AI Exposure, Job Insecurity, and Household Financial Risk-Taking

Abstract

We study how generative artificial intelligence (GenAI) affects household financial decisions through labor market insecurity. Exploiting the release of ChatGPT in November 2022 as a plausibly exogenous technological shock, we combine occupation-level and county-level AI exposure measures with individual labor outcomes and household financial data. Workers in more AI-exposed occupations face greater job insecurity: they experience greater layoffs, lower voluntarily quit rates, reduced remote work, and greater resilience on multiple jobs. Households respond by reducing financial risk exposure. In high GenAI exposure counties, households lower equity holdings, especially investments in high-beta stocks, and deleverage by reducing total debt, credit card balances, and mortgage debt. High-exposure counties also experience weaker growth in home-purchase and home-improvement mortgage applications. Together, these findings suggest that AI-driven labor market uncertainty induces conservative financial decisions among exposed households.

JEL classification: D14, G11, G51, J23, O33.

Keywords: Generative AI, job security, household financial decision, equity investment.

I. Introduction

The public release of ChatGPT in November 2022 represents a major shock to the workplace. Generative artificial intelligence can perform, complement, or accelerate a wide range of cognitive tasks, including writing, coding, summarizing, customer support, and data analysis. What makes this technology especially important is not only its scope, but also its speed of diffusion. According to UBS research, ChatGPT reached 100 million monthly active users within two months of its launch, a pace that far exceeds the adoption speed of many earlier transformative technologies. The rapid spread of generative AI suggests that its economic effects may materialize quickly in labor markets and, in turn, in household financial decisions.

The labor market consequences of generative AI are potentially distinct from those of earlier automation technologies. Industrial robots and other forms of automation often affected routine, manual, or production-intensive tasks (Acemoglu and Restrepo, 2020). By contrast, large language models directly affect many nonroutine and white-collar tasks performed by highly educated and relatively high-income workers. Prior research shows that large language models have the potential to affect a substantial share of the U.S. workforce. Eloundou, Manning, Mishkin and Rock (2024) estimate approximately 80 percent of workers may see at least 10 percent of their tasks affected with greater exposure concentrated in higher-wage, higher-education occupations. Other studies document that AI adoption changes firm workforce composition (Babina, Fedyk, He and Hodson, 2023a,b), affects firm valuations and labor demand Eisfeldt, Schubert and Zhang (2023), and reduces demand for some online freelancers (Hui, Reshef and Zhou, 2024; Demirci, Hannane and Zhu, 2025). These findings suggest that generative AI can change workers' expected employment prospects even before all realized employment effects are fully observed.

In this paper, we examine whether job insecurity caused by generative AI affects household financial decisions. Household financial decisions are shaped not only by realized labor income, but also by expectations about future labor income risk. A household anticipating a

higher probability of layoff and weaker bargaining power may reduce its exposure to financial risk even before current income declines. Such behavior can manifest itself through lower stock market exposure, a shift away from risky securities, lower debt balances, delayed mortgage borrowing, or fewer discretionary investments in housing. The relevant mechanism is therefore not only passive deterioration after an income shock, but also active precautionary adjustment in response to greater job insecurity.

The distinction is central to our analysis. Much of the existing evidence in household finance studies shows how unemployment, income losses, or adverse cash-flow shocks affect consumption, borrowing, default, and portfolio choices after households experience a direct economic setback. By contrast, generative AI provides a setting in which households can respond to perceived labor market risk before a real income loss is observed. If workers believe that AI exposure increases the probability of job displacement or weakens future labor income prospects, they can reduce financial risk exposure by changing both the asset side and the liability side of the household balance sheet. This paper provides evidence on this anticipatory channel.

To study this question, we construct measures of AI exposure in occupations and local labor markets at the county level. Following Eloundou et al. (2024) and Eisfeldt et al. (2023), we use task-level information from the O*NET database to evaluate the extent to which generative AI can affect occupational tasks. We then link these occupational exposures to households, and aggregate these exposures to counties using county-industry employment composition. This approach captures how workers and local labor markets are exposed to the disruption of generative AI. Similar to the findings in the literature, we confirm that high-exposure occupations and counties are concentrated in knowledge-intensive sectors such as finance, information, and professional services, while low-exposure areas rely more heavily on physical-labor-intensive sectors such as construction and manufacturing.

Our empirical strategy uses the release of ChatGPT in November 2022 as a quasi-experimental technological shock. We compare changes in labor-market outcomes and house-

hold financial decisions before and after ChatGPT's release across households and counties with different levels of AI exposure. The difference-in-differences framework allows us to compare more exposed and less exposed households (and counties) while absorbing time trends that affect all households or counties. We examine individual labor-market outcomes, household equity and debt decisions using SIPP data, and mortgage application behavior using county-level HMDA data over the 2021-2023 period.

We begin by documenting that AI exposure increases job insecurity. Workers in more AI-exposed occupations become significantly more likely to experience layoffs after ChatGPT's release. A higher exposure to GenAI is associated with an increase in the probability of layoff, and the effect is economically meaningful. For example, the workers in the top-tercile and above-median GenAI exposure groups experience 1.7 and 1.5 percentage points higher layoff probabilities, respectively. Consistent with an increase in perceived job risk, exposed workers are less likely to voluntarily quit for a new job or self-employment. This decline in voluntary quits suggests that workers are less willing to leave existing employment when their outside options become more uncertain.

Other labor market adjustments reinforces this job insecurity interpretation. We find that workers in high-exposure occupations spend less time working from home after the ChatGPT shock, consistent with employees attempting to maintain visibility or strengthen attachment to their employers when job insecurity rises. The economic magnitude is meaningful. For example, the above-median exposure indicator implies a 3.8 percentage point decline. Exposed workers are also more likely to hold multiple jobs simultaneously, suggesting a precautionary labor-supply response to heightened employment uncertainty. Workers may seek additional sources of income or insurance against the possibility of future job loss. Taken together, these findings indicate that the exposure to GenAI is associated with heightened job insecurity rather than only actual layoff in employment.

We then examine whether households respond to this job insecurity by reducing financial risk. Focusing first on the asset side, we find that households with greater AI exposure reduce

equity investment following the release of ChatGPT.[1] Importantly, this portfolio adjustment hold among households with no observed layoffs, higher exposure to GenAI is associated with a significant decline in stock holdings. This indicates that the portfolio adjustment is not merely a response to realized income loss. Instead, households appear to reduce risky assets in anticipation of greater uncertainty about future labor income.

While the SIPP data provides information on the value of household stock investment, they do not identify specific securities. To further examine portfolio composition under job insecurity risk, we draw on SEC 13F filings from local investment advisors, following Pan, Pikulina, Siegel and Wang (2024) in using local advisor portfolios as a proxy for the equity choices of local households. We find that equity investment growth in high-beta stocks decreases by roughly 3 percentage points in more GenAI-exposed areas after ChatGPT's release. In contrast, investment growth in low-beta stocks increases by about 2 percentage points. These findings are consistent with a shift away from high-risk equity positions rather than a mechanical reduction in all investment activity. Households and local investors do not merely reduce equity exposure; they also tilt portfolios away from stocks with greater systematic risk. This evidence supports that AI-driven job insecurity leads exposed households to reduce overall financial risk-taking from equity investment.

We next examine household liability, and find that households also reduce leverage. Among households without observed layoffs, higher exposure to GenAI is associated with lower total debt after the release of ChatGPT. The reduction is sizable across important categories of household borrowing. Credit card balances decline for more exposed households. Mortgage debt against the primary residence also declines, indicating that the deleveraging response is not limited to small or short-term borrowing that carries higher interest rates. These findings show that exposed households adjust not only their investment portfolios but also their debt positions, lowering balance-sheet risk when labor income becomes more uncertain.

[1]The SIPP survey reports the total value of stock holdings at the end of each reference period. We deflate these values to 2021 dollars using S&P 500 returns.

We complement the household-level evidence with county-level mortgage application data from HMDA. In counties subject to greater GenAI exposure, growth in home-purchase mortgage applications declines significantly after the release of ChatGPT. A one percentage point increase in county-level AI exposure is associated with a 0.81 percentage point decline in home-purchase mortgage application growth. The decline is even stronger for home-improvement loan applications. Since home purchases and home improvements are large and often debt-financed commitments, lower application growth suggests that households become less willing to take on new long-term financial obligations. This result complements the household level evidence on lower mortgage balances and indicates that GenAI-driven job insecurity leads households to avoid new housing-related debt.

Overall, the evidence is consistent with a proactive risk-reduction channel. The exposure to GenAI increases perceived labor market risk, and households respond by lowering financial risk, i.e., reducing equity holdings, shifting away from high-beta stocks, deleveraging by lowering credit card and mortgage debt, and becoming less likely to seek new home-purchase or home-improvement loans. These responses are difficult to explain as passive consequences of realized income losses, especially given that several results are estimated among households with no observed layoffs. This pattern instead suggests that households actively adjust financial choices when the disruption of GenAI raises job insecurity, even before that insecurity materializes into actual job loss.

This paper makes several contributions. First, it contributes to the growing literature on generative AI and labor markets. Existing studies document that AI affects firm valuations, workforce composition, labor demand, freelancers, and occupational exposure (Eisfeldt et al., 2023; Babina et al., 2023a,b; Hui et al., 2024; Eloundou et al., 2024; Demirci et al., 2025). We build on this work by showing that the consequences of generative AI extend beyond firms and employment outcomes to household financial choices, and identify the household balance sheet as an additional channel through which AI-driven labor market risk is transmitted.

Second, we contribute to household finance by showing that job insecurity can lead house-

holds to actively reduce financial risk. A large literature links labor income risk and income uncertainty to household portfolio choice and balance-sheet outcomes (Guiso, Jappelli and Terlizzese, 1996; Heaton and Lucas, 2000; Cocco, Gomes and Maenhout, 2005; Ganong and Noel, 2019). Our paper contributes to this literature by identifying an anticipatory channel: households reduce financial risk in response to perceived job insecurity arising from AI exposure before any income loss is realized.

Third, our study contributes to the broader literature on technological change and economic risk. Prior work shows that trade, automation, robots, and AI reshape local labor markets, employment, and wages (Autor, Dorn and Hanson, 2013; Acemoglu and Restrepo, 2020; Acemoglu, Autor, Hazell and Restrepo, 2022; Acemoglu and Restrepo, 2022). Recent studies of generative AI further show that the new technology affects productivity, online labor demand, firm value, and occupational exposure (Noy and Zhang, 2023; Brynjolfsson, Li and Raymond, 2025; Hui et al., 2024; Eisfeldt et al., 2023; Demirci et al., 2025; Eloundou et al., 2024). We add to this literature by showing that generative AI can also affect household balance sheets. Because generative AI diffuses rapidly and directly penetrates many white-collar occupations, households may adjust financial decisions before the long-run labor market consequences are fully known. By linking the exposure to GenAI to individual job insecurity, household balance sheets, and local mortgage demand, we show that technological shocks can propagate from the workplace to household finance through precautionary behavior.

More broadly, our findings suggest that generative AI is not only a productivity shock or a labor demand shock. It is also a household financial risk shock. Even if generative AI raise productivity in the long run, its short-run diffusion can increase uncertainty for workers and lead households to adopt more conservative financial positions. These adjustments may reduce household vulnerability at the individual level, but they may also have broader implications for local credit demand, housing investment, and aggregate risk-taking in financial markets.

The rest of the paper proceeds as follows. Section II discusses the related literature. Section III describes the data and construction of the key measures. Section IV and Section V report the main empirical evidence. Section VI concludes.

II. Literature Review and Background

The rapid adoption of generative AI has attracted significant attention from researchers, with early studies focusing primarily on the implication of AI adoption on labor demand and firm-level outcomes. A burgeoning literature has investigated how AI adoption affects firm-level workforce composition and organizational structure. Babina et al. (2023a) examine firm-level AI investment, using employee resumes to measure firm-level AI investment. They find a spark increase in AI investments across sectors. The research reveals that AI-investing firms experience higher growth in sales, employment, market valuation, and product innovation. However, the same authors also find that AI-investing firms face higher systematic risks (Babina et al., 2023a).

Eisfeldt et al. (2023) examine how generative AI affects firm-level productivity and market valuations. They document significant positive effects on firm performance for companies with high AI exposure, suggesting that early adopters may gain competitive advantages. More directly addressing workforce restructure, Babina et al. (2023b) show that firms investing in AI tend to hire more educated workers, increase junior-level positions and decrease the share of workers in middle-management and senior roles. This research suggests that the investment of AI at workplace has profound organizational changes beyond simply job displacement, which could have significant impact on labor market dynamics.

Recent empirical evidence has begun to document the realized effects of AI adoption on employment outcomes. Using online freelancing platform, Demirci et al. (2025) document a 21% reduction in automation-prone jobs within eight months of ChatGPT's introduction, suggesting that AI's labor market effects materialize more rapidly than previous technological

shocks. Similarly, Hui et al. (2024) also focus on online freelancers labor market and find that generative AI introduction has adverse effects on freelancers' employment opportunities, particularly in tasks involving writing, translation, and basic programming. Acemoglu et al. (2022) study the effect of artificial intelligence on labor markets using establishment-level data on online job vacancies. They observe the growth of AI-related job postings but find that AI-adopted establishments reduce hiring in non-AI positions.

While the research on AI's firm-level and labor market impacts has grown rapidly, evidence on household and local economy effects remains limited. Previous literature has studied how international trade shock affect labor market and local economy. For example, Autor et al. (2013) examine local exposure to international trade and find the exposure leads to employment decline and the persistent effects on employment and wages extended from directly-affected workers to local communities. This research demonstrates that economic shocks in labor market can have far-reaching effects on household and local communities. Other studies examine how automation technological shock affect labor market. For instance, Acemoglu and Restrepo (2020) investigate the impact of industrial robot on local labor market. They find negative automation effects on local employment and wages. Importantly, they find that automation effects are concentrated among workers without college education and in routine manual occupations. The same authors find that automation decreases wage (Acemoglu and Restrepo, 2018) and increases income inequality in the U.S. (Acemoglu and Restrepo, 2022).

The rapid adoption of generative AI in the workplace may influence employment and household finances through two competing channels. On the one hand, displacement effects may reduce earnings, increase income volatility, and heighten financial stress. On the other hand, productivity effects may enhance job performance and strengthen household balance sheets. The net impact therefore remains an open question.

A large literature establishes that labor income risk shapes household financial decisions, providing the theoretical and empirical foundation for our analysis. Portfolio choice models

treat human capital as a nontradable asset and predict that greater labor income uncertainty reduces households' willingness to bear financial risk (Campbell and Viceira, 2002; Viceira, 2001). In these frameworks, workers with riskier labor income optimally hold less equity to avoid compounding human capital risk with financial risk. Consistent with this view, empirical studies find that households facing greater earnings volatility hold fewer risky assets and adopt more conservative portfolios (Angerer and Lam, 2009; Betermier, Jansson, Parlour and Walden, 2012; d'Astous and Shore, 2024). The effects of income uncertainty extend beyond portfolio choice to liability management. Precautionary saving models predict that households facing uninsurable income risk accumulate buffer stocks of safe assets and avoid large financial commitments (Aiyagari, 1994; Carroll, 1997). Consistent with these predictions, empirical work finds that income uncertainty affects debt exposure, mortgage borrowing, and credit access (Campbell, 2006; Matsa, Melzer and Zator, 2024). Taken together, this literature suggests that shocks to perceived labor income risk can lead to broad adjustments in the household balance sheet. Our paper builds on this foundation by identifying generative AI exposure as a novel source of labor income risk and tracing its effects on household portfolio and debt decisions.

Despite the growing literature on AI's economic impacts, existing research focuses primarily on firm-level outcomes and labor market effects, with limited attention to household and local economies. This represents an overlook in the literature, given that household financial stability is central for both economic stability and policy design. This paper aims to fill the gap by examining how generative AI affects household financial decisions, exploiting the ChatGPT's release as a quasi-experimental shock.

III. Data and Methodology

A. *Household Survey Data*

We obtain household financial and individual employment data from the Survey of Income and Program Participation (SIPP), a longitudinal survey that follows households over multiple years through periodic interviews.[2] Our analysis focuses on the 2021–2023 reference period, which includes waves two to four in the 2021 Panel and waves one to three in the 2022 panel. To measure exposure to artificial intelligence, we assign each individual an occupation-level AI exposure measure based on their pre-ChatGPT occupation, prior to the release of ChatGPT in late 2022. We restrict the sample to individuals and households with at least one observation before the ChatGPT shock. When an individual holds multiple jobs, we assign AI exposure based on the highest-earning job. The final sample consists of 5,660 individuals and 3,800 households.

B. *AI Exposure*

In this paper, we construct AI exposure at different levels, i.e., individual, household, and county, in the following steps. Our AI exposure measures are based on occupation-level AI exposure, which is derived from task-level AI exposure.

B.1. Evaluating Occupation-level AI Exposure

Occupational tasks First, we measure the task-level Generative AI exposure using the Occupational Information Network (O*NET 28.1 database), following Eisfeldt et al. (2023). Each occupation has a list of tasks involved. We evaluate whether these tasks can be done more efficiently with the assistance of ChatGPT and other large language models (LLMs) or the applications that can be developed by the LLMs. Details of the evaluation methodology are provided in Appendix section C.

[2]The SIPP survey data is available here: https://www.census.gov/programs-surveys/sipp.html

Evaluating tasks' AI exposure using OpenAI Specifically, we use Open AI's GPT 4.0 Preview model[3] to evaluate the task's AI exposure. Following Eloundou et al. (2024); Eisfeldt et al. (2023), we classify the exposure level into four categories. The classification of the categories are based on whether ChatGPT or other LLMs can shorten the time to complete the task by at least half for the same quality as human work. Specifically, *E0* indicates no exposure as the tool is either insufficiently useful for this task, or cannot be brought to bear as a result of the intrinsic nature of the task. *E1* implies that half reduction in time to compete is feasible with the existing LLM interfaces. *E2* indicates that such a productivity gain is feasible but only when the current capabilities of the model can be deployed through applications with further inputs (e.g., access to internet or proprietary databases), or if it is trained on domain-specific issues or data. *E3* requires the access to both the LLM and a system that could view, caption, and create images as well as any systems powered by the LLM (those in E2 above) that would reduce the time it takes to complete the task by half. It is important to note that among the four categories, *E1* indicates the highest Generative AI exposure as half of time can be reduced to complete the task with the existing LLM interfaces.

We submit the task description and the occupation title to Open AI. Importantly, we do not only require it returns a category of AI exposure but also provide explanations, which allows us to manually audit whether GPT is interpreting the task description correctly and check the creditworthiness of the job. After evaluating task-level AI exposure, we aggregate tasks' AI exposure to the occupation level. Specifically, we calculate the share of tasks for each occupation that can be completed in a shortened time with the help of Generative AI. Following Eloundou et al. (2024); Eisfeldt et al. (2023), we focus on the tasks that affected by Generative AI, which are tasks labeled as E1, E2, or E3. We then construct the occupation-level AI exposure measure as follows:

[3]We performed the task in OpenAI during March 30, 2024 to April 2, 2024.

$$\omega_j = \frac{N_{1,j} + 0.5 \times (N_{2,j} + N_{3,j})}{N_j}, \tag{1}$$

j represents occupation, ω_j is the AI exposure of occupation j, $N_{0,j}$ is the number of tasks that are labeled as *E0*, $N_{1,j}$ is the number of tasks that are labeled as *E1*, $N_{2,j}$ is the number of tasks that are labeled as *E2*, $N_{3,j}$ is the number of tasks that are labeled as *E3*, N_j equals $N_{0,j} + N_{1,j} + N_{2,j} + N_{3,j}$, is the total number of tasks for occupation j.

Occupations' AI exposure We aggregate the task-level exposure to the occupation level (6-digit SOC code) to match the industry-occupation matrix data from the Bureau of Labor Statistics. Table A1 reports the top and bottom 20 AI exposure occupations. Occupations with the highest exposure include clerks, proofreaders, computer programmers, interpreters, and translators. In contrast, occupations with the lowest exposure include roofers, construction laborers, tire repairers and changers, and maids and housekeeping cleaners.

B.2. Evaluating Individual and Household-level AI Exposure

In the individual-level analysis of job security and employment outcomes, we use each individual's occupation-level AI exposure measure. An individual's AI exposure is determined by the AI exposure associated with their occupation.

In the household financial decision analysis, we employ two household-level Generative AI exposure measures constructed as follows. First, we calculate the value-weighted average AI exposure of household income earners, where the weights are based on each individual's monthly earnings within the household. An individual's AI exposure is determined by the AI exposure of their occupation. Second, we construct an alternative measure based solely on the highest earner in the household, using the occupation-level AI exposure associated with the highest earner's occupation.

B.3. Measuring County-level AI Exposure

Next, we create an industry-level AI exposure measure using industry employment patterns from the Bureau of Labor Statistics (BLS). Specifically, we aggregate the occupational-level Generative AI exposure to the industry-level using the 2022 BLS Industry-occupation matrix data. Lastly, we aggregate the industry-level Generative AI exposure to the county-level using the 2022 county-level industry[4] employment data from the Bureau of Economic Analysis. Figure 1 illustrates how we construct the county-level AI exposure step by step from O*NET occupation-task database.

To measure county-level AI exposure, we start with the industry-level first and then use the county-industry employment data to construct the county-level AI exposure measure. We measure the industry-level AI exposure using value-weighted occupational exposure as in Equation 2. The weight is calculated using the 2022 industry-occupation matrix from the Bureau of Labor Statistics.

$$AI_exposure_n = \sum_{j=1}^{J} \omega_j \times \frac{Emp_{j,n,2022}}{Emp_{n,2022}}, \tag{2}$$

where j represents occupation, ω_j is the AI exposure of occupation j, n represents industry (NAICS 2-digit), $Emp_{j,n,2022}$ is for occupation j's employment in industry n in 2022, $Emp_{n,2022}$ is industry n's total employment across all occupations in 2022. Table A3 presents the industry-level AI exposure. The top AI exposure industries include finance and insurance, management of companies, and information, and professional, scientific, and technical services. The bottom AI exposure industries are those industries that require more physical works, for example, manufacturing, mining, oil and gas extraction, construction, and accommodation and food services.

County-level AI exposure We proxy county-level AI exposure using the value-weighted industry exposure, which is determined by county-industry employment matrix in 2022,

[4]Here, industry is defined by NAICS 2-digit.

which is obtained from the Bureau of Economic Analysis. Equation 3 illustrates how we aggregate the industry-level AI exposure to the county-level using the weight of employment of industries in a county.

$$AI_exposure_i = \sum_{n=1}^{N} AI_exposure_n \times \frac{Emp_{n,i,2022}}{Emp_{i,2022}}, \tag{3}$$

where i represents county, n represents industry (NAICS 2-digit), $Emp_{n,i,2022}$ is for industry n's employment in county i in 2022, $Emp_{i,2022}$ is county i's total employment across all industries in 2022.[5] Table A2 presents the list of top and bottom 20 AI exposure counties in the U.S. Figure 2 provides the heatmap of the county-level AI exposure. From the figure, we can infer that the high AI exposure counties widely spread across the country instead of clustering in certain areas.

C. Household Equity Portfolio Choices

In the SIPP household survey, respondents report the value of their stock investments but not the specific stocks they hold. To examine whether Generative AI exposure affects households' portfolio choices, we therefore utilize 13F filings of local financial advisers, which report the names and market values of securities held in advisers' portfolios. Local financial advisers primarily serve local households and therefore provide a useful proxy for county-level household equity investment choices.

First, we follow Pan et al. (2024) to identify local financial advisers. We collect all U.S. advisers data using Form ADV filings from the SEC between 2021 and 2023. We identify advisers that primarily serve individual investors by requiring the percentage of individual clients to be no less than 50% of the adviser's client base or AUM by individual client to be no less than 50% of total AUM in a given year. To focus on advisers who serve locally, we

[5]For robustness, we also construct county-level AI exposure using the value-weighted NAICS 3-digit industry AI exposure. The weight is determined by the personal income of an industry in a county, which is obtained from the Bureau of Economic Analysis "Personal Income by Major Component and Earnings by NAICS Industry."

exclude advisers who report office locations in more than one county.

Next, we manually match the local adviser data with asset managers from Thomson Reuters 13F filings database for 2021-2023, which covers advisers with more than $100 million in Section 13F securities, such as domestic stocks and exchange-traded funds (ETFs). Following Pan et al. (2024), we require our sample advisers to those whose 13F holdings are at least 50% but no more than 110% of their total AUM reported in Form ADV to exclude advisers with large short positions in their portfolio. We identify holding records for about 27% local advisers in our sample. Among 3,057 counties in our original sample, there are 374 counties with identified local advisers and matched 13F equity holding data.

D. Empirical Strategy

We employ a difference-in-differences framework to examine how exposure to generative AI affects household financial decisions, using the release of ChatGPT in November 2022 as the event.

$$Y_{i,t} = \beta GenAIexposure_i \times PostChatGPT_t + \alpha_i + \lambda_t + \epsilon_{i,t}, \tag{4}$$

where index i denotes an individual, household, or county, and t denotes year. The dependent variable, $Y_{i,t}$, represents employment-related or financial outcomes, such as layoff indicators, stock investment, household debt, etc. α_i denotes individual-, household-, or county-level fixed effects, while λ_t represents year fixed effects that capture aggregate macroeconomic shocks.

Our main explanatory variable is the interaction between Generative AI exposure at the individual, household, or county level and an indicator for the post-ChatGPT period. We measure Generative AI exposure using both the continuous exposure measure and indicator variables for whether an individual, household, or county belongs to the top tercile or above-median exposure group. To account for serial correlation in the error terms, we cluster

standard errors at the individual, household, or county level, depending on the unit of observation.

E. Summary Statistics

Our sample spans from 2021 to 2023, covering two years before and one year after the ChatGPT shock in November 2022. The final sample combines multiple data sources on household financial behaviors at the individual-, household-, and county-level.

First, the Survey of Income and Program Participation (SIPP) provides individual occupation and employment information, as well as detailed household financial data. Second, to measure local household portfolio choices, we follow Pan et al. (2024) to identify financial advisors operating primarily within individual counties using SEC Form ADV filings, and track their equity holdings through Thomson Reuters 13F data. This approach allows us to capture how local financial advisors adjust their portfolio allocations following the introduction of ChatGPT. Because local advisors primarily serve local investors, their portfolio choices likely reflect the preferences and expectations of local clients. Third, we supplement these financial data with county-level mortgage application data from the Home Mortgage Disclosure Act (HMDA).

The final SIPP survey sample includes 5,660 individuals and 3,800 households. Among the surveyed individuals, 2% reported experiencing a layoff within a given year. Approximately 5% reported voluntarily leaving their jobs for new employment opportunities or transitioning to self-employment. In addition, 29% reported working from home at least one day per week, while 20% reported beginning to hold multiple jobs simultaneously.

The SIPP surveyed households hold an average of $30,587 in stock investments (measured in 2021 value). To ensure comparability across years, we adjust stock values in 2022 and 2023 to constant 2021 value using the growth rate of the S&P 500 index. In addition, the average household carries $131,157 in total debt, including credit card debt, mortgages, educational loans, and vehicle loans. Among the debt, mortgage is the largest and accounts for 82% of

total debt. The average education debt balance is $9,024, while the average vehicle debt balance is $8,417. Average credit card debt amounts to $4,194. On average, each household has two income-earning individuals.

IV. Individual Employment and Job Security

This section reports the empirical results. We first examine whether workers in AI-exposed occupations face greater job insecurity after the release of ChatGPT. We then study whether households respond by reducing financial risk exposure on the equity investment side and the liability side of the balance sheet.

If the household results were driven primarily by realized layoffs, we would expect financial adjustments to be concentrated among workers who actually lose jobs and should primarily reflect liquidity constraints or financial distress. The precautionary-motive interpretation makes a different prediction. Once AI exposure raises expected labor-income risk, households should reduce financial risk before an actual layoff is observed. We therefore emphasize specifications that exclude households with observed layoffs and outcomes that reflect deliberate risk reduction: lower equity exposure, lower high-beta investment, reduced leverage, and weaker demand for new mortgage credit.

A. AI Exposure and Layoffs

We first study whether greater exposure to GenAI is associated with a higher probability of job loss. The conjecture is that workers in occupations with greater GenAI exposure will face higher layoff risk after the release of ChatGPT, because exposed jobs become more susceptible to automation or substitution. We estimate individual-level regressions using the SIPP data from 2021 to 2023. The dependent variable is an indicator variable equal to one if the worker experiences a layoff. The specifications include individual fixed effects and year fixed effects, so the estimates compare changes in layoff risk for the same worker before

and after the release of ChatGPT across different levels of occupational GenAI exposure. Standard errors are clustered at the individual level. Table 2 reports the results.

The results show that workers in more exposed occupations become more likely to experience layoffs after the release of ChatGPT. In Column (1), the coefficient on the interaction term (GenAI×Post-ChatGPT) is positive and statistically significant at the 5% level, suggesting that worker subject to greater GenAI exposure are associated with an increase in the probability of layoff after the shock.

Columns (2) and (3) use indicator variables for high-exposure occupations to examine robustness. The coefficient on the top-tercile exposure indicator interacted with the post-ChatGPT dummy is 1.7% and statistically significant at the 1% level. The coefficient on the above-median exposure indicator is 1.5% and also statistically significant at the 1% percent level. Thus, workers in highly exposed occupations experience 1.5 to 1.7 percentage points higher layoff probabilities after the release of ChatGPT, relative to workers in less exposed occupations. The coefficients correspond to 75% to 85% of the sample mean. Therefore, the economic magnitude is meaningful relative to the average layoff probability in the sample.

The evidence in Table 2 provides the first piece of evidence for the job insecurity channel. The release of ChatGPT does not merely change public attention to AI; it was followed by a meaningful increase in layoff risk among workers whose occupational tasks are more exposed. If workers perceive that AI exposure raises the probability of job loss, they may reduce financial risk before any actual income loss occurs. In this sense, layoffs serve as evidence of the underlying shock to labor-income risk, rather than as the only channel through which household financial choices may change.

B. AI Exposure and Job Security

In this section, we examine whether GenAI exposure affects other aspects of job security. Layoffs capture realized job loss, but workers could change labor market behavior even before a layoff occurs. We predict that exposed workers will respond to heightened job security risk

and behave more cautiously. In particular, they are less likely to voluntarily quit, less likely to work from home as workplace visibility may become more valuable for signaling productivity, and more likely to hold multiple jobs as a form of income insurance.

Table 3 reports the results. The dependent variables are dummy variables for voluntarily quit for new jobs or self-employed, working from home, and starting working for more jobs simultaneously. These outcomes are useful because layoffs alone do not fully capture job insecurity. Workers can face higher expected job loss or weaker outside options without experiencing an actual layoff during the sample period. The expected signs are therefore negative for voluntary quits and work from home, and positive for multiple-job holding. The regressions use the same individual-level SIPP sample as in Table 2, include individual and year fixed effects, and cluster standard errors at the individual level.

Panel A reports the results for voluntary quits. The coefficient on the GenAI exposure measure interacted with the post-ChatGPT indicator is -0.053 and statistically significant at the 1% percent level. The top-tercile and above-median exposure indicators imply a decline of 1.6 and 2.1 percentage points, respectively. These estimates indicate that workers in more exposed occupations become less likely to leave their current jobs voluntarily after ChatGPT shock. This pattern is consistent with an increase in perceived job insecurity. Workers appear less willing to move to a new job or self-employment when AI raises uncertainty about the value of their skills or the availability of outside opportunities.

Panel B shows that AI exposure is also associated with less work from home. The coefficient on the GenAI exposure interaction is -0.086 and statistically significant at the 1% level. The top-tercile and above-median indicators imply declines of 2.2 and 3.8 percentage points. This result is consistent with workers responding to job insecurity by increasing workplace attachment or visibility. While the data do not allow us to observe workers' motives directly, the decline in remote work is consistent with the broader pattern of more cautious labor-market behavior among exposed workers.

Panel C examines whether workers start working multiple jobs simultaneously. The

coefficient on the continuous exposure interaction is 0.076 and statistically significant at the 5% level. The top-tercile indicator implies a 2.2 percentage point increase, and the above-median indicator implies a 3.1 percentage point increase. The rise in multiple-job holding suggests a precautionary labor-supply response. Workers exposed to generative AI may seek additional income sources to insure against future job loss or weaker earnings in their primary occupation.

Taken together, Table 2 and Table 3 show that GenAI exposure increases job insecurity along several dimensions. Exposed workers are more likely to be laid off, less likely to voluntarily quit, less likely to work from home, and more likely to hold multiple jobs. These patterns are difficult to reconcile with a pure productivity-enhancement interpretation in the short run. Instead, the evidence suggests that the release of ChatGPT increased both realized and perceived employment risk among workers in exposed occupations.

V. Household Financial Decision

A. AI Exposure and Household Equity Investment

We now test whether households exposed to generative AI reduce financial risk. We conjecture that households will reduce financial risk even when the worker has not yet experienced an actual layoff. We examine equity investment because it exposes households to aggregate financial risk and can be adjusted more readily than many real assets. While a decline in equity holdings among exposed households is consistent with a precautionary motive, an alternative channel operating through realized job loss would predict the same pattern. To isolate the effect from the precautionary-motive channel, we restrict households hold stock equities and have not experienced a layoff in the current year, which helps separate precautionary portfolio adjustment from the direct effect of realized job loss.

Table 4 reports the results. The dependent variable is the natural logarithm of the household's stock holdings, adjusted by the S&P 500 return to 2021 values to facilitate comparison

across years. . All specifications include household fixed effects and year fixed effects, and standard errors are clustered at the household level. The household fixed effects absorb time-invariant differences across households, and the year fixed effects absorb aggregate changes in market conditions.

The results show a sizable decline in equity holdings among more exposed households. In Column (1), the interaction between value-weighted household GenAI exposure and the post-ChatGPT indicator is -2.163 and statistically significant at the 5% level. A one-standard-deviation increase in GenAI exposure is associated with a approximately a 34.6% reduction in household equity holdings. Column (2) uses the GenAI exposure of the highest-earning household member and obtains a similar coefficient of -2.015, which is roughly about 32% reduction with a one-standard-deviation increase in GenAI exposure. The coefficient is also statistically significant at the 5% level. The dummy variables (top tercile and above median) yield the same conclusion.

It is worth noting that this test is conducted among households without actual layoffs. Households are not simply selling stocks after experiencing layoffs. Instead, exposed households appear to reduce risky assets when they anticipate income loss due to job insecurity arising from the disruption of generative AI. This result is consistent with standard portfolio choice intuition (Chang, Hong and Karabarbounis, 2018; Gálvez and Paz-Pardo, 2026). Labor income is an important component of household wealth. When future labor income becomes riskier, households are expected to reduce exposure to risky assets in order to lower total household financial risk. Table 4 provides direct evidence of such a response, i.e., the portfolio adjustment occurs among households for which the ChatGPT shock operates through perceived labor-income risk rather than realized job loss, consistent with a precautionary rather than distress-driven mechanism.

B. AI Exposure and Equity Portfolio Risk

We next extend the analysis to ask whether exposed households reduce not only equity holdings but also the riskiness of their portfolio composition. If households reduce financial risk for precautionary reasons, investment will decrease more for high-beta stocks and could shift toward lower-beta stocks. This test helps distinguish between alternative interpretations. If households experience wealth reduction following realized job losses, they reduce equity holdings overall without a clear prediction regarding portfolio risk composition. A precautionary-risk channel, by contrast, predicts a sharper reduction in high-beta stocks and potential reallocation toward safer equity positions. Table 5 therefore examines not only whether households invest less in equity, but whether equity investment becomes more conservative. Since the SIPP survey data do not report the specific equity holdings, we utilize local investors' portfolio holdings from 13F filings, following Pan et al. (2024). The data frequency for this test is therefore quarterly.

We report results in Table 5. The dependent variable is county-level equity investment growth by local investment advisors, measured separately for high-beta and low-beta stocks. To isolate changes in investment quantities from price appreciation, we compute equity investment growth using prior-period prices. We classify stocks as high- or low-beta based on the sample median beta of approximately 0.98, and examine each group separately: Panel A reports results for high-beta stocks and Panel B for low-beta stocks. The main specifications include county and year fixed effects, and standard errors are clustered at the county level.

Panel A shows that investment growth in high-beta stocks declines in counties with more AI-exposed population after ChatGPT's release. The coefficient on the GenAI exposure interaction is negative and statistically significant. The top-tercile exposure indicator implies a 2.7 to 3.0 percentage point decline in high-beta stock investment growth. The above-median indicator is negative but not statistically significant. Overall, the evidence indicates that high-exposure areas experience lower investment growth in stocks with greater systematic risk.

Panel B provides an interesting contrast. Investment growth in low-beta stocks increases in counties with more AI-exposed population in several specifications. The coefficient on GenAI interaction term is positive and statistically significant, and the top-tercile indicator is positive and statistically significant in the baseline specification. The low-beta results suggest that the portfolio adjustment is not simply a uniform withdrawal from all equity positions. Rather, local investors appear to reallocate away from high-beta stocks and toward safer equity positions.

The results in Table 4 and Table 5 therefore point to two related forms of risk reduction. At the household level, exposed households reduce overall stock holdings. At the portfolio level, households in high-exposure areas tilt away from high-beta stocks. Both patterns are consistent with households lowering financial risk when labor-market uncertainty rises. The composition result in Table 5 is especially difficult to reconcile with a pure wealth-depletion channel, as it reflects a shift in the risk profile of the portfolio rather than a reduction in overall investment alone.

C. AI Exposure and Household Leverage

Debt is central to the risk reduction because leverage amplifies future income shocks. We next examine whether households reduce leverage in response to AI-induced job insecurity. When labor income becomes more uncertain, forward-looking households may reduce debt in order to lower fixed repayment obligations and preserve financial flexibility. This prediction differs from a distress channel after actual unemployment, in which balances might rise through borrowing or missed repayment. A decline in debt among households without observed layoffs is consistent with precautionary deleveraging driven by heightened employment uncertainty. Therefore, the prediction is that households with greater GenAI exposure will lower their debt balance to lower financial risk and repayment obligations.

We report the results in Table 6. The dependent variable is the natural logarithm of household total debt. As in the equity analysis, the sample focuses on households without

actual layoffs, which helps distinguish proactive deleveraging from the direct consequences of realized unemployment. The regressions include household fixed effects and year fixed effects, and standard errors are clustered by county.

The estimates show that more exposed households reduce total debt after the release of ChatGPT. In Column (1), the coefficient on value-weighted household GenAI exposure interacted with the post-ChatGPT dummy is negative and statistically significant at the 1% level. A one-standard-deviation increase in GenAI exposure is associated with a 14.7% reduction in total household debt following the ChatGPT shock. Column (2), which uses the exposure of the highest-earning household member, produces a similar coefficient estimate and is also statistically significant at the 1% level. The results are robust to indicator variables. Top tercile and above-meidan GenAI exposure households reduce total debt by approximately 23% relative to less exposed households.

These magnitudes suggest economically meaningful deleveraging. The results imply that households exposed to greater AI-related labor-market risk reduce debt balances even in the absence of observed layoffs. This behavior is consistent with precautionary deleveraging among households facing greater employment uncertainty associated with GenAI exposure. When expected labor income becomes more uncertain, households may prefer to reduce fixed payment obligations and lower the probability that a future income shock leads to financial distress.

The leverage results complement the equity results. Households do not respond only by changing equity investment. They also lower liabilities. This joint reduction in risky assets and debt exposure is consistent with a broad decline in household financial risk-taking following the increase in employment uncertainty associated with GenAI exposure. Importantly, the decline in debt among households without observed layoffs suggests that the response is an attempt to reduce vulnerability to a possible future income shock.

D. AI Exposure and Debt Composition

The composition of debt is important for interpreting the underlying mechanism. We next examine which components of household debt drive the decline in leverage. Credit card debt is unsecured and can be adjusted relatively quickly, making it particularly sensitive to short-run precautionary behavior. Mortgage debt is larger and closely tied to housing investment and long-term financial commitments. Education debt, in contrast, is generally less flexible and more difficult to adjust over a short horizon. If households actively manage balance-sheet risk in response to heightened employment uncertainty, we would expect stronger reductions in credit card and mortgage debt than in education debt.

We report the results in Table 7. The dependent variables are the natural logarithms of credit card balances, debt against the primary residence, education debt, and auto loan. The regressions include household and year fixed effects, use the sample of households without observed layoffs, and cluster standard errors by county.

Panel A reports results for credit card balances. The coefficients on the interaction term are negative and statistically significant (columns 1 and 2), indicating that households subject to greater GenAI exposure reduce their credit card balances following the ChatGPT shock. The magnitude is economically meaningful. A one-standard-deviation increase in GenAI exposure_vw is associated with a 10.9% reduction in credit card balance. Households in the above-median GenAI exposue reduce credit card balances by 29.3% (column 4). These estimates suggest that exposed households reduce unsecured consumer debt after the ChatGPT shock.

Panel B documents a parallel deleveraging response in mortgage debt. The coefficient on value-weighted household GenAI exposure is negative and statistically significant at the 5% level. A one-standard-deviation increase in GenAI exposure_vw is associated with a 13% reduction in mortgage balance. Results are robust to an alternative exposure measure based on the highest-earning household member. Specifications employing median and top-tercile exposure cutoffs yield negative but statistically insignificant. Taken together, the evidence

suggests that the deleveraging response extends to the largest component of many household balance sheets: debt secured by the primary residence.

The debt composition results reinforce the interpretation that households are actively reducing financial risk exposure. Exposed households reduce unsecured credit card debt and mortgage debt. The pattern is more consistent with active deleveraging than with a passive deterioration in household finances. If the results were driven only by current income losses, one might expect higher debt accumulation or delinquency rather than lower balances among households with no observed layoffs. Instead, households appear to reduce obligations that would increase their exposure to financial distress in the event of future income loss, consistent with a forward-looking precautionary motive.[6]

E. AI Exposure and Mortgage Applications

To supplement the household debt analysis, we collect home mortgage application data from the Home Mortgage Disclosure Act (HMDA) to examine whether GenAI exposure affects household mortgage applications. The prediction is that if AI exposure increases job insecurity, households should be less willing to initiate new long-term debt commitments, especially discretionary borrowing such as home-improvement loans.

This test helps distinguish precautionary behavior from realized distress. Existing balances can change through repayment, refinancing, or amortization. Mortgage applications, by contrast, measure households' willingness to initiate new debt-financed commitments. If AI exposure raises expected labor-income risk, households should be less willing to purchase homes or finance home improvements. The prediction is particularly strong for home-improvement loans, which are more discretionary and easier to postpone than home purchases.

We report the results in Table 8. The dependent variables are county-level growth rates in the number of home-purchase mortgage applications and home-improvement mortgage

[6]In untabulated results, we examine the response of education and vehicle debt and do not find significant changes in those balances.

applications from HMDA. These outcomes capture new borrowing intentions rather than existing debt balances. The regressions include county fixed effects and year fixed effects, and standard errors are clustered at the county level.

Panel A reports a significant decline in home-purchase mortgage application growth. A one-percentage-point increase in county-level AI exposure is associated with a 0.81 percent point decline in home-purchase mortgage application growth, statistically significant at the 1% level. The categorical measures show a similar pattern. Counties in the top tercile of the GenAI exposure experience 1.8% lower growth, and counties above the median experience 1.2% lower growth.

Panel B shows an even larger response for home-improvement mortgage applications. A one percentage point increase in county-level AI exposure is associated with a 3.96 percentage point decline in home-improvement application growth. The coefficient is statistically significant at the 1% level. The indicator measures imply a decline of 8.1% and 5.3% for top-tercile and above-median exposure counties, respectively.

The mortgage application results are consistent with the debt-balance evidence in Table 6 and Table 7. Households in AI-exposed areas reduce existing leverage and also become less likely to initiate new housing-related borrowing. The larger decline in home-improvement applications is particularly informative. Home improvement projects are often discretionary and easier to postpone than home purchases. A sharp decline in these applications is therefore consistent with households delaying optional, debt-financed investment when future labor income becomes more uncertain.

Collectively, the evidence in Table 2 through Table 8 points to a consistent mechanism: greater GenAI exposure increases labor-market insecurity which, in turn, leads households to reduce financial risk along multiple dimensions. Exposed households lower overall stock holdings, reallocate away from high-beta stocks, reduce total debt, decrease credit card and mortgage balances, and become less likely to apply for new housing-related loans. Importantly, these patterns are observed in samples that exclude households with observed layoffs,

suggesting that the financial adjustments are not merely passive outcomes of realized income losses. Instead, households appear to make forward-looking financial decisions in response to AI-induced job insecurity.

VI. Conclusion

The rapid diffusion of ChatGPT and other large language models represents a labor market shock with potentially important implications for household finance. This paper studies whether the exposure to generative AI affects household financial decisions through job insecurity. We construct measures of GenAI exposure across occupations, households, and counties, and use the release of ChatGPT in November 2022 as a technological shock.

The overall evidence indicates that GenAI exposure increases labor market insecurity. Workers in more exposed occupations become more likely to experience layoffs, less likely to voluntarily quit, less time to work from home, and more likely to hold multiple jobs after the release of ChatGPT. These labor-market responses are accompanied by more conservative household financial decisions. Exposed households reduce equity holdings, shift away from high-beta stocks, lower total debt, reduce credit card and mortgage balances, and become less likely to apply for home-purchase and home-improvement mortgages.

Collectively, the results are consistent with a precautionary risk-reduction channel. The household financial response is difficult to explain purely as a consequence of realized income loss, as the effects persist among households with no observed layoffs. Instead, the findings suggest that households actively reduce financial risk exposure when exposure to generative AI raises expected labor-income risk.

These findings highlight the importance of household balance sheets in evaluating the economic consequences of generative AI. Even if AI raises productivity in the long run, its short-run diffusion may increase job insecurity and lead households to reduce financial risk-taking, leverage, and housing-related credit demand.

REFERENCES

Acemoglu, D., Autor, D., Hazell, J., Restrepo, P., 2022. Artificial intelligence and jobs: Evidence from online vacancies. Journal of Labor Economics 40, S293–S340.

Acemoglu, D., Restrepo, P., 2018. The race between man and machine: Implications of technology for growth, factor shares, and employment. American Economic Review 108, 1488–1542.

Acemoglu, D., Restrepo, P., 2020. Robots and jobs: Evidence from US labor markets. Journal of Political Economy 128, 2188–2244.

Acemoglu, D., Restrepo, P., 2022. Tasks, automation, and the rise in US wage inequality. Econometrica 90, 1973–2016.

Aiyagari, S.R., 1994. Uninsured idiosyncratic risk and aggregate saving. The Quarterly Journal of Economics 109, 659–684.

Angerer, X., Lam, P.S., 2009. Income risk and portfolio choice: An empirical study. The Journal of Finance 64, 1037–1055.

Autor, D.H., Dorn, D., Hanson, G.H., 2013. The China syndrome: Local labor market effects of import competition in the united states. American Economic Review 103, 2121–2168.

Babina, T., Fedyk, A., He, A.X., Hodson, J., 2023a. Artificial intelligence and firms' systematic risk. Working paper, Columbia University, NBER, and CEPR .

Babina, T., Fedyk, A., He, A.X., Hodson, J., 2023b. Firm investments in artificial intelligence technologies and changes in workforce composition 31325.

Betermier, S., Jansson, T., Parlour, C., Walden, J., 2012. Hedging labor income risk. Journal of Financial Economics 105, 622–639.

Brynjolfsson, E., Li, D., Raymond, L., 2025. Generative AI at work. The Quarterly Journal of Economics 140, 889–942.

Campbell, J.Y., 2006. Household finance. The journal of finance 61, 1553–1604.

Campbell, J.Y., Viceira, L.M., 2002. Strategic asset allocation: portfolio choice for long-term investors. Clarendon Lectures in Economic.

Carroll, C.D., 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. The Quarterly journal of economics 112, 1–55.

Chang, Y., Hong, J.H., Karabarbounis, M., 2018. Labor market uncertainty and portfolio choice puzzles. American Economic Journal: Macroeconomics 10, 222–62.

Cocco, J.F., Gomes, F.J., Maenhout, P.J., 2005. Consumption and portfolio choice over the life cycle. The Review of Financial Studies 18, 491–533.

d'Astous, P., Shore, S.H., 2024. Human capital risk and portfolio choices: Evidence from university admission discontinuities. Journal of Financial Economics 154, 103793.

Demirci, O., Hannane, J., Zhu, X., 2025. Who is AI replacing? the impact of generative ai on online freelancing platforms. Management Science .

Eisfeldt, A.L., Schubert, G., Zhang, M.B., 2023. Generative AI and firm values. Journal of Finance, forthcoming .

Eloundou, T., Manning, S., Mishkin, P., Rock, D., 2024. GPTs are GPTs: Labor market impact potential of llms. Science 384, 1306–1308.

Ganong, P., Noel, P., 2019. Consumer spending during unemployment: Positive and normative implications. American Economic Review 109, 2383–2424.

Guiso, L., Jappelli, T., Terlizzese, D., 1996. Income risk, borrowing constraints, and portfolio choice. The American Economic Review , 158–172.

Gálvez, J., Paz-Pardo, G., 2026. Richer earnings dynamics, consumption and portfolio choice over the life cycle. Journal of Financial Economics 176, 104206.

Heaton, J., Lucas, D., 2000. Portfolio choice and asset prices: The importance of entrepreneurial risk. The journal of finance 55, 1163–1198.

Hui, X., Reshef, O., Zhou, L., 2024. The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. Organization Science 35, 1977–1989.

Matsa, D.A., Melzer, B.T., Zator, M., 2024. Dual credit markets: Income risk, household debt, and consumption. National Bureau of Economic Research .

Noy, S., Zhang, W., 2023. Experimental evidence on the productivity effects of generative artificial intelligence. Science 381, 187–192.

Pan, Y., Pikulina, E., Siegel, S., Wang, T.Y., 2024. Political divide and the composition of households' equity portfolios. Working Paper, University of Utah .

Viceira, L.M., 2001. Optimal portfolio choice for long-horizon investors with nontradable labor income. The Journal of Finance 56, 433–470.

Figure 1: The Construction of County-level Generative AI Exposure

This figure shows the steps to construct the county-level Generative AI exposure.

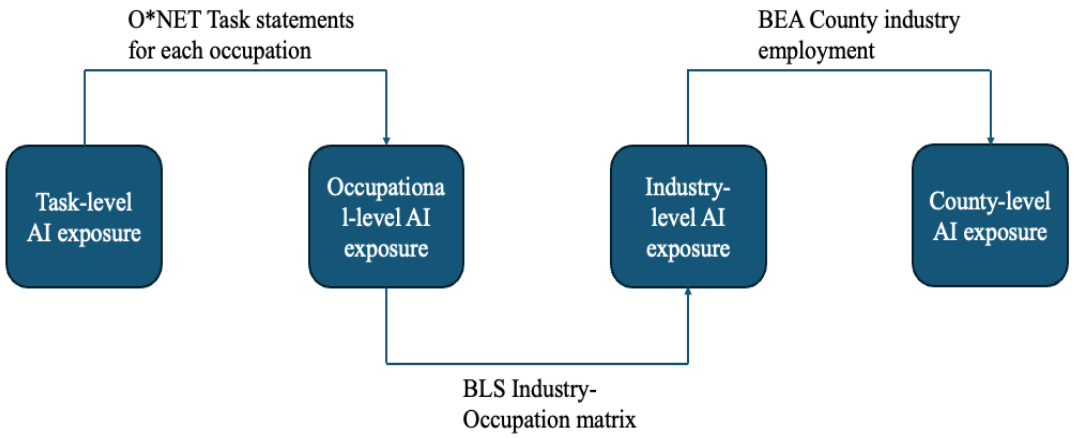


Figure 2: The county-level AI Exposure

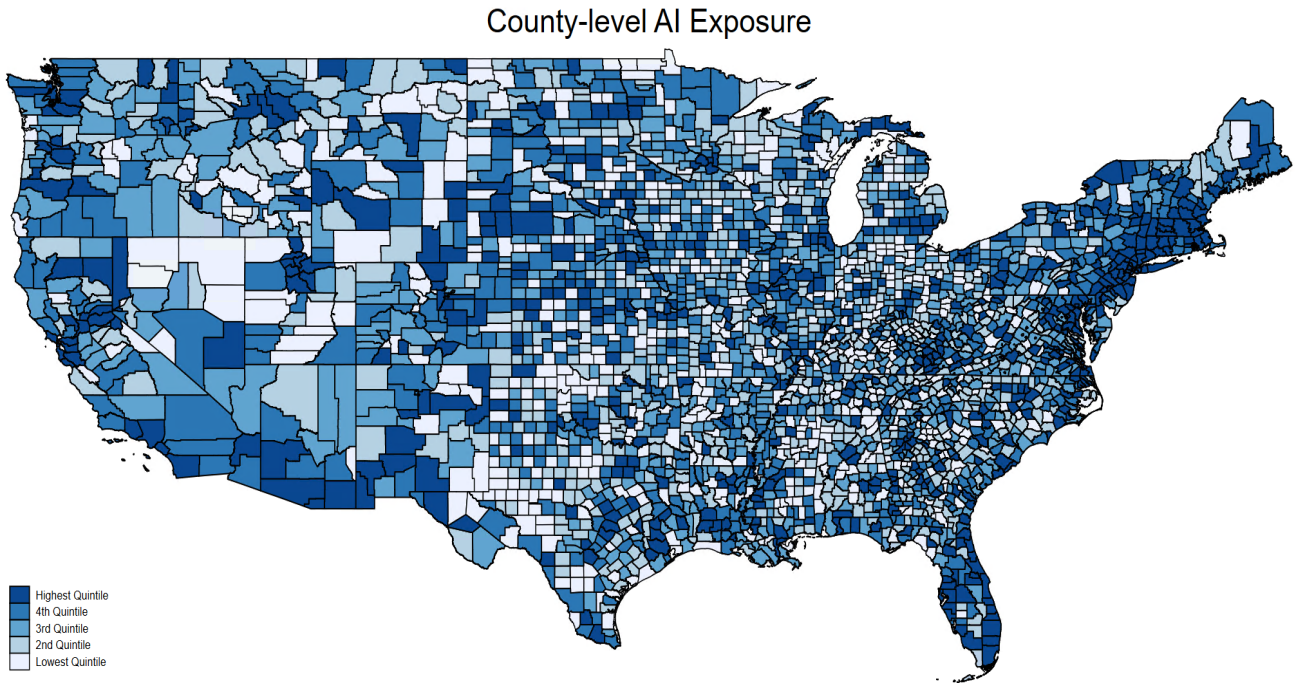


Table 1: Summary Statistics

VARIABLES	N	Mean	STD	P25	P50	P75
SIPP Data						
Individual Level Data						
GenAI exposure	14,994	0.26	0.17	0.12	0.25	0.39
Layoff dummy	14,994	0.02	0.15	0	0	0
Voluntarily quit for new job or switch to self-employed	14,994	0.05	0.21	0	0	0
Work from home dummy	14,656	0.29	0.45	0	0	1
Start working for more jobs	14,316	0.20	0.40	0	0	0
Household Level Data						
GenAI exposure_vw	10,272	0.27	0.16	0.14	0.26	0.38
GenAI exposure (highest earned person)	10,272	0.27	0.17	0.13	0.26	0.40
Household S&P 500 adjusted value of stocks	10,272	30587	128095	0	0	249
Total debt	10,272	131157	173452	2000	52700	205000
Credit card debt	10,272	4194	8929	0	0	4000
Home mortgage	10,272	107889	162073	0	0	175000
Number of earning people	10,272	2	1	1	1	2
County Level Data						
GenAI exposure	9,152	0.24	0.01	0.23	0.24	0.25
Home purchase mortgage growth	9,152	-0.04	0.20	-0.18	-0.07	0.08
Home improvement mortgage growth	8,842	0.21	0.41	-0.09	0.14	0.42
In counties with identified local investment advisors						
High-beta stocks investment growth	3,976	0.01	0.22	-0.05	0.00	0.05
Low-beta stocks investment growth	3,978	0.07	0.23	-0.02	0.03	0.10

Table 2: AI Exposure and Layoffs

This table reports the individual-level OLS regression results. The data is from the Census SIPP survey data. The dependent variable is layoff dummy. The sample period is 2021 to 2023. The variable definition is in the Appendix section D. The fixed effects of the regressions are reported at the bottom of the table. T-statistics are in parentheses. Standard errors are clustered at the county level. Coefficients statistically significant at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Lay off dummy

VARIABLES	(1)	(2)	(3)
GenAI exposure×Post-ChatGPT	0.039**		
	(2.46)		
GenAI exposure(top tercile)×Post-ChatGPT		0.017***	
		(3.25)	
GenAI exposure(above median)×Post-ChatGPT			0.015***
			(3.03)
Constant	0.019***	0.020***	0.020***
	(12.08)	(31.72)	(22.70)
Observations	14,994	14,994	14,994
R-squared	0.429	0.429	0.429
Individual FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Cluster at individual	Yes	Yes	Yes

Table 3: AI Exposure and Job Security

This table reports the individual-level OLS regression results. The data is from the Census SIPP survey data. The dependent variables are three dummy variables for voluntarily quit for new jobs or self-employed, work from home, and start working for multiple jobs. The sample period is 2021 to 2023. The variable definition is in the Appendix section D. The fixed effects of the regressions are reported at the bottom of the table. T-statistics are in parentheses. Standard errors are clustered at the county level. Coefficients statistically significant at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Panel A: Voluntarily quit for new job or self-employed

VARIABLES	(1)	(2)	(3)
GenAI exposure×Post-ChatGPT	-0.053***		
	(-2.73)		
GenAI exposure(top tercile)×Post-ChatGPT		-0.016**	
		(-2.31)	
GenAI exposure(above median)×Post-ChatGPT			-0.021***
			(-3.02)
Constant	0.053***	0.050***	0.052***
	(27.91)	(57.37)	(42.10)
Observations	14,994	14,994	14,994
R-squared	0.431	0.431	0.431
Individual FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Cluster at individual	Yes	Yes	Yes

Panel B: Work from Home

VARIABLES	(1)	(2)	(3)
GenAI exposure×Post-ChatGPT	-0.086***		
	(-3.15)		
GenAI exposure(top tercile)×Post-ChatGPT		-0.022**	
		(-2.04)	
GenAI exposure(above median)×Post-ChatGPT			-0.038***
			(-4.08)
Constant	0.300***	0.294***	0.298***
	(111.71)	(225.50)	(174.76)
Observations	14,546	14,546	14,546
R-squared	0.784	0.784	0.784
Individual FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Cluster at individual	Yes	Yes	Yes

Panel C: Start working for more jobs simultaneously

VARIABLES	(1)	(2)	(3)
GenAI exposure×Post-ChatGPT	0.076**		
	(2.06)		
GenAI exposure(top tercile)×Post-ChatGPT		0.022*	
		(1.66)	
GenAI exposure(above median)×Post-ChatGPT			0.031**
			(2.44)
Constant	0.188***	0.193***	0.190***
	(51.60)	(116.73)	(81.60)
Observations	14,149	14,149	14,149
R-squared	0.494	0.494	0.494
Individual FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Cluster at individual	Yes	Yes	Yes

Table 4: AI Exposure and Equity Investment

This table reports the OLS regression results of the household equity investment. The dependent variable is the natural logarithm of household stock investment value. To facilitate comparisons across years, we adjust stock investment values to constant 2021 dollars using the growth rate of the S&P 500 index. The control variables are county-level education, Ln(total population), Ln(Median_age), Ln(household income), and GDP Growth. Variable definitions are provided in the Appendix section D. T-statistics are in parentheses. Standard errors are clustered at the county level. Coefficients statistically significant at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Ln(household S&P 500 adjusted value of stocks)

VARIABLES	(1)	(2)	(3)	(4)
GenAI exposure_vw×Post-ChatGPT	-2.163**			
	(-2.00)			
GenAI exposure (highest earned person)×Post-ChatGPT		-2.015**		
		(-2.07)		
GenAI exposure_vw(top tercile)×Post-ChatGPT			-0.529*	
			(-1.67)	
GenAI exposure_vw(above median)×Post-ChatGPT				-0.552*
				(-1.68)
Number of earning people	0.200	0.202	0.201	0.203
	(0.86)	(0.87)	(0.87)	(0.87)
Constant	6.610***	6.593***	6.449***	6.489***
	(16.01)	(16.04)	(16.30)	(16.25)
Observations	3,734	3,734	3,734	3,734
R-squared	0.494	0.494	0.494	0.494
Household FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Cluster at household	Yes	Yes	Yes	Yes

Table 5: AI Exposure and Equity Investment

This table reports the OLS regression results of the county-level quarterly growth rate of equity investment.We classify stocks into high-beta and low-beta groups based on the median beta of stocks held by local investment advisors, which is 0.98. We follow Pan et al. (2024) to identify local investment advisors using the SDC form ADV filings. The sample contains 374 counties with identified local investment advisors. The equity holding data comes from Thomson Reuters 13F. The dependent variable is the value-weighted growth rate of the equity investment in the high-beta stocks by the local advisors (quarterly or annually). The sample period is 2021-2023. Variable definitions are provided in the Appendix section D. The fixed effects of the regressions are reported at the bottom of the table. T-statistics are in parentheses. Standard errors are clustered at the county level. Coefficients statistically significant at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Panel A: High-beta Stocks Investment Growth Rate

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
GenAI exposure×Post-ChatGPT	-1.092*			-1.035*		
	(-1.96)			(-1.70)		
GenAI exposure(top tercile)×		-0.027**			-0.030**	
Post-ChatGPT		(-2.09)			(-2.20)	
GenAI exposure(above median)×			-0.012			-0.012
Post-ChatGPT			(-0.94)			(-0.86)
GenAI exposure	0.128					
	(0.26)					
GenAI exposure(top tercile)		0.011				
		(1.17)				
GenAI exposure(above median)			0.001			
			(0.07)			
Education	0.000	0.000	0.000			
	(0.83)	(0.35)	(0.71)			
lntot_population	0.004	0.003	0.003			
	(1.13)	(0.84)	(1.04)			
lnMedian_age	-0.038	-0.040	-0.039			
	(-1.16)	(-1.22)	(-1.19)			
lnHousehold income	-0.005	-0.002	-0.004			
	(-0.28)	(-0.13)	(-0.20)			
county GPD growth	-0.000	-0.000	-0.000			
	(-0.14)	(-0.12)	(-0.05)			
Constant	0.206	0.145	0.142	0.106*	0.015***	0.014***
	(1.00)	(0.76)	(0.75)	(1.91)	(9.11)	(5.41)
Observations	3,928	3,928	3,928	3,974	3,974	3,974
R-squared	0.090	0.090	0.089	0.151	0.151	0.150
FIPS FE	No	No	No	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Cluster at FIPS	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Low-beta Stocks Investment Growth Rate

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
GenAI exposure×Post-ChatGPT	1.100**			1.119**		
	(2.16)			(2.13)		
GenAI exposure(top tercile)×Post-ChatGPT		0.023*			0.022*	
		(1.89)			(1.77)	
GenAI exposure(above median)×Post-ChatGPT			0.019			0.020
			(1.60)			(1.50)
GenAI exposure	-0.529					
	(-0.98)					
GenAI exposure(top tercile)		-0.007				
		(-0.69)				
GenAI exposure(above median)			-0.018*			
			(-1.82)			
Education	-0.000	-0.000	-0.000			
	(-0.65)	(-1.11)	(-0.31)			
lntot_population	0.002	0.001	0.003			
	(0.48)	(0.33)	(0.83)			
lnMedian_age	0.028	0.027	0.031			
	(0.81)	(0.77)	(0.90)			
lnHousehold income	-0.004	-0.003	-0.007			
	(-0.22)	(-0.15)	(-0.39)			
county GPD growth	0.001	0.001	0.001			
	(1.13)	(1.11)	(0.98)			
Constant	0.035	0.001	0.007	-0.030	0.069***	0.068***
	(0.15)	(0.00)	(0.04)	(-0.64)	(44.47)	(28.01)
Observations	3,930	3,930	3,930	3,976	3,976	3,976
R-squared	0.156	0.156	0.156	0.228	0.228	0.228
FIPS FE	No	No	No	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Cluster at FIPS	Yes	Yes	Yes	Yes	Yes	Yes

Table 6: AI Exposure and Household Leverage

This table reports the individual-level OLS regression results. The data is from the Census SIPP survey. The dependent variable is the natural logarithm of the total amount of credit card balance, home mortgage balance, education debt, and auto loan. The sample period is 2021 to 2023. Variable definitions are provided in the Appendix section D.The fixed effects of the regressions are reported at the bottom of the table. *T*-statistics are in parentheses. Standard errors are clustered at the county level. Coefficients statistically significant at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Household Total Debt

VARIABLES	(1)	(2)	(3)	(4)
GenAI exposure_vw×Post-ChatGPT	-0.863**			
	(-2.24)			
GenAI exposure (highest earned person)×Post-ChatGPT		-0.858**		
		(-2.49)		
GenAI exposure_vw(top tercile)×Post-ChatGPT			-0.231*	
			(-1.96)	
GenAI exposure_vw(above median)×Post-ChatGPT				-0.228**
				(-2.00)
Number of earning people	0.345***	0.346***	0.344***	0.343***
	(3.69)	(3.70)	(3.68)	(3.67)
Constant	8.289***	8.288***	8.233***	8.249***
	(54.00)	(54.59)	(54.62)	(54.54)
Observations	9,800	9,800	9,800	9,800
R-squared	0.812	0.812	0.812	0.812
Household FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Cluster at year	Yes	Yes	Yes	Yes

Table 7: AI Exposure and Household Categorized Debts

This table reports the individual-level OLS regression results. The data is from the Census SIPP survey data. The dependent variables are the natural logarithm of the credit card balance or home mortgage balance, respectively. The sample period is 2021 to 2023. Variable definitions are provided in the Appendix section D. The fixed effects of the regressions are reported at the bottom of the table. T-statistics are in parentheses. Standard errors are clustered at the county level. Coefficients statistically significant at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Panel A: Credit Card

VARIABLES	(1)	(2)	(3)	(4)
GenAI exposure_vw×Post-ChatGPT	-0.650*			
	(-1.66)			
GenAI exposure (highest earned person)×Post-ChatGPT		-0.731**		
		(-2.03)		
GenAI exposure_vw(top tercile)×Post-ChatGPT			-0.126	
			(-0.95)	
GenAI exposure_vw(above median)×Post-ChatGPT				-0.293**
				(-2.33)
Number of earning people	0.352***	0.353***	0.352***	0.351***
	(3.36)	(3.36)	(3.35)	(3.34)
Constant	3.329***	3.336***	3.281***	3.321***
	(19.54)	(19.68)	(19.55)	(19.76)
Observations	9,800	9,800	9,800	9,800
R-squared	0.712	0.712	0.712	0.712
Household FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Cluster at year	Yes	Yes	Yes	Yes

Panel B: Home Mortgage

VARIABLES	(1)	(2)	(3)	(4)
GenAI exposure_vw×Post-ChatGPT	-0.765**			
	(-2.29)			
GenAI exposure (highest earned person)×Post-ChatGPT		-0.822***		
		(-2.71)		
GenAI exposure_vw(top tercile)×Post-ChatGPT			-0.182	
			(-1.56)	
GenAI exposure_vw(above median)×Post-ChatGPT				-0.153
				(-1.41)
Number of earning people	0.081	0.082	0.081	0.080
	(0.94)	(0.95)	(0.93)	(0.92)
Constant	5.873***	5.877***	5.821***	5.828***
	(41.23)	(41.43)	(41.69)	(41.79)
Observations	9,800	9,800	9,800	9,800
R-squared	0.892	0.892	0.892	0.892
Household FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Cluster at year	Yes	Yes	Yes	Yes

Table 8: AI Exposure and Mortgage Applications (County-level)

This table reports the county-level OLS regression results. The dependent variables are county-level home purchase mortgage growth rate (in fraction) in Panel A and home improvement mortgage growth rate (in fraction) in Panel B. The growth rate is based on number of applications. The data is obtained from Consumer Financial Protection Bureau (CFPB), Home Mortgage Disclosure Act (HMDA) data. The sample period is 2021 to 2023. Variable definitions are provided in the Appendix section D. The fixed effects of the regressions are reported at the bottom of the table. Home purchase mortgage growth rate is winsorized at 1% and 99% level. Home improvement mortgage growth rate is winsorized at 5% and 95% level because home improvement loans are substantially smaller and less frequent, making their growth rates more susceptible to extreme outliers and measurement noise. T-statistics are in parentheses. Standard errors are clustered at the county level. Coefficients statistically significant at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Panel A: Home purchase mortgage growth rate

VARIABLES	(1)	(2)	(3)
GenAI exposure×Post-ChatGPT	-0.807***		
	(-2.68)		
GenAI exposure(top tercile)×Post-ChatGPT		-0.018***	
		(-2.67)	
GenAI exposure(above median)×Post-ChatGPT			-0.012*
			(-1.75)
Constant	0.029	-0.034***	-0.034***
	(1.20)	(-44.79)	(-29.73)
Observations	9,147	9,147	9,147
R-squared	0.531	0.531	0.531
FIPS FE	Y	Y	Y
Year FE	Y	Y	Y
S.E. Clustered by County	Y	Y	Y

Panel B: Home improvement mortgage growth rate

VARIABLES	(1)	(2)	(3)
GenAI exposure×Post-ChatGPT	-3.955***		
	(-5.39)		
GenAI exposure(top tercile)×Post-ChatGPT		-0.081***	
		(-4.56)	
GenAI exposure(above median)×Post-ChatGPT			-0.053***
			(-3.02)
Constant	0.527***	0.218***	0.218***
	(8.95)	(109.24)	(73.47)
Observations	8,783	8,783	8,783
R-squared	0.312	0.311	0.310
FIPS FE	Y	Y	Y
Year FE	Y	Y	Y
S.E. Clustered by County	Y	Y	Y

Appendix A. Generative AI Exposure

Table A1: Top and Bottom 20 AI Exposure Occupations

Panel A: Top 20 AI exposure occupations

ONET_SOC6	Occupation	number of tasks	AI_exp	AI_exp_rank
43-4021	Correspondence Clerks	20	0.825	1
43-9081	Proofreaders and Copy Markers	11	0.818	2
15-1251	Computer Programmers	17	0.765	3
27-3091	Interpreters and Translators	17	0.765	3
13-2071	Credit Counselors	23	0.739	5
43-3011	Bill and Account Collectors	15	0.733	6
13-1131	Fundraisers	28	0.714	7
27-3043	Writers and Authors	36	0.708	8
27-3043	Poets, Lyricists and Creative Writers	36	0.708	8
15-1242	Database Administrators	18	0.694	10
19-3022	Survey Researchers	16	0.688	11
15-1243	Data Warehousing Specialists	43	0.686	12
15-1243	Database Architects	43	0.686	12
43-6011	Executive Secretaries and Executive Administrative Assis	22	0.682	14
43-9111	Statistical Assistants	14	0.679	15
15-2051	Business Intelligence Analysts	54	0.667	17
15-2051	Data Scientists	54	0.667	17
15-2051	Clinical Data Managers	54	0.667	17
15-1254	Web Developers	29	0.655	19
13-1151	Training and Development Specialists	20	0.650	20
25-9031	Instructional Coordinators	30	0.650	20

Panel B: Bottom 20 AI exposure occupations

ONET_SOC6	Occupation	number of tasks	AI_exp	AI_exp_rank
37-2012	Maids and Housekeeping Cleaners	25	0	924
47-2044	Tile and Stone Setters	25	0	924
47-3013	Helpers–Electricians	25	0	924
49-3021	Automotive Body and Related Repairers	25	0	924
47-2051	Cement Masons and Concrete Finishers	26	0	924
47-2053	Terrazzo Workers and Finishers	26	0	924
47-2081	Drywall and Ceiling Tile Installers	26	0	924
47-4061	Rail-Track Laying and Maintenance Equipment Operators	26	0	924
49-3031	Bus and Truck Mechanics and Diesel Engine Specialists	26	0	924
49-3093	Tire Repairers and Changers	26	0	924
51-7042	Woodworking Machine Setters, Operators, and Tenders, Exc	26	0	924
47-2121	Glaziers	27	0	924
51-6041	Shoe and Leather Workers and Repairers	27	0	924
29-9093	Surgical Assistants	28	0	924
47-2181	Roofers	28	0	924
51-6021	Pressers, Textile, Garment, and Related Materials	28	0	924
51-2011	Aircraft Structure, Surfaces, Rigging, and Systems Assem	30	0	924
47-2061	Construction Laborers	32	0	924
51-4031	Cutting, Punching, and Press Machine Setters, Operators,	32	0	924
49-9063	Musical Instrument Repairers and Tuners	37	0	924

Table A2: Top and Bottom 20 AI Exposure Counties

This table reports the top and bottom 20 AI exposure counties in the U.S.

Panel A: Top 20 AI exposure counties

FIPS	County	AI_exp
36061	New York, NY	0.304
6075	San Francisco, CA	0.293
51013	Arlington, VA	0.293
32017	Lincoln, NV	0.291
19049	Dallas, IA	0.290
51919	Fairfax, VA	0.290
8014	Broomfield, CO	0.290
25025	Suffolk, MA	0.289
47187	Williamson, TN	0.285
6081	San Mateo, CA	0.284
48085	Collin, TX	0.283
13121	Fulton, GA	0.282
51099	King George, VA	0.281
25017	Middlesex, MA	0.280
11001	District of Columbia	0.280
34035	Somerset, NJ	0.280
34027	Morris, NJ	0.280
29510	St. Louis, MO	0.280
8035	Douglas, CO	0.280
42029	Chester, PA	0.279

Panel B: Bottom 20 AI exposure counties

FIPS	County	GenAI exposure
2068	Denali Borough, AK	0.177
22023	Cameron, LA	0.184
48301	Loving, TX	0.193
2060	Bristol Bay Borough, AK	0.199
30057	Madison, MT	0.200
32029	Storey, NV	0.201
6051	Mono, CA	0.201
2013	Aleutians East Borough, AK	0.201
48237	Jack, TX	0.203
18117	Orange, IN	0.204
8047	Gilpin, CO	0.205
48235	Irion, TX	0.207
51036	Charles City,	0.207
48283	La Salle, TX	0.208
13319	Wilkinson, GA	0.208
18165	Vermillion, IN	0.209
40007	Beaver, OK	0.209
48383	Reagan, TX	0.209
18051	Gibson, IN	0.210
48311	McMullen, TX	0.210

Table A3: Industry (NAICS-2 digit) AI Exposure

This table reports the industry-level AI exposure.

NAICS-2	Description	AI_exp	AI_exp(core tasks only)
52	Finance and Insurance	0.404	0.408
55	Management of Companies and Enterprises	0.389	0.398
51	Information	0.383	0.387
54	Professional, Scientific, and Technical Services	0.382	0.391
61	Educational Services	0.324	0.321
42	Wholesale Trade	0.282	0.297
92	Public Administration	0.271	0.282
53	Real Estate and Rental and Leasing	0.266	0.271
62	Health Care and Social Assistance	0.261	0.255
44-45	Retail Trade	0.254	0.256
81	Other Services (except Public Administration)	0.243	0.239
48-49	Transportation and Warehousing	0.230	0.238
71	Arts, Entertainment, and Recreation	0.227	0.226
22	Utilities	0.225	0.246
56	Administrative and Support and Waste Management and Remediation Services	0.220	0.228
11	Agriculture, Forestry, Fishing and Hunting	0.199	0.215
31-33	Manufacturing	0.187	0.191
21	Mining, Quarrying, and Oil and Gas Extraction	0.160	0.174
23	Construction	0.144	0.149
72	Accommodation and Food Services	0.124	0.114

Table A4: Occupation and Tasks Examples

Panel A: Computer programmers

Tasks	GPT_label
Write, analyze, review, and rewrite programs, using workflow chart and diagram, and applying knowledge of computer capabilities, subject matter, and symbolic logic.	E1
Correct errors by making appropriate changes and rechecking the program to ensure that the desired results are produced.	E1
Perform or direct revision, repair, or expansion of existing programs to increase operating efficiency or adapt to new requirements.	E1
Write, update, and maintain computer programs or software packages to handle specific jobs such as tracking inventory, storing or retrieving data, or controlling other equipment.	E1
Consult with managerial, engineering, and technical personnel to clarify program intent, identify problems, and suggest changes.	E0
Conduct trial runs of programs and software applications to be sure they will produce the desired information and that the instructions are correct.	E1
Prepare detailed workflow charts and diagrams that describe input, output, and logical operation, and convert them into a series of instructions coded in a computer language.	E1
Compile and write documentation of program development and subsequent revisions, inserting comments in the coded instructions so others can understand the program.	E1
Consult with and assist computer operators or system analysts to define and resolve problems in running computer programs.	E1
Perform systems analysis and programming tasks to maintain and control the use of computer systems software as a systems programmer.	E1
Write or contribute to instructions or manuals to guide end users.	E1
Investigate whether networks, workstations, the central processing unit of the system, or peripheral equipment are responding to a program's instructions.	E0
Assign, coordinate, and review work and activities of programming personnel.	E0
Train subordinates in programming and program coding.	E1
Develop Websites.	E1
Train users on the use and function of computer programs.	E1
Collaborate with computer manufacturers and other users to develop new programming methods.	E0

Panel B: Software developers

Tasks	GPT_label
Analyze information to determine, recommend, and plan installation of a new system or modification of an existing system.	E2
Analyze user needs and software requirements to determine feasibility of design within time and cost constraints.	E2
Confer with data processing or project managers to obtain information on limitations or capabilities for data processing projects.	E0
Confer with systems analysts, engineers, programmers and others to design systems and to obtain information on project limitations and capabilities	E0
Consult with customers or other departments on project status, proposals, or technical issues, such as software system design or maintenance.	E0
Coordinate installation of software system.	E0
Design, develop and modify software systems, using scientific analysis and mathematical models to predict and measure outcomes and consequences of design.	E1
Determine system performance standards.	E2
Develop or direct software system testing or validation procedures, programming, or documentation.	E1
Modify existing software to correct errors, adapt it to new hardware, or upgrade interfaces and improve performance.	E1
Monitor functioning of equipment to ensure system operates in conformance with specifications.	E0
Obtain and evaluate information on factors such as reporting formats required, costs, or security needs to determine hardware configuration.	E2
Prepare reports or correspondence concerning project specifications, activities, or status.	E1
Recommend purchase of equipment to control dust, temperature, or humidity in area of system installation.	E2
Specify power supply requirements and configuration.	E2
Store, retrieve, and manipulate data for analysis of system capabilities and requirements.	E1
Supervise and assign work to programmers, designers, technologists, technicians, or other engineering or scientific personnel.	E0
Supervise the work of programmers, technologists and technicians and other engineering and scientific personnel.	E0
Train users to use new or modified equipment.	E1

Appendix B. Selected GPT explanations for some tasks

Computer programmers

Task 1: Write, analyze, review, and rewrite programs, using workflow chart and diagram, and applying knowledge of computer capabilities, subject matter, and symbolic logic.

Label: E1

Explanation: The task falls under the E1 category. The task involves writing, analyzing, reviewing, and rewriting programs, which directly aligns with the capabilities of the LLM as described in the E1 category. Specifically, the LLM can assist in writing and transforming text and code according to complex instructions, providing edits to existing text or code following specifications, and writing code that can help perform a task that used to be done by hand. Additionally, the task of applying knowledge of computer capabilities, subject matter, and symbolic logic can be supported by the LLM's ability to generate explanations and provide feedback on documents or code. The use of workflow charts and diagrams might not be directly supported in terms of visual creation, but the LLM can help in planning and describing the logic that would go into such charts and diagrams, thus reducing the time it takes to complete the task with equivalent quality by at least half.

Task 2: Consult with managerial, engineering, and technical personnel to clarify program intent, identify problems, and suggest changes.

Label: E0

Explanation: The task falls under the E0 category. The primary reason for this classification is that the task involves a high degree of human interaction, specifically consulting with managerial, engineering, and technical personnel to clarify program intent, identify problems, and suggest changes. This process requires nuanced communication, understanding of complex technical and organizational contexts, and often real-time problem-solving and negotiation skills that are beyond the current capabilities of an LLM. While an LLM could assist in preparing documents, summarizing discussions, or even drafting potential solutions or changes to code based on specifications, the core of the task—effective communication and consultation with a variety of stakeholders to solve problems—relies heavily on human skills and intuition. Therefore, even though an LLM could support some aspects of the task, it would not significantly reduce the time it takes to complete the core components of the task with high quality by at least half.

Software developers

Task 1: Consult with customers or other departments on project status, proposals, or technical issues, such as software system design or maintenance.

Label: E0

Explanation: The task falls under the E0 category. Consulting with customers or other departments on project status, proposals, or technical issues, such as software system design or maintenance, involves a high degree of human interaction and understanding of specific, potentially up-to-date contexts that an LLM alone cannot provide. While an LLM can assist in generating reports, summarizing documents, or even drafting responses to

technical queries, the nuanced understanding required for effective consultation, especially in dynamic and complex fields like software development, goes beyond text generation. This task requires understanding the specific needs, concerns, and contexts of the project or issue at hand, which can vary greatly and may involve up-to-date information not available to the LLM. Additionally, the task involves decision-making based on these consultations, which is explicitly mentioned as a limitation in the E0 category. Therefore, despite the LLM's capabilities in assisting with certain aspects of the task, the core requirement for human interaction and decision-making based on up-to-date, specific contexts means that the LLM would not significantly reduce the time it takes to complete the task with high quality by at least half.

Task 2: Determine system performance standards.

Label: E2

Explanation: The task of determining system performance standards falls under the E2 category. This task involves analyzing data, understanding the specific requirements of a system, and making recommendations based on that analysis. While the LLM can assist in providing general knowledge about system performance and possibly help in drafting documents or guidelines based on general principles, the task requires a deeper understanding of specific system requirements, data analysis, and the ability to make tailored recommendations. These aspects may not be fully addressed by the LLM alone but could be significantly aided by additional software developed on top of the LLM that specializes in analyzing system performance data and making recommendations. Such software could leverage the LLM's capabilities to process and summarize information, but it would need to integrate more specialized knowledge and data analysis tools to effectively determine system performance standards.

Painters, Construction and Maintenance

Task 1: remove old finishes by stripping, sanding, wire brushing, burning, or using water or abrasive blasting.

Label: E0

Explanation: The task involves physical labor and manual skills such as stripping, sanding, wire brushing, burning, or using water or abrasive blasting to remove old finishes, which are tasks that cannot be performed by a language model or any software tool. These tasks require hands-on work, precise measurements, and often a degree of human judgment based on the condition of the material being worked on. Therefore, this task falls under the E0 category, as the LLM and systems built on top of it cannot perform physical tasks, make decisions that impact human livelihood, or replace the need for human labor in tasks that require physical interaction with materials.

Fashion designer

Task 1: Research the styles and periods of clothing needed for film or theatrical productions.

Label: E1

Explanation: The task of researching styles and periods of clothing needed for film or theatrical productions can be significantly expedited by using a large language model (LLM) like ChatGPT. The LLM can provide detailed information on historical clothing styles, the cultural context of different periods, and even suggest sources for further research. This task involves writing and transforming text according to complex instructions, summarizing documents, and providing feedback on documents, all of which are capabilities listed under the E1 category. The LLM can help fashion designers by quickly generating comprehensive overviews of fashion trends from different periods, which would otherwise require extensive manual research. This use of the LLM can reduce the time it takes to complete the research with equivalent quality by at least half, as it streamlines the information gathering and synthesis process.

Appendix C. OpenAI Methodology Notes

GPT prompt for exposure scoring

The following prompt structure was based on the rubric language by Eloundou et al. (2024) and Eisfeldt et al. (2023). Below is the instruction prompts submitted before asking GPT-4.0 Preview to classify each task statement (using the version as of March 30th to April 2nd, 2024). The GPT temperature is set to 0.

systemprompt="Consider the most powerful OpenAI large language model (LLM). This model can complete many tasks that can be formulated as having text input and text output where the context for the input can be captured in 2000 words. The model also cannot draw up-to-date facts (those from $<$ 1 year ago) unless they are captured in the input. Assume you are a worker with an average level of expertise in your role trying to complete the given task. You have access to the LLM as well as any other existing software or computer hardware tools mentioned in the task. You also have access to any commonly available technical tools accessible via a laptop (e.g. a microphone, speakers, etc.). You do not have access to any other physical tools or materials. You are a helpful research assistant who wants to label the given tasks according to the rubric below. Equivalent quality means someone reviewing the work would not be able to tell whether a human completed it on their own or with assistance from the LLM. If you aren't sure how to judge the amount of time a task takes, consider whether the tools described exposed the majority of subtasks associated with the task.

rubric:

E1 - Direct exposure: Label tasks E1 if direct access to the LLM through an interface like ChatGPT or the OpenAI playground alone can reduce the time it takes to complete the task with equivalent quality by at least half. This includes tasks that can be reduced to: - Writing and transforming text and code according to complex instructions, - Providing edits to existing text or code following specifications, - Writing code that can help perform a task that used to be done by hand, - Translating text between languages, - Summarizing medium-length documents, - Providing feedback on documents, - Answering questions about a document, - Generating questions a user might want to ask about a document, - Writing questions for an interview or assessment, - Writing and responding to emails, including ones that involve refuting information or engaging in a negotiation (but only if the negotiation is via written correspondence), - Maintain records of written data, - Prepare training materials based on general knowledge, or - Inform anyone of any information via any written or spoken medium.

E2 - Exposure by LLM-powered applications: Label tasks E2 if having access to the LLM alone may not reduce the time it takes to complete the task by at least half, but it is easy to imagine additional software that could be developed on top of the LLM that would reduce the time it takes to complete the task by half. This software may include capabilities such as: - Summarizing documents longer than 2000 words and answering questions about those documents, - Retrieving up-to-date facts from the Internet and using those facts in combination with the LLM capabilities, - Searching over an organization's existing knowledge, data, or documents and retrieving information, - Retrieving highly specialized domain knowledge, - Make recommendations given data or written input, - Analyze written information to inform decisions, - Prepare training materials based on highly specialized knowledge, - Provide counsel on issues, and - Maintain complex databases.

E3 - Exposure given image capabilities: Suppose you had access to both the LLM and a system that could view, caption, and create images as well as any systems powered by the LLM (those in E2 above). This system cannot take video as an input and it cannot produce video as an output. This system cannot accurately retrieve very detailed information from image inputs, such as measurements of dimensions within an image. Label tasks as E3 if there is a significant reduction in the time it takes to complete the task given access to a LLM and these image capabilities: - Reading text from PDFs, - Scanning images, or - Creating or editing digital images according to instructions. The images can be realistic but they should not be detailed. The model can identify objects in the image but not relationships between those options.

E0 - No exposure: Label tasks E0 if none of the above clearly decrease the time it takes for an experienced worker to complete the task with high quality by at least half. Some examples: - If a task requires a high degree of human interaction (for example, in-person demonstrations) then it should be classified as E0. - If a task requires precise measurements then it should be classified as E0. - If a task requires reviewing visuals in detail then it should be classified as E0. - If a task requires any use of a hand or walking then it should be classified as E0. - Tools built on top of the LLM cannot make any decisions that might impact human livelihood (e.g.hiring, grading, etc.). If any part of the task involves collecting inputs to make a final decision (as opposed to analyzing data to inform a decision or make a recommendation) then it should be classified as E0. The LLM can make recommendations. - Even if tools built on top of the LLM can do a task, if using those tools would not save an experienced worker significant time completing the task, then it should be classified as E0. - The LLM and systems built on top of it cannot do anything that legally requires a human to perform the task. - If there is existing technology not powered by an LLM that is commonly used and can complete the task then you should mark the task E0 if using an LLM or LLM-powered tool will not further reduce the time to complete the task. When in doubt, you should default to E0.

Appendix D. Variable definitions

Variables	Definitions	Sources
GenAI exposure (task-level)	The Generative AI exposure of tasks determined by GPT-4-0125 Preview Model on OpenAI April 2, 2024.	OpenAI
GenAI exposure (occupation-level)	The weighted average task-level Generative AI exposure aggregated to the occupation-level Generative AI exposure.	O*NET
GenAI exposure (county-level)	The weighted average industry-level Generative AI exposure aggregated to county-level. The weight is determined by the number of jobs of each NAICS-2 industry in a county. The county-industry jobs data is from the BEA 2022 employment by NAICS industries.	Bureau of Economic Analysis (BEA)
GenAI exposure (state-level)	The weighted average county-level Generative AI exposure aggregated to the state-level using the 2022 county population data.	As above
GenAI exposure (industry-level)	The weighted average occupation-level AI exposure aggregated to industry -level. The weight is determined by number of jobs of each occupation in an industry using the 2022 BLS Occupation-industry matrix data.	Bureau of Labor Statistics (BLS)
GenAI exposure (top tercile)	A binary variable that equals 1 if a county or household's AI exposure is in top tercile.	
GenAI exposure (above median)	A binary variable that equals 1 if a county or household's AI exposure is above the sample median.	
Post-ChatGPT	A dummy variable that equals 1 for 2023 and 0 for 2021 and 2022.	
Layoff dummy	A dummy variable that equals 1 if an individual reported at least one month of layoff in a reference year.	SIPP
voluntarily quit for new job or self-employed	A dummy variable that equals 1 if an individual reported to quit previous job and switch to self-employed or get a new job without observable layoff.	As above
Work from home	A dummy variable that equals 1 if an individual reported to work from home for at least one day in a week on average for the reference year.	As above
Working for more jobs	A dummy variable that equals 1 if an individual reported to start working for	As above

Continued on the next page

Variables	Definition	
	more jobs simultaneously in a reference year.	
Household stock investment	The natural logarithm of household stock investment value adjusted by the S&P 500 growth rate (in 2021 value).	SIPP
High-beta equity growth (quarterly)	The value-weighted quarterly growth rate of high-beta stocks (beta is above median of 0.98) of all local advisors in a county.	Thomas Reuters 13F data
Low-beta equity growth (quarterly)	The value-weighted quarterly growth rate of high-beta stocks (beta is below median of 0.98) of all local advisors in a county. Local advisor is defined as investment managers whose offices are located in no more than one county.	As above
Household total debt	The natural logarithm of the total amount of household debt including credit card balance, home mortgage balance, education debt, and vehicle debt.	SIPP
Credit card debt	The natural logarithm of household credit card debt balance.	As above
Home mortgage	The natural logarithm of household home mortgage balance.	As above
Home purchase mortgage applications growth	The number of home purchase mortgage applications growth rate in a county.	Home Mortgage Disclosure Act
Home improvement mortgage applications growth	The number of home improvement mortgage applications growth rate in a county.	As above
GDP growth	The county GDP growth rate of real Gross Domestic Product (Percent change from last year same month).	BEA
Education	The percentage of population over 25 years old holding bachelor degree or higher in 2022.	U.S. Census
Polulation	The natural logarithm of county total population in 2022.	As above
Median age	The natural logarithm of county median age in 2022.	As above
Household income	The natural logarithm of household median income in a county in 2022.	As above