of AI (Aldasoro et al (2024)).

On average, US households are generally positive about the impact of AI on their job prospects. Respondents assess the chance of higher productivity or more job opportunities because of gen AI to average between 27% and 30% (Graph 2.A). In stark contrast, they see only a 13% chance of losing their job or seeing a decline in their salary because of AI.


The expected job market benefits of AI differ across demographics (similar to the use of AI), but there are no differences in expected risks. Men, those younger than 60 and households with higher educational attainment and income expect the largest possible benefits from gen AI (Graph 2.B). Differences are more pronounced for productivity than for job opportunities.⁴ Somewhat surprisingly, all demographic groups deem the risks of job or wage loss from AI to be low (Graph 2.C).

Consumers’ trust in gen AI

Gen AI converses through everyday language and shows almost uncannily human-like capabilities in content creation. An increasing number of companies are already experimenting with chat bots that interact directly with end users. But how much do consumers trust gen AI compared to humans?

In gen AI we (do not) trust

Graph 3



Panel A reports the average responses to the following question: “In the following areas, would you trust artificial intelligence (AI) tools less or more than traditional human-operated services? For each item, please indicate your level of trust on a scale from 1 (much less trust than in a human) to 7 (much more trust).” Panel B reports average trust levels and 95% confidence intervals for the questions combined by household group. Panel C reports average scores and 95% confidence intervals to the question: “How much do you trust the following entities to safely store your personal data when they use artificial intelligence tools? For each of them, please indicate your level of trust on a scale from 1 (no trust at all in the ability to safely store personal data) to 7 (complete trust).”

Source: Federal Reserve Bank of New York, *Survey of Consumers Expectations*; authors’ calculations.

³ See Brynjolfsson et al (2023), Noy and Zhang (2023), and Peng et al (2024).

⁴ Compared with women, men report a 7 percentage point (pp) higher chance for productivity gains and a 2.5 pp higher chance for job opportunities. The respective differences are 12 pp and 11 pp for younger vs older respondents, 14 pp and 10 pp for those with higher vs lower education, and 7 pp and 3 pp for those with higher vs lower incomes.


In general, users have lower trust in gen AI tools than their fellow human beings. Respondents were asked to rate their trust in AI tools over traditional human-operated services on a scale from 1 (much less trust than in a human) to 7 (much more trust) in different areas. While relative trust in gen AI services is considerably lower than in humans in each area (Graph 3.A), trust is particularly low in banking, public policy and health. These patterns suggest that the adoption of gen AI tools, at least in the short run, is more likely to succeed in the fields of education and information. Differences across demographic groups are small, with the exception that women report significantly lower trust (Graph 3.B). This pattern aligns with their lower use and knowledge of gen AI and could be related to concerns about security and privacy when dealing with companies online (Armantier et al (2021)).

There are marked differences in the trust households place in how AI tools store their personal data depending on which institutions provide such tools (Graph 3.C). Respondents report the highest trust in traditional financial institutions, eg to store their bank transaction history, geolocation or social media data. The median respondent chose a (relatively low) value of 3. Trust in a government agency (ie federal and local governments) is slightly lower. Trust was the lowest for big techs (eg large technology companies such as Amazon, Apple, Meta or Google). The median respondent assigned a value of 2, and four fifths of respondents selected a value between 1 and 3. Values are similar across demographic groups.

Data are a prime concern for US households, who invariably deem regulation of AI necessary. When asked about their specific concerns, respondents appear equally worried about the risk of data breaches and the abuse of data for unintended purposes (Graph 4.A). Moreover, households clearly state that regulation on the use of AI tools is necessary (Graph 4.B). These assessments hold irrespective of demographics and of the type of agent engaging with AI (eg individuals, financial institutions, non-financial companies, and doctors and scientists).

US households' perspectives on risks and the need for rules

Graph 4



Panel A reports the average scores to the following questions where respondents were asked to reply by indicating a value from 1 (lowest score) to 7 (highest score): 1 "Do you think that sharing your personal information with artificial intelligence tools will decrease or increase the risk of data breaches (that is, your data becoming publicly available without your consent)?", 2 "Are you concerned that sharing your personal information with artificial intelligence tools could lead to the abuse of your data for unintended purposes (such as for targeted ads)?", 3 "Are you concerned that an increased reliance on artificial intelligence will have negative effects on human interactions or relationships?" Panel B reports average scores to the question "To what extent do you agree that there should be rules or restrictions on how individuals and firms can use artificial intelligence tools? Please indicate your level of agreement on a scale from 1 (I totally disagree) to 7 (I totally agree)."

Source: Federal Reserve Bank of New York, *Survey of Consumers Expectations*; authors' calculations.

Policy implications

Our results inform the debate on the implications of the rise of gen AI for labour markets and inequality, as well as for data privacy.

Observers worry that AI might increase inequality, as it benefits some workers, eg those whose tasks require logical reasoning, while making the tasks of other workers obsolete. Think of nurses who, with the assistance of AI, can more accurately interpret X-ray images. In contrast, tasks requiring only intermediate levels of cognitive abilities, such as text summarisation, might be done by gen AI. The associated displacement of some workers could lead to declines in employment and have profound implications for economic inequality. While the jury is still out on whether AI will indeed benefit better-educated workers, the survey evidence presented in this Bulletin suggests that men and better-educated households, which tend to have higher incomes, have better knowledge of AI, use it more and expect greater benefits from it. If this self-assessment turned out to be correct, gen AI would deepen the digital divide and exacerbate existing inequalities.

The use of gen AI tools also raises important concerns about users' privacy and the value of data. Personal data lie at the heart of the digital economy. Without training on abundant data, gen AI could have never reached its current performance. But companies collect and analyse the personal data shared by users when employing gen AI, often without consumers' explicit consent or full understanding. As our survey results have shown, consumers value their privacy and are concerned about the abuse and misuse of data. Consistent with these concerns, households agree with the proposition of having rules or restrictions on the use of AI. These considerations pose a trade-off for policymakers, who need to balance the improvement of efficiency through greater use of data with the protection of users' right to privacy.

References

Aldasoro, I, S Doerr, L Gambacorta and D Rees (2024): "The impact of artificial intelligence on output and inflation", *BIS Working Papers*, no 1179, April.

Armantier, O, S Doerr, J Frost, A Fuster and K Shue (2021): "Whom do consumers trust with their data? US survey evidence", *BIS Bulletin*, no 42, May.

Armantier, O, S Doerr, J Frost, A Fuster and K Shue (2024): "Nothing to hide? Gender and age differences in willingness to share data", *BIS Working Papers*, forthcoming.

Briggs, J and D Kodnani (2024): "The AI transition one year later: on track, but macro impact still several years off", *Goldman Sachs Global Economic Analysis*, April.

Brynjolfsson, E, D Li and L Raymond (2023): "Generative AI at work", *NBER Working Papers*, no 31161, November.

Chen, S, S Doerr, J Frost, L Gambacorta and H S Shin (2023): "The fintech gender gap", *Journal of Financial Intermediation*, vol 54, no 101026.

Doerr, S, J Frost, L Gambacorta and H Qiu (2022): "Population ageing and the digital divide", *SUERF Policy Briefs*, no 270, February.

Noy, S and W Zhang (2023): "Experimental evidence on the productivity effects of generative artificial intelligence", *Science*, vol 381, no 6654, pp 187–92.

Peng, S, W Swiatek, A Gao, P Cullivan and H Chang (2024): "AI revolution on chat bot: evidence from a randomized controlled experiment", working paper.

Previous issues in this series

No 85 28 March 2024	From financial inclusion to financial health	Carlos Cantú, Jon Frost, Tirupam Goel and Jermy Prenio
No 84 23 January 2024	Artificial intelligence in central banking	Douglas Araujo, Sebastian Doerr, Leonardo Gambacorta and Bruno Tissot
No 83 04 January 2024	Testing the cognitive limits of large language models	Fernando Perez-Cruz and Hyun Song Shin
No 82 20 December 2023	The contribution of monetary policy to disinflation	Pongpitch Amatyakul, Fiorella De Fiore, Marco Lombardi, Benoit Mojon and Daniel Rees
No 81 13 December 2023	Interest rate risk of non-financial firms: who hedges and does it help?	Ryan Banerjee, Julián Caballero, Enisse Kharroubi, Renée Spigt and Egon Zakrajšek
No 80 23 November 2023	Monetary policy, financial conditions and real activity: is this time different?	Fernando Avalos, Deniz Igan, Cristina Manea and Richhild Moessner
No 79 2 November 2023	Lessons from recent experiences on exchange rates, capital flows and financial conditions in emerging market economies	Pietro Patelli, Jimmy Shek and Ilhyock Shim
No 78 3 October 2023	Mapping the realignment of global value chains	Han Qiu, Hyun Song Shin and Leanne Si Ying Zhang
No 77 13 September 2023	Margins and liquidity in European energy markets in 2022	Fernando Avalos, Wenqian Huang and Kevin Tracol
No 76 7 September 2023	The oracle problem and the future of DeFi	Chanelle Duley, Leonardo Gambacorta, Rodney Garratt and Priscilla Koo Wilkens
No 75 19 May 2023	Disinflation milestones	Benoit Mojon, Gabriela Nodari and Stefano Siviero
No 74 13 April 2023	The changing nexus between commodity prices and the dollar: causes and implications	Boris Hofmann, Deniz Igan and Daniel Rees
No 73 11 April 2023	Stablecoins versus tokenised deposits: implications for the singleness of money	Rodney Garratt and Hyun Song Shin
No 72 11 April 2023	The tokenisation continuum	Iñaki Aldasoro, Sebastian Doerr, Leonardo Gambacorta, Rodney Garratt and Priscilla Koo Wilkens

All issues are available on our website www.bis.org.