Generative AI at work:

Survey evidence from Central Banks*

Patryk Perkowski
Yeshiva University

Aleš Maršál
National Bank of Slovakia

September 12, 2024

Abstract

Generative AI promises to transform the modern workplace, but understanding how workers use these tools is crucial for unlocking their productivity benefits. In this paper, we present the results of three workplace surveys at central banks: a survey of employees across the entire National Bank of Slovakia and one focused on PhD economists at both the Czech National Bank and the European Central Bank. While the vast majority of employees have heard of generative AI tools and believe using these tools would increase their productivity, less than half of employees use these tools at work. A significant barrier is the lack of clear use cases, identified by 95% of central bankers, and managerial guidance, identified be 88% respondents. The strongest predictors of adoption is the belief that colleagues are using these tools. Despite fears of automation, 79% of respondents see generative AI as a complement to their work, not a replacement. Our findings also uncover gender disparities and strong peer effects: women are 15% less likely to use generative AI which can be explained by the perceived differences in productivity benefits. Overall, our results highlight several opportunities and tensions for widespread generative AI adoption in central banks and other modern high-skilled workplaces.

Keywords: Generative AI, Large Language Models, Central Banking, Technology Adoption, Worker Productivity
JEL Classification: J24, M15, E58, O33

*Patryk Perkowski (patryk.perkowski@yu.edu); Aleš Maršál (ales.marsal@nbs.sk). We are grateful to Pavel Gertler and Jan Klacso for their support in running these surveys. The views expressed herein are those of the authors and do not necessarily represent the views of the National Bank of Slovakia, the European Central Bank, or the Czech National Bank. IRB approval for this study was obtained from WCG IRB under protocol number 20241892.

1 Introduction

In this paper, we examine the current use of generative AI in central banking, exploring both its existing applications and the potential for greater adoption. We partnered with three central banks – the European Central Bank (ECB), the National Bank of Slovakia (NBS), and the Czech National Bank (CNB) – to collect novel survey data on generative AI use at work and employee beliefs regarding the technology. Our goal was to understand now only how central bankers are currently using generative AI in the workplace, but also potential barriers to broader adoption. For the former, we investigate the types of tasks where central bankers use generative AI, the generative AI tools they use, and the amount of time of work they spend interacting with the technology. For the latter, we investigate why employees may be unlikely to use the technology at work and ways to increase adoption.

Since the launch of ChatGPT in late 2022, generative AI (GAI) technologies have generated widespread interest due to their potential to augment the productivity of modern workers. A growing body of literature has shown that working alongside generative AI can lead to significant productivity gains across various professional settings, including midlevel professional writing tasks (Noy and Zhang, 2023), customer support agents (Brynjolfsson et al., 2023), consultants (Dell'Acqua et al., 2023), and central bankers (Maršál and Perkowski, 2024). These findings suggest that many high-skill, high-wage jobs will be impacted by generative AI (Eloundou et al., 2023; Felten et al., 2023).

Central bankers, as white collar professionals working in a quantitatively driven, unstructured data-intensive environment, are obvious candidates to benefit substantially from generative AI. Indeed, central banks have been early adopters of machine learning techniques in areas such as financial supervision and forecasting (Doerr et al., 2021; Araujo et al., 2022; Aldasoro et al., 2024). Although central bankers believe that adopting generative AI at work could lead to productivity gains of over 18%, only 45% of them currently use these technologies.

There is striking agreement across respondents[1] on the primary reasons for not harnessing the full potential of GAI. Almost 95% of central bankers identifies the lack of clear use cases as main culprit. Indeed, results from our survey also finds that usage is much lower for specialized workplace tasks versus more generalist ones, highlighting the need for demonstration of use cases for more specialized central banking tasks. In addition, while over 50% of central bank employees agree that bank leadership

[1] Respondents agree across occupations at the central bank as well across different central banks.

supports the use of GAI, 88% of them also identify insufficient bank support as a reason for not expanding their use of generative AI at work. This highlights a significant gap between the potential productivity gains from GAI and its actual implementation within central banks, where employees remain uncertain about how to effectively integrate these tools into their workflows.

Wide spread believe is that the low utilization comes from fears and worries regarding the use of generative AI in this field. One reason why central banks hesitate to fully integrate generative AI is concern over data privacy. Central banks handle sensitive data, and there is a legitimate fear that using AI models like ChatGPT, which rely on vast external datasets, could inadvertently expose confidential information. Our survey findings, however, suggest that these concerns, while valid, are not the primary obstacles preventing widespread adoption of generative AI. Our results suggest that that central bank employees most benefit from using gen AI for task such as such as coding (62% at NBS and 30% at ECB of all employees), literature reviews, and technical modeling. These activities, vital to day-to-day operations, present a relatively safe space for generative AI implementation without risking sensitive data breaches. On the other hand, generative AI is used significantly less in policy-related (12%) work where privacy concerns might be more pressing.

We further rule out that the lack of adoption is driven by fears of job displacement. Despite reports of automation fears in the public and business press (Leffer, 2023; Ernst & Young LLP, 2023), we find that central bankers do not perceive AI as a threat to their job security but rather a tool to enhance their existing roles. 79% of the respondents disagree that their job is at risk of automation and only 11% believe that advances in GAI will lead to massive unemployment. Importantly, 77% of respondents perceive GAI as complement to their work rather than substitute. Thus, these fears do not drive adoption decisions in our domain.

What predicts the adoption of GAI within central banks? One of the strongest predictors of generative AI usage among central bankers is their belief about how much their colleagues are using these technologies. Given that the adoption of new technologies often spreads through workplace or industry networks, which reduces the costs of learning and adaptation (see, for example, Rastogi (2023); Békés and Harasztosi (2020); Nadiri (1993)), broader adoption in this sector will likely be driven by increased visibility of generative AI among co-workers.

We also find evidence that the mix of job tasks and beliefs regarding others' usage are related to generative AI usage at work. Using the task classification system from Autor et al. (2003) and Autor

and Handel (2013), we show that central bankers whose roles feature a higher percentage of analytical, non-routine cognitive, and leadership and management tasks are more likely to use generative AI at work. However, task composition alone accounts for only 12 percent of the variation in generative AI usage at work.

Our results also point to some broader patterns regarding differences in generative AI use, particularly for gender disparities in generative AI adoption. Humlum and Vestergaard (2023) shows that women are 20 percentage points less likely to use Chat GPT at work than men, even within the same occupations. Interestingly, our findings reveal a mixed picture. While women at the National Bank of Slovakia (NBS) were 15% less likely to use generative AI compared to men, this gender gap does not persist among PhD economists at the European Central Bank (ECB), where men and women use generative AI at similar rates. We find that this difference is driven by believes in the productivity benefits of the technology. At the NBS, men believed that the productivity gains from generative AI were 26% higher than those reported by women. This is consistent with findings in the information systems literature (see, for example, Venkatesh and Morris (2000), where men are more strongly influenced by their perceptions of usefulness. More broadly, workplace users of generative AI believed the productivity benefits were 60% higher than non-users, suggesting that experience with the technology likely shapes perceptions and influence adoption.

Our results indicate that users tend to rely heavily on a single generative AI tool, 91 percent report using ChatGPT, across all workplace applications. This is probably from the large part driven by the fact that from the user perspective ChatGPT largely outperforms competing models[2] (i.e Laskar et al. (2023), Giannakopoulos et al. (2023)). This situation presents tough decision for central banks: Should they continue using more secure, yet less powerful, in-house models or should they embrace the more powerful external models, despite the inherent risks.

Overall, our manuscript contributes to three complementary literatures.

First, our work contributes to the expanding literature on generative AI usage at work. Papers across various settings have found large productivity benefits of generative AI across various workplace settings.[3] Indeed, estimates suggest that a large scope of workplace tasks and jobs are likely to be impacted through generative AI use (Eloundou et al., 2023; Felten et al., 2023), particularly for higher-

[2]Another likely explanation is a first mover advantage for ChatGPT versus other generative AI tools.

[3]See, for example, (Noy and Zhang, 2023) for midlevel professional writing tasks, (Brynjolfsson et al., 2023) for customer support agents, (Dell'Acqua et al., 2023) for consultants, and (Maršál and Perkowski, 2024) for central bankers.

skilled and higher-wage workers. Humlum and Vestergaard (2023) provide survey evidence on how workers in 11 exposed occupations use ChatGPT at work. We complement this paper by examining the use of generative AI in central banking. We provide novel survey data on not only workplace usage in this setting, but also worker beliefs regarding the technology's impact on productivity and labor markets more generally, their trust and knowledge of generative AI, and reputational considerations associated with generative AI usage at work, among others.

Second, we contribute to the growing literature on the use of digital technologies in central banking. Central banks have recently adopted new forms of technology for their core functions of monetary policy and financial stability. For example, in the realm of financial supervision, central banks have used machine learning techniques to manage large datasets and enhance analytical capabilities (Doerr et al., 2021; Araujo et al., 2022). Neural networks and deep learning have been applied in forecasting and sentiment analysis to enhance the precision and depth of central bank communication (Hagos et al., 2024; Sleiman, 2023). Central banks have also recently adopted genAI tools for cybersecurity (Aldasoro et al., 2024). Our manuscript contributes to this agenda by highlighting how genAI tools are used throughout central banking. We pay special attention to the tools, tasks, and time spent using the technology, for both PhD economists and non-PhD economists, across three different central banks.

Lastly, our paper contributes to the literature on barriers to technology adoption within organizations. In the context of soccer-ball producers in Pakistan, Atkin et al. (2017) show that misaligned incentives between employees and owners is an important barrier to technology adoption. Goehring et al. (2024) show that career concerns affected director adoption of digital cinematography in the U.S. motion picture industry. Gupta et al. (2020) provides evidence that language marriers limited the adoption of modern agricultural technologies for Indian farmers. We contribute to this literature by examining barriers to generative AI adoption in central banks. We find that in addition to the task make-up of jobs, employee beliefs regarding automation fears and reputational considerations constrain the adoption of generative AI in our setting.

Our paper proceeds as follows. Section 2 discusses the three surveys that were run. Section 3 examines how employees use generative AI at work, focusing on its use across various workplace tasks, across different generative AI tools, and across time. Section 4 looks at employee attitudes regarding generative AI, including beliefs regarding the labor market impact of generative AI, reputational

considerations, trust and knowledge of generative AI, and organizational practices to boost workplace adoption. Section 5 explores various factors that may drive the lack of generative AI adoption observed in the sample. Section 6 concludes.

2 Survey

In the spring of 2024, we partnered with three central banks to run surveys on generative AI in the workplace: the National Bank of Slovakia, the Czech National Bank, and the European Central Bank. The National Bank of Slovakia and the Czech National Bank are the central banks of Slovakia and the Czech Republic, respectively, overseeing monetary policy, financial stability, and banking supervision within each country. Meanwhile, the European Central Bank oversees monetary policy for all of the countries in the Eurozone.

The survey was designed to understand key questions regarding generative AI use within these samples. We focused on large language models since these are the generative AI tool that is most applicable to the work done in central banking. The survey module consisted of five modules.

- **Background data**: We asked participants various demographic and background questions, including their tenure at the organization, whether they manage any workers, and the language(s) they use at work. We also asked them to detail the amount of hours they spend across various categories of work such as routine cognitive tasks.
- **Exposure to Generative AI**: We asked participants several questions regarding their exposure to and knowledge of generative AI. Participants reported whether they had heard of generative AI tools before, which tools they have heard of, and whether they have used generative AI before.
- **Workplace Usage of Generative AI**: We asked participants about their use of generative AI across various workplace tasks. Both samples required the same 10 general tasks where generative AI can be helpful: (i) writing emails or reports, (ii) generating code or debugging, (iii) creating content (articles, blogs, etc), (iv) data analysis and visualization, (v) educational purposes or learning new skills, (vi) customer service, (vii) research and information gathering; (ix) design and creative work, (x) other. In addition, the ECB sample included workplace tasks specifically related to the work done by PhD economists. These tasks included: (i)

conducting literature reviews, (ii) cleaning data and preprocessing, (ii) exploratory data analysis, (iv) statistical modeling and econometric analysis, (v) generating economic models, simulations, or forecasts, (vi) writing research papers or reports, (vii) drafting policy recommendations, (viii) explaining economic concepts or theories, (ix) answering research questions, (x) peer review or providing feedback, (xi) research ideation and brainstorming, (xii) coding or programming tasks; (xiii) creating presentations or teaching materials; (xiv) project or task management. For each of the general tasks, participants reported whether they have used generative AI for the task, which generative AI tool(s) they have used for the task, and how many hours in an average week they used generative AI tools for the task.[4]

- **Attitudes toward Generative AI**: We asked participants a series of questions to gauge their attitudes toward generative AI. These questions spanned several subthemes, including (i) fears of automation, (ii) reputational considerations at work, (iii) trust in generative AI, (iv) workplace redesign with generative AI, (v) technical knowledge of generative AI, and (vi) barriers to workplace adoption. These questions were collected using a 7-point Likert scale from strongly disagree to strongly agree.

- **Prediction**: We asked participants to participate in several prediction exercises. We asked participants to predict the percentage of employees who have previously used generative AI and who have used it in the workplace. We also asked participants to predict the impact of generative AI use on their own productivity.

This survey data from all three samples complements one another to provide a comprehensive look at how generative AI is used in central banking. On the one hand, the sample from the National Bank of Slovakia was collected alongside a forthcoming field experiment on generative AI in the workplace (Maršál and Perkowski, 2024), and covers employees throughout the entire organizational hierarchy. This includes not only those in non-routine cognitive roles where generative AI may be especially helpful, but also those in jobs with more manual tasks. On the other hand, our sample from the Czech National Bank and the European Central Bank focuses on a specific subgroup: PhD economists in central banks. These surveys were conducted as a part of a seminar on how economists can use generative AI, and thus permits a deeper look at how this specialized group uses generative AI at

[4]Given the addition of the survey questions for the economic tasks for the ECB sample, we had to drop some survey questions for this sample. We have data on whether each participant in the ECB sample used generative AI across the various economic tasks plus how many hours a week they spent using generative AI on these tasks, but not which tools were used. This was done to prevent participant fatigue during the survey.

work. We leverage the unique benefits of each sample to provide a holistic view of how generative AI is currently used in central banking.

Table 1 displays summary statistics for the entire sample of participants across both pools. The sample is highly well-educated: over 90 percent of participants have a masters degree or higher. 37 percent of the sample is female, and 15 percent are managers. The average respondent is 41 years old and has spent 9 years at their organization. On average participants work for 45 hours a week. Almost a third of their work time is spent on analytical tasks, a quarter on non-routine cognitive tasks, and just under 20 percent on routine cognitive tasks. The remaining quarter of their time is spent primarily on leadership and management tasks, and interpersonal and communication tasks.

3 Generative AI at work

3.1 Rates of exposure to and usage of generative AI

Figure 1 examines average exposure to generative AI and workplace usage. Our results indicate high exposure to generative AI but lower workplace usage. Almost 90 percent of the sample has heard of large language models before, and 76 percent of participants have used a large language model before. However, only 45 percent of the sample has used a large language model in the workplace, versus over 60 percent for personal use.

In Appendix section A.1.1 we display summary statistics across these samples to examine demographic and workplace differences in those who use generative AI at work, versus those who do not. Those who use generative AI at work are on average younger (37.8 years old vs 43 years old), are more likely to be male (30 percent of generative AI users are female, while between 40 and 45 percent of non-users are female), and spend more of their time at work using the English language (36 percent vs between 20 and 32 percent for non-work users). These results indicate that there are important demographic differences in who is likely to use generative AI at work.

3.2 The use of generative AI across tasks

We next examine the distribution of generative AI usage across workplace tasks. We are interested in understanding the types of tasks where generative AI is used, which generative AI tools are used, and the amount of time workers spend interacting with these tools at work.

We first examine the distribution of generative AI usage across workplace tasks. We are interested in both the general workplace tasks that were asked for both samples, plus the economist-specific tasks for the ECB sample. Those who reported using generative AI at work were asked to select which workplace tasks they used generative AI.

Figure 2 examines generative AI use across various workplace tasks for participants who reported using generative AI at work. Panel A of Figure 2 focuses on the general workplace tasks, displaying the average percent of generative AI users at work, who report using generative AI for each task. Meanwhile, Panel B of Figure 2 focuses on economist tasks. Both samples are limited to respondents who reported using generative AI at work.

The results in Panel A of Figure 2 indicate that generative AI is used across a wide variety of general workplace tasks. Over half of workplace users reported using generative AI for coding or debugging, writing email and reports, and educational purposes or learning new skills. Around 44 percent used it for translation. Another third reported using generative AI for research and information gathering, creating content, and data analysis. In Appendix section A.2.1, we examine the distribution of the number of workplace tasks where generative AI is used, among respondents who use generative AI at work. The results indicate that usage is not siloed into specific tasks: the median respondent reported using generative AI for three workplace tasks. Only 13 percent of those who use generative AI in the workplace only use it for one task, while around 40 percent use generative AI for four or more job tasks. These results suggest that users are well-versed in the various ways to use generative AI in their general workplace tasks.

While our sample uses generative AI across many general workplace tasks, usage for more specialized economics tasks is much more limited. The results in Panel B of Figure 2 shows that usage rates are much lower for economist tasks. The most common uses here are coding, writing papers, and conducting econometric analysis, though these rates are much lower than in Panel A. Indeed, in Appendix section A.2.1, we also examine the number of economist tasks where one uses generative AI.

The median respondent reported using generative AI for only one economics task. Meanwhile, over 40 percent of respondents reported using it for no economics tasks. These individuals use generative AI for general workplace tasks, but do not use it for more specialized tasks corresponding to their roles.

Overall, these results indicate that while generative AI tools are used for a variety of general workplace tasks, adoption lags for more specialized, domain-specific tasks. This suggests that there may be barriers in fully embracing these tools for tasks that require a higher level of domain expertise or specialized knowledge. sbnuqsuqs

3.3 The use of generative AI across tools

We next examine the type of generative AI tools that participants use for each general workplace task. For each general workplace task, participants indicated whether they use OpenAI's ChatGPT, Google Gemini, Github Copilot, Meta's LLaMa, Anthropic's Claude, or another generative AI tool. Participants were allowed to select multiple options from the list. We use this data to understand how participants decide which generative AI tools to use by task.

In Table 2, we examine the use of various generative AI tools across the sample. The first column displays the average number of generative AI tools used for each general task, while the second column displays the percent of users who reported using ChatGPT for this task. The sample is limited to those who report using generative AI for the given general task.

The results in Table 2 indicate that ChatGPT was by far the most common generative AI tool used across all general tasks. 87 percent of users who use generative AI tools for emails report using ChatGPT, 91 percent of those who use generative AI for coding, 92 percent for those who use generative AI for learning new skills, and 91 percent for those who use generatie AI for research and information gathering. Moreover, participants are unlikely to use multiple tools at once: the averages in column 1 are all close to one. Taken together, these results indicate that users tend to focus on one generative AI across all workplace tasks, and that tool is ChatGPT.

These results suggest that one potential roadblock for unlocking the productivity benefits of generative AI may be understanding how workers select tools for given tasks. In our sample, workers are most likely to use ChatGPT regardless of job task. However, other tools may be more appropriate for a

given task. For example, DeepL is a specialty generative AI tool used specifically for translation.[5] This suggests there may be productivity benefits from workers targeting the right specialized tools to the right tasks. Further research on how workers trade-off the benefits versus costs of various tools, and how to incentive worker exploration across these AI tools, will be helpful in this regard.

3.4 The use of generative AI across time

Lastly, we examine the distribution of generative AI tools across time. Figure 3 examines how often users spend interacting with generative AI tools at work. Panel A of Figure 3 displays a histogram of the percent of one's workweek that is spent interacting with generative AI tools, while Panel B displays the average number of hours spent across each general workplace sample. The sample in both figures is limited to those who report using generative AI at work.

The results in Figure 3 indicate substantial variation in the amount of time that is spent interacting with generative AI. The median user reports spending 9 percent of their work week interacting with generative AI at work. Assuming a 40 hour work week, this is between 3 and 4 hours a week. However, there are some "super-generative AI" users, who report spending more than 20 percent of their work week alongside this technology. These individuals are rare, though: most users report less than 5 percent of their work week interacting with generative AI. This indicates that across users, there is substantial heterogeneity in the extent to which the technology is integrated into their workday. Meanwhile, Panel B indicates that the most time is spent interacting with generative AI for coding, for learning new schools, for conducting research, and for translation. Overall, these results indicate that even among those who use generative AI at work, generative AI usage does not make up a lot of their work time. This suggests that there may be further opportunities to increase usage, perhaps through the display of use cases.

4 Attitudes toward generative AI

In addition to the data on generative AI usage at work, our survey also examined respondent attitudes toward generative AI more generally. In this section, we investigate these beliefs.

[5] https://www.deepl.com/en/translator

4.1 Beliefs regarding the labor market impact of generative AI

Our survey included several questions on respondent beliefs regarding the labor market impact of generative AI. Since the launch of ChatGPT in late 2022, generative AI technologies have generated much interest in their capabilities and concern regarding their labor market impact. Indeed, recent estimates suggest that a large scope of workplace tasks and jobs are likely to be impacted through generative AI use (Eloundou et al., 2023; Felten et al., 2023), particularly for higher-skilled and higher-wage workers. The capabilities of these technologies have led to fears of automation and workforce changes. Articles in both the popular and business press have discussed "AI anxiety", or the fear that increasing adoption of AI technologies will lead to negative impacts on jobs such as automation (see, for example Leffer (2023) and Ernst & Young LLP (2023)). Such fears are not new, dating back to at least the Luddites during the Industrial Revolution (Merchant, 2023; Jones, 1996). As technology journalist Merchant has written, "Now, with nearly half of Americans worried about how AI will affect jobs, Luddism has blossomed" (Merchant, 2024).

To understand these beliefs,, our survey collected six items regarding generative AI in the workplace, and fears of automation and workplace changes more broadly. Participants responded using a 7-point Likert scale to the extent to which they agreed or disagreed with various statements. Our survey included four items to capture fears of automation, including worrying about one's skills being less valuable with generative AI, beliefs in whether generative AI will substitute or complement labor, beliefs in whether generative AI will lead to massive unemployment, and beliefs whether their job is at risk of automation. We also included two questions regarding negative workplace changes, particularly among reduced workplace automation and increased workplace monitoring.

Table 3 displays participant responses for these six items. Column 1 displays the average rating given to each question. The rating ranges from -3 (average response of strongly disagree) to 3 (average response of strongly agree), with an average response of 0 as indifferent between the two. The aggregated "agree" response combines respondents who replied "strongly agree", "agree", and "somewhat agree", while the aggregated "disagree" response combines respondents who replied "strongly disagree", "disagree", and "somewhat disagree". The remaining columns give average ratings across various subsamples, including those with economics PhDs vs not (columns 4 and 5), gender (columns 6 and 7), team role (columns 8 and 9), and GAI workplace usage (columns 10 and 11).

The results indicate that respondents are less pessimistic about the labor market impacts of generative AI. Only 11 percent of respondents believe that advances in generative AI will lead to massive unemployment (versus 80 percent who disagree), and just 12 percent worried their job is at risk of automation (versus 79 percent who disagree). Instead, respondents believe generative AI will complement labor (77 percent agree versus 14 percent disagree). While respondents are skeptical that generative AI will reduce their autonomy at work (9 percent agree versus 80 percent disagree), there are some concerns regarding workplace monitoring: almost a quarter of responses believe that generative AI will increase workplace monitoring.

While the respondents as a whole are not worried about the negative labor market impacts of generative AI, there are some noteworthy differences across the various subsamples. Non-PhD economists are twice as likely than PhD economists to believe that advances in generative AI will lead to massive unemployment (15 percent versus 7 percent), and almost twice as likely to believe that generative AI will increase workplace monitoring (28 percent versus 15 percent). Despite these differences in the aggregate impact of generative AI on the labor market and workplace, there are smaller differences in participant belief's regarding the individual impact of the technology. If anything, PhD economists are more likely to believe their skills will be less valuable with generative AI (20 percent versus 14 percent) and worry their job is at risk of automation (12 percent versus 9 percent). These results line up with recent experimental evidence that finds a flattening of the skills distribution through generative AI (Brynjolfsson et al., 2023; Dell'Acqua et al., 2023).

4.2 Reputational concerns with generative AI

Our survey was also interested in understanding how employees think about reputational considerations with generative AI use. A large literature in economics and management posits that employees must maintain a positive reputation from their peers, managers, and other firms, and that this need for a positive reputation influences behavior in organizations (Gibbons and Murphy, 1992; Holmstrom, 1999). For example, in the context of the US motion picture industry, Goehring et al. (2024) find that career concerns affected director adoption of digital cinematography. In the context of generative AI, some workers may worry that using assistance from generative AI can make them appear less skilled or less intelligent to their manager or co-workers, and this may limit adoption.

Our survey collected five items regarding reputational concerns and generative AI usage, collected

again using a 7-point Likert scale. We included two questions regarding respondent comfort with others knowing they used generative AI at work. We asked this question for both their managers and also their coworkers. We also included two other questions on reputational considerations, asking both if they fear their coworkers would judge their work negatively if they used generative AI, and also if they themselves would view their coworkers less favorably if the coworker used generative AI for a workplace task. Lastly, we asked whether participants use generative AI to ask questions they are shy to ask at work.

Table 4 displays participant responses to these questions. The results in Table 4 provide some evidence of pluralistic ignorance in the reputational consequences of generative AI. On the one hand, respondents themselves do not judge their coworkers less favorably if their coworker used a large language model for their work (84 percent versus 5 percent). However, a larger percentage of respondents believe that their coworkers and managers would have negative evaluations of their work if use of generative AI was revealed. For example, 20 percent of respondents believe their co-workers would negatively judge their work if they used generative AI. Interestingly, the magnitude is much smaller for manager evaluations, where only 8 percent would be uncomfortable with their manager knowing they used generative AI for their work. Thus, it appears that individuals believe there will be negative evaluations from co-workers, not managers. Moreover, almost 30 percent of generative AI workplace users use the technology to ask questions they are shy to ask about.

Interestingly, these reputational considerations are larger for PhD economists than non-economists. PhD economists are 25 percentage points less likely to be comfortable with their manager knowing they used a large language model for their work (61 percent versus 86 percent), and more than twice as likely to not want their coworker to know about their generative AI usage at work (15 versus 7 percent).

4.3 Trust and knowledge of generative AI

Our survey also included questions on trust and knowledge of generative AI. We included three questions regarding trust: an aggregate trust question, plus questions on the two-subcomponents of trust: trust in competence and trust in motives (Twyman et al., 2008). Our trust in competence question asked the extent to which employees trusted the output of generative AI, while the trust in motives question gauged the extent to which participants believed that the development of generative

AI prioritized ethical considerations and fairness concerns.

Our survey also included four questions to measure participant knowledge of generative AI. Three of these questions had correct answers, and we wanted to measure if participants understand that large language models are statistical models of text prediction, that they do not "understand" the meaning or context of their output, and that current large language models do not dynamically update their knowledge bases. The last question was an attitudinal question, which asked participants the extent to which they agree they have no idea how generative AI works. Lastly, we asked participants to report how much more or less productive they would be working with a large language model at the bank, versus working without using one at the bank?

Table 5 displays the average responses for these questions. No consistent pattern emerges in the trust questions. Although participants are more likely to trust large language models (53 percent) versus not trust them (28 percent), they seem skeptical about the output of large language models and their development.

Meanwhile, participants seem decently informed about large language models. 59 percent disagree that they have no idea how large language models work (versus 22 percent who agree). 80 percent understand that large language models are statistical models of prediction, and over half understand that these models do not understand the meaning or context behind their output. One area of misinformed beliefs is regarding the knowledge base of these tools: over 80 percent of participants wrongly believe that current large language models dynamically update their knowledge bases. In terms of the subsamples, PhD economists and generative AI workplace users are more informed on generative AI than non-PhD economists and non-workplace users. Lastly, on average, employees believe that they would be almost 14 percent more productive if they used a large language model at work.[6] Men believed that the productivity benefits of generative AI were 26 percent $(= (14.0 - 11.1)/11.1)$ higher than woman, and generative AI workplace users believed the productivity benefits were around 60 $(= (16.6 - 10.3)/10.3)$ percent higher than non-workplace users.

[6]See Appendix Figure A.4 for the distribution of responses for this question, both on average and then for various subgroups.

4.4 Organizational readiness and workplace re-design

Our survey also included various items regarding organizational readiness for generative AI and workplace redesign. Indeed, a large literature emphasizes that new technologies often require complementary changes within organizations, including the redesign of business processes, changes to the product and service mix of the organization, and changes to human resource management practices. (Hitt and Brynjolfsson, 1997; Bresnahan et al., 2002; Bartel et al., 2007; Aral et al., 2012; Brynjolfsson et al., 2021a,b) Although it is likely that widespread adoption of generative AI will require fewer complementary organizational changes than general-purpose technologies of the past (McAfee et al., 2023), there are still important changes regarding employee training and privacy rules, among others.

Our survey included various items to measure employee beliefs regarding these changes. One survey item asked the extent to which participants agreed that effectively using generative AI will require structural changes to the organization. We also asked the extent to which respondents believe that there is organizational support for generative AI from both bank executives and team managers. We asked two questions regarding awareness of the organization's rules for using generative AI at work, and how easy it is to follow these rules. Lastly, we included three survey items to gauge which organizational practices would be most likely to increase the use of generative AI.

Table 6 displays the average responses for the survey items related to organizational readiness and workplace re-design. Overall, the sample was split on whether effectively using generative AI would require structural changes at the bank (43 percent agree versus 34 percent disagree), with limited differences across samples. There was more agreement on executive and managerial support for generative AI, with over 50 percent agreeing versus less than 10 percent disagreeing. This also included interesting differences, whereby employees are 9 percentage points more likely to agree that team managers support generative AI usage versus bank executives. Participants were split on their knowledge and awareness of the organization's rules regarding generative AI at work, with an average score of 0.1 and 0.2. Instead, 37 percent were uncertain about the simplicity and ease of following the bank's rules on generative AI. Lastly, employees gave the most support for successful use cases in increasing generative AI use at work, followed by clear support from bank executives.

5 What predicts generative AI usage at work?

In this section, we explore various factors that may drive the lack of generative AI adoption observed in the sample. Our goal is to compare outcomes for those who use generative AI versus those who do not, to shed light on potential ways for organizations to increase the adoption of generative AI in the workplace. Understanding these differences and their source is key to increasing workplace adoption and targeting the right organizational interventions. For example, showcasing opportunities for upskilling using generative AI will likely be effective in increasing usage among those with higher fears of automation, while managerial training will likely be more effective for those whose reputational considerations hamper adoption.

5.1 Workplace tasks and generative AI usage

We first investigate whether there is a relationship between workplace tasks and a worker's propensity to use generative AI at work. Our analysis here is based upon the task-based approach, which posits that jobs are made up of tasks that require different types of skills (Autor et al., 2003; Autor, 2013; Autor and Handel, 2013). To the extent that these job tasks differ to the extent which generative AI is useful, we would expect differences in generative AI usage. While the task-based approach has been useful for studying skill-biased technological change (see, for example, Autor et al. (2003) and Acemoglu and Autor (2012), we use it here to examine the relationship between workplace tasks and generative AI usage at work.

Our survey used the taxonomy introduced from the task-based approach, and asked respondents to report the number of hours they spend on various task categories. The task categories included: (i) routine cognitive work, (ii) non-routine cognitive work; (iii) routine manual work; (iv) non-routine manual work; (v) analytical work; (vi) leadership and management work; and (vii) interpersonal and communication work. As illustrated in Table A.1, 28 percent of the sample's working time is spent on analytical tasks, 27 percent on non-routine cognitive tasks, 19 percent on routine cognitive tasks, and 16 percent on leadership and management tasks.

Figure 4a displays the results of regressions that predict various measures of generative AI usage on the task composition of work. We use three measures of generative AI usage: (i) a binary indicator if the respondent has ever used a large language model; (ii) a binary indicator whether they use a large

language model at work; and (iii) a continuous measure that captures the percent of their workweek that is spent using large language models.

The results indicate a strong relationship between workplace tasks and generative AI usage at work. Respondents who spend a larger percent of their workweek on analytical, non-routine cognitive, and leadership and management tasks are more likely to use generative AI tools at work. The results further suggest a positive relationship between generative AI usage and interpersonal and communication tasks, though our estimates are noisier here. Meanwhile, there was no relationship between generative AI usage and the proportion of work that was routine, either as routine cognitive or routine manual, nor in non-routine manual work.

These results suggest that workplace task composition is one reason why some workers are likely to use generative AI while others are not. However, this is not the entire story: we find that workplace task composition accounts for just 12 percent of the variation in generative AI usage at work.

5.2 Worker beliefs and generative AI usage

We additionally investigate whether worker beliefs can explain differences in workplace usage in our sample. To do so, we aggregate our belief survey items into four belief scores: one on automation fears, one on trust in generative AI, one on reputational concerns with generative AI, and one on productivity beliefs. Each index is normalized to have mean zero and standard deviation one. We then regress our various measures of generative AI usage at work, on these indices. Given the strong prognostic ability of task make-up documented in section 5.2, we additionally control for this in our regressions.

The results in Figure 4b show a strong relationship between participant beliefs regarding the productivity impact of generative AI and likelihood of using generative AI at work. Conditional on the same job type, participants who have larger beliefs in the productivity impact of generative AI are much more likely to use generative AI at work. This lines up with the results in Table 5, which shows that generative AI users report an average productivity impact of 16.6 percent versus 10.3 percent for non-workplace users. The results in Figure 4b control for the task make-up of one's job and suggest that the presentation of successful use cases may be especially useful in boosting generative AI adoption in the workplace.

The results in Figure 4b also highlight the role that reputational considerations play in adoption

decisions. Those who report higher reputational fears associated with generative AI usage are less likely to use generative AI at work. This suggests that fostering a culture of acceptance regarding generative AI usage may be key to increasing adoption. Interestingly, we see no evidence of a relationship between trust in generative AI and usage at work.

Overall, generative AI workplace adoption seems to be driven by the task make-up of work, individuals' beliefs regarding the productivity impact of generative AI, and reputational concerns regarding generative AI usage at work.

6 Conclusion

Overall, the results in this paper provide novel evidence on how generative AI is being used in central banking. The results suggest several opportunities and challenges in unlocking the full productivity benefits of generative AI in the workplace. Among current workplace users of generative AI, there are barriers in extending usage from general to more complex and domain-specific applications that require higher levels of expertise. Additionally, workplace users exhibit a heavy reliance on a single generative AI tool, ChatGPT, across virtually all workplace tasks rather than exploring and leveraging more specialized tools. This singular focus on ChatGPT may limit productivity gains from more strategic task-tool matching with generative AI. For non-users of generative AI at work, the results point to several possible explanatory factors like the task make-up of work, their beliefs regarding the productivity impact of generative AI, and reputational risks. Understanding and addressing these perceived barriers is key to driving broader adoption.

Looking ahead, there are several important opportunities for researchers and practitioners to not only examine the various ways that workers use generative AI on the job, but also systematically test interventions and change management strategies to enable more productive human-machine augmentation with generative AI. Targeted guidance on optimal task-tool matching and alleviating automation anxiety through upskilling may help organizations better realize the transformative potential of generative AI. This study provides an important empirical foundation for such efforts going forward.

References

Acemoglu, Daron and David Autor, "Chapter 12," in "Handbook of Labor Economics," Vol. 4, Elsevier, 2012, pp. 1043–1171.

Aldasoro, I. et al., "Generative Artificial Intelligence and Cyber Security in Central Banking," BIS Papers 145, Bank for International Settlements 2024.

Aral, Sinan, Erik Brynjolfsson, and Lynn Wu, "Three-Way Complementarities: Performance Pay, Human Resource Analytics, and Information Technology," *Management Science*, 2012, *58* (5), 913–931.

Araujo, D. et al., "Machine Learning in Central Banking," IFC Bulletins 57, Irving Fisher Committee on Central Bank Statistics 2022.

Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K. Khandelwal, and Eric Verhoogen, "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan," *The Quarterly Journal of Economics*, August 2017, *132* (3), 1101–1164.

Autor, David and Michael Handel, "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages," *Journal of Labor Economics*, 2013, *31* (S1), S59 – S96.

Autor, David H, "The "task approach" to labor markets: an overview," *Journal of Labour Market Research*, 2013, *46* (3), 185–199.

Autor, David H., Frank Levy, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 11 2003, *118* (4), 1279–1333.

Bartel, Ann, Casey Ichniowski, and Kathryn Shaw, "How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills*," *The Quarterly Journal of Economics*, 11 2007, *122* (4), 1721–1758.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt, "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *The Quarterly Journal of Economics*, 2002, *117* (1), 339–376.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, January 2021, *13* (1), 333–72.

_ , Danielle Li, and Lindsey R Raymond, "Generative AI at Work," Working Paper 31161, National Bureau of Economic Research April 2023.

_ , Wang Jin, and Kristina McElheran, "The power of prediction: predictive analytics, workplace complements, and business performance," *Business Economics*, October 2021, *56* (4), 217–239.

Békés, Gábor and Péter Harasztosi, "Machine imports, technology adoption, and local spillovers," *Review of World Economics (Weltwirtschaftliches Archiv)*, May 2020, *156* (2), 343–375.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," Working Paper No. 24-013, Harvard Business School Technology & Operations Mgt. Unit September 2023.

Doerr, S. et al., "Big Data and Machine Learning in Central Banking," BIS Working Papers 930, Bank for International Settlements 2021.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," 2023.

Ernst & Young LLP, "AI Anxiety in Business Survey," 2023.

Felten, Edward W., Manav Raj, and Robert Seamans, "Occupational Heterogeneity in Exposure to Generative AI," April 2023. Available at SSRN: https://ssrn.com/abstract=4414065 or http://dx.doi.org/10.2139/ssrn.4414065.

Giannakopoulos, Kostis, Argyro Kavadella, Anas Aaqel Salim, Vassilis Stamatopoulos, and Eleftherios G Kaklamanos, "Evaluation of the Performance of Generative AI Large Language Models ChatGPT, Google Bard, and Microsoft Bing Chat in Supporting Evidence-Based Dentistry: Comparative Mixed Methods Study," *J Med Internet Res*, Dec 2023, *25*, e51580.

Gibbons, Robert and Kevin J Murphy, "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *Journal of Political Economy*, June 1992, *100* (3), 468–505.

Goehring, Grant, Filippo Mezzanotti, and S. Abraham Ravid, "Technology Adoption and Career Concerns: Evidence from the Adoption of Digital Technology in Motion Pictures," June 2024. Working Paper.

Gupta, Apoorv, Jacopo Ponticelli, and Andrea Tesei, "Language Barriers, Technology Adoption and Productivity: Evidence from Agriculture in India," Working Paper 27192, National Bureau of Economic Research May 2020.

Hagos, D.H. et al., "Recent Advances in Generative AI and Large Language Models: Current Status, Challenges, and Perspectives," *IEEE Transactions on Artificial Intelligence*, 2024.

Hitt, Lorin M. and Erik Brynjolfsson, "Information Technology and Internal Firm Organization: An Exploratory Analysis," *Journal of Management Information Systems*, 1997, *14* (2), 81–101.

Holmstrom, Bengt, "Managerial Incentive Problems: A Dynamic Perspective," NBER Working Papers 6875, National Bureau of Economic Research, Inc January 1999.

Humlum, Anders and Emilie Vestergaard, "The Adoption of ChatGPT," Technical Report 16992, IZA Discussion Paper 2023.

Jones, Steven E., *Rebels Against The Future: The Luddites And Their War On The Industrial Revolution: Lessons For The Computer Age*, Reading, MA: Addison-Wesley Publishing Company, Apr 1996.

Laskar, Md Tahmid Rahman, M Saiful Bari, Mizanur Rahman, Md Amran Hossen Bhuiyan, Shafiq Joty, and Jimmy Huang, "A Systematic Study and Comprehensive Evaluation of ChatGPT on Benchmark Datasets," in Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, eds., *Findings of the Association for Computational Linguistics: ACL 2023*, Association for Computational Linguistics Toronto, Canada July 2023, pp. 431–469.

Leffer, Lauren, "'AI Anxiety' Is on the Rise—Here's How to Manage It," *Scientific American*, October 2023.

Maršál, Aleš and Patryk Perkowski, "The impact of generative AI in central banking: Evidence from the National Bank of Slovakia," *Work in progress*, 2024.

McAfee, Andrew, Daniel Rock, and Erik Brynjolfsson, "How to capitalize on generative AI," *Harvard Business Review*, 2023, *101* (6), 42–48.

Merchant, Brian, *Blood in the Machine: The Origins of the Rebellion Against Big Tech*, Little Brown and Company, Sep 2023.

_ , "The New Luddites Aren't Backing Down," *The Atlantic*, February 2024.

Nadiri, M. Ishaq, "Innovations and Technological Spillovers," Working Paper 4423, National Bureau of Economic Research August 1993.

Noy, Shakked and Whitney Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, 2023, *381* (6654), 187–192.

Rastogi, Ananya, "Impact of spillovers on technological changes," *Nature Computational Science*, 2023, *3* (11), 914–914.

Sleiman, J.P., "Generative Artificial Intelligence and Large Language Models for Digital Banking: First Outlook and Perspectives," *Journal of Digital Banking*, 2023, *8* (2), 102–117.

Twyman, Matt, Nigel Harvey, and Clare Harries, "Trust in motives, trust in competence: Separate factors determining the effectiveness of risk communication," *Judgment and Decision Making*, 2008, *3* (1), 111–120.

Venkatesh, Viswanath and Michael G. Morris, "Why Don't Men Ever Stop to Ask for Directions? Gender, Social Influence, and Their Role in Technology Acceptance and Usage Behavior," *MIS Quarterly*, 2000, *24* (1), 115–139.

Tables

Table 1: **Summary statistics**

	Full Sample N=345		NBS sample N=162		Economist sample N=183	
	Mean	SD	Mean	SD	Mean	SD
Education, less than high school	0.00	0.00	0.00	0.00	0.00	0.00
Education, high school or associate's degree	0.03	0.16	0.03	0.16	0.00	0.00
Education, university	0.04	0.19	0.04	0.19	0.01	0.11
Education, masters degree or higher	0.90	0.31	0.90	0.31	0.98	0.15
Age, in years	40.71	8.98	40.71	8.98	39.62	10.88
Female	0.37	0.48	0.37	0.48		
Tenure at bank, in years	8.89	10.33	8.89	10.33	9.07	7.78
Manages or leads team	0.10	0.30	0.10	0.30	0.25	0.43
% of work in English	0.30	0.26	0.30	0.26		
Total weekly hours	47.18	11.67	47.18	11.67	42.42	7.60
Percent of work week, routine cognitive	0.19	0.13	0.19	0.13	0.15	0.10
Percent of work week, non-routine cognitive	0.27	0.16	0.27	0.16	0.21	0.12
Percent of work week, routine manual	0.03	0.06	0.03	0.06	0.03	0.05
Percent of work week, non-routine manual	0.02	0.06	0.02	0.06	0.02	0.04
Percent of work week, analytical	0.28	0.19	0.28	0.19	0.36	0.18
Percent of work week, leadership and management	0.16	0.11	0.16	0.11	0.08	0.10
Percent of work week, interpersonal and communication	0.04	0.07	0.04	0.07	0.15	0.08

Notes: This table displays summary statistics for the respondents. The full sample columns pool respondents from all three samples. The NBS sample captures respondents across the entire National Bank of Slovakia, while the economist sample captures responses from PhD economists at the ECB and CNB.

Table 2: **Generative AI tools used across tasks**

	Percent that have used each GAI tool						Average number of GAI tools used
	ChatGPT	Gemini	Co-Pilot	LLama	Claude	Other	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
General sample							
Email and reports	0.87	0.11	0.21	0.03	0.00	0.05	1.26
Code and debugging	0.91	0.12	0.21	0.06	0.00	0.00	1.29
Creating content	0.95	0.16	0.21	0.11	0.05	0.00	1.47
Data analysis	1.00	0.00	0.00	0.00	0.00	0.11	1.11
Educational purposes	0.92	0.21	0.26	0.03	0.03	0.00	1.45
Research	0.91	0.30	0.35	0.04	0.04	0.04	1.70
Design and creative work	0.67	0.11	0.11	0.00	0.00	0.11	1.00
Translation	0.79	0.12	0.26	0.06	0.03	0.12	1.38
PhD economists							
Workplace or personal tasks	0.99	0.13	0.13	0.12	0.04	0.07	1.42

Notes: This table examines the generative AI tools used across the sample. The first six columns display the percent of respondents who have used each GAI tool for each task, while the last column displays the average number of GAI tools for each task. Each row limits the sample to respondents who have reported using generative AI for the given task. The data on tool usage across workplace tasks was collected for the general sample, while the PhD econonmist sample reported their tool use at a more general level (whether they used the tool for any workplace or personal task).

Table 3: **Survey responses on automation fears and workplace changes related to generative AI**

	Overall sample			Average rating, by subsamples							
	Average rating	% disagree	% agree	Background		Gender		Team role		GAI workplace	
				Econ PhD	Other	Women	Men	Manager	Non-manager	Users	Nonusers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
I worry my skills will be less valuable with GAI	-1.2	0.70	0.18	-1.0	-1.4	-1.4	-1.5	-1.3	-1.2	-1.4	-1.1
GAI will complement labor, not substitute for it	1.3	0.14	0.77	1.4	1.2	1.4	1.1	1.0	1.4	1.3	1.3
Advances in GAI will lead to massive unemployment	-1.4	0.79	0.11	-1.5	-1.4	-1.6	-1.2	-1.3	-1.5	-1.5	-1.4
I worry my job is at risk of automation	-1.6	0.79	0.12	-1.5	-1.7	-1.8	-1.6	-1.7	-1.5	-1.6	-1.6
GAI will reduce my workplace autonomy	-1.5	0.80	0.09	-1.5	-1.6	-1.7	-1.5	-1.4	-1.5	-1.7	-1.3
GAI will increase workplace monitoring	-0.7	0.51	0.23	-0.9	-0.4	-0.5	-0.4	-0.7	-0.7	-0.8	-0.5

Notes: This table displays survey responses for questions related to automation fears and workplace changes related to generative AI. Column 1 displays the average rating given to each question. The rating ranges from -3 (average response of strongly disagree) to 3 (average response of strongly agree), with an average response of 0 as indifferent between the two. The aggregated "agree" response combine respondents who replied "strongly agree", "agree", and "somewhat agree", while the aggregated "disagree" response combine respondents who replied "strongly disagree", "disagree", and "somewhat disagree". The remaining columns give average ratings across various subsamples, including those with economics PhDs vs not (columns 4 and 5), gender (columns 6 and 7), team role (columns 8 and 9), and GAI workplace usage (columns 10 and 11).

Table 4: **Survey responses on reputational considerations associated with generative AI workplace usage**

	Overall sample			Average rating, by subsamples							
	Average rating	% disagree	% agree	Background		Gender		Team role		GAI workplace	
				Econ PhD	Other	Women	Men	Manager	Non-manager	Users	Nonusers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
I would be comfortable with my manager knowing I used a LLM for my work	1.5	0.08	0.80	1.1	1.9	1.7	2.0	1.3	1.6	1.5	1.5
My coworkers would judge my work negatively if I used a LLM	-1.0	0.61	0.20	-0.9	-1.1	-1.0	-1.1	-0.8	-1.0	-1.1	-0.9
I would not want my coworker to know if I used a LLM	-1.3	0.75	0.13	-1.0	-1.6	-1.5	-1.7	-1.1	-1.4	-1.4	-1.2
I would view my coworker less favorably if they used a LLM	-1.8	0.84	0.05	-1.7	-2.0	-1.9	-2.1	-1.6	-1.9	-2.0	-1.6
I use GAI to ask questions I am shy to ask to people	-0.8	0.53	0.25	-0.5	-1.0	-0.9	-1.1	-1.0	-0.7	-0.7	-0.9

Notes: This table displays survey responses for questions on reputational considerations associated with generative AI workplace usage. Column 1 displays the average rating given to each question. The rating ranges from -3 (average response of strongly disagree) to 3 (average response of strongly agree), with an average response of 0 as indifferent between the two. The aggregated "agree" response combine respondents who replied "strongly agree", "agree", and "somewhat agree", while the aggregated "disagree" response combine respondents who replied "strongly disagree", "disagree", and "somewhat disagree". The remaining columns give average ratings across various subsamples, including those with economics PhDs vs not (columns 4 and 5), gender (columns 6 and 7), team role (columns 8 and 9), and GAI workplace usage (columns 10 and 11).

Table 5: **Survey responses on trust and knowledge of generative AI**

	Overall sample			Average rating, by subsamples							
	Average rating	% disagree	% agree	Background		Gender		Team role		GAI workplace	
				Econ PhD	Other	Women	Men	Manager	Non-manager	Users	Nonusers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
I do not trust LLMs	-0.6	0.53	0.28	-0.4	-0.7	-0.8	-0.7	-0.4	-0.6	-0.7	-0.4
I trust the output of LLMs	-0.4	0.50	0.32	-0.2	-0.6	-0.6	-0.6	-0.5	-0.4	-0.3	-0.5
LLM development prioritizes ethical considerations, fairness, and the well-being of users	-0.4	0.42	0.22	-0.8	0.0	0.2	-0.1	-0.7	-0.4	-0.6	-0.3
LLMs are statistical models of text prediction	1.5	0.08	0.80	1.7	1.3	1.1	1.4	1.3	1.5	1.8	1.1
LLMs "understand" the meaning or context behind their output	-0.6	0.55	0.30	-1.0	-0.3	-0.3	-0.3	-0.9	-0.6	-0.7	-0.6
Current LLMs dynamically update their knowledge bases	1.4	0.11	0.81	1.1	1.7	1.8	1.7	1.3	1.4	1.3	1.4
I have no idea how LLMs work	-0.8	0.59	0.22	-0.9	-0.7	-0.3	-1.0	-0.7	-0.8	-1.2	-0.4
Belief of productivity impact of LLM	13.6			14.2	13.0	11.1	14.0	12.8	13.8	16.6	10.3

Notes: This table displays survey responses for questions related to trust and knowledge of generative AI. Column 1 displays the average rating given to each question. The rating ranges from -3 (average response of strongly disagree) to 3 (average response of strongly agree), with an average response of 0 as indifferent between the two. The aggregated "agree" response combine respondents who replied "strongly agree", "agree", and "somewhat agree", while the aggregated "disagree" response combine respondents who replied "strongly disagree", "disagree", and "somewhat disagree". The remaining columns give average ratings across various subsamples, including those with economics PhDs vs not (columns 4 and 5), gender (columns 6 and 7), team role (columns 8 and 9), and GAI workplace usage (columns 10 and 11).

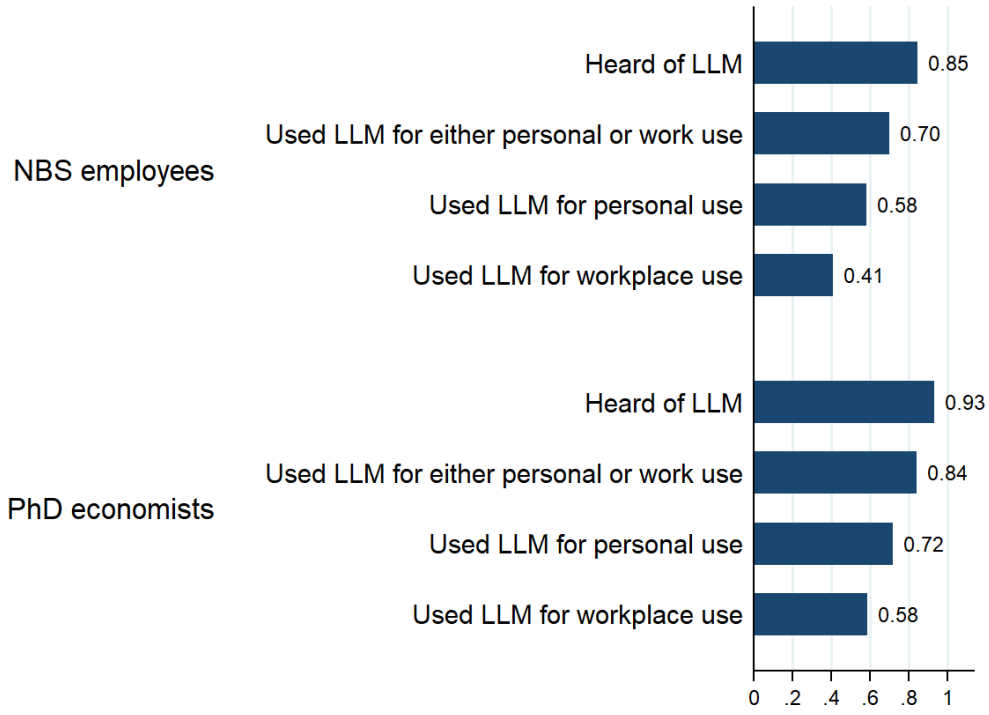
Table 6: **Survey responses on organizational policies regarding generative AI at work**

	Overall sample			Average rating, by subsamples							
	Average rating	% disagree	% agree	Background		Gender		Team role		GAI workplace	
				Econ PhD	Other	Women	Men	Manager	Non-manager	Users	Nonusers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Effectively using GAI requires structure changes at the bank	0.1	0.34	0.45	0.1	0.0	-0.1	0.0	0.1	0.0	-0.1	0.2
Bank executives support GAI	0.6	0.11	0.49	0.8	0.4	0.4	0.5	0.7	0.6	0.6	0.6
My manager supports me using GAI at work	0.9	0.06	0.58	0.9	0.9	0.9	0.9	1.0	0.9	1.1	0.6
I am aware of the bank's rules regarding LLMs at work	0.2	0.35	0.53	0.4	0.1	-0.3	0.3	0.4	0.2	0.8	-0.4
It is simple and easy to follow the bank's rules regarding LLMs at work	0.1	0.28	0.35	-0.1	0.3	-0.1	0.5	0.0	0.1	0.3	-0.1
More info about benefits and usages would increase my use of LLMs at work	1.6	0.05	0.83	1.6	1.5	2.0	1.3	1.5	1.6	1.5	1.7
Demonstrations of use cases would increase my use of LLMs at work	2.0	0.03	0.94	2.0	1.9	2.0	1.9	1.9	2.0	2.0	2.0
Clear support from the bank would increase my use of LLMs at work	1.7	0.02	0.88	1.6	1.9	1.9	1.8	1.6	1.7	1.7	1.7

Notes: This table displays survey responses for questions related to organizational policies regarding generative AI at work. Column 1 displays the average rating given to each question. The rating ranges from -3 (average response of strongly disagree) to 3 (average response of strongly agree), with an average response of 0 as indifferent between the two. The aggregated "agree" response combine respondents who replied "strongly agree", "agree", and "somewhat agree", while the aggregated "disagree" response combine respondents who replied "strongly disagree", "disagree", and "somewhat disagree". The remaining columns give average ratings across various subsamples, including those with economics PhDs vs not (columns 4 and 5), gender (columns 6 and 7), team role (columns 8 and 9), and GAI workplace usage (columns 10 and 11).

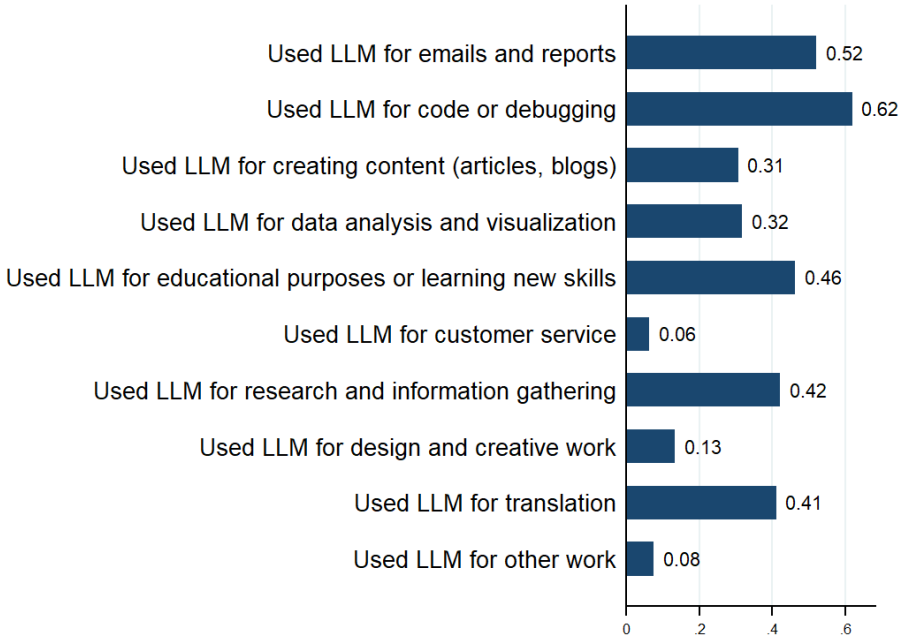
Figures

Figure 1: **Generative AI exposure and use**

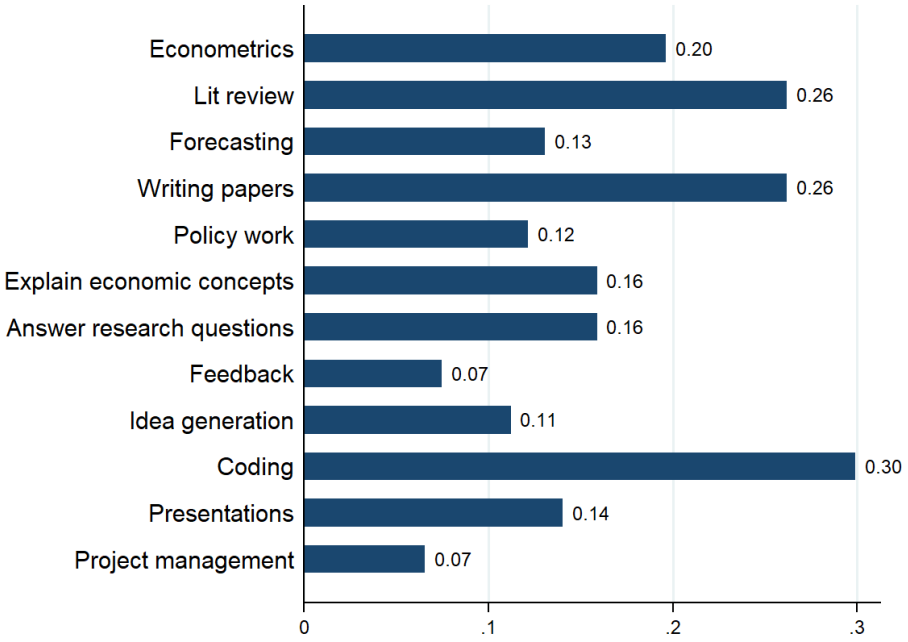


Notes: This figure displays exposure to and use of generative AI tools in our sample.

Figure 2: **Generative AI usage by task, among those who use generative AI at work**

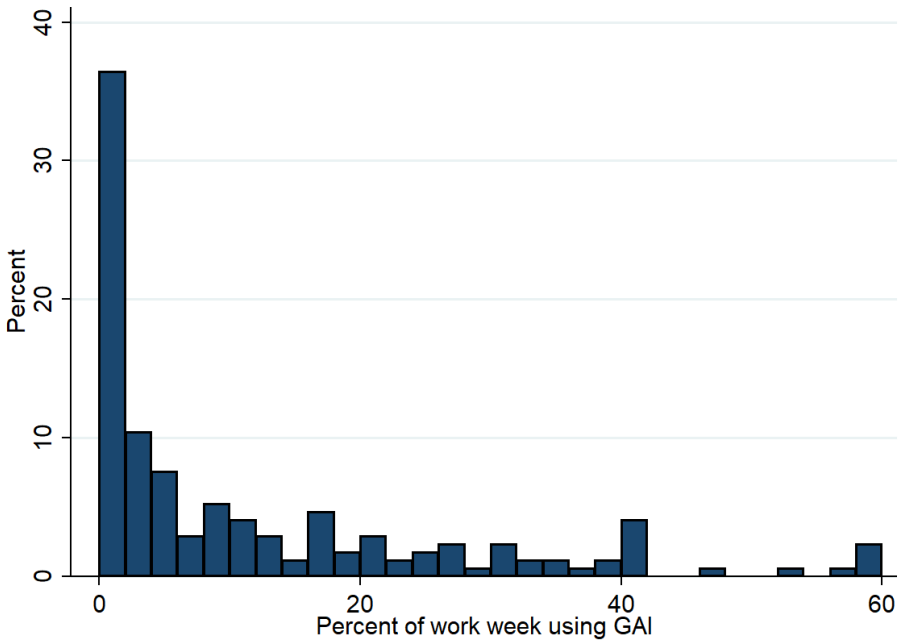


(a) Percent of generative AI workplace users who report using generative AI for each general task

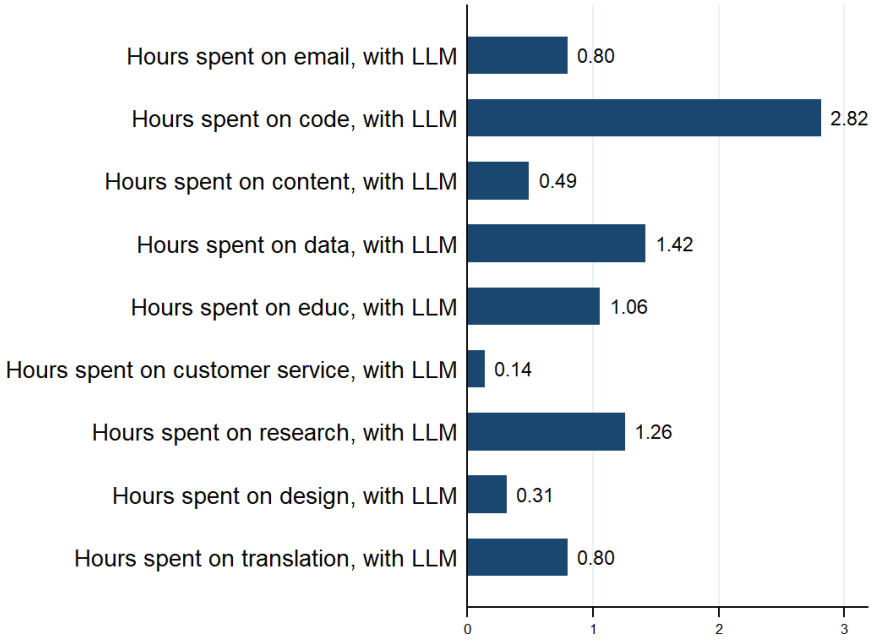


(b) Percent of generative AI workplace users who report using generative AI for each economist task

Figure 3: **Generative AI usage by time, among those who use generative AI at work**

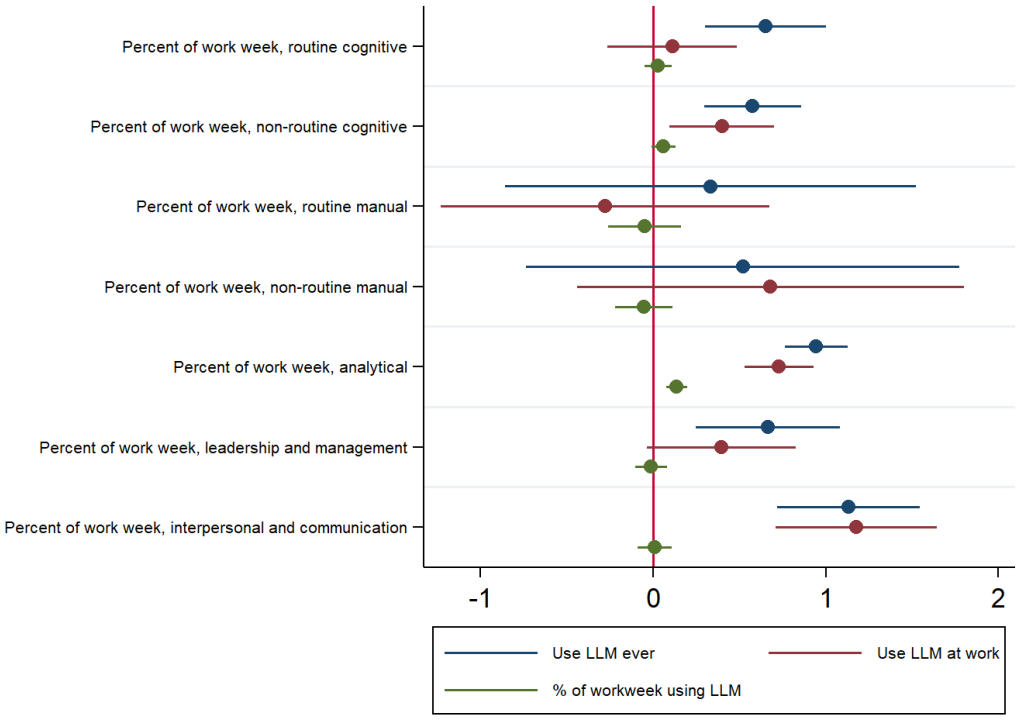


(a) Percent of work week spent interacting with generative AI tools, among those who use generative AI at work

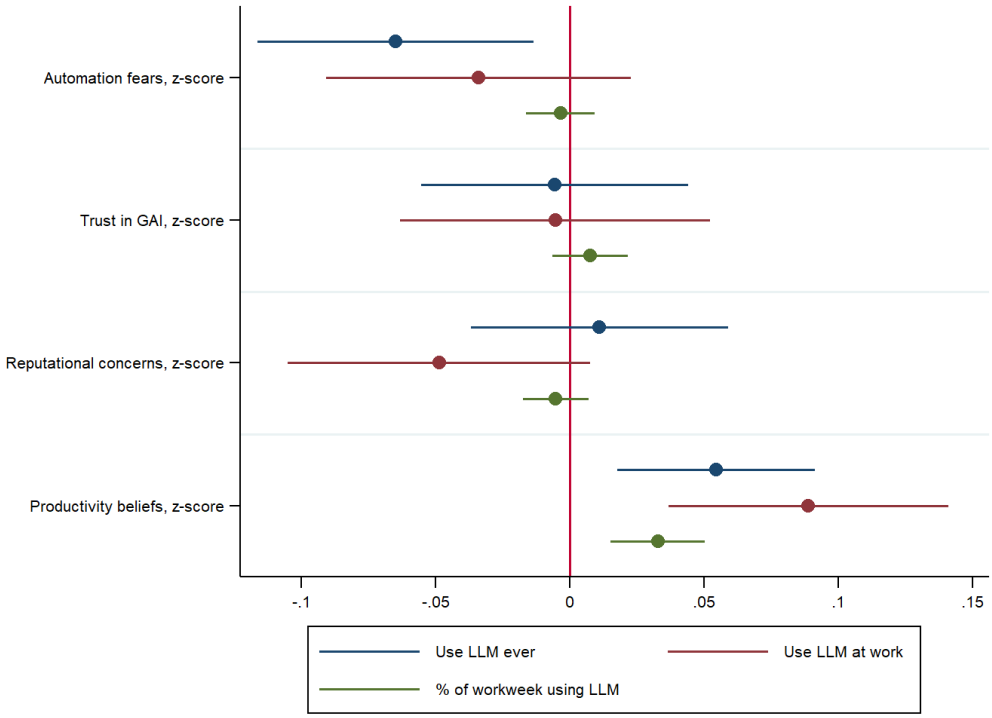


(b) Mean number of hours spent interacting with generative AI tools, by task

Figure 4: **Predictors of generative AI usage at work**



(a) Task make-up



(b) Beliefs

Appendix: For Online Publication Only

A Additional survey results

A.1 Rates of exposure to and workplace usage of generative AI

A.1.1 Summary statistics by usage of generative AI

Table A.1 displays summary statistics depending on the usage of the participant.. Column 1 displays summary statistics for respondents who have never used generative AI, column 2 for those who have used generative AI but not at work, and column 3 for those who use generative AI at work. Meanwhile, column 4 displays summary statistics across the entire sample of respondents.

The results in Table A.1 indicate some demographic and workplace differences in those who use generative AI at work, versus those who do not. Those who use generative AI at work are on average younger (37.8 years old vs 43 years old), are more likely to be male (30 percent of generative AI users are female, while between 40 and 45 percent of non-users are female), and spend more of their workload using English (36 percent vs between 20 and 32 percent for non-work users).

Table A.1: **Summary statistics**

	Never used LLM before	Used LLM, not at work	Used LLM, at work	Entire Sample
	1	2	3	Total
Education, high school	0.04	0.00	0.03	0.03
Education, college	0.04	0.05	0.03	0.04
Education, more than college	0.90	0.89	0.90	0.90
Age	42.98	42.66	37.79	40.66
Female	0.40	0.45	0.30	0.37
Tenure	8.95	9.71	8.19	8.81
Manager	0.16	0.05	0.07	0.10
Has laptop	0.90	0.92	0.97	0.94
Percent of workload in English	0.20	0.32	0.36	0.30
Total work hours reported	50.01	46.22	45.69	47.21
Percent of work, routine- cognitive tasks	0.22	0.22	0.15	0.19
Percent of work, non-routine cognitive tasks	0.27	0.26	0.27	0.27
Percent of work, routine manual tasks	0.05	0.04	0.02	0.03
Percent of work, non-routine cognitive tasks	0.03	0.02	0.02	0.02
Percent of work, analytical tasks	0.24	0.26	0.32	0.28
Percent of work, leadership and management tasks	0.15	0.16	0.17	0.16
Percent of work, interpersonal and communication tasks	0.04	0.02	0.05	0.04

Notes: This table displays summary statistics by usage of LLMs in the workplace.

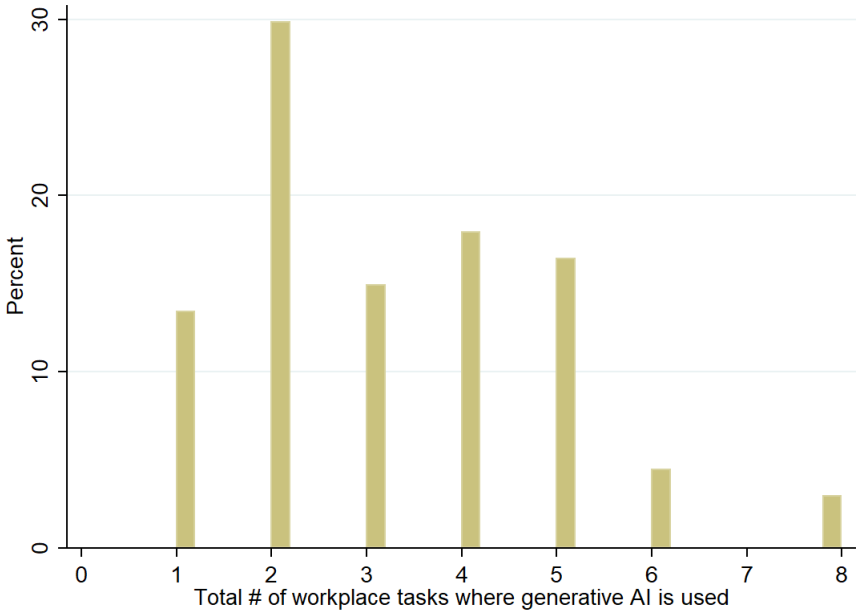
A.2 Workplace usage of generative AI

A.2.1 Distribution of number of tasks where generative AI is used

Figure A.1 displays a histogram of the number of tasks where generative AI is used, for those who use generative AI at work. The median respondent reported using generative AI for three workplace tasks. Around 13 percent of workplace users only use generative AI for one task. Around 40 percent

of users use generative AI for four or more job tasks.

Figure A.1: Distribution of number of general workplace tasks where generative AI is used, for those who use generative AI at work



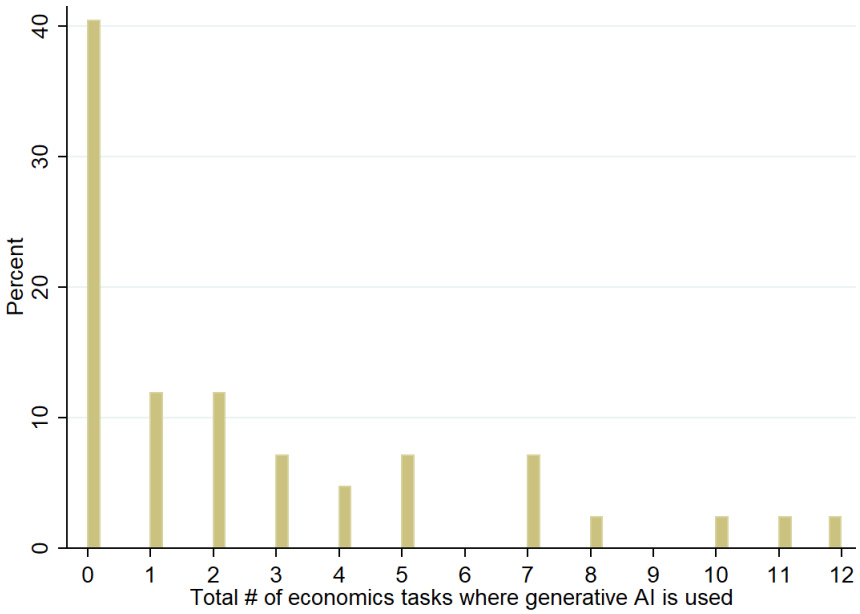
Notes: This figure displays a histogram of the number of workplace tasks where generative AI is used. The sample is limited to respondents who reported using generative AI in the workplace.

Figure A.2 displays the corresponding histogram of the number of economics tasks where generative AI is used, for those who use generative AI at work. The median respondent reported using generative AI for only one economics task. Meanwhile, over 40 percent of respondents reported using it for no economics tasks. These participants report using generative AI for general workplace tasks, but do not use it for more specialized economics tasks.

A.2.2 Distribution of hours spend using generative AI at work

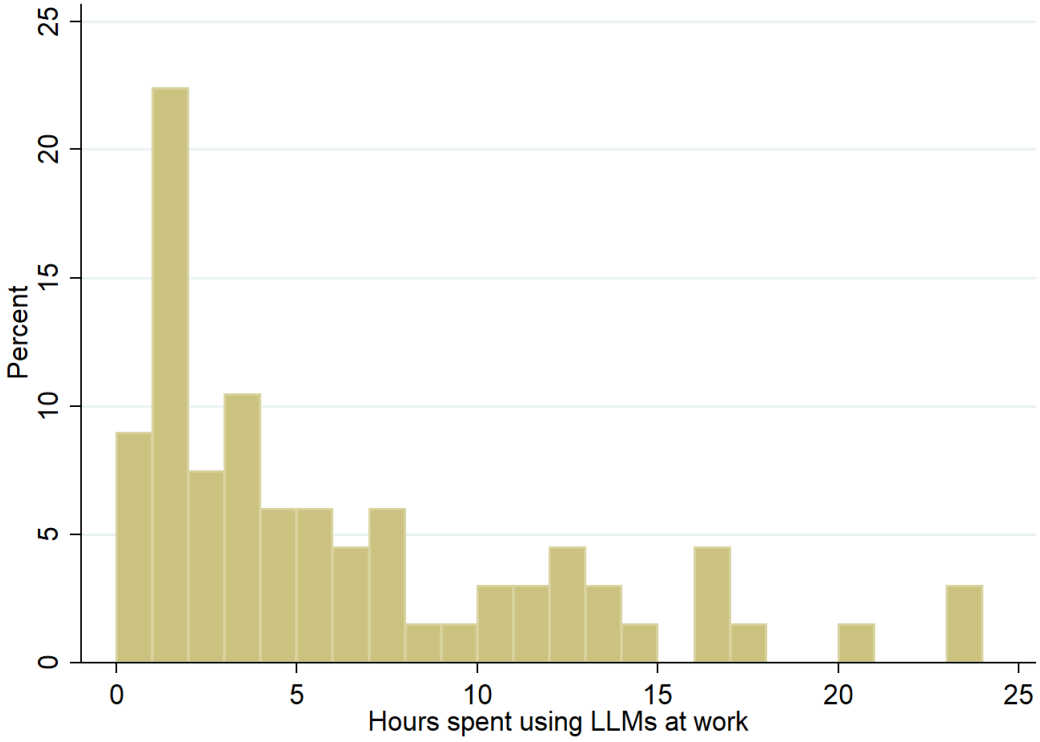
Figure A.3 displays a histogram of the number of hours where generative AI is used, for those who use generative AI at work. The average respondent reported using generative AI for around six hours a week, and the median respondent reported using generative AI for four hours a week. The bottom quarter of users reported using generative AI for up to 1.5 hours a week.

Figure A.2: Distribution of number of economics workplace tasks where generative AI is used, for those who use generative AI at work



Notes: This figure displays a histogram of the number of economics workplace tasks where generative AI is used. The sample is limited to respondents who reported using generative AI in the workplace.

Figure A.3: Distribution of workplace hours where generative AI is used, for those who use generative AI at work

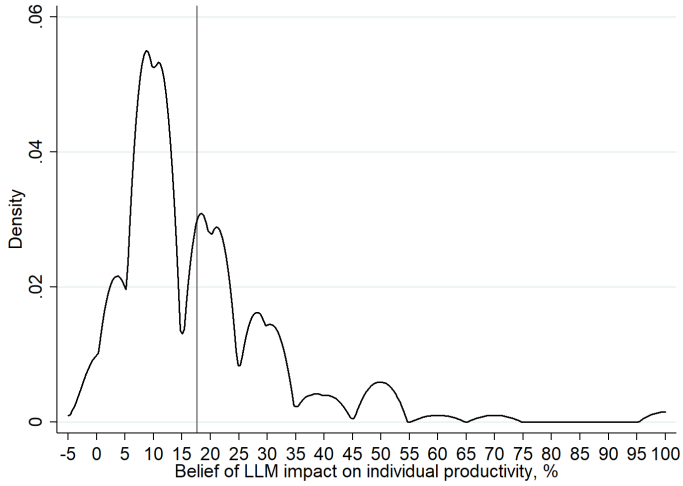


Notes: This figure displays a histogram of the number of workplace hours where generative AI is used. The sample is limited to respondents who reported using generative AI in the workplace.

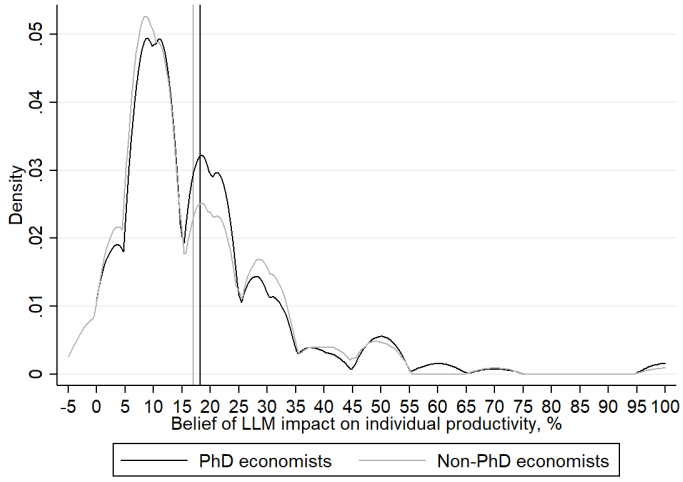
A.3 Productivity impacts of generative AI

In Figure A.4, we display density plots of participant beliefs regarding the productivity impact of generative AI. Panel A displays the overall distribution, Panel B compares PhD economists vs non-PhD economists, and Panel C compares generative AI workplace users versus non-users.

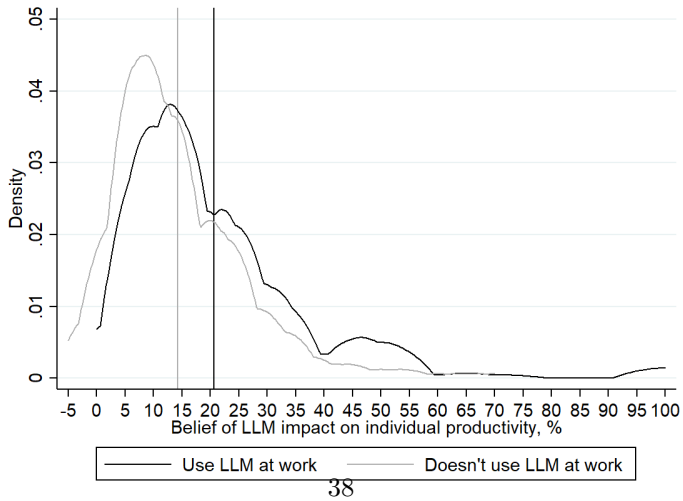
Figure A.4: **Beliefs in the productivity impacts of generative AI, overall and by subgroups**



(a) All participants



(b) PhD economists vs non-PhD economists



(c) Workplace users vs non-workplace users