

Organizational Technology Ladders: Remote Work and Generative AI Adoption

Gregor Schubert

UCLA Anderson *

June 22, 2025

Abstract

What drives inequality in technology adoption among firms? In this study, I propose that there is an “organizational technology ladder”: firms’ human and organizational capital adapts when firms incorporate new technologies and this adaptation can make it easier to adopt later technologies, leading to path-dependent divergence in technological capabilities. To provide empirical evidence, I exploit the sequential exposure of firms to two large changes in organizational technology: the rise of remote work 2020-2021, and the proliferation of generative AI technology since November 2022. I use detailed job posting data to study how firms’ labor demand transformed in response to these changes in technology. I develop an IV approach to show that firms that had to adopt remote work later demand more generative AI-related skills. I also provide evidence from a synthetic difference-in-differences estimation that firms that were more exposed to generative AI technology reduced their demand for remote workers after ChatGPT was released. I explore the mechanism for these effects: When firms adopt remote work technology, they invest in skills that enable a more rapid generative AI adoption. Firms that are less able to accommodate remote work because they have lower managerial, communication, or decision-making capabilities are more likely to invest in generative AI to reduce their reliance on remote workers. I rationalize these results using a task-based model of firm investments in new technology.

*I would like to thank participants at the UNC CREDA Real Estate Research Symposium, the UCLA Global Economics and Management Brown Bag Lunch, the UCLA Finance Brown Bag Seminar, the Hong Kong University Finance Brown Bag, as well as Lindsey Raymond and Anna Stansbury for helpful comments and discussion. I thank the UCLA Ziman Center for Real Estate’s Rosalinde and Arthur Gilbert Program in Real Estate, Finance and Urban Economics for generous funding. Email: gregor.schubert@anderson.ucla.edu. Website: <https://sites.google.com/view/gregorschubert>.

“As a general rule, it now looks like AI may be able to replace human labour in many virtual settings, meaning that if a task can be done remotely, it can also be potentially automated.”

Frey and Osborne (2024)

“The big way to protect yourself [from AI] as an individual is to be in a role that requires some in-person interaction, even if that’s every other month...To meet co-workers, manage, or mentor every other month creates an activity that AI cannot do.”

Nick Bloom (Business Insider, 2023)

Organizations routinely confront major technological shifts that require substantial adaptation in order for firms to benefit from them. The speed and effectiveness of firms’ responses to these changes often depend critically on their existing organizational and technological capabilities. Most recently, the release of ChatGPT in November 2022, and the subsequent rapid improvements in generative AI technology, are propelling a transformation of how many firms operate. Generative AI has the potential to increase productivity in a large share of tasks across occupations (Eloundou, Manning, Mishkin, and Rock, 2023), and field experiments have found large productivity increases when workers are given access to generative AI-based tools (Noy and Zhang, 2023; Brynjolfsson, Li, and Raymond, 2023). As a result, firms have adopted this new technology at an unprecedented pace compared to previous technology waves (Bick, Blandin, and Deming, 2024).

However, the pace of adoption has been highly uneven, with large differences in generative AI use emerging between different workers, often driven by factors internal to firms (Humlum and Vestergaard, 2025). Evidence by Bloom, Hassan, Kalyani, Lerner, and Tahoun (2021) suggests that early adopters of previous disruptive technologies retained persistent advantages over time. Given the large potential productivity impacts and labor market consequences of generative AI tools, it is therefore of first-order importance to understand why some firms lag behind while others quickly embrace generative AI technology.

In this paper, I argue that different technology waves are interconnected in their effect on firms: each adopted innovation changes whom firms hire, how decisions in the organization are made, and what technical infrastructure workers can access. These changes, in turn, alter the firm’s readiness for the next rung of technological innovation. This creates a path dependence between the technologies that a firm chooses to adopt, which I call an “organizational technology ladder.”

To provide evidence for this dynamic, I take advantage of the natural experiment created by the technological adaptation necessitated by the COVID-19 pandemic 2020-2021: virtually overnight, lockdowns forced many firms to adopt technology enabling large-scale remote work (Barrero, Bloom, and Davis, 2023). I argue that the associated shifts in work

processes, investments in digital infrastructure, and hiring patterns permanently changed firms' readiness for the emergence of generative AI in late 2022.


My empirical strategy exploits variation in *how much* remote work different firms were forced to accommodate, to show that greater remote work adoption caused firms to also adopt generative AI at a faster pace. Using detailed job posting data, I find that remote work led to organizational adaptation that eases the use of generative AI tools. At the same time, I show that technological inter-dependence does not just run one way: organizational changes resulting from generative AI adoption also alter firms' use of remote work. Moreover, I provide a theoretical framework that rationalizes the empirically observed organizational changes associated with these technology interactions.

As a first step in this analysis, I motivate the idea of links between different technologies by providing new descriptive evidence that remote work and generative AI are likely to affect one another at the level of firms and occupations. I document the following facts: (1) The same industries that have a high remote work prevalence are also adopting generative AI at a higher rate, which is illustrated in Figure 1: the log of the remote share explains 43% of the log generative AI adoption rate at the industry sector level—and similar patterns hold at the firm-level; (2) Generative AI adoption rates are higher in remote jobs; (3) Higher rates of remote hiring in an occupation during the pandemic predict more rapid generative AI adoption since the release of ChatGPT; (4) The rapid rise in remote work shares in hiring that began with the Covid pandemic reversed around the same time when ChatGPT was released and the reversal was greatest among occupations that adopted generative AI at a higher rate; (5) The firms that hired more for data management skills before the pandemic, adopt generative AI more rapidly. These facts show that there is large potential for the two technologies to compound or mitigate one another's effects. Moreover, the time pattern is suggestive of both remote work complementing the adoption of generative AI, and generative AI adoption substituting for remote work demand.

Second, the main part of the paper provides empirical estimates of the causal effects of a higher rate of remote work at a firm on its adoption of generative AI, and of the effect of generative AI exposure on remote work after the release of ChatGPT. I develop an IV strategy that identifies exogenous variation in remote work based on firms' exposure to labor market pressure to provide remote work options for their workers. I develop two different instruments based on this idea: One instrument is based on the firm's exposure to labor markets with competition from *other firms* that are more likely be offering remote work as a benefit for their workers, interacted with a firm's own ability to offer such a perk. The other instrument uses a firm's exposure to labor markets with long commutes as an exogenous variation in the benefit of remote work during the pandemic, interacted with the teleworkability of

Figure 1:
Remote work and generative AI adoption by industry sector

This figure plots the share of job postings in the YTD Sep. 2024 period that are for jobs that mention generative AI relative to those that are for remote jobs. The job postings data are from Lightcast and are aggregated into 2-digit NAICS industry sectors. The red line indicates a linear best fit.



the firm's jobs. I also construct instruments to identify within-firm variation in remote work across different occupations by further interacting these proxies for labor market pressure at the firm level with national trends in remote work across different occupations.

Using this IV approach, I find that, across firms, a 10 pp higher remote work share causes firms to increase their hiring for generative AI skills by at least 0.6 pp by 2024. Within firms, a 10 pp difference in remote work across occupations is associated with a 2.2 pp increase in generative AI mentions when hiring. These effects are robust to controlling for a flexible set of control variables that capture firms' fundamental suitability for remote work and generative AI automation, their technology capabilities before the pandemic, as well as industry sector fixed effects. Moreover, within-firm estimates control for both firm and occupation fixed effects.

Next, I estimate the effect of exposure to generative AI on remote hiring after the release of ChatGPT. I use a synthetic difference-in-differences design following Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021), which compares the hiring behavior of firms with high (top decile) generative AI automation potential to firms with comparable remote work trends before the release. I find that high generative AI exposure firms significantly reduce their remote hiring share after the ChatGPT release relative to the control group.

Greater exposure to generative AI led to a reduction in remote hiring by about 19% for the most exposed firms relative to a matched synthetic control group. This effect results from a reduction in the level of hiring across all positions, with a disproportionate decline of remote jobs.

Third, I develop a simple conceptual framework that captures the potential channels for *why* remote work and generative AI technology adoption are connected and *how* firms change their organizations when they adopt technological innovations. This model builds on the task-based approach to modeling automation (Acemoglu and Restrepo, 2022; Thompson and Autor, 2024) but incorporates the organizational effects of the two technologies of interest: I model the effect of remote work as an increase in productivity by the worker that comes at the cost of lower decision-making quality. Generative AI is modeled as an automation technology that reduces the tasks done by humans, and which leaves tasks to the human that are more sensitive to good decision-making.

In the model, each firm optimally divides workers' total time between routine production and decision-making, and the production boost from better decision-making depends both on the worker's own time investment and on the decision support from managers at the firm. The firm also decides whether to adopt new technologies that build on one another. Because the costs of adopting each technology depend on the existing levels of information technology at the firm, an investment in remote work lowers the incremental investment needed to implement generative AI. In this setting, firms differ because their tasks may be more suitable for remote work, or because they have greater generative AI automation potential, or because its output is more sensitive to a high quality of decision-making.

Under these basic assumptions, I can rationalize the observed positive relation between remote work adoption and subsequent generative AI adoption, as well as the negative effect of generative AI adoption on remote work. Intuitively, as long as the complementarities in technical infrastructure are large enough, remote work adoption always makes generative AI easier to implement. At the same time, generative AI adoption increases the importance of good decision-making. This reduces firm incentives to allow remote work, because it reduces the ability of firm management to support workers in their decisions. In short, the tasks that remain for humans to do after AI automation need more guidance from others, and thus benefit more from being done in person in the office. Importantly, this framework also results in testable predictions for how different organizational characteristics affect the magnitude of these effects, and how firms adapt to the adoption of remote work: if firms are unable to manage remote workers to be productive, automation with generative AI provides greater benefits.

Fourth, I validate these predictions by exploring the empirical evidence for the mecha-

nism through which the technologies are linked: I find that remote work adoption during the pandemic reshapes organizations as it causes firms to increase hiring for roles that are associated with higher decision-making skills, leadership skills, and managerial responsibility. At the same time, remote work adoption causes firms to increase their investment in technology skills, such as data management and machine learning. Within firms, the specific occupations that hire remotely during the pandemic also see an increase in mentions of communication and decision-making skills, and an emphasis on worker independence, but the firm’s technology skill investments are more likely to be driven by non-remote positions. When allowing for heterogeneity in the effects of remote work on generative AI adoption, I find that firms with a greater emphasis on communication-intensive roles and decision-making, and with more managerial skills, are less likely to adopt generative AI in response to remote work, while firms that previously hired for more technology skills show stronger “technology ladder” effects.

Similarly, generative AI exposure is less likely to result in a decline in remote hiring for firms with better managerial skills, better communication, and workers with more individual decision-making skills, in line with the idea that organizations that can better overcome the communication barriers associated with remote work are then less likely to embrace generative AI tools—compared to firms that struggle with remote processes. As a case study for this effect, I show that firms that issued return-to-office mandates—as a proxy for perceiving remote work to have low benefits—have a significantly higher likelihood of investing in generative AI skills in response to higher remote work shares.

These results suggest the existence of an “organizational technology ladder”: firms that adopt one technology—remote work in this case—then benefit more from subsequent technology waves—here, generative AI—that build on the same investments in digital infrastructure and require similar changes in work processes. The results on changes in skill demand suggest that this mechanism in part operates through the way that remote work induces firms to increase their level of technology skill, and shifts their hiring towards more experienced workers in decision-making intensive roles. These changes in organizational characteristics increase their ability to benefit from the adoption of generative AI tools. This path dependency can further exacerbate differences in productivity among firms as the effects of the two biggest technological changes in firms’ work processes during the last decade amplify one another in some organizations, and are dampened in others.

Related Literature. This paper builds on, and aims to link, two growing, but separate, literatures on the effects of remote work and generative AI on firms.¹ While this paper is,

¹For an overview of the remote work literature, see, e.g. Barrero et al. (2023), while an early review of the emerging evidence regarding firm-level effects of generative AI can be found in Eisfeldt and Schubert

to my knowledge, the first study to draw a direct link between remote work and generative AI adoption, previous studies have found that information and communication technology investments may complement work-from-home productivity: Boeri, Crescenzi, and Rigo (2024) show that, while the overall effect of work-from-home on productivity was negative for firms in Italy, this effect was less negative for larger firms, those that had previously invested in laptops and server setups before the pandemic, and for knowledge-intensive firms. Also, in line with the findings in this paper, some studies have found that one of the reasons why remote work benefits from existing investments in communication technology may be that remote work creates more communication overhead (Bao, Li, Xia, Zhu, Li, and Yang, 2022; Gibbs, Mengel, and Siemroth, 2023). My finding that greater remote work may be associated with a greater investment in potentially automating technologies is also consistent with a literature that finds that remote work changes the attachment between firms and their employees, leading to a decline in the quality of employee engagement with co-workers (Gibbs et al., 2023; Akan, Barrero, Bloom, Bowen, Buckman, Davis, Pardue, and Wilkie, 2024; Emanuel and Harrington, 2024).

Moreover, my IV estimation approach relies on the argument that work-from-home is a workplace amenity that matters for labor market competition among employers. There is now a growing literature that shows that work-from-home is considered equivalent to a sizeable wage increase for workers in terms of the attractiveness of a job. Maestas, Mullen, Powell, Von Wachter, and Wenger (2023) find that the opportunity to telecommute is equivalent to a 4.1% wage increase for workers. Similarly, Powell and Wenger (2024) show that hiring managers consider providing opportunities to work from home as equivalent to a 9.4% wage increase for their workers. Other studies have found that offering hybrid work options increases worker retention (Bloom, Han, and Liang, 2024), and that remote workers are more productive as a result of shorter commute times and better work-life balance (Choudhury, Foroughi, and Larson, 2021).

The generative AI-related insights in this paper relate to an emerging literature that studies the effect of generative AI adoption on firms. For example, Gulati, Marchetti, Puranam, and Sevchenko (2025) show that generative AI adoption in particular roles within firms is associated with higher demand for cognitive and social skills in job postings for those roles.

The proposed notion of a “technology ladder” that operates through organizational adaptation builds on earlier studies which showed that technology can affect the organization of firms and thereby change wage inequality (Garicano and Rossi-Hansberg, 2006), and that technologies differ in their effect on the centralization of decision-making within firms

(Forthcoming).

(Bloom, Garicano, Sadun, and Van Reenen, 2014). Moreover, there is existing evidence that the effect of technologies on productivity depends on the degree to which complementary organizational changes are made (Brynjolfsson and Hitt, 2000), including developing managerial experience or worker skills (Brynjolfsson, Rock, and Syverson, 2021). Bresnahan, Brynjolfsson, and Hitt (2002) argue organizational adaptation can reinforce the effect of technologies on labor demand. The contribution in this paper is to link the effects of adopting one technology to the likelihood and effectiveness of adopting another.

This paper has the following structure: the next section discusses the data that I use. Section II presents the empirical approach to estimating the causal effects of remote work adoption on generative AI and vice versa. Then, I provide a set of new stylized facts regarding the relation between remote work and generative AI in Section III.A; I explore the determinants of remote work adoption in Section III.B; I estimate the causal effect of remote work on generative AI adoption in Section III.C, and the effect of generative AI exposure on remote work in Section III.D. To provide an explanation for the empirical findings, Section IV provides a simple conceptual framework, and in Section V, I provide additional evidence on the heterogeneity of the effects and the changes in hiring that result from remote work and generative AI adoption. Finally, Section VI draws out the implications of these results.

I. Data

Occupational wage and employment data. I obtain data on employment and average wages by occupation and MSA from the Bureau of Labor Statistics' Occupational Employment Statistics. I crosswalk all occupational data to SOC 2010 codes and crosswalk New England City and Town Areas (NECTAs) into Core-Based Statistical Area codes (CBSAs) for the few geographies where the OES data is not provided at the level of CBSAs.

Job postings. Both hiring activity and investments in Generative AI skills and other characteristics of new jobs are measured using job postings data from Lightcast. This data consists of the near-universe of online job postings in the U.S. from Jan. 2010 to Sep. 2024. I filter job postings in the following way: (1) Drop job postings by staffing companies. (2) Drop all jobs flagged as internships. (3) Retain only full-time jobs, dropping part-time jobs. Moreover, I create variables that represent particular key characteristics based on the job postings data: most importantly, the classification of *remote* jobs used in most of the analyses below represents jobs that indicate that the position is fully remote. However, some jobs may allow only for hybrid remote working, so, where applicable, I will explicitly note where *hybrid remote* jobs are included. I also create dummy variables for the minimum level of required education distinguishing whether a high school, Associate, Bachelor, Master, or Ph.D.

degree are required. Similarly, I create dummies for different required experience levels in years. When job postings are aggregated into firm or occupation or firm-by-occupation level panels, the counts reflect the number of job postings *posted* in each category in the time period, which does not necessarily reflect actual hiring, but rather proxies for hiring demand.

Generative AI adoption. While I cannot observe the *usage* of Generative AI within firms, the job posting data allows me to see when firms are either requiring Generative AI-related skills from new workers, or describing the work activities in a job as involving Generative AI tools. Therefore, I will use mentions of Generative AI and related technologies in job postings as a proxy for the degree to which those jobs involve using Generative AI-related tools. Labels for "Generative AI" skills mentioned in a job posting are supplied by Lightcast based on the job posting text.

Occupation characteristics. To evaluate whether the estimated effects differ between occupations with different characteristics, I use O*Net data on the characteristics of different occupations to construct summary measures that assess the prevalence of different skills in different occupations. I construct the following measures of occupation characteristics: To capture the importance of “decision-making” in a job, I average the level of related work activities involved in an occupation according to O*Net, based on Deming (2021).² Indicators of coordination required, interaction, social skills and social skills interacted with analytical skills are constructed as described in Deming (2017). I construct a measure of job inflexibility based on Goldin (2014), and a measure of leadership skills based on Schubert, Stansbury, and Taska (2019), which is an average of the work activities for which a high value best predicts managerial leadership positions.

Skill demand. I also construct firm- and occupation-by-firm measures of skill demand directly from the job posting data by using tags provided by Lightcast for job postings that require a college degree, and advanced degree, or different skills. I focus on teamwork, communication, decision-making, data management, data science, machine learning, deep learning, and independence skills. I also construct an indicator for jobs that require more than 4 years of experience.

Commuting cost. To measure difference in commuting cost across different MSAs, I use county-level data on average travel time to work by county from IPUMS NHGIS (Manson, Schroeder, Van Riper, Knowles, Kugler, Roberts, and Ruggles, 2024), which is based on the 2015-2019 American Community Survey data. I aggregate this data to the MSA level using population weights from the 2000 Census.

Return-to-Office policies. I obtain information on firms’ return-to-office (RTO) policies

²These activities are: Making decisions and solving problems; Developing objectives and strategies; Organizing, Planning and prioritizing work.

by collecting crowdsourced public information from the website Flex Index. I match companies listed on Flex Index to companies in Compustat using company names and locations. The final dataset includes data on 1,336 companies and their office presence requirements. I classify a company as having provided a “Return-to-Office” mandate, if its RTO policy requires either a full-time presence, or specifies a non-zero minimum number of days in the office.³

II. Empirical Approach

To test whether there is in fact a “technology ladder effect”, where one technology’s adoption impacts the use of the other by firms, we want to estimate the effect β of having a greater share of remote jobs on firms’ investment in generative AI skills. This effect is estimated from regressions of the form

$$\text{GenAIJobShare}_i = \alpha + \beta \text{RemoteWorkShare}_i + \text{Controls}_i + \varepsilon_i, \quad (1)$$

where the unit of observation i can be a firm or an occupation-by-firm unit. The Controls_i are detailed below when discussing how I address identification concerns. While the main estimation will be cross-sectional, using job posting data for the 12 months ending in, and including, September 2024, to construct the dependent variable, some of the control variables are constructed in earlier periods to capture past characteristics of the firm, occupation, or location.

A. Remote work and Generative AI: identification issues

Why would having many workers in remote positions *cause* a firm to adopt generative AI at a faster rate? To structure the empirical approach, it is important to first clarify the potential causal channels of interest and empirical challenges in identifying them.

The effects of remote work on generative AI adoption are potentially ambiguous: generative AI tools could complement remote workers more than in-office workers, for example building on their existing familiarity with digital workflows, or enabling better remote work tools. However, if generative AI is more likely to automate remote work tasks entirely, the use of these tools can also substitute for remote workers. It is therefore an empirical question whether a greater use of remote workers is likely to lead to more or less investment in generative AI skills by firms. The organizational technology ladder mechanism where

³See Appendix B for more information on the Flex Index data cleaning and the method for matching to Compustat.

remote work adoption *causally* affects generative AI adoption can arise from a number of different channels:

1. *Augmentation*: Remote work jobs might benefit directly from applying generative AI tools, and this greater benefit of adoption drives higher demand for generative AI skills.
2. *Digital infrastructure*: the digital infrastructure, data processing, and other technical capacity improvements that result from implementing a remote work-friendly organization, make it easier to adopt generative AI technologies.
3. *Workflow restructuring*: by making workflows modular, asynchronous, decentralized, and not reliant on real-time in-person input, while incorporating mechanisms for remote quality validation, remote-adapted firms are more suitable for inserting automated AI tools into the value chain.
4. *Automation*: For some firms remote workflows are less suitable, but they find themselves unable to enforce a post-pandemic return to in-person work. In that case, automation through generative AI can provide a channel for organizational adjustment that allows firms to reduce the share of their work done remotely.

On the other hand, there are also possible mechanisms through which remote work and generative AI adoption may coincide in the same firms and occupations, but without a causal link from one to the other:

4. *Task characteristics*: it is possible that the tasks that are amenable to being done remotely also happen be the ones for which generative AI is useful. In that case, greater adoption of both technologies is driven by greater suitability for the technologies being correlated.
5. *Innovation orientation / tech-savviness*: some firms may have organizational cultures and leadership that are more open to, and capable at, embracing new technologies and being at the leading edge of innovation. This trait may make these firm more likely to experiment with, and adopt, both remote work and generative AI as the newest waves of technological progress.

One key empirical issue is trying to distinguish these causal and non-causal channels. To see that this is a relevant concern, I compare a proxy for the remote work suitability of an occupation—the “teleworkability” measure by Dingel and Neiman (2020)—to the measure of generative AI exposure at the occupation level from Eisefeldt, Schubert, Zhang, and Taska

(2023). Figure 3, Panel A, shows that there is a positive relationship between occupations having tasks that are suitable for remote work and tasks that are exposed to generative AI capabilities.

The ideal setting for estimating the effect of remote work adoption on generative AI adoption would therefore require comparing groups of jobs that have similar *suitability* for both remote work and generative AI deployment, and also are at companies with similar “tech-friendliness” or innovative capacity. The ideal experiment to identify causal effects would then require one group of these jobs to experience greater prevalence of remote work for an exogenous reason, so that we can compare generative AI adoption rates to see what the magnitude of the causal channel operating *through* remote work adoption is.

I approximate this natural experiment by controlling for potential confounders and instrumenting for remote job prevalence using exogenous variation that is plausibly unrelated to firm-level differences in unobserved characteristics that might be driving generative AI adoption. As control variables, I include Dingel and Neiman (2020) teleworkability and Eisefeldt et al. (2023) generative AI exposure scores to capture the intrinsic nature of the tasks performed in different firms and occupations that might make them more or less suitable for adopting these technologies. Moreover, the most stringent specifications include fixed effects at the level of the firm or occupation that capture general differences in tech-savvy, or in the tendency to adopt either of the technologies of interest, as well as controls for the average education level of a firm’s hiring that might capture the ability of employees to implement innovations.

B. Remote work instruments

To identify exogenous variation in the remote job share at the firm level, I exploit a novel source of variation in whether a firm adopts remote work at a high rate in a particular location: the interaction between *labor market pressure* and the *ability* to let employees work remotely. That is, many studies (e.g. Maestas et al. (2023), Powell and Wenger (2024)) have found that workers consider the ability to work remotely as a sizable non-monetary benefit.

As a result, similar to the way that workers’ outside options encourage firms to match wages on offer elsewhere (Schubert, Stansbury, and Taska, 2024), labor market competition induces firms to offer remote work options. This pressure should be larger in labor markets where *other* employers are likely offering this perk, or where commuting is more burdensome and workers therefore likely perceive a larger benefit of working from home (Flynn, Ghent, and Nair, 2024). Based on this intuition, I construct two novel instruments

for remote work.

Competition instrument. While the actual adoption of remote work by other employers in particular labor markets is both difficult to observe (as remote job locations are, by definition, not well defined in job postings), and might be simultaneously determined with post-pandemic choices by the firm of interest, the *average pre-pandemic ability* to work remotely in the labor markets that a firm hires in is both observable and unlikely to be driven by the focal firm’s ex post adoption behavior. Thus, I construct a firm-level instrument for remote work adoption based on labor market competition as

$$\begin{aligned}
 Z_f^T &= \left(\sum_m \phi_{fm,2019} T_{m,2019} \right) \times T_{f,2019} \\
 &= \underbrace{\text{Avg. Labor Market Teleworkability}_{f,2019}}_{\substack{\text{Remote work adoption potential} \\ \text{of the firm's labor market pre-Covid}}} \times \underbrace{\text{Firm Teleworkability}_{f,2019}}_{\substack{\text{Firm remote work adoption} \\ \text{potential pre-Covid}}},
 \end{aligned}$$

where $T_{m,2019}$ is the average Dingel and Neiman (2020) teleworkability among job postings in a particular MSA based on 2019 job postings, and the MSAs are weighted by the share $\phi_{fm,2019}$ of all firm f hiring done in each location m as of 2019. $T_{f,2019}$ is the average teleworkability among job postings by firm f in 2019.

Commuting cost instrument. To proxy for the pressure to permit remote work during and after the pandemic that arises from workers’ perceived cost of commuting, I use estimates of the pre-pandemic average travel time to work in each MSA, based on ACS data from 2015-2019. I aggregate these geographic proxies to a firm level measure of commuting cost using the firm’s share of hiring in different locations in 2019. As the desire by employees to work remotely will only translate into remote jobs if positions at a firm are teleworkable, I again interact this measure of remote work pressure with the firm’s teleworkability when constructing the commuting cost instrument:

$$\begin{aligned}
 Z_f^C &= \left(\sum_m \phi_{fm,2019} C_{m,15-19} \right) \times T_{f,2019} \\
 &= \underbrace{\text{Avg. Labor Market Commuting Cost}_{f,2019}}_{\substack{\text{Travel time to work} \\ \text{in the firm's labor market pre-Covid}}} \times \underbrace{\text{Firm Teleworkability}_{f,2019}}_{\substack{\text{Firm remote work adoption} \\ \text{potential pre-Covid}}},
 \end{aligned}$$

where $C_{m,15-19}$ is the average travel time by MSA, and the MSAs are weighted by the share $\phi_{fm,2019}$ of all firm f hiring done in each location m as of 2019.

Occupation-by-firm instruments. As I also want to be able to estimate exogenous variation across occupations *within firms*, I construct occupation-by-firm-level instruments based on a similar intuition. For occupations where the remote work share was higher at

the national level post-Covid (but before the ChatGPT release) in 2021-2022, employers who are able to let their workers work remotely likely face more labor market pressure to grant this perceived perk. Thus, adoption of remote work in an occupation-by-firm cell can vary in response to this labor market channel, independent of overall firm incentives to adopt generative AI. To be specific, I construct occupation-by-firm instruments by interacting the average labor market teleworkability or average commuting cost in a firm's hiring markets as of 2019 with the remote work share in the occupation during the late pandemic years 2021-2022:

$$Z_{fo}^T = \left(\sum_m \phi_{fm,2019} T_{m,2019} \right) \times \text{RemoteWorkShare}_{o,21-22}$$

$$Z_{fo}^C = \left(\sum_m \phi_{fm,2019} C_{m,15-19} \right) \times T_{f,2019} \times \text{RemoteWorkShare}_{o,21-22}$$

The commuting cost version of this instrument also includes teleworkability of the firm, as this seems to empirically lead to a stronger predictive relationship with remote work adoption.

Exclusion restriction. These shift-share instruments rely only on pre-Covid job compositions in different locations, firms and occupations, and exogenous national trends in remote work before generative AI technology was salient and are therefore unlikely to be correlated with endogenous generative AI adoption choices by a firm, other than through the remote work channel. Formally, while the relevance condition in this setting requires that the *interaction* between labor market suitability for remote work (due to competition or commuting costs) and a firm's ability to let workers go remote drives remote work adoption in a way that is not explained by either overall labor market suitability for remote work, or firm suitability on their own. The exclusion restriction then requires, for example, that if we compare two firms that have similar teleworkability, similar remote shares pre-pandemic, similar exposure to generative AI, are in the same industry, and have similar observable hiring for technology skills, then the one that finds itself in a labor market with longer commutes does not systematically differ from the other firm in a way that drives generative AI adoption *except* through greater remote work adoption during the pandemic in the high commuting cost firm.

At the occupation-by-firm level, the exclusion restriction is even more narrow and requires that greater pressure to adopt remote work, arising from high national remote work shares in an occupation, leads to differential generative AI adoption for some occupations, but only within firms that face higher labor market pressure to allow remote work, and that this effect operates only through a remote work channel. Moreover, in this setting we are

able to condition on a firm’s overall tendency to adopt generative AI, as well as an occupation’s overall tendency to use generative AI.

As a firm’s or occupation-by-firm’s teleworkability or exposure to particular labor markets on their own might be correlated with particular unobservable firm characteristics, the IV estimations will always control for the uninteracted components of the instrument.

Unit of observation. Depending on the particular dimension of the variation in generative AI adoption that we are interested in, a different unit of observation is appropriate. One question is whether the *same firms* are seeing higher remote work adoption and generative AI adoption, which matters for inequality in productivity across the economy as some firms might see compounding effects of technology, and might also provide evidence about the degree to which there is variation in corporate strategy across firms. This question is best studied using firm-level evidence. At the same time, there are questions of whether generative AI adoption is more prevalent in the same occupations where remote work was adopted, which speaks to the degree to which these technologies complement or substitute for one another *within* particular jobs. I will use a firm-by-occupation sample to study the latter questions, such that I can identify differences in effects across occupations within the same firm.

C. Estimating the effect of generative AI adoption on remote hiring

The stylized fact pattern shown in Panel B of Figure 2, and described in Section III.A, shows that higher later generative AI adoption is associated with a strong reversal in remote work shares around the time when ChatGPT was released. I want to test more formally whether these differential trend breaks in remote work shares around the ChatGPT release period might be driven at least in part by generative AI innovations.

Identifying whether this association in trends is likely to represent a causal “technology substitution effect” requires an empirical design that can address the identification issues mentioned before. Given the surprising release of ChatGPT in November 2022 and rapid rise in adoption afterwards (see Figure 2), an event study methodology, such as differences-in-differences might be appropriate in this case. However, estimating a causal effect using a differences-in-differences approach normally requires a parallel trend assumption to hold (see, e.g. Rambachan and Roth (2023)), which assumes that treatment and control group can be expected to behave similarly in the counterfactual scenario where no treatment occurs. As the descriptive analysis shows, generative AI adoption is clearly not randomly assigned with regard to remote work shares—in fact, I argue in this paper that higher remote work adoption eventually *causes* greater generative AI adoption. Moreover, there is

a positive task level correlation in exposure to the two technologies—as shown in Figure 3—which means that adoption will also be correlated. As a result, firms that have higher later generative AI adoptions likely also had steeper run-ups in remote work earlier in the pandemic. Any causal effect estimation of the effect going in the other direction—from generative AI adoption to remote work—therefore needs to try and hold remote work adoption trends constant when comparing firms around the time of the ChatGPT release event.

In traditional difference-in-differences designs, one would use control variables to try and capture the common variation in the confounding variable, which would mean trying to incorporate all the control variables that can capture potential drivers of differential remote work share trends at the firm level. As there is not necessarily a definitive list of these drivers, one would worry that the control and treatment group are not fully comparable in terms of their remote work trends even after including suitable observables.

Synthetic difference-in-differences. A more appropriate methodology in this case is therefore the synthetic difference-in-differences (SDID) design proposed by Arkhangelsky et al. (2021): instead of ex ante selecting which control variables might align remote work trends in the pre-event period between firms, this methodology first estimates which untreated firms have remote work trends that are most similar to the pre-trend of the treated units, and uses a weighted average of the most-comparable firms to create a “synthetic” control group. The estimation of post-event relative trends between the treated and control groups then proceeds analogous to a difference-in-differences estimation, but weighting potential control group firms based on the optimal weights, and also applying time weights that optimally put higher weights on pre-periods that are more comparable to post-event periods for the control group.

Conceptually, this corresponds to determining the treatment effect τ of being in a high generative AI exposure group in a regression equation of the form

$$\text{RemoteWorkShare}_{it} = \mu + \alpha_i + \gamma_t + \tau \mathbb{1}[\text{Post-ChatGPT}]_t \times \mathbb{1}[\text{High GenAI Exposure}]_i + \varepsilon_{it}, \quad (2)$$

which includes firm fixed effects α_i and time fixed effects γ_t , which absorb level differences between firms and national trends in the remote work share. The treatment effect τ is estimated with SDID by solving

$$(\hat{\tau}, \hat{\mu}, \hat{\alpha}, \hat{\gamma}) = \underset{\tau, \mu, \alpha, \gamma}{\text{argmin}} \left\{ \sum_{i=1}^N \sum_{t=1}^T (\text{RemoteWorkShare}_{it} - \mu - \alpha_i - \gamma_t - \tau \text{Treated}_{it})^2 \hat{\omega}_i^{\text{sdid}} \hat{\lambda}_t^{\text{sdid}} \right\} \quad (3)$$

where Treated_{it} stands for the interaction $\mathbb{1}[\text{Post-ChatGPT}]_t \times \mathbb{1}[\text{High Gen. AI Exposure}]_i$

between the post-release period indicator and the high generative AI exposure group assignment. The difference to conventional difference-in-differences comes from the optimal weights $\hat{\omega}_i^{\text{sdid}}$ and $\hat{\lambda}_t^{\text{sdid}}$ that are estimated in a first step to create a control group that matches the treatment group’s trend in the remote work share (as opposed to differences in levels which are absorbed into the firm-level fixed effects).⁴

Event study data. In order to use the SDID estimator, I need a balanced panel of company-by-period observations. I retain all firms that have at least 10 job postings in all quarters from Q1 2021 to Q3 2024, which is the sample period for the event study estimation, consisting of 7 quarters before and after the ChatGPT release in Q4 2022.

Treatment group. To define the treated group, note that *actual* adoption of generative AI is endogenous with regard to a firm’s characteristics. In contrast, a firm’s *exposure* to generative AI, based on its employment structure before ChatGPT was released, more plausibly represents a quasi-exogenous shock as it interacts with the exogenous timing of the release. I use the Eisefeldt et al. (2023) measure of occupation exposure to generative AI, and compute a firm’s 2022 exposure in its hiring as a job posting weighted average over the exposure of the occupations that it is hiring for. I sort firms by the continuous exposure variable and assign the top decile of firms by exposure to the “treated” group. Note that this means that the estimated effect does not compare firms with high exposure to firms with zero exposure, but rather to a synthetic control group that has *some* exposure and similar remote work trends to the highest exposure groups.

III. Remote Work and Generative AI Adoption

In this section, I first provide descriptive facts about the relation between remote work and generative AI that motivate the causal analysis. Second, I explore what firm characteristics are associated with greater adoption of remote work. Then, I provide evidence of the “technology ladder effect:” remote work adoption has a positive causal effect on the subsequent investment in generative AI skills. Finally, I estimate the effect of generative AI adoption on a firm’s hiring of remote workers.

A. Stylized facts: remote work and generative AI

To understand whether firms are more or less likely to change their investments in hiring for generative AI skills if they already have a remote workforce, or might change their

⁴I implement this estimation procedure using the `sdid` package in Stata, which is described in Clarke, Pailanir, Athey, and Imbens (2023).

remote hiring with the emergence of generative AI tools, I start by providing descriptive evidence that these two technologies are likely to be connected. In this section, I document a number of novel facts about remote work and generative AI skill demand.

Fact #1: Firms that hire remotely also hire for generative AI skills. Is there a cross-sectional relationship between occupational remote work adoption and generative AI use? Figure 1 shows that the industry-level correlation between the adoption of remote work and of generative AI is strongly positive. This does not necessarily imply that the two technologies should be positively correlated in their prevalence at the firm level. For instance, if they represent different solutions to a similar organizational problem but involve some fixed costs or are not compatible with one another, firms might choose to deploy one or the other, but rarely both. In that case, a positive correlation at the industry level could coincide with a negative correlation at the firm level. However, in Figure A.1 I show in job postings data for the year-to-date (YTD) Sep. 2024 that the relation is also positive across firms, with the log of the remote share explaining 13% of the variation in the log generative AI adoption rate across firms.

Fact #2: Generative AI adoption is higher in remote jobs. To provide some intuition for which occupations are adopting the technologies of interest, Table I ranks 6-digit occupations by their prevalence of remote jobs (panel A) and generative AI (panel B) in job postings, as of YTD Sep. 2024. Each table shows the top 20 occupations by technology prevalence that have at least 5,000 job postings in the sample.

The lists show that remote work is highly prevalent among occupations that vary substantially in the degree of social interactions and education required, including Actuaries, Telemarketers, and Mental Health Counselors, for example. High generative AI adoption is more prevalent among technical or writing-heavy occupations, but also includes Animators, as well as Marketing Managers. In panel C, I explore whether the list of top generative AI adopters looks different *conditional* on being a remote worker (including only occupations with at least 100 remote job postings). Notably, many of the top adopters in general are also heavily featuring generative AI in remote job postings—there is substantial overlap between the lists panels B and C—but the adoption rates among remote workers are substantially larger. For example, Technical Writer job postings mention generative AI in 6.1 percent of all job postings, but in 18 percent of remote job postings. This shows that the occupations that adopt generative AI at high rates are similar between remote and non-remote workers, but generative AI adoption rates are higher among remote workers.

Fact #3: Higher remote hiring during the pandemic predicts rapid generative AI adoption. How has the adoption of generative AI and remote work changed over time? Both of these technologies experienced a key time period that accelerated their adoption: for

remote work, the onset of the Covid pandemic in Q1 2020 played an important role as many positions went remote during the first year of the pandemic. For generative AI, the release of ChatGPT in November 2022 launched the technology into public awareness.

Figure 2, Panel A, shows how adoption rates of generative AI have varied over time and by remote work prevalence: panel A shows that generative AI adoption surged after ChatGPT was released, with 0.3-0.4% of all job postings mentioning the technology on average, or 21K jobs per quarter as of Q3 2024. Importantly, this generative AI adoption rate has been much higher in occupations that had high rates of remote work prevalence before ChatGPT was released. The occupations in the highest quartile of remote work adoption have generative AI adoption rates around 4 times the national average. Figure A.2, panel A shows that the same positive correlation holds in changes: occupations that adopted remote work at a faster rate from 2019 to 2022 also increased hiring more for generative AI-related skills from 2022-2024.

Fact #4: Generative AI adoption coincides with post-pandemic declines in remote hiring. What happened to remote work *after* the release of ChatGPT? Panel B of Figure 2 shows that the remote work share in job postings experienced a surge starting in Q2 2020, and continued to rise after that, all the way *until ChatGPT was released*. After Q4 2022, remote work jobs started to decline again. It is possible that this timing is a coincidence and results from the release coinciding with changes in the hiring cycle for industries with large shares of remote workers. For instance, early 2023 saw large layoffs in the tech sector in the U.S. that were likely unrelated to the ChatGPT release and may have impacted remote work hiring patterns.⁵ Panel B of Figure A.2 shows that generative AI adoption from the ChatGPT release to Q3 2024 is negatively correlated with changes in remote work during the same time period among the largest occupations. The rise of remote work reversed exactly when ChatGPT was released, and this reversal was greatest in occupations that ended up adopting generative AI at a higher rate.

B. Determinants of remote work adoption

What type of firm is more likely to adopt remote work? The conceptual framework suggests that firms with greater technology skills should be more likely to invest in remote work technology. I test this prediction and the effect of other firm characteristics often used in the finance literature in regressions of the form

$$\text{RemoteWorkShare}(21-22)_i = \alpha_{ind} + \beta \text{FirmCharacteristic}(2019)_i + \text{Controls}_i + \varepsilon_i,$$

⁵For a list of reasons for the layoffs, see, e.g.: <https://stratechery.com/2023/the-four-horsemen-of-the-tech-recession/>

where all independent variables are standardized and the controls include the 2019 value of the dependent variable, so the coefficients can be interpreted as the effect of a standardized difference in the characteristic on changes in the remote work share. To proxy for the effect of simply hiring for roles that are more suitable for remote work, the control variables also include firm-level teleworkability in 2019 and 2021/2022, as well as the company’s share of jobs requiring a college education in 2019, and NAICS 2-digit sector fixed effects.

The results are shown in Figure 4: The first row shows the effect of a higher share of firm hiring in computer occupations on remote work adoption: firms that hire more for computer-related positions are more likely to go remote. The coefficients in rows 2-5 capture various dimensions of technology skill prevalence in pre-pandemic hiring of the firm and show that hiring for data management, deep learning, machine learning, and data science skills predicts greater remote work adoption. While these findings are not causal, they are consistent with the predictions regarding a technological capability mechanism for greater remote work adoption in the model.

While the model does not make explicit predictions for other organizational and financial characteristics, these can provide context for understanding why some firms adopt remote work: the figure shows that firms that are larger or less R&D-intensive, or that have a higher labor share or labor intensity in production⁶ are more likely to adopt remote work. With regard to financials, measures of higher earnings, such as gross profitability, ROA, and ROE, all predict higher remote work. However, as these measures require Compustat data, the sample size for these estimations is much smaller than for the main job posting sample and these results should therefore be considered only suggestive.

C. Remote Work Impact on Generative AI Adoption

To test the key hypothesis that there exists a “technology ladder effect,” this section estimates the causal effect of remote work on Generative AI adoption both across firms, and across occupations within firms. Then, I explore heterogeneity with regard to which occupations are the ones that adopt Generative AI technologies in response to higher remote work prevalence at a firm, in order to better understand the mechanism for the effect.

First-stage. To visualize the relationship that underlies the variation in the instrument, Figure 5 plots the relationship between the competition IV instruments and the remote work shares at the firm and occupation-by-firm level, including a full set of control variables. As the figure shows, the instruments have a strongly positive relationship with the endogenous remote work prevalence (see also Table II, column (5) for the firm-level

⁶The latter is defined as $\ln(\text{Employment}/\text{PP\&E})$.

first-stage coefficient).

Firm-level results. To what degree is the correlation between remote work shares at firms and their tendency to invest in hiring workers with generative AI skills causal? I use the IV approach to estimate the equation

$$100 \times \text{GenAIJobShare}('23-'34)_i = \beta \text{RemoteWorkShare}('21-'22)_i + \alpha_{ind} + \text{Controls}_i + \varepsilon_i$$

for the generative AI share for the YTD Sep. 2024 period and the remote share during 2021-2022 (excl. Q4 2022). The results using the competition instrument are shown in Table II. Column (1) shows the OLS results, and columns (2)-(4) control for increasingly stringent measures of the underlying firm characteristics that might be jointly driving both remote work adoption and generative AI investments. The coefficient corresponds to a % elasticity between the remote share and the generative AI skill share in the firm's job postings, such that the most stringent specification in column (4) suggests that a 10 pp higher remote share causes about a 0.6 pp higher generative AI skill share. The control variables include proxies for the task-level teleworkability (Dingel and Neiman, 2020) and generative AI exposure (Eisfeldt et al., 2023) at the firm, the average commuting time in a firm's labor markets, as well as industry sector fixed effects, the level of education required for average job postings at the firm, and the firm's pre-pandemic history of hiring for technology skills as a proxy for its tech-savviness. These controls mean that the effect is unlikely to be driven by a sorting of jobs that are more suitable for both technologies into the labor markets that are affected by the exogenous remote work variation induced by labor market competition. Note that the IV estimates are substantially larger than the OLS estimates, suggesting that firms that endogenously choose to offer remote work tend to be less likely to adopt generative AI.

Robustness to choice of instrument and controls. In the estimation in Table II, I used the instrument based on labor market competition from other teleworkable firms interacting with a firm's own teleworkability. However, there may be residual concerns that this instrument still selects for firms that are unobservably different due to their sorting into labor markets with particular competitors. To allay this concern, I repeat the same estimation with the instrument based on average commuting time in its labor markets interacting with a firm's teleworkability. If one believes that this is more plausibly excluded from the estimation equation, conditional on the other controls, then the relevant estimates are shown in Table A.1: the effect estimates also significant and positive and I cannot statistically reject the hypothesis that they are identical to those in Table II. As the instruments yield very similar estimates, but the first-stage Kleibergen-Paap F-statistics suggest the competition IV is a stronger instrument than the commuting time IV and yields more pre-

cise estimates, I use the former as my baseline for all other estimates below unless stated otherwise. However, all results are robust to using the commuting time instrument instead.

To evaluate the sensitivity of the results, we can also consider the changes in the estimated coefficient between columns (2), (3), and (4) in Table II: As I control for a more detailed set of firm characteristics that are plausibly correlated with a firm’s overall ability to embrace innovation or new technology, the causal effect estimate changes very little and, if anything, increases. This suggests that the estimated effect is not driven by channels correlated with a firm’s industry, or hiring for technology skills or advanced degrees before the pandemic. The robustness of the effect is further supported by the within-firm estimates.

Within-firm results. Is the coincidence between remote work and generative AI driven only by firm-level dynamics, or is there also a causal effect of one on the other across occupations *within* firms? I estimate these within-firm causal effects using the IV approach that exploits occupation-level differences in the tendency to adopt remote work during the pandemic, interacted with one of the two instruments that capture the firm’s labor market pressure to adopt remote work.

Table III shows the results for this within-firm effect using the competition IV: columns (2)-(3) correspond to adding occupation fixed effects and similar control variables as were already included in Table II, and column (4) adds firm fixed effects and occupation fixed effects. This means that the latter effect is only identified off within-firm variation across occupations. The estimated effect suggests that higher remote work prevalence has a large positive causal effect across occupations within a firm (and across firms within an occupation), with a size in column (4) about thrice that of the across-firm effect. This provides additional confirmation that the across-firm results are not driven by some unobserved firm-level attitude towards innovation—as such an effect would be eliminated by the firm fixed effects included here. Appendix Table A.2 shows again that the effect estimates are statistically indistinguishable when using the commuting time instrument instead.

These results also mean that, while differences in corporate processes and capabilities as a result of firm-level changes in response to remote work can perhaps explain part of the puzzle of why the adoption of these two technologies is correlated, this is not the full story. The inclusion of occupation and firm fixed effects means that the effects cannot be explained purely by a change in the composition of hiring by firms towards occupations that tend to use generative AI. The large within-firm causal effects suggest that *particular jobs* become more suitable for generative AI if those jobs themselves adopted higher rates of remote work at that firm.

D. *Generative AI Effects on Remote Work*

What happens to remote work after a firm adopts generative AI? In this section, I provide firm-level evidence on how remote hiring at firms evolves in response to generative AI adoption.

The first question I explore whether there is systematic evidence supporting the stylized fact pattern shown in Panel B of Figure 2, and described in Section III.A, of higher later generative AI adoption being associated with a strong reversal in remote work shares at exactly the moment when ChatGPT was released. That is, I want to know whether this association reflects a causal “technology substitution effect.”

Firm-level trends. To show the variation in the raw data that underlies this estimation and to verify that the timing of generative AI innovations is associated with remote work prevalence even at the firm level, I do a similar analysis to that in Panel B of Figure 2, but for firms. In Figure 8, I sort firms by their quartile of generative AI exposure. The reason to use generative AI exposure rather than later generative AI adoption here is that we already know that actual generative AI adoption will be affected by remote work trends, so this might lead to reverse causality in the estimation. In contrast, the firm’s task-based exposure to generative AI, measured before the ChatGPT release, is unlikely to be shaped by the differential remote work trends after the release that we’re interested in. Comparing the firm-level remote work trends by generative AI exposure shows similar patterns to the occupation-level results: firms with higher exposure have steeper run-ups in remote shares that reverse in the quarter when ChatGPT is released.

Generative AI exposure and generative AI adoption. While the use of generative AI exposure based on the task composition of jobs, instead of actual generative AI adoption avoids the endogeneity issues associated with realized firm choices to adopt the new technology, it raises the question of whether the Eisefeldt et al. (2023) measure of firm-level exposure to generative AI actually predicts firm-level adoption of generative AI. One way to validate this “first-stage” relationship is in the cross-section: Figure 9 shows how the share of all job postings (panel A) and remote job postings (panel B) at the firm level in YTD Sep. 2024 varies with the firm’s generative AI exposure. As the figures show, there is a monotonically increasing relation between the predicted exposure measure and the actual adoption across firms.

I can also estimate this relationship in the same SDID event study setting as the remote work effects that we are interested in. Column (1) of Table IV shows the estimated post-ChatGPT effect of being in the top exposure decile on the generative AI share of job postings. The estimate shows that the generative AI share increases by 0.2pp on average more in the top decile than in the matched control group with lower exposure.

Generative AI effects on remote work shares. How much does the remote job share change for high generative AI exposure firms after the ChatGPT release? Columns (2) and (3) of Table IV show that the estimated effect on remote work shares is a reduction by 3.4 pp. This is a substantial effect: the treated group has a level of remote work of 18.2% in the pre-period (7.3% for the control group), so the post-release impact of being in the high exposure group corresponds to a 19% reduction on average in remote work prevalence in job postings.

To verify that the synthetic difference-in-differences matching procedure actually results in a suitable control group, Figure 10, Panel (A) shows details on the raw time variation in remote work across the treated and control groups. As the graph shows, the SDID weights select control units that have a very similar run-up in remote work before the release of ChatGPT, but they differ in their behavior in the post-period that is untargeted by the synthetic control weights. Note that any differences in remote work levels between treatment and control groups are eliminated by the firm fixed effects included in the estimation. Panel (B) shows the standard event study plot commonly used in difference-in-differences designs, which illustrates the treatment effect for each period. As the graph shows, there are minimal differences between the treated and control groups in the pre-period, but the estimated gap becomes large and consistently negative starting in the ChatGPT release quarter.

Generative AI effects on the level of job postings. Is the estimated effect on the composition of hiring accompanied by a decline in the *level* of hiring by firms? Columns (3) and (4) of Table IV show the relative decline in total job postings per firm in the treated group after the release of ChatGPT, for both remote and total job postings. The estimates show that overall hiring declines for both remote and all jobs, but more than proportionally so for remote jobs. This suggests that generative AI adoption does not merely lead to a reallocation of positions from remote work to non-remote work, but also involve a substitution of the automation tools for human labor.

In order to understand *why* remote work and generative AI adoption impact one another in the observed ways, I provide a conceptual framework below that proposes a mechanism that can explain these results.

IV. Conceptual framework: organizational technology ladders

Why would having many workers in remote positions *cause* a firm to adopt generative AI at a faster rate? And why would generative AI adoption reduce remote hiring? In this

section, I provide a conceptual framework for how remote work and generative AI adoption can be linked that is based on a firm’s decision to invest in different technologies, which can then have effects on the ability, and cost, of investing in another technology. That is, from an organizational management point of view, one technology might represent a “ladder” to adopting another technology, which introduces path dependence into firm investment choices, and interdependence between the technologies in their impact on firms, workers, and the broader economy.

Model overview. Below, I develop a parsimonious partial equilibrium model of the firm choice of investment in technology and its consequences for organizational structure. The model builds on the task-based approach to modeling automation (Acemoglu and Restrepo, 2022; Thompson and Autor, 2024) but explicitly incorporates firm tech-savviness, investments in new technologies, the decision-making effort of the worker and variable decision-making sensitivity of tasks. The firm divides workers’ total time between routine production and decision-making, and adjusts this allocation in response to the availability of support from managers at the firm and the degree to which an occupation benefits from better decision-making. This framework captures all of the mechanisms discussed in Section II.A. To adopt new technologies, like remote work or generative AI, the firm needs to invest in sufficient information and communication technology (ICT) capital.

In this setting, I model remote work adoption as a trade-off between greater effective work hours for the worker and less effective support from managers in making effective decisions. In contrast, Generative AI adoption expands the frontier of tasks that are automated.

Based on the two assumptions that (i) remote work reduces the quality of centralized support for worker decision-making; and (ii) the tasks that are still done by the worker after automating some tasks with generative AI tend to be more sensitive to the quality of decision-making, this framework can explain both the increase in generative AI adoption for firms that go remote, and the reduction in remote hiring for firms that adopt generative AI.

Importantly, this framework results in testable predictions for how different organizational characteristics affect the magnitude of these effects, and how firms adapt to the adoption of remote work. These predictions allow me to explain the heterogeneity across firms in the relationship between remote work and generative AI, and suggest empirical tests to validate the mechanism captured in the model.

A. Occupational output

For simplicity, assume that each firm consists only of one production occupation and one production worker, with potentially some non-producing “managers”—although the results would be the same if a firm has multiple workers in the same occupation. A worker at firm f in occupation j , which, in this simplified framework also indexes the type of firm that employs the worker, produces output by executing the non-automated subset of the tasks $x \in [0, j]$ that are sorted by automatability. Output y_{fj} given by the following expression:

$$\ln y_{fj} = \underbrace{k_{fj}}_{\text{Automated production}} + \underbrace{\int_{k_{fj}}^j \ln\left(\frac{1-D_f}{j-k_{fj}}\right) dx}_{\text{Non-decision tasks}} + \underbrace{\int_{k_{fj}}^j \delta(x) Q_f dx}_{\text{Decision-making}} + \underbrace{\int_{k_{fj}}^j r_j \times \mathbb{1}[\text{Remote}_f] dx}_{\text{Remote time savings}} \quad (4)$$

Here, the first term is the automated product $k_{fj} < j$, which represents how much of the task spectrum $[0, j]$ is automated (or does not require labor) at firm f . In the second term, the integral $\int_{k_j}^j$ can be seen as a stylized representation of the “remaining tasks” from k_j to j that must be performed by labor of firm f . D_f is the time spent on making independent decisions by workers at firm f rather than producing output, so the second term represents the output from the share of tasks that are still done by the worker with the time remaining after accounting for time spent on decision-making tasks.

The third term can be thought of as either a productivity “boost” from good decision-making or as the output from a separate planning task that accompanies each part of the worker’s job, and which varies in importance across occupations. Its output depends of two terms: The term $\delta(x)$ captures the importance of decision-making for tasks of different complexity that the worker does. Importantly, I assume that $\delta'(x) > 0$, i.e. when we sort tasks by automatability, less automatable tasks tend to benefit more from good decision-making by the worker. For tractability, I will assume $\delta(x) = e^x$ below, but the results only require that δ increases in x . The second term is Q_f , which is the worker’s decision-making quality. Decision quality is defined as

$$Q_f = \ln \left[\underbrace{\eta_f D_f}_{\text{Worker input}} + \underbrace{\rho_f \mathbb{1}[\text{Remote}] M_f}_{\text{Firm input}} \right],$$

where η_f scales the importance of independent decision-making relative to firm-level coordination. M_f is the normal decision support from the firm (e.g. from better central management), which I assume to increase with the decision intensity of the job, and $\rho_f < 1$ represents a barrier to communication and coordination that reduces the effectiveness of

this decision support. Here, I explicitly make this barrier a function of remote work, but it could also be used to model the coordination effects of a firm’s technology choices more generally, e.g. when thinking about offshoring. Moreover, ρ_f can vary across firms, as some firms having better leadership or communication skills that can bridge the barrier and therefore have a higher ρ_f (a smaller discount factor on decision support when remote).

The last term in equation 4 represents the productivity benefit of working remotely, which may differ by occupation. This allows for some occupations where working from home is easily accomplished and productive, while in other occupations it may either be physically infeasible (e.g. for an emergency room nurse) or undesirable for other reasons. It is a function of r_j , the remote suitability of the occupation, and an indicator for whether the firm is remote.

B. Technology adoption

We will work backwards to derive the firm’s optimal choices for technology investment. After a firm chooses whether to adopt a technology, it determines the efficient amount of time for its workers to spend on decision-making. There is a trade-off, as time spent on decisions D_f is time spent not producing task output, so spending more time on decisions will only be optimal for firms that have tasks are sensitive to good decision-making, i.e. where $\delta(x)$ is high on average.

Remote work and decision quality. Firms that work remotely invest more worker time to make good decisions in order to undo the reduction in central decision support due to communication issues. What does this mean for the overall decision quality in remote firms, all else equal? In Appendix C, I show that an decrease in central decision support is not going to be fully offset by an increase in local decision intensity. This implies that *remote firms have lower overall decision quality even though they invest more worker effort in local decision-making*. That is, $D_f^R > D_f^{NR}$ and $Q_f^R < Q_f^{NR}$, where R and NR denote optimal choices conditional on going remote or staying non-remote.

Taking into account the optimal allocation of worker time to decision-making that results, the firm decides in an earlier period t whether to adopt a new technology τ . Technologies are assumed to arrive at random and new technologies are not anticipated by the firm. The firm’s technology status—the last technology it adopted—when making its period t adoption decision is τ_f .

Technology implementation costs. Each new technology τ that the firm adopts—where, for concreteness, I will focus on $\tau \in \{\text{Remote}, \text{GenAI}\}$ —requires updating the firm’s

level of information technology to $I(\tau)$, which has cost

$$c_f(\tau) = \frac{1}{T_f} (I(\tau) - I(\tau_f))^2,$$

where T_f is a measure of the firm's innovation orientation or tech-savviness, which reduces its investment costs. That is, new investment is costlier if the firm did not adopt previous technologies, such that there is a bigger technology gap $I(\tau) - I(\tau_f)$ between the firm's level of technology infrastructure and the level required for the new technology. I assume that later technologies always require a higher level of information technology, and that remote work is the last technology wave that precedes generative AI.

Optimal technology adoption. Each firm adopts a new technology if the benefits of doing so outweigh the benefits *or* if the firm is forced to do so by policy, for example, having to adopt remote work as a result of pandemic lockdowns. Generally, the firm adopts new technologies if the increase in output is greater than the cost of adopting the technology, i.e.

$$\ln y_{fj}(\tau) - \ln y_{fj}(\tau_f) > c_f^\tau.$$

Note that a firm of a given type j will be more likely to adopt a given technology if it is more tech-savvy (higher T_f), or if it recently invested in other advanced technology, i.e. when the technology gap $I(\tau) - I(\tau_f)$ is small.

Different technologies will change different aspects of the production technology, potentially affecting the automation of tasks, the benefits from remote work, the cost of communicating with managers, etc. While the model so far has been kept deliberately general, in this paper I will focus on a particular sequence of technology waves that reflects the large changes experienced during the 2020-2024 period: First, a change in the attractiveness of remote work technology as a result of the Covid pandemic. Second, a change in the potential for automation of particular tasks as a result of the proliferation of generative AI based tools.

C. Technology Wave #1: Remote work

While there were certainly firms that allowed for remote work before the Covid pandemic in 2020, the sudden need to continue work when forced to work-from-home during lockdowns led to both a large increase in the adoption of remote work in “teleworkable” occupations, as well as large investments, and rapid improvements, in the technology tools enabling work-from-home, e.g. virtual meeting and collaboration software.

Here, remote work is assumed to provide an additional productivity boost r_j to all the

tasks done in an occupation, as shown in equation 4. This is motivated by empirical observations that remote workers save time on grooming and commuting, and have more flexibility in their time use over the day, and that some of these time savings are used to work more (Barrero, Bloom, and Davis, 2020; Pabilonia and Vernon, 2022). This productivity boost can vary across occupations, such that there might be “teleworkable” jobs with a high r_j and other jobs where r_j is small or even negative (Dingel and Neiman, 2020). At the same time, remote work reduces the quality of decision-making support from management through triggering the communication barrier $\rho_f^{\mathbb{1}[\text{Remote}]} < 1$. This assumption is empirically supported by the findings in Choudhury, Espinosa, Makridis, Schirmann, et al. (2025) that workers coordinate less with managers when they are not co-located.

Based on the model above, a firm will adopt remote work (indexed by R) rather than stay non-remote (NR) if $\ln y_{fj}^R - \ln y_{fj}^{NR} > c_f^R$, which can be written as

$$\underbrace{(j - k_{fj})r_j}_{\text{Productivity effect}} + \underbrace{(j - k_{fj})\ln\left(\frac{1 - D_f^R}{1 - D_f^{NR}}\right)}_{\text{Output loss from } \Delta \text{ Decision-making effort}} + \underbrace{(e^j - e^{k_{fj}})\ln\left(\frac{Q_f^R}{Q_f^{NR}}\right)}_{\text{Output loss from } \Delta \text{ Decision-making quality}} > c_f^R$$

where R and NR again indicate optimal decisions conditional on going remote or not going remote. Note that we have previously derived that $D_f^R > D_f^{NR}$ and $Q_f^R < Q_f^{NR}$, such that the last two terms on the left are negative—remote work requires more time spent on decisions by the remote worker, and also decreases overall decision quality, which both represent a cost to be weighed against any potential productivity benefits.

Empirical predictions. Note that this model above predicts particular patterns of remote work adoption among firms:

- P1. More teleworkable occupations and firms (higher remote work benefit r_j) are more likely to adopt remote work.
- P2. More tech-savvy or innovative firms (higher T_f) are more likely to adopt remote work for a given occupation.
- P3. Firms with a higher level of existing technology investments (a smaller technology gap $I(\tau) - I(\tau_f)$) are more likely to adopt remote work
- P4. Higher adoption of remote work leads to an increase in demand for individual decision-making skills—“upskilling”—in the firm’s workers

Some of these predictions are already supported by existing findings in the remote work literature, e.g. the teleworkability measure by Dingel and Neiman (2020) has been shown

to predict at least part of the variation in actual remote work adoption (Hansen, Lambert, Bloom, Davis, Sadun, and Taska, 2023), which confirms P1. In the empirical sections of this paper, I provide novel evidence for predictions P2, P3, and P4.

D. Technology Wave #2: Generative AI and the technology ladder effect

I model the effect of generative AI tool availability as pushing the automation boundary k_j upward. That is, if tasks are sorted by j , then the frontier of automatable tasks expands to

$$k_{fj}^G = k_{fj} + \underbrace{\text{GenAIPotential}_j \times \mathbb{1}[\text{GenAIAdoption}_f]}_{\Delta k_{fj} \text{ with Gen. AI}},$$

where G indexes values conditional on the adoption of generative AI. Here, $\text{GenAIPotential}_j \geq 0$ measures how amenable occupation j is to the new technology, i.e. what share of its tasks are exposed to generative AI automation. Whether this exposure leads to automation at the firm depends on $\text{GenAIAdoption}_f \in \{0, 1\}$ which indicates whether a firm invests in generative AI adoption. If a firm invests in the necessary information technology capital $I(\tau = G)$, it moves more tasks into the automatable set. Hence, the fraction of tasks requiring actual worker effort shrinks from $[k_{fj}, j]$ to $[k_{fj}^G, j]$. Note that when k_j increases, more tasks become automated, so fewer tasks remain for labor—but because $\delta'(x) > 0$ the tasks that remain benefit more from higher-level decision-making.

The firm adopts generative AI if $\ln y_{fj}^G - \ln y_{fj}^{\text{NG}} > c_f^G$. Note that the cost c_f^G is lower for firms with high existing technology capabilities, so this directly implies another prediction which I can confirm in the data:

P5. *Firms with higher levels of technology skills are more likely to adopt generative AI.*

Remote work effect on generative AI adoption. How does the adoption of remote work affect the decision of whether to adopt generative AI technology?

The difference in the net benefits from generative AI adoption between remote and non-remote firms is

$$\left[\ln y_{fj}^G - \ln y_{fj}^{\text{NG}} - c_f^G \right]_R - \left[\ln y_{fj}^G - \ln y_{fj}^{\text{NG}} - c_f^G \right]_{NR} \quad (5)$$

$$= \underbrace{-(k^G - k_{fj})r_j}_{\text{Loss of remote productivity (-)}} - \underbrace{(k^G - k_{fj}) \ln \left(\frac{1 - D_f^R}{1 - D_f^{NR}} \right) - (e^{k_{fj}^G} - e^{k_{fj}}) \ln \left(\frac{Q_f^R}{Q_f^{NR}} \right)}_{\Psi_f: \text{Relative human productivity effect (+)}} + \underbrace{(c_f^G|_{NR} - c_f^G|_R)}_{\text{Cost advantage (+)}}. \quad (6)$$

As $(k^G - k_{fj})r_j > 0$, the first term implies that remote firms are disincentivized from adopting generative AI if they perceive large benefits from work-from-home, which have been assumed to scale with the share of tasks done by humans. The second and third terms—the relative size of the human output loss Ψ_f due to generative AI adoption—are positive.⁷ That is, as remote workers produce output with lower decision quality, there is less of a loss if generative AI produces the output in their stead. The last term reflects the fact that the cost of investing in generative AI c_f^{Gf} will be lower (and thus the adoption of generative AI more beneficial) if the firm previously adopted remote work, as I assume that remote work is the most recent technology wave before generative AI and therefore firms that adopt remote work have a higher level of information technology capital than those which do not. Intuitively, this equation captures the trade-off that the work quality of remote workers is lower, so replacing it with generative AI is more profitable, and remote firms also already have some of the necessary infrastructure in place—but this has to be weighed against remote workers' longer effective work hours.

I make the assumption that the work-from-home productivity boost for the automated tasks is not too large relative to the technology investment cost savings for generative AI adoption when already having implemented remote work and the output quality differences among remote and non-remote workers, i.e. that

$$c_f^G|_{NR} - c_f^G|_R + \Psi_f > (k^G - k_{fj})r_j. \quad (7)$$

If this holds, then we get **the technology ladder effect**:

P6. Firms that adopt remote work have a higher incentive to also adopt generative AI.

Note that it is an empirical question whether the underlying assumption is realistic. Testing whether this effect exists is one of the key contributions in the empirical section below. If there is heterogeneity across firms, equation 6 also implies additional predictions with regard to when equation 7 is more likely to hold. The difference in output quality between non-remote and remote workers is larger when ρ_f is smaller, i.e. centralized decision support declines more with remote work. This means that firm organizational characteristics that can overcome these communication barriers (increase ρ_f) lower the quality loss when going remote, which then reduces the relative benefits of generative AI adoption between remote and non-remote firms.

If we assume that better management and greater communication skills allow firms to better overcome remote communication barriers, i.e. they have a larger remote effectiveness

⁷Here, we can use the envelope theorem to ignore reoptimizations of the decision intensity in the remote and non-remote state when considering generative AI adoption.

ρ_f , the model predicts the following heterogeneity for the technology ladder effect:

P7. *The effect of remote work on generative AI adoption is less positive for firms with greater communication skills and managerial or leadership skills.*

Similarly, the degree to which remote work incurs an output penalty due to a change in centralized decision-making support also depends on the initial relative importance of individual decision-making: if D_f is large relative to M_f , then changes in the effectiveness of managerial support will only lead to a small relative decline in decision quality, and a smaller benefit of automating tasks in remote jobs:

P8. *The effect of remote work on generative AI adoption is smaller for firms with greater individual decision-making skills.*

Generative AI effects on remote work. Conversely, does generative AI adoption impact a firm's incentive to continue using remote workers? The change in output from allowing remote work can be written as

$$\ln y_{fj}^R - \ln y_{fj}^{NR} = (j - k_{fj}) \left(r_j + \ln \left(\frac{1 - D_f^R}{1 - D_f^{NR}} \right) + \frac{e^j - e^{k_{fj}}}{j - k_{fj}} \ln \left(\frac{Q_f^R}{Q_f^{NR}} \right) \right), \quad (8)$$

where it can be shown that

$$\frac{\partial}{\partial k_{fj}} \left(\frac{e^j - e^{k_{fj}}}{j - k_{fj}} \right) > 0,$$

or, intuitively, that the average decision-making sensitivity of tasks increases with automation, as lower-complexity tasks are more likely to be automated. Because $\ln \left(\frac{Q_f^R}{Q_f^{NR}} \right) < 0$, this means that equation 8 is decreasing in k_{if} for firms that initially prefer remote work, i.e. the overall benefit from remote work declines with greater automation. This means that the model predicts a **technology substitution effect**:

P9. *Higher adoption of generative AI leads to a decline in the use of remote work technology, if it was previously adopted.*⁸

Note that the size of this effect of generative AI adoption on remote work is scaled by the decision quality loss due to remote work adoption. This means that the same factors that reduce the impact of remote work on decision quality will also reduce the technology substitution effect:

⁸Note that I am assuming that for the purpose of the decision of whether to continue letting workers work remotely, the technical infrastructure is not relevant, as there is no way to “recoup” the original outlay for information technology capital to enable work-from-home when later eliminating remote work.

P10. *Adoption of generative AI leads to a smaller decline in the use of remote work for firms with greater communication, managerial or leadership skills.*

P11. *Adoption of generative AI leads to a smaller decline in the use of remote work for firms with greater individual decision-making skills.*

In the next section, I will further explore the empirical evidence to see whether it aligns with the predictions from this conceptual framework that go beyond the technology ladder effect and technology substitution effect that are supported by the results in Sections III.C and III.D.

V. Mechanism: organizational adaptation

The empirical results in in Sections III.C and III.D showed that greater remote work adoption leads companies to embrace generative AI, while the adoption of generative AI is followed by a decline in remote hiring. The conceptual framework proposed in Section IV suggests that these patterns can be explained with a model where workers’ quality of decision-making declines when they work remotely, and the use of generative AI tools tends to increase the importance of strategic decision-making on the part of the human worker. This framework also generates additional predictions about how firms adapt their organizations when they adopt remote work, and how these effects should vary across firms. In this section, I provide evidence on how empirical adoption patterns in different organizations and changes in hiring for particular skills align with these predictions.

A. *Heterogeneity of remote work effects on generative AI*

In this section, I estimate whether particular firm characteristics are associated with higher or lower impacts of remote work adoption on generative AI adoption. I estimate interacted regression models of the form

$$100 \times \text{GenAIJobShare}('23-'34)_i = \beta \text{RemoteWorkShare}('21-'22)_i + \alpha_{ind} + \text{Controls}_i \\ + \gamma \mathbb{1}[\text{High}_i] \times \text{RemoteWorkShare}('21-'22)_i + \varepsilon_i$$

where the “High” indicator is computed for different firm characteristics and indicates that the firm is above the average across firms for the measure as of 2022.

The results of estimating these interacted models are shown in Figure 6, where I show both the estimated baseline effect of remote work on generative AI adoption for the “low” group (lower panel), and the estimated interaction coefficient that captures the difference

in remote work effects between the high and the low group for each characteristic (upper panel).

I find patterns that align with the proposed conceptual framework of technology adoption: On the one hand, the “low” group coefficients confirm that the significant positive effect of remote work adoption on generative AI adoption is not driven by any particular sub-group of firms, as the baseline effects remain significant in all of the interactions.

On the other hand, there is systematic variation in the size of the effect with particular firm characteristics: (a) firms with greater advanced degree hiring are more likely to adopt generative AI when workers are remote; (b) firms that normally hire for occupations that are more social skill-, communication- or interaction-intensive, or which explicitly mention communication skills are less responsive in their generative AI adoption to the remote work share; (c) firms that normally hire workers with more decision-making skills are less likely to increase generative AI adoption when their remote work shares are high; (d) firms that more frequently mention specialized tech skills in any role are significantly more responsive; and (e) firms with more leadership skills and independence in the workers they hire are less likely to adopt generative AI after going remote.

We can also estimate whether particular *occupation* characteristics are associated with higher or lower impacts of remote work adoption on generative AI adoption when holding overall firm and occupation effects constant. Again, I define “High” indicators at the national level using O*Net data for all measures that indicate that occupations are above the job posting-weighted average for the measure, or average Lightcast characteristics of the hiring in that occupation-by-firm cell.

The results, shown in Figure 7 show very similar patterns to those found with regard to firm characteristics, which suggests that a similar mechanism makes it more or less desirable to adopt generative AI in response to remote work for particular jobs within firms as across firms.

These findings directly support the predictions of the conceptual framework: occupations and firms that experience a smaller loss in output quality from remote work because they have better communication, better management, or workers that have more individual decision-making capability are less likely to want to replace those remote workers with automation tools. Conversely, firms and occupations that already have a higher level of technology skills find it easier to adopt the new technology.

The next section explores whether firms *change* their organization or capabilities as a result of remote work.

B. Remote Work Impact on Hiring Characteristics

This section tests how the adoption of remote work changed firm hiring patterns with regard to skills. These results help us understand through which mechanism the technology ladder effects on subsequent generative AI adoption might be operating. I estimate specifications of the form

$$\text{Skill}(2022)_i = \alpha + \beta \text{RemoteWorkShare}('21-'22)_i + \gamma \text{Skill}(2019)_i + \text{Controls}_i + \varepsilon_i, \quad (9)$$

where the dependent variable captures different measures of the composition of the firm's hiring with regard to indicators of skill in its job postings, or the average characteristics (derived from O*Net scores) of the occupations that the firm is hiring for. These regressions also control for the past level of the dependent variable pre-pandemic, so the estimated coefficients capture the effect on changes in the skill composition of job postings between the period before and after the lockdown period of the pandemic.

Figure 11 shows the results of the IV estimation (using the competition IV) at the firm level (and Appendix Figure A.3 shows the OLS estimates), which additionally controls for a rich set of pre-pandemic hiring characteristics to capture potential confounders in hiring skill trends.⁹ These results should be considered jointly with the findings at the occupation-by-firm level in Figure 12 for the skill characteristics derived from job postings. In these regressions, I am able to include firm and occupation fixed effects and can therefore distinguish whether the effect is driven by changes in the skills demanded in the occupations that went remote themselves or by changes in the non-remote parts of a firm's organization. However, the set of occupation-by-firm skill outcomes is more limited as not some skill characteristics vary only at the occupation level.

Upskilling. First, I explore whether firms change the formal requirements for workers' experience or education when increasing the remote share of their workers, i.e. whether there is upskilling as a result of the technology adoption. The results shown in rows 1-3 of Figure 11 show that firms that adopt remote work at a higher rate increase their hiring of workers with at least a college degree and with at least 5 years of experience. This shows that firms are *upskilling* their hiring when a larger share of their workers go remote. Row 4 shows that firms also shift their hiring away from occupations that involve repetitive tasks,

⁹The firm-level regressions include the following control variables: NAICS 2-digit fixed effects; company's remote work share in 2019, the company's uninteracted exposure to MSA teleworkability in 2019, uninteracted firm-level teleworkability in 2019, the company's share of jobs requiring a college education and the share requiring an advanced degree in 2019, the share of the company's job postings in 2019 that was for computer occupations or manager positions; the log of total job postings in 2019 and in 2022; the company's labor market exposure to MSA remote shares in 2019.

making the firm’s jobs less routine. These patterns look similar in the occupation-by-firm results in Figure 12, where the first 3 rows show that firms also increase hiring for college educated and experienced workers within the firm more in occupations that have higher remote shares, suggesting that these are characteristics that firms are looking for in the remote workers themselves.

Communication skills. Which particular skills are firms more likely to hire for when they adopt remote work? In rows 5-9 of Figure 11 I explore whether firms are more or less likely to hire for jobs or skills associated with more communication. The results show that the firm does not shift its hiring towards roles that tend to involve more interactiveness or “inflexible” client responsibilities as defined by (Goldin, 2014). I do find evidence that roles requiring greater social skills as defined by Deming (2017) are more likely to appear in a firm’s hiring when remote work is adopted. At the firm level there is not a significant shift towards mentioning more teamwork or communication skills in job postings. However, Figure 12 shows that, within firms, the jobs that go remote to a larger degree are significantly more likely to mention teamwork and communication. This shows that these are skills considered useful for remote workers themselves, rather than changes in the overall firm culture.

Decision-making. In contrast, I find evidence of a significant shift in the composition of firm hiring towards roles requiring individual decision-making: rows 10-12 of Figure 11 show that firms increase the share of job postings that hire for roles that have high decision-making intensity as defined by Deming (2021) and also are more likely to involve decision-making without supervision. While the share of a firm’s jobs asking for decision-making skills on the part of the worker does not increase significantly overall, Figure 12 shows that the more remote occupations within firms significantly increase their likelihood of mentioning decision-making skills, in line with the model’s prediction that workers in remote jobs are less able to rely on someone else making decision for them, which leads firms to optimally increase the effort devoted to individual decision-making by remote workers.

Technology skills. There is even stronger evidence of a shift in the composition in hiring towards computer-related occupations in firms that went remote to a larger degree. Moreover, the share of job postings mentioning data management, data science, machine learning, or deep learning skills increases significantly in firms where remote work increases, as shown in rows 13-17 of Figure 11. Importantly, Figure 12 shows that this is not because remote roles themselves become more technical: with the exception of data management skills, the more remote occupations within a firm are less likely to be increasing their mentions of technical skills. This suggests that firms are investing in their digital infrastructure in general when they have more remote workers, but the remote workers

themselves only need to have the data skills to operate within digital workflows.

Management skills. Do firms also increase their demand for managerial skills to deal with the increased difficulty of supervising remote workflows? Alekseeva, Azar, Giné, and Samila (2024) show that non-generative AI adoption was associated with greater hiring for managerial roles and demand for managers with more interpersonal and cognitive skills. The last 3 rows of Figure 11 show a similar pattern for remote firms, which hire more for roles that involve greater leadership skills, and also for managerial roles. Figure 12 shows that, within firms, remote work adopting roles were more likely to ask for “independence” when hiring, in line with the assumption in the model that remote work lowers dependence on centralized decision-making.

C. *Heterogeneity in generative AI Adoption by hiring characteristics*

To close the loop on whether the characteristics that are engendered by greater remote work adoption then make it more likely that a firm in fact adopts generative AI, I split my sample based on the characteristics of a firm’s hiring in 2022 to see if particular hiring patterns lead to a stronger adoption of generative AI when a firm (or an occupation within a firm) has generative AI exposure.

I estimate cross-sectional specifications of the form

$$100 \times \text{GenAIJobShare}('23-'34)_i = \beta \text{GenAIExp}('21-'22)_i + \gamma \text{GenAIExp}('21-'22)_i \times \mathbb{1}[\text{High}_i] \\ FE_s + \text{Controls}_i + \varepsilon_i,$$

where the unit of observation i is either a firm, or an occupation-by-firm. Here, $\mathbb{1}[\text{High}_i]$ indicates whether a firm or an occupation-by-firm cell is above average in the characteristic in the sample. Some occupation characteristics are based on time-invariant 2018 O*Net data, which is the case for repetitiveness, social skills, interactiveness, inflexibility, decision-making intensity, decision-making without supervision, computer occupation, manager occupation, and leadership skills. The remaining occupation-by-firm level characteristics are means of hiring characteristics in 2022 Lightcast data. Firm-level and occupation-by-firm level hiring characteristics are computed as job posting-weighted means over job postings in 2022.

The results at the firm level are shown in Figure 13, where the top panel shows estimates of $\hat{\gamma}$ that correspond to the difference in effects between the high and the low group, and the bottom panel shows the baseline effect $\hat{\beta}$ for the low group.

I find that firms that demand higher skills as proxied by education levels in their hiring and hire for occupations with less repetitive tasks are more likely to adopt generative

AI if they have productivity potential from the technology. The same is true of firms that hired for more teamwork skills or for occupations that require social skills before ChatGPT was released. Firms that hire for jobs that normally require problem-solving and decision-making, or un-supervised decision-making, or for which the firm requires decision-making skills, are significantly more likely to adopt generative AI as well. I find the most pronounced difference in adoption between firms that have low and high technical skills such as machine learning or deep learning, with firms that hire more for these skills substantially more likely to adopt generative AI. Similarly, firms that hired more for managerial positions in 2022 were more likely to adopt generative AI.

Note that the types of positions that are associated with greater generative AI adoption are exactly the types of hiring that become more prevalent at the firm level with remote work adoption as shown in Figure 11. Together with the evidence in Figure 13, this provides one potential mechanism for the technology ladder effect, which is that firms change their organizational structure and skill mix in response to remote work adoption, and this skill mix then proves useful for generative AI adoption. I discuss this connection in more detail in Section VI.

The corresponding results for within-firm differences in generative AI adoption in response to generative AI exposure at the occupation-by-firm level are shown in Figure 14. The upper panel of the figure shows that, controlling for firm and occupation fixed effects, I find very similar patterns with regard to *which* occupations within the firm are more likely to adopt in response to exposure to generative AI: the only differences to the across-firm associations are that, within firms, the jobs requiring greater social skills or leadership skills are less likely to respond disproportionately to exposure, which aligns with the idea that these are positions that enable functioning communication and governance of technology-adopting firms, but in support of other roles that are more likely to be the ones directly working with the technology. Interestingly, the within-firm results also show that computer occupations are not more responsive to generative AI exposure, which suggests that the outsized adoption associated with technical skills is driven by other technical positions within the firm, e.g. more research-oriented analytical roles.

D. Return-to-office policies and generative AI adoption

The previous results suggest that firms that have existing organizational characteristics that make remote work less disruptive, are less likely to adopt generative AI. I explore this hypothesis further by considering firms that have publicly stated that remote work is not working out for them. I use data on firms that have stated return-to-office (RTO)

policies, collected from the crowd-sourced Flex Index website. Firms that have a publicly known return-to-office mandate are likely to consider remote work to be less productive. To test this idea, I create an indicator for whether a firm appears in the Flex Index data as requiring a minimum presence in the office. All other firms (whether they appear in Flex Index or not) are labeled as not having an RTO mandate.¹⁰

Then, I estimate IV regression specifications of the form

$$100 \times \text{GenAIJobShare}('23-'34)_i = \beta \text{RemoteWorkShare}('21-'22)_i + FEs + \text{Controls}_i \\ + \gamma \mathbb{1}[\text{RTOMandate}_i] \times \text{RemoteWorkShare}('21-'22)_i + \varepsilon_i$$

both at the firm and at the occupation-by-firm level. For the purposes of the estimation, I treat RTO mandate status as exogenously assigned, conditional on the control variables, which are the same as in the baseline effect estimations in Tables II and Table III. The results are shown in Table VI: I find that firms that have an RTO mandate are twice as likely to adopt generative AI for each percentage point of remote hiring during the pandemic (column (1)). Moreover, in line with the idea that this remote work aversion leads to a firmwide investment in generative AI, rather than a targeted adoption only in remote jobs, column (2) shows that there is no evidence that the within-firm effect of remote work on generative AI is different in RTO firms.

E. Remote work effect heterogeneity by hiring characteristics.

Which firm characteristics make it more likely that a firm responds to generative AI adoption with a reduction in remote hiring? I again split the sample based on the characteristics of a firm's hiring in 2022 and repeat the synthetic diff-in-diff estimation for each subgroup. The results are shown in Figure 15: Comparing the above-average firms in the upper panel to the below-average firms in the lower panel, some clear patterns emerge: high generative AI-exposure firms hiring for *fewer* social skills or communication-intensive roles (which includes managers) and that hire for *more* technical skills are more likely to reduce their remote work share after the release of ChatGPT.

The evidence with regard to educational background is more mixed: exposed firms that are hiring more for experienced workers, but also those hiring more for repetitive tasks or for fewer college workers see a larger reduction in remote hiring after the ChatGPT release.

The heterogeneity with regard to decision-making skills aligns closely with the mechanism captured in the model in Section IV: firms that have more jobs that require decision-

¹⁰See Appendix B for details on the Flex Index data.

making—for which remote communication barriers are presumably more problematic—are more likely to reduce remote hiring if they have generative AI exposure. However, firms that hire for more decision-making skills, i.e. where workers can be expected to better cope with going remote, are less likely to see generative AI exposure lead to a reduction in remote work.

F. Generative AI effects on the characteristics of remote work

What happens to the skill composition of hiring at firms with high exposure to generative AI after ChatGPT is released? I repeat the same estimation as in equation 3, but with the share of *remote jobs* or *non-remote jobs* that have particular characteristics as the dependent variable. This analysis captures whether the decline in remote hiring shown in Table IV is concentrated in particular types of remote work and whether the exposure to generative AI further changes the skill mix of a firm’s overall hiring. The non-remote job skill composition results are shown in panel A of Table V and the remote work skill composition in panel B.

Columns (1)-(3) of the table show that firms with the highest exposure to generative AI post-ChatGPT shift demand away from workers with mid-level experience (who represent the largest group of job postings), leaving entry-level demand unchanged, and significantly increasing hiring of more experienced workers. At the same time, columns (5) and (6) show that the most exposed firms saw no substantial change in the demand for college or advanced degrees among remote workers.

Interestingly, there is strong evidence in column (8) of Table V that generative AI adoption increases demand for decision-making skills among all workers. This provides evidence in favor of the assumption in the model that generative AI adoption is associated with a shift towards more decision-intensive tasks for human workers.

There is also evidence for an increase in communication skills for all workers in column (9), while demand for data management skills among remote workers is unaffected, but increases among non-remote workers, as column (7) shows.

VI. Discussion

The results above tell the following story of how the technology ladder between remote work adoption and generative AI adoption operates: First, many firms adopted remote work in response to greater availability of related tools and a physical need to adopt them during the Covid pandemic, which then changed optimal work processes and firms’ skill mix. I find that the change in work processes that resulted from this first technology shock was

associated with hiring more experienced workers, consistent with the finding by Emanuel and Harrington (2024) that the productivity effects of remote work are more negative for less experienced workers.

I also find evidence that remote work caused firms to hire for skills that would allow them to better adapt to the new work style: increasing data management skills and individual decision-making skills, and also hiring additional managers. This is consistent with the increased demands on managers as a result of a reduced ability to train and supervise remote workers.

Demand for technical skills related to information processing, such as machine learning, also increases substantially in more remote firms. However, this is *not* because of greater demand for these skills in remote jobs themselves, but rather as part of a firm's non-remote hiring. This is in line with other studies that have found that firms increased the share of investments going to IT equipment and computers during the pandemic (Barrero, Bloom, and Davis, 2021), and suggests that they are investing in broader digital capabilities and technical infrastructure to support remote workers.

Additionally, I find some evidence that firms increased their hiring for social skill-intensive roles and for teamwork skills both at more remote firms, and in remote jobs in particular. Other research has found that remote work increases the time and effort devoted to communication (Gibbs et al., 2023) and leads to a decrease in the amount of information sharing across collaboration networks (Yang, Holtz, Jaffe, Suri, Sinha, Weston, Joyce, Shah, Sherman, Hecht, et al., 2022), which firms might seek to counter-act by investing in hiring for management and team coordination roles that can overcome these issues.

After ChatGPT was released, remote work then facilitated the adoption of generative AI. While the causal effect on generative AI adoption is robust across different sub-samples, the heterogeneity tests reveal that technology skills in particular differentiate which firms implemented generative AI at a faster pace in response to initially going remote. Together with the evidence on changes in hiring as a result of remote work, this suggests a direct mechanism through which organizational adaptation increases the benefits from the subsequent generative AI technology wave: technology skills enable remote work, but firms also invest more in technology skills in response. Generative AI adoption then builds on the technical infrastructure and skill investments made for remote work.

While remote work adoption enables the use of generative AI tools, the use of the latter also displaces hiring for the former. That is, the event study design around the release of ChatGPT shows that firms that are more exposed to generative AI reduce their hiring for remote workers. But the size of this effect depends on an organization's ability to accommodate remote work: the same characteristics—such as hiring more for communication,

individual decision-making, or managerial skills—that lead to a lower adoption of generative AI in response to remote shares also make it less likely that generative AI exposure leads to a reduction in remote hiring, in line with the quote at the beginning of the paper that interacting with colleagues and customers makes employees harder to automate.

This suggests that firms that incur a smaller quality penalty from going remote are less likely to demand generative AI automation, in line with the latter representing a substitute for “bad” remote work. In contrast, firms which more actively mitigate the management and coordination issues arising from remote work are less eager to adopt generative AI. A salient example of this dynamic are firms which have issued return-to-office mandates to overcome the perceived disadvantages of remote work for their organizations: these firms see a much larger adoption of generative AI tools relative to their remote hiring than other firms.

Finally, organizations also change as a result of generative AI exposure: exposed firms upskill further and hire more for decision-making skills, as generative AI tools complement the human ability to structure and evaluate analytical work output. They also invest more in technology skills, which are associated with a greater tendency for the firm to substitute generative AI for remote workers.

The dynamic of this “organizational technology ladder” raises important issues for policymakers and businesses responding to technological changes. It suggests that differences in the ability to innovate and to transform an organization in response to an initial technology shock can compound into broader competitive advantages over time if the ability to benefit from later technological waves depends on how eagerly an organization embraced the former. This path dependency means that an initial heterogeneity in technology exposure can cascade into some firms and worker groups being systematically affected by later changes in a way that would be difficult to anticipate based on studying these technology shocks in isolation. As I argue in this study, remote work and generative AI technologies are intimately connected in their effect on firms and local labor markets, both by coinciding in the same jobs and by mutually affecting firms’ ability and willingness to invest in the respective other technology.

As both technologies continue to evolve and firms continue to adapt, understanding these technological complementarities is crucial for predicting how labor markets and firm productivity will evolve, and for policymakers and researchers to be able to design policies that harness these changes for the benefit of society. In addition, the effects on skill demand and heterogeneity of effects for different roles highlight that particular worker groups—particularly those in interactive or technical roles, or in organizations with better communication and managerial capabilities—might be less negatively impacted by generative AI,

even if their position is remote.

REFERENCES

- Acemoglu, Daron, and Pascual Restrepo, 2022, Tasks, automation, and the rise in us wage inequality, *Econometrica* 90, 1973–2016.
- Akan, Mert, Jose Maria Barrero, Nicholas Bloom, Tom Bowen, Shelby Buckman, Steven J. Davis, Luke Pardue, and Liz Wilkie, 2024, Americans now live farther from their employers, Unpublished manuscript.
- Alekseeva, Liudmila, José Azar, Mireia Giné, and Sampsa Samila, 2024, Ai adoption and the demand for managerial expertise, *FEB Research Report MSI_2412* 1–60.
- Arkhangelsky, Dmitry, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager, 2021, Synthetic difference-in-differences, *American Economic Review* 111, 4088–4118.
- Bao, Lingfeng, Tao Li, Xin Xia, Kaiyu Zhu, Hui Li, and Xiaohu Yang, 2022, How does working from home affect developer productivity?—a case study of baidu during the covid-19 pandemic, *Science China Information Sciences* 65, 142102.
- Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis, 2020, 60 million fewer commuting hours per day: How americans use time saved by working from home, *University of Chicago, Becker Friedman Institute for Economics Working Paper* .
- Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis, 2021, Why working from home will stick, Technical report, National Bureau of Economic Research.
- Barrero, José María, Nicholas Bloom, and Steven J Davis, 2023, The evolution of work from home, *Journal of Economic Perspectives* 37, 23–49.
- Bick, Alexander, Adam Blandin, and David J Deming, 2024, The rapid adoption of generative ai, Technical report, National Bureau of Economic Research.
- Bloom, Nicholas, Luis Garicano, Raffaella Sadun, and John Van Reenen, 2014, The distinct effects of information technology and communication technology on firm organization, *Management Science* 60, 2859–2885.
- Bloom, Nicholas, Ruobing Han, and James Liang, 2024, Hybrid working from home improves retention without damaging performance, *Nature* 1–6.
- Bloom, Nicholas, Tarek Alexander Hassan, Aakash Kalyani, Josh Lerner, and Ahmed Tahoun, 2021, The diffusion of disruptive technologies, Technical report, National Bureau of Economic Research.
- Boeri, Filippo, Riccardo Crescenzi, and Davide Rigo, 2024, Work from home and firm productivity: The role of ict and size, Working paper, The London School of Economics.
- Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt, 2002, Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence, *The quarterly journal of economics* 117, 339–376.

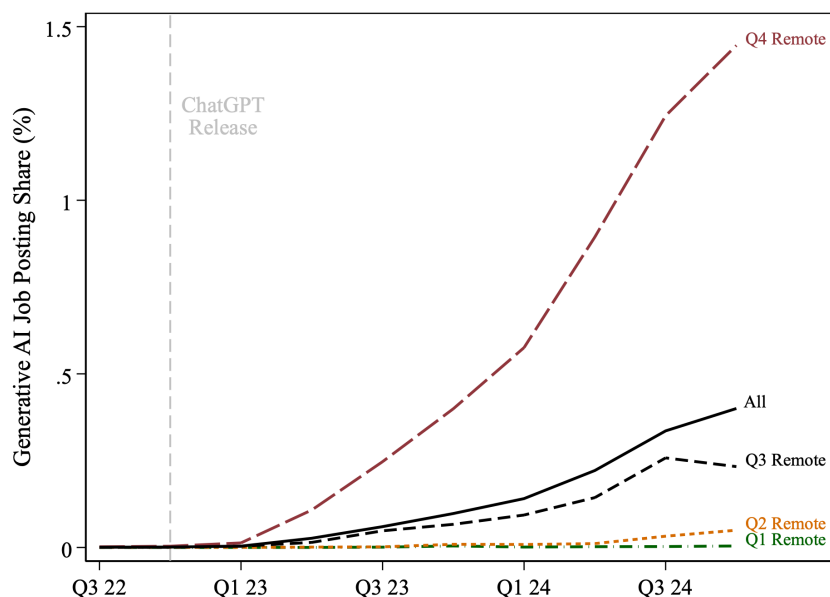
- Brynjolfsson, Erik, and Lorin M Hitt, 2000, Beyond computation: Information technology, organizational transformation and business performance, *Journal of Economic perspectives* 14, 23–48.
- Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond, 2023, Generative ai at work, Technical report, National Bureau of Economic Research.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, 2021, The productivity j-curve: How intangibles complement general purpose technologies, *American Economic Journal: Macroeconomics* 13, 333–372.
- Business Insider, 2023, Fully remote workers are at highest risk of losing their jobs to ai like chatgpt, Accessed: 2025-01-10.
- Choudhury, Prithwiraj, Miguel Espinosa, Christos Makridis, Kyle Schirmann, et al., 2025, Internal communication and remote work, *Christos and Schirmann, Kyle, Internal Communication and Remote Work (January 21, 2025)* .
- Choudhury, Prithwiraj, Cirrus Foroughi, and Barbara Larson, 2021, Work-from-anywhere: The productivity effects of geographic flexibility, *Strategic Management Journal* 42, 655–683.
- Clarke, Damian, Daniel Pailanir, Susan Athey, and Guido Imbens, 2023, On synthetic difference-in-differences and related estimation methods in stata, *Stata Journal* .
- Deming, David J, 2017, The growing importance of social skills in the labor market, *The quarterly journal of economics* 132, 1593–1640.
- Deming, David J, 2021, The growing importance of decision-making on the job, Technical report, National Bureau of Economic Research.
- Dingel, Jonathan I, and Brent Neiman, 2020, How many jobs can be done at home?, *Journal of Public Economics* 189, 104235.
- Eisfeldt, Andrea L, and Gregor Schubert, Forthcoming, Gnerative ai and finance, *The Annual Review of Financial Economics* .
- Eisfeldt, Andrea L, Gregor Schubert, Miao Ben Zhang, and Bledi Taska, 2023, The labor impact of generative ai on firm values, *Available at SSRN 4436627* .
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, 2023, Gpts are gpts: An early look at the labor market impact potential of large language models, *arXiv preprint arXiv:2303.10130* .
- Emanuel, Natalia, and Emma Harrington, 2024, Working remotely? selection, treatment, and the market for remote work, *American Economic Journal: Applied Economics* 16, 528–559.
- Flynn, Sean, Andra C Ghent, and Vasudha Nair, 2024, Determinants and consequences of return to office policies, *Available at SSRN* .
- Frey, Carl Benedikt, and Michael Osborne, 2024, Generative ai and the future of work: a reappraisal, *Brown Journal of World Affairs* 30.
- Garicano, Luis, and Esteban Rossi-Hansberg, 2006, Organization and inequality in a knowledge economy, *The Quarterly journal of economics* 121, 1383–1435.
- Gibbs, Michael, Friederike Mengel, and Christoph Siemroth, 2023, Work from home and productiv-

- ity: Evidence from personnel and analytics data on information technology professionals, *Journal of Political Economy Microeconomics* 1, 7–41.
- Goldin, Claudia, 2014, A grand gender convergence: Its last chapter, *American economic review* 104, 1091–1119.
- Gulati, Piyush, Arianna Marchetti, Phanish Puranam, and Victoria Sevchenko, 2025, Generative ai adoption and higher order skills, *arXiv preprint arXiv:2503.09212* .
- Hansen, Stephen, Peter John Lambert, Nicholas Bloom, Steven J Davis, Raffaella Sadun, and Bledi Taska, 2023, Remote work across jobs, companies, and space, Technical report, National Bureau of Economic Research.
- Humlum, Anders, and Emilie Vestergaard, 2025, The unequal adoption of chatgpt exacerbates existing inequalities among workers, *Proceedings of the National Academy of Sciences* 122, e2414972121.
- Maestas, Nicole, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger, 2023, The value of working conditions in the united states and implications for the structure of wages, *American Economic Review* 113, 2007–2047.
- Manson, Steven, Jonathan Schroeder, David Van Riper, Katherine Knowles, Tracy Kugler, Finn Roberts, and Steven Ruggles, 2024, Ipums national historical geographic information system: Version 19.0 [dataset], <http://doi.org/10.18128/D050.V19.0>.
- Noy, Shakked, and Whitney Zhang, 2023, Experimental evidence on the productivity effects of generative artificial intelligence, *Available at SSRN 4375283* .
- Pabilonia, Sabrina Wulff, and Victoria Vernon, 2022, Telework and time use, *Handbook of Labor, Human Resources and Population Economics* 1–26.
- Powell, David, and Jeffrey B. Wenger, 2024, Workplace amenities and the costs to firms in the united states, Working paper, RAND Corporation.
- Rambachan, Ashesh, and Jonathan Roth, 2023, A more credible approach to parallel trends, *Review of Economic Studies* 90, 2555–2591.
- Schubert, Gregor, Anna Stansbury, and Bledi Taska, 2019, Employer concentration and outside options, Technical report.
- Schubert, Gregor, Anna Stansbury, and Bledi Taska, 2024, Employer concentration and outside options, *Available at SSRN 3599454* .
- Thompson, Neil, and David Autor, 2024, Does automation replace experts or augment expertise? the answer is yes, Schumpeter Lecture, European Economic Association, MIT CSAIL.
- Yang, Longqi, David Holtz, Sonia Jaffe, Siddharth Suri, Shilpi Sinha, Jeffrey Weston, Connor Joyce, Neha Shah, Kevin Sherman, Brent Hecht, et al., 2022, The effects of remote work on collaboration among information workers, *Nature human behaviour* 6, 43–54.


Figures and Tables

Figure 2:
Remote work and occupational generative AI adoption trends.

This figure shows the share of job postings in each quarter that are for jobs that mention generative AI (panel A) or for remote jobs (panel B). In each panel, the occupations are aggregated into job posting-weighted quartiles of adoption of the other technology: panel A shows generative AI adoption by the occupation's quartile of national remote work adoption in 2021-2022 (excl. Q4 2022), and panel B shows remote work share by generative AI demand as of YTD Sep. 2024. The grey drop line indicates the period (Q4 2022) when ChatGPT was released.



(A) Gen. AI share by remote work quartile



(B) Remote work share by Gen. AI adoption quartile

Figure 3:
Remote Work Suitability and Generative AI Exposure by Occupation

This figure shows the relation between Generative AI exposure and remote work suitability. Generative AI exposure and teleworkability are measured at the SOC 2010 6-digit occupation level based on data from Eisfeldt et al. (2023) and Dingel and Neiman (2020), and aggregated across detailed occupations based on employment weights as of 2022 from the Occupational Employment Statistics. The line of best fit in red is estimated with employment weights.




Figure 4:
Firm characteristics and remote work adoption

This figure shows OLS estimates of coefficients for the effect of standardized firm characteristics in 2019 on the firm's remote work prevalence in 2021/2022 (excl. Q4 2022) job postings in a regression of the form

$$\text{RemoteWorkShare}('21-'22)_i = \alpha_{ind} + \beta \text{FirmCharacteristic}(2019)_i + \text{Controls}_i + \varepsilon_i,$$

where the controls in all regressions include the 2019 value of the dependent variable as a control variable, so the coefficients can be interpreted as the effect of a standardized difference in the characteristic on changes in the remote work share. The control variables also include NAICS 2-digit fixed effects, firm-level teleworkability in 2019 and 2021/2022; and the company's share of jobs requiring a college education in 2019. The 95% confidence intervals shown are based on heteroskedasticity-robust standard errors clustered at the NAICS 2-digit sector level.




Figure 5:
Firm-level and occupation-by-firm level IV first stage

This figure shows binscatters of the first-stage relationships underlying the IV results in Table II, column 4, and Table III, column 4. The first panel shows the effect of the interaction between a firm’s exposure to MSA level teleworkability and firm-level teleworkability (both measured in 2019) on the horizontal axis, and the firm’s job posting remote work prevalence in the YTD as of Sep. 2024 on the vertical axis. The second panel shows the effect of the interaction between a firm’s exposure to MSA level teleworkability and the occupation-level national remote work share in 2021-2022 on the horizontal axis, and the occupation-by-firm’s job posting remote work prevalence in the YTD as of Sep. 2024 on the vertical axis. The values on both axes in both graphs are residualized with regard to the same control variables as the IV results in the tables.




Figure 6:
Remote work effects on Gen. AI adoption by firm characteristics

This figure shows coefficients estimated using IV for the effect of remote work prevalence in a firm's 2021/2022 (excl. Q4 2022) job postings on the prevalence of generative AI mentions in a firm's job postings in 2023/2024 (excl. Q4 2024) in a regression of the form:

$$100 \times \text{GenAIJobShare}('23-'24)_i = \alpha_{ind} + \beta \text{RWS}('21-'22)_i + \gamma \text{RWS}('21-'22)_i \times \mathbb{1}[\text{High}] + \text{Controls}_i + \varepsilon_i,$$

where *RWS* is the RemoteWorkShare and the dependent variable has been scaled by 100 for better readability. So, a coefficient of 10 indicates that a 10 pp change in remote work causes a 1pp change in generative AI adoption. $\mathbb{1}[\text{High}]$ indicates whether a firm is above average in the characteristic. The characteristics of jobs are measured in 2022 in Lightcast data, or are time-invariant based on 2018 O*Net data, and are computed from 2022 averages over a firm's hiring composition. The instrument consists of the interaction between firm-level exposure to MSA teleworkability through its hiring labor markets and firm-level teleworkability (all measured in 2019). All regressions also include the following control variables: NAICS 2-digit fixed effects; company's remote work share in 2019, the company's uninteracted exposure to MSA teleworkability in 2019, uninteracted firm-level teleworkability in 2019, the company's share of jobs requiring a college education and the share requiring an advanced degree in 2019; the company's labor market exposure to MSA remote shares in 2019; the company's teleworkability of 2023/2024 hiring, and the company's generative AI exposure in 2023/2024 hiring. The 95% confidence intervals shown are based on heteroskedasticity-robust standard errors clustered at the firm level.

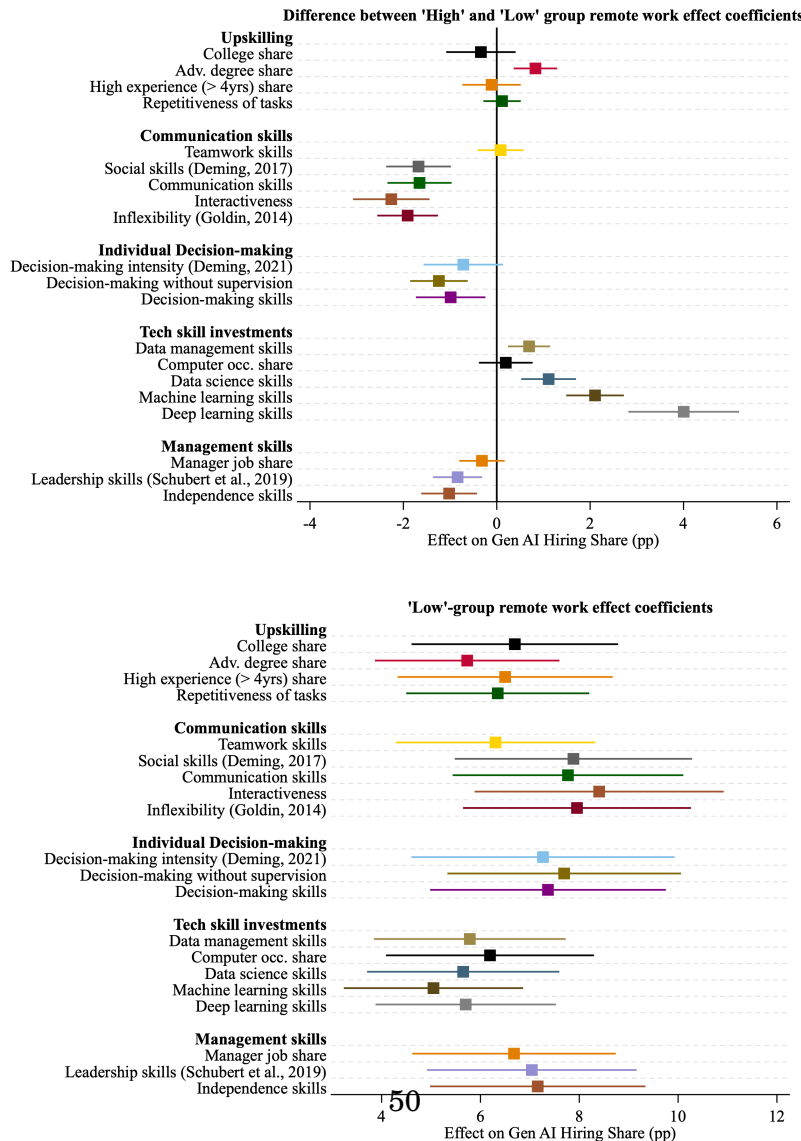


Figure 7:
Remote work effects on Gen. AI adoption by occupation characteristics

This figure shows coefficients estimated using IV for the effect of remote work prevalence in an occupation-by-firm's 2021/2022 (excl. Q4 2022) job postings on the prevalence of generative AI mentions in an occupation-by-firm's job postings in 2023/2024 (excl. Q4 2024) in a regression of the form:

$$100 \times \text{GenAIJobShare}('23-'24)_i = \alpha_{\text{firm}} + \alpha_{\text{occ}} + \beta \text{RWS}('21-'22)_i + \gamma \text{RWS}('21-'22)_i \times \mathbb{1}[\text{High}_i] + \text{Controls}_i + \varepsilon_i,$$

where RWS is the RemoteWorkShare and the dependent variable has been scaled by 100 for better readability. So, a coefficient of 10 indicates that a 10 pp change in remote work causes a 1pp change in generative AI adoption. $\mathbb{1}[\text{High}_i]$ indicates whether a firm is above average in the characteristic. The characteristics are either measured in 2024 in Lightcast data, if they vary at the occupation-by-firm level, or are time-invariant based on 2018 O*Net data if they are occupation-level characteristics, which is the case for: repetitiveness, social skills, interactiveness, inflexibility, decision-making intensity, decision-making without supervision, computer occupation, manager occupation, and leadership skills.

The baseline instrument consists of the interaction between firm-level exposure to MSA teleworkability through its hiring labor markets (measured in 2019) and the national remote work share in the occupation in 2022, and the interacted regressions also include the interaction between the $\mathbb{1}[\text{High}_i]$ indicator and the baseline instrument as an additional instrument. All regressions also include the following control variables: 6-digit occupation fixed effects, firm fixed effects, and the occupation-by-firm's share of jobs requiring a college education and the share requiring an advanced degree in 2023/24. The 95% confidence intervals shown are based on heteroskedasticity-robust standard errors clustered at the firm level.

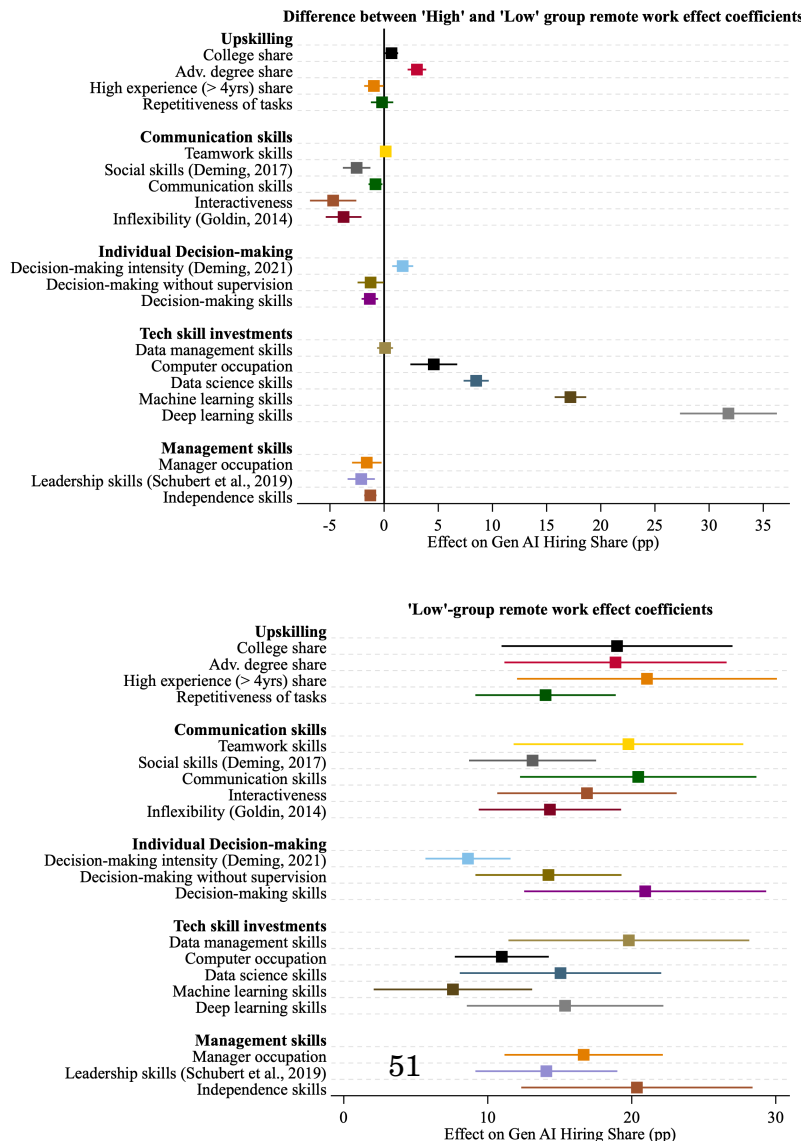


Figure 8:
Remote work trends by firm-level generative AI exposure

This figure shows the share of job postings in each quarter that are for jobs that are for remote jobs. In each panel, the jobs are aggregated into job posting-weighted quartiles of the firms' Eisefeldt et al. (2023) generative AI exposure, based on job postings in 2021-2022 (excl. Q4 2022). The grey drop line indicates the period (Q4 2022) when ChatGPT was released.


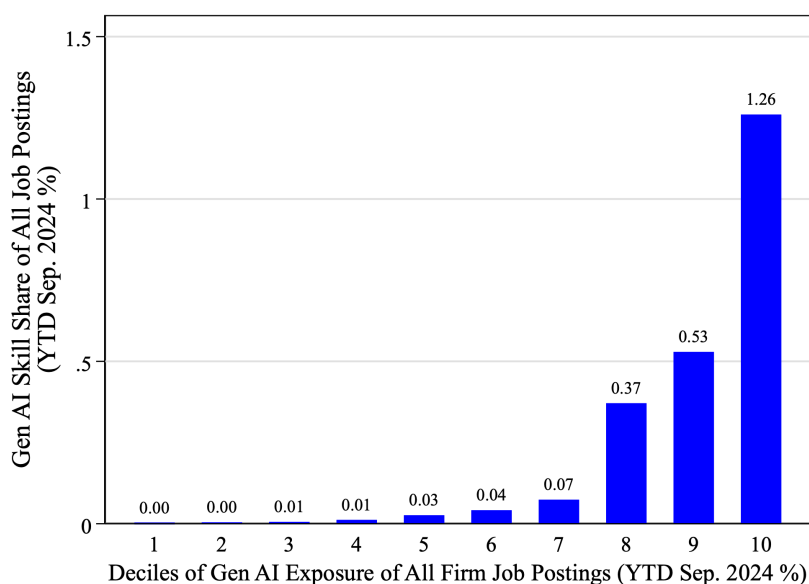



Figure 9:
Generative AI exposure and generative AI adoption by firms

These graphs show the relationship between the share of job postings at the firm level that mention generative AI skills in the YTD Sep. 2024 period, and the Eisfeldt et al. (2023) measure of generative AI exposure of the firm's job postings. In panel A the vertical axis shows the share of generative AI skill mentions in all job postings, and panel B shows the share of generative AI skill mentions only in remote job postings. In each panel, firms are sorted into deciles by exposure, weighted by total job postings and each bar shows the total job posting-weighted mean share of all (panel A) or remote (panel B) posted jobs at firms in that decile of exposure that mention generative AI.



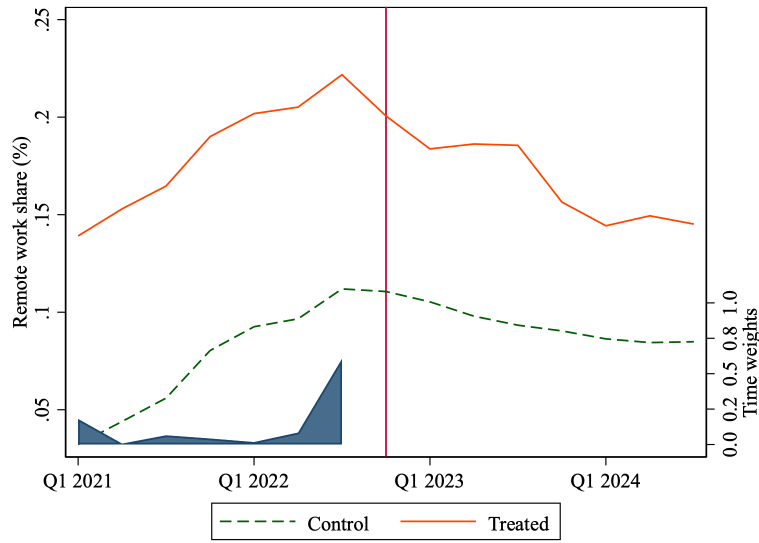
(A) Generative AI exposure and adoption: all jobs




(B) Generative AI exposure and adoption: remote jobs

Figure 10:
Synthetic diff-in-diff event study results

These graphs show the differences in remote work share trends between the treated and control groups that underlie the synthetic diff-in-diff effect estimates reported in Table IV. Panel A shows the remote work trends over time in the synthetic control group and the treatment group. The treatment group is defined as firms that are in the top decile of generative AI exposure based on the composition of their 2021/2022 job postings (excl. Q4 2022). The control group is determined by weighting firms outside the top decile to find the best fit to the pre-period trends in the treatment group, as described in Arkhangelsky et al. (2021). The right axis and the solid area in the graph shows the optimal time weights applied to different pre-periods in the estimation. Panel B shows the estimated difference in remote work shares between the control and treatment group in each quarter (blue dots), as well as the 95% confidence interval based on block bootstrap standard errors computed from 100 resamples, using the procedure detailed in Clarke et al. (2023).



(A) Remote Work Trends: Synthetic Control vs. Treatment Group



(B) Dynamic Treatment Effects of Top Decile Gen. AI Exposure

Figure 11:
Effects of remote work adoption on firm hiring characteristics

This figure shows coefficients estimated using IV for the effect of remote work prevalence in a firm's 2021/2022 (excl. Q4 2022) job postings on the standardized characteristics of the firm's job postings in 2022 in a regression of the form

$$\text{Skill}(2022)_i = \alpha + \beta \text{RemoteWorkShare}('21-'22)_i + \gamma \text{Skill}(2019)_i + \text{Controls}_i + \varepsilon_i,$$

where the controls in all regressions include the standardized 2019 value of the dependent variable as a control variable, so the coefficients can be interpreted as the effect of changes in remote work shares on changes in the composition. The instrument consists of the interaction between firm-level exposure to MSA teleworkability through its hiring labor markets and firm-level teleworkability (all measured in 2019). All regressions also include the following control variables: NAICS 2-digit fixed effects; company's remote work share in 2019, the company's uninteracted exposure to MSA teleworkability in 2019, uninteracted firm-level teleworkability in 2019, the company's share of jobs requiring a college education and the share requiring an advanced degree in 2019, the share of the company's job postings in 2019 that was for computer occupations or manager positions; the log of total job postings in 2019 and in 2022; the company's labor market exposure to MSA remote shares in 2019. The 95% confidence intervals shown are based on heteroskedasticity-robust standard errors clustered at the NAICS 2-digit sector level.




Figure 12:
Effects of remote work adoption on occupation-by-firm hiring characteristics

This figure shows coefficients estimated using IV for the effect of remote work prevalence in an occupation-by-firm's 2021/2022 (excl. Q4 2022) job postings on the standardized characteristics of the occupation-by-firm's job postings in 2022 in a regression of the form

$$\text{Skill}(2022)_i = \alpha + \beta \text{RemoteWorkShare}('21-'22)_i + \gamma \text{Skill}(2019)_i + \text{Controls}_i + \varepsilon_i,$$

where the controls in all regressions include the standardized 2019 value of the dependent variable as a control variable, so the coefficients can be interpreted as the effect of changes in remote work shares on changes in the composition. The instrument consists of the interaction between firm-level exposure to MSA telework-ability through its hiring labor markets and occupation-level national remote work shares in 2021-2022. All regressions also include occupation-level and firm-level fixed effects. The 95% confidence intervals shown are based on heteroskedasticity-robust standard errors clustered at the firm level.




Figure 13:
Across-firm Gen. AI exposure effects on Gen. AI adoption by firm characteristics

This figure shows the the cross-sectional estimates of the differential effect of generative AI exposure in 2021/2022 on generative AI adoption in 2023/24 based on the hiring characteristics of a firm in 2022. Each row in the figures shows coefficient estimates corresponding to a different specification of the form

$$100 \times \text{GenAIJobShare}('23-'24)_i = \alpha_{ind} + \beta \text{GenAIExp}('21-'22)_i + \gamma \text{GenAIExp}('21-'22)_i \times \mathbb{1}[\text{High}_i] + \text{Controls}_i + \varepsilon_i,$$

where i is an occupation-by-firm unit and the dependent variable has been scaled by 100 for better readability. So, a coefficient of 10 indicates that a 10 pp change in generative AI exposure causes a 1pp change in generative AI adoption. $\mathbb{1}[\text{High}_i]$ indicates whether an occupation-by-firm cell is above average in the characteristic in the sample. The characteristics are either measured as means of the hiring characteristics of 2021/2022 Lightcast data, if they vary at the occupation-by-firm level, or are time-invariant based on 2018 O*Net data if they are occupation-level characteristics, which is the case for: repetitiveness, social skills, interactivens, inflexibility, decision-making intensity, decision-making without supervision, computer occupation, manager occupation, and leadership skills. The top panel shows estimates of $\hat{\gamma}$ that correspond to the difference in effects between the high and the low group, and the bottom panel shows the baseline effect $\hat{\beta}$ for the low group. The 95% confidence intervals shown are based on heteroskedasticity-robust standard errors clustered at the firm level.




Figure 14:
Within-firm Gen. AI exposure effects on Gen. AI adoption by occupation characteristics

This figure shows the the cross-sectional estimates of the differential effect of generative AI exposure in 2021/2022 on generative AI adoption in 2023/24 for different subsets of occupation-by-firm cells based on static characteristics of the occupation or the hiring characteristics of the occupation-by-firm in 2022. Each row in the figures shows coefficient estimates corresponding to a different specification of the form

$$100 \times \text{GenAIJobShare}('23-'24)_i = \alpha_{firm} + \beta \text{GenAIExp}('21-'22)_i + \gamma \text{GenAIExp}('21-'22)_i \times \mathbb{1}[\text{High}_i] + \text{Controls}_i + \varepsilon_i,$$

where i is an occupation-by-firm unit and the dependent variable has been scaled by 100 for better readability. So, a coefficient of 10 indicates that a 10 pp change in generative AI exposure causes a 1pp change in generative AI adoption. $\mathbb{1}[\text{High}_i]$ indicates whether an occupation-by-firm cell is above average in the characteristic in the sample. The characteristics are either measured as means of the hiring characteristics of 2021/2022 Lightcast data, if they vary at the occupation-by-firm level, or are time-invariant based on 2018 O*Net data if they are occupation-level characteristics, which is the case for: repetitiveness, social skills, interactivens, inflexibility, decision-making intensity, decision-making without supervision, computer occupation, manager occupation, and leadership skills. The top panel shows estimates of $\hat{\gamma}$ that correspond to the difference in effects between the high and the low group, and the bottom panel shows the baseline effect $\hat{\beta}$ for the low group. All regressions control for firm fixed effects and the share of hiring at the occupation-by-firm level in 2023/24 that requires a college degree or an advanced degree. The 95% confidence intervals shown are based on heteroskedasticity-robust standard errors clustered at the firm level.

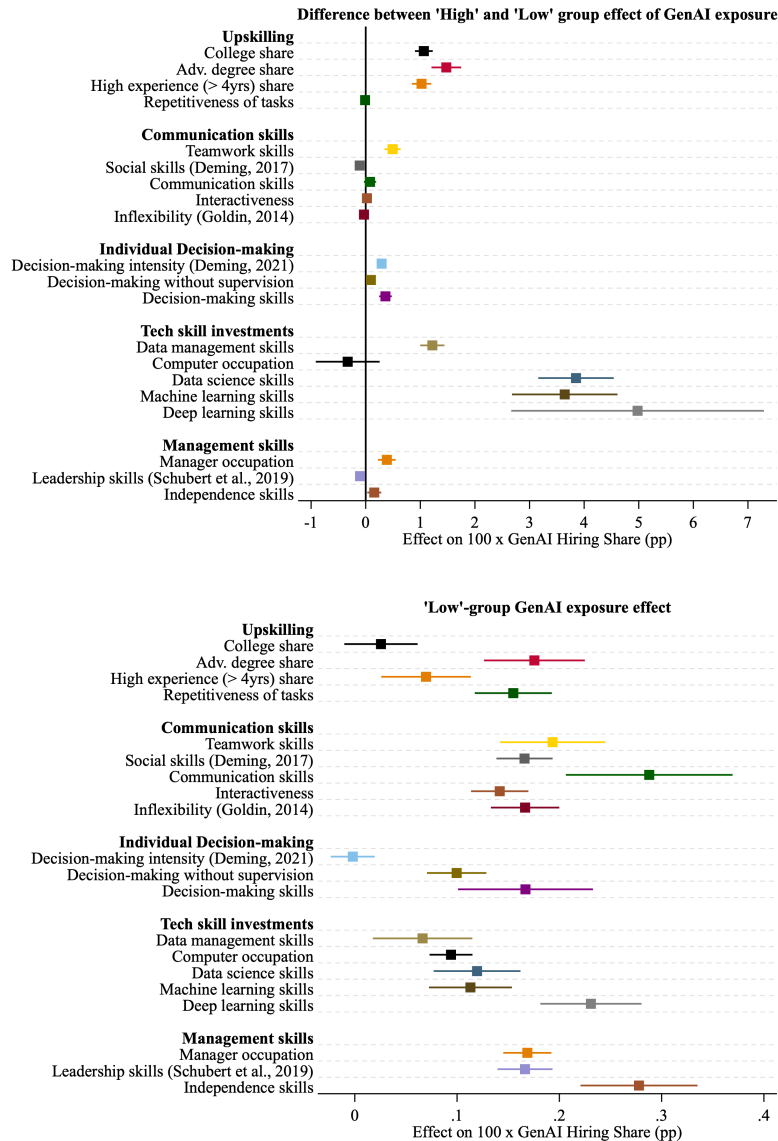


Figure 15:
Gen. AI exposure effects on remote work by firm characteristics

This figure shows the synthetic diff-in-diff effect estimates of the difference between the treated and control groups in remote work shares after the release of ChatGPT for different subsets of firms based their hiring characteristics. The treatment group is defined as firms that are in the top decile of generative AI exposure based on the composition of their 2021/2022 job postings (excl. Q4 2022). The control group is determined by weighting firms outside the top decile to find the best fit to the pre-period trends in the treatment group, as described in Arkhangelsky et al. (2021). The subgroups are defined as firms that are above-median (square red markers) or below-median (grey triangles) in terms of each characteristic of firm hiring in 2022. The 95% confidence intervals shown are based on block bootstrap standard errors computed from 50 resamples, using the procedure detailed in Clarke et al. (2023).




Table I:
Remote work and generative AI prevalence by occupation.

This table shows the share of all job postings in the YTD ending Sep. 2024 in each 6-digit SOC 2010 occupation that are for remote jobs (panel A), require generative AI-related skills (panel B), and are both remote and require generative AI skills (panel C). Each panel shows the top 20 occupations, ranked by the measure of interest, which have at least 5,000 job postings during the measurement period. The last column of the table shows the total job postings in the sample for that occupation.

(a) A: Remote shares

SOC	Occupation title	Remote	Jobs
15-2011	Actuaries	40	9,871
41-9041	Telemarketers	30	6,026
21-1014	Mental Health Counselors	28	60,334
27-3042	Technical Writers	28	19,584
41-3041	Travel Agents	28	5,208
13-2053	Insurance Underwriters	26	19,385
13-1031	Claims Adjusters, Examiners, and Investigators	24	58,753
43-9041	Insurance Claims and Policy Processing Clerks	23	20,677
41-9031	Sales Engineers	21	13,462
27-3043	Writers and Authors	21	27,545
19-3094	Political Scientists	20	6,688
21-1022	Healthcare Social Workers	20	25,264
15-1121	Computer Systems Analysts	20	78,807
15-1134	Web Developers	20	69,144
13-1075	Labor Relations Specialists	20	5,305
19-3031	Clinical, Counseling, and School Psychologists	19	25,873
23-1011	Lawyers	19	103,553
21-1013	Marriage and Family Therapists	19	26,145
13-1111	Management Analysts	18	125,206
15-2041	Statisticians	18	6,128

(b) B: Generative AI shares

SOC	Occupation title	Gen. AI	Jobs
15-1131	Computer Programmers	8.5	19,768
27-3042	Technical Writers	6.1	19,584
27-3043	Writers and Authors	5.8	27,545
15-2099	Mathematical Science Occupations, All Other	4	187,460
27-3091	Interpreters and Translators	3	22,998
27-3041	Editors	1.9	16,550
11-2021	Marketing Managers	1.7	197,468
15-1133	Software Developers, Systems Software	1.7	218,581
15-1132	Software Developers, Applications	1.7	218,581
11-3021	Computer and Information Systems Managers	1.4	17,145
15-1134	Web Developers	1.4	69,144
15-1141	Database Administrators	1.2	133,219
27-1011	Art Directors	1.1	9,699
43-9011	Computer Operators	1.1	305,659
27-1014	Special Effects Artists and Animators	1.1	6,281
25-9031	Instructional Coordinators	.98	28,773
25-3099	Teachers and Instructors, All Other	.95	38,229
11-9041	Architectural and Engineering Managers	.95	65,851
27-1021	Commercial and Industrial Designers	.89	18,390
15-2041	Statisticians	.78	6,128

(c) C: Generative AI share in remote jobs

SOC	Occupation title	Gen. AI Remote	Jobs
27-3091	Interpreters and Translators	21	22,998
27-3042	Technical Writers	18	19,584
15-1131	Computer Programmers	11	19,768
27-3043	Writers and Authors	9.7	27,545
27-1014	Special Effects Artists and Animators	6.4	6,281
25-9021	Farm and Home Management Educators	4.8	20,921
15-2099	Mathematical Science Occupations, All Other	3.9	187,460
17-1011	Architects, Except Landscape and Naval	3.5	9,934
25-9031	Instructional Coordinators	3.4	28,773
15-2041	Statisticians	2.9	6,128
25-2031	Sec. School Teachers (exc. Special/Tech. Ed.)	2.8	131,514
11-9041	Architectural and Engineering Managers	2.5	65,851
11-9032	Education Administrators, K-12	2.2	54,092
19-4021	Biological Technicians	2	7,372
29-9011	Occupational Health and Safety Specialists	2	47,336
15-1134	Web Developers	1.9	69,144
11-2021	Marketing Managers	1.6	197,468
13-1151	Training and Development Specialists	1.4	66,980
15-1132	Software Developers, Applications	1.3	218,581
15-1133	Software Developers, Systems Software	1.3	218,581

Table II:
Remote work and Gen. AI Adoption: Firm-level (Competition IV)

This table shows estimates of the remote share effect at the firm level in a specification of the form

$$100 \times \text{GenAIJobShare}('23-'34)_i = \beta \text{RemoteWorkShare}('21-'22)_i + \alpha_{ind} + \text{Controls}_i + \varepsilon_i$$

for the generative AI share for the YTD Sep. 2024 period and the remote share during 2021-2022 (excl. Q4 2022). Here, the instrument consists of the interaction between firm-level exposure to MSA teleworkability through its hiring labor markets and the firm's own teleworkability (measured in 2019). For comparison, see the results in Table A.1 using an instrument based on commuting time. The control variables include proxies for the task-level teleworkability (Dingel and Neiman, 2020) and generative AI exposure (Eisfeldt et al., 2023) at the firm, 2019 remote work shares at the firm, exposure to MSA remote work levels and commuting time through the firm's hiring labor markets, hiring share of different technology skills in 2019, as well as industry sector fixed effects and the level of education required for average job postings at the firm. T-test statistics based on heteroskedasticity-robust standard errors clustered at the company level in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

<i>Dependent variable:</i>	100 × Generative AI Job Share (%)				Remote Share (%)
<i>Estimation:</i>	OLS	IV	IV	IV	OLS
	(1)	(2)	(3)	(4)	(5)
Remote Job Share (%)	0.51*** (10.91)	5.65*** (11.21)	5.96*** (8.52)	6.31*** (6.64)	
MSA Telew. ₁₉ × Firm Telew. ₁₉					21.50*** (12.15)
Observations	140,540	101,191	101,179	88,472	88,472
1st-stage KP F-stat.		560	256	148	
Firm remote work & telework. (2019)		X	X	X	X
Firm exposure to MSA remote work & telework. (2019)		X	X	X	X
Firm exposure to MSA commuting time ('15-'19)		X	X	X	X
Firm Gen. AI potential			X	X	X
Firm Teleworkability			X	X	X
Firm tech skill hiring: AI, DL, ML, Data Mgmt. (2019)			X	X	X
Firm adv. education requirements			X	X	X
2-dig. Industry FEs				X	X

Table III:**Remote work effects on Gen. AI Adoption: Occ. × Firm-level (Competition IV)**

This table shows estimates of the remote share effect at the occupation-by-firm level in a specification of the form

$$100 \times \text{GenAIJobShare}('23-'34)_i = \beta \text{RemoteWorkShare}('21-'22)_i + \alpha_{\text{firm}} + \alpha_{\text{occ}} + \text{Controls}_i + \varepsilon_i$$

for the generative AI share for the YTD Sep. 2024 period and the remote share during 2021-2022 (excl. Q4 2022). The instrument consists of the interaction between firm-level exposure to MSA teleworkability through its hiring labor markets (measured in 2019) and occupation-level remote work adoption as of 2021-2022. T-test statistics based on heteroskedasticity-robust standard errors double-clustered at the company and 6-digit occupation level in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

<i>Dependent variable:</i>	100 × Generative AI Job Share (%)			
<i>Estimation:</i>	IV (1)	IV (2)	IV (3)	IV (4)
Remote Job Share (%)	4.40*** (5.74)	7.78*** (4.80)	8.78*** (4.05)	22.13** (2.45)
Observations	2,166,999	2,166,992	1,953,368	2,166,368
1st-stage KP F-stat.	533	96	56	14
Firm remote work & telework. (2019)	X	X	X	
Firm exposure to MSA remote work & telework. (2019)	X	X	X	
Firm exposure to MSA commuting time ('15-'19)	X	X	X	
Firm Gen. AI potential			X	
Firm Teleworkability			X	
Firm tech skill hiring: AI, DL, ML, Data Mgmt. (2019)			X	
Firm × Occ. adv. educ. requirements			X	X
2-dig. Industry FEs			X	
Occupation FEs		X	X	X
Firm FEs				X

Table IV:
Synthetic diff-in-diff event study: effects of GenAI exposure

This table shows estimates of the effect of being in the high generative AI exposure group post-ChatGPT release, estimated using the synthetic diff-in-diff estimator in equation 3. The treatment group is defined as firms that are in the top decile of generative AI exposure based on the composition of their 2021/2022 job postings (excl. Q4 2022). The control group is determined by weighting firms outside the top decile to find the best fit to the pre-period trends in the treatment group, as described in Arkhangelsky et al. (2021). The dependent variable is the share of all jobs that mention generative AI skills in column 1, and the share of all jobs that are remote in columns 2 and 3. T-test statistics based on standard errors from a block bootstrap procedure with 50 resamples at the company level in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

<i>Dependent variable:</i>	Share of All Jobs		Total Jobs	
<i>Job category:</i>	Gen. AI (1)	Remote (2)	Remote (3)	All (4)
$\mathbb{1}[\text{Post-ChatGPT}] \times \mathbb{1}[\text{High Gen. AI Exposure}]$	0.002*** (14.09)	-0.034*** (-8.04)	-10.711*** (-3.05)	-29.727** (-2.55)
Observations	222,555	222,555	222,555	222,555
<i>Estimation:</i>	SDID	SDID	SDID	SDID
Firm FEs	X	X	X	X
Time FEs	X	X	X	X

Table V:
Change in firm-level job characteristics after GenAI adoption

This table shows estimates of the effect of being in the high generative AI exposure group post-ChatGPT release, estimated using the synthetic diff-in-diff estimator in equation 3. The treatment group is defined as firms that are in the top decile of generative AI exposure based on the composition of their 2021/2022 job postings (excl. Q4 2022). The control group is determined by weighting firms outside the top decile to find the best fit to the pre-period trends in the treatment group, as described in Arkhangelsky et al. (2021). The dependent variables represent shares of different job categories among all non-remote worker (panel A) and remote worker (panel B) job postings at the firm, in each panel only including only firms that have non-remote or remote job postings, respectively, in all quarters included in the sample (Q1 2021-Q3 2024). T-test statistics based on standard errors from a block bootstrap procedure with 50 resamples at the company level in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

<i>Dep. var. job category:</i>	High Experience (1)	Medium Experience (2)	Low Experience (3)	College (4)	Advanced degree (5)	Data mgmt. (6)	Decision making (7)	Communi- cation (8)
<u>Panel A: Share of Non-Remote Jobs</u>								
$\mathbb{1}[\text{Post}] \times \mathbb{1}[\text{High Gen. AI Exp.}]$	0.017*** (7.07)	-0.011*** (-4.08)	-0.004 (-1.64)	0.005 (1.56)	0.000 (0.47)	0.003*** (2.70)	0.016*** (4.58)	0.009** (2.49)
Observations	220,980	220,980	220,980	220,980	220,980	220,980	220,980	220,980
<u>Panel B: Share of Remote Jobs</u>								
$\mathbb{1}[\text{Post}] \times \mathbb{1}[\text{High Gen. AI Exp.}]$	0.011** (2.12)	-0.012 (-1.59)	0.001 (0.09)	0.005 (0.69)	-0.001 (-0.58)	0.005 (1.21)	0.023*** (4.20)	0.014* (1.71)
Observations	33,900	33,900	33,900	33,900	33,900	33,900	33,900	33,900
<i>Estimation:</i>	SDID	SDID	SDID	SDID	SDID	SDID	SDID	SDID
Firm FEs	X	X	X	X	X	X	X	X
Time FEs	X	X	X	X	X	X	X	X

Table VI:
Return-to-Office firms have higher remote work effects on Gen. AI Adoption

This table shows estimation results for the following regression specifications: The specification in column 1 is identical to the one in column 4 of Table II and the specification in column 2 is identical to that in column 4 of Table III, except for the following differences: the endogenous remote share is interacted with an indicator for a company that has a return-to-office (RTO) policy as of March 2025, and so is the instrument for the endogenous variable; an indicator for a firm having an RTO policy is added as an additional control variable where this is not collinear with fixed effects. T-test statistics based on heteroskedasticity-robust standard errors clustered at the firm level (column 1) or double-clustered at the company and 6-digit occupation level (column 2) in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

<i>Dependent variable:</i>	100 × Generative AI Job Share (%)	
<i>Estimation:</i>	IV (1)	IV (2)
Remote Job Share (%)	6.22*** (6.56)	22.19** (2.44)
Remote Job Share (%) x 1[Return-to-Office]	6.48*** (3.39)	-0.94 (-0.40)
Observations	88,472	2,166,368
1st-stage KP F-stat.	74	7
Firm remote work & telework. (2019)	X	
Firm exposure to MSA remote work & telework. (2019)	X	
Firm exposure to MSA commuting time ('15-'19)	X	
Firm Gen. AI potential	X	
Firm Teleworkability	X	
Firm tech skill hiring: AI, DL, ML, Data Mgmt. (2019)	X	
Firm × Occ. adv. educ. requirements	X	X
2-dig. Industry FEs	X	
Occupation FEs		X
Firm FEs		X
"Firm has RTO Policy" Indicator	X	

A Appendix Figures & Tables

Figure A.1:
Remote work and generative AI at the firm level

This figure plots the share of job postings in the YTD Sep. 2024 period that are for jobs that mention generative AI relative to those that are for remote jobs. The job postings data are from Lightcast and are aggregated by company. The red line indicates a linear best fit. The analysis only includes companies that have at least 10 job postings in the sample period shown and have non-zero job postings in all quarters from Q1 2021 to Q3 2024.

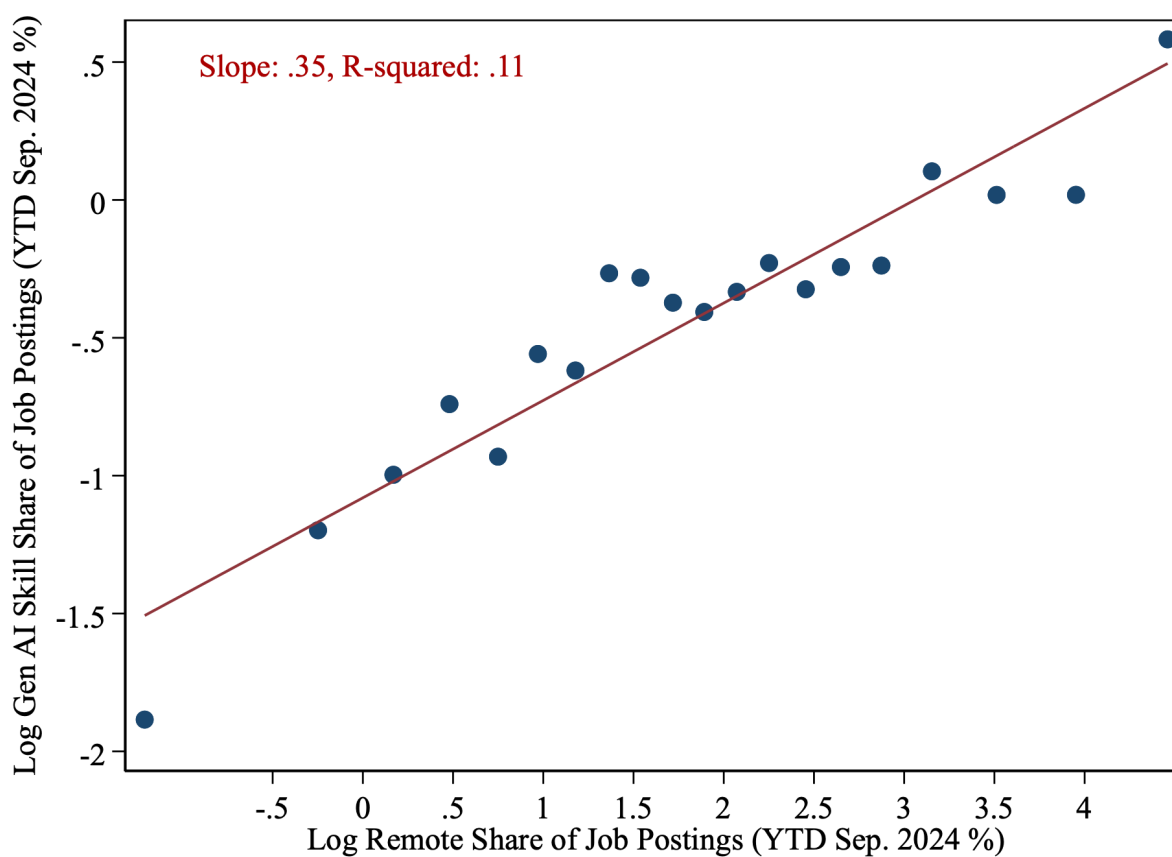

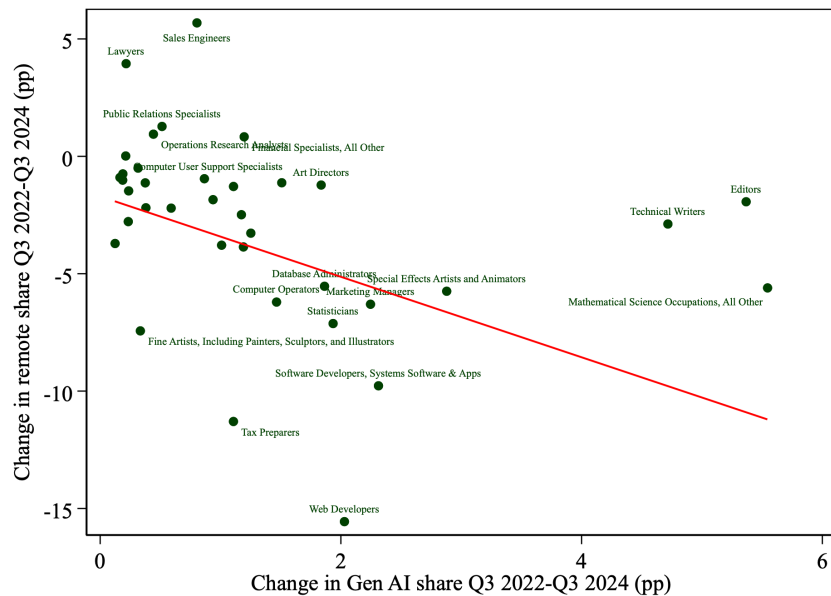


Figure A.2:
Pandemic remote hiring and generative AI adoption

These figures focus on the top quartile of remote work adopting occupations (based on 2021/2022 job postings). Panel A shows the change in hiring for generative AI skills after the release of ChatGPT (Q3 2022-Q3 2024) as a function of remote work changes during the pandemic (Q3 2019-Q3 2022). Panel B shows remote work changes after the ChatGPT release as a function of generative AI adoption over the same period. Both graphs only include occupations with > 1K job postings in Q3 2022 with some Q3 2024 Gen AI adoption (> 0.1% of jobs). Two outlier occupations, 'Computer Programmers' and 'Writers & Authors' are not included for better visibility, but follow similar patterns. The data are Lightcast job postings where “Gen AI share” is the share of jobs mentioning Gen. AI-related keywords.



(A) Δ Remote Work ('19-'22) and Gen. AI adoption ('22-'24)



(B) Gen. AI adoption ('22-'24) and Δ Remote Work ('22-'24)

Figure A.3:
Effects of remote work adoption on firm hiring characteristics (OLS)

This figure shows coefficients estimated using OLS for the effect of remote work prevalence in a firm's 2021/2022 (excl. Q4 2022) job postings on the standardized characteristics of the firm's job postings in 2022 in a regression of the form

$$\text{Skill}(2022)_i = \alpha + \beta \text{RemoteWorkShare}('21-'22)_i + \gamma \text{Skill}(2019)_i + \text{Controls}_i + \varepsilon_i,$$

where the controls in all regressions include the standardized 2019 value of the dependent variable as a control variable, so the coefficients can be interpreted as the effect of changes in remote work shares on changes in the composition. All regressions also include the following control variables: NAICS 2-digit fixed effects; company's remote work share in 2019, the company's uninteracted exposure to MSA teleworkability in 2019, uninteracted firm-level teleworkability in 2019, the company's share of jobs requiring a college education and the share requiring an advanced degree in 2019, the share of the company's job postings in 2019 that was for computer occupations or manager positions; the log of total job postings in 2019 and in 2022; the company's labor market exposure to MSA remote shares in 2019. The 95% confidence intervals shown are based on heteroskedasticity-robust standard errors clustered at the NAICS 2-digit sector level.




Figure A.4:
Remote work effects on Gen. AI adoption by occupation characteristics: excl. deep learning intensive occupation

This figure shows coefficients corresponding to the same analysis as Figure 7, but omitting the top 10 occupations by average deep learning skill share in job postings in 2021/22 (which are also all the occupations with average deep learning skill shares above 1%). The excluded occupations are: Computer and Information Research Scientists; Mathematical Science Occupations, All Other; Sociologists; Physical Scientists, All Other; Semiconductor Processing Technicians; Social Science Research Assistants; Atmospheric, Earth, Marine, and Space Sciences Teachers, Postsecondary; Physicists; Statisticians; Biological Scientists, All Other. Sample sizes for the regressions shown below range from 1.72M to 2.15M, and first-stage Kleibergen Paap F-statistics are above 18 for all IV regressions shown.

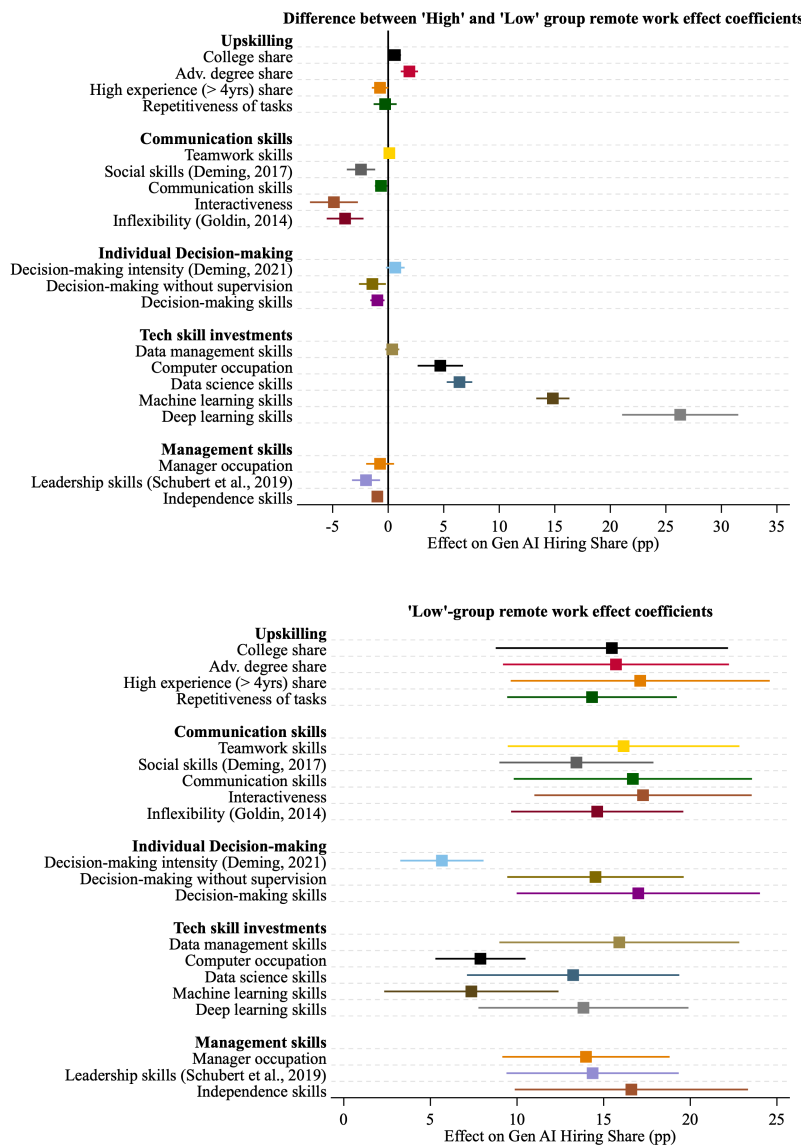


Figure A.5:
SDID Gen. AI exposure effects on Gen. AI adoption by firm characteristics

This figure shows the synthetic diff-in-diff effect estimates of the difference between the treated and control groups in generative AI adoption after the release of ChatGPT for different subsets of firms based their hiring characteristics. The treatment group is defined as firms that are in the top decile of generative AI exposure based on the composition of their 2021/2022 job postings (excl. Q4 2022). The control group is determined by weighting firms outside the top decile to find the best fit to the pre-period trends in the treatment group, as described in Arkhangelsky et al. (2021). The subgroups are defined as firms that are above-median (top panel estimates) or below-median (bottom panel estimates) in terms of each characteristic of firm hiring in 2022. The 95% confidence intervals shown are based on block bootstrap standard errors computed from 50 resamples, using the procedure detailed in Clarke et al. (2023).




Figure A.6:
Firm data management capabilities and generative AI adoption

This figure shows the share of job postings in each quarter that are for jobs that mention generative AI. The job posting trend is shown by quartile of firms based on the firm's data management skill hiring share in 2019. Moreover, the sample is restricted to job postings where the occupation is in the top quartile by remote share in 2021/2022 (excl. Q4 2022), using 2019 hiring as weights. The grey drop line indicates the period (Q4 2022) when ChatGPT was released.




Table A.1:
Remote Work Effects on Gen. AI Adoption: Firm-level (Commuting Cost IV)

This table shows estimates of the remote share effect at the firm level in a specification of the form

$$100 \times \text{GenAIJobShare}('23-'34)_i = \beta \text{RemoteWorkShare}('21-'22)_i + \alpha_{ind} + \text{Controls}_i + \varepsilon_i$$

for the generative AI share for the YTD Sep. 2024 period and the remote share during 2021-2022 (excl. Q4 2022). Here, the instrument consists of the interaction between firm-level exposure to MSA commuting costs through its hiring labor markets and the firm's own teleworkability (measured in 2019). For comparison, see the results in Table II using an instrument based on MSA teleworkability. The control variables include proxies for the task-level teleworkability (Dingel and Neiman, 2020) and generative AI exposure (Eisfeldt et al., 2023) at the firm, 2019 remote work shares at the firm, exposure to MSA remote work levels and commuting time through the firm's hiring labor markets, hiring share of different technology skills in 2019, as well as industry sector fixed effects and the level of education required for average job postings at the firm. T-test statistics based on heteroskedasticity-robust standard errors clustered at the company level in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

<i>Dependent variable:</i>	100 × Generative AI Job Share (%)				Remote Share (%)
<i>Estimation:</i>	OLS	IV	IV	IV	OLS
	(1)	(2)	(3)	(4)	(5)
Remote Job Share (%)	0.51*** (10.91)	6.08*** (9.64)	6.60*** (7.28)	6.80*** (5.83)	
MSA Commute Time ₁₉ × Firm Telew. ₁₉					0.28*** (9.25)
Observations	140,540	101,191	101,179	88,472	88,472
1st-stage KP F-stat.		275	129	86	
Firm remote work & telework. (2019)		X	X	X	X
Firm exposure to MSA remote work & telework. (2019)		X	X	X	X
Firm exposure to MSA commuting time ('15-'19)		X	X	X	X
Firm Gen. AI potential			X	X	X
Firm Teleworkability			X	X	X
Firm tech skill hiring: AI, DL, ML, Data Mgmt. (2019)			X	X	X
Firm adv. education requirements			X	X	X
2-dig. Industry FEs				X	X

Table A.2:
Remote work effects on Gen. AI Adoption YTD Sep. 2024: Occ. × Firm-level
(Commuting Cost IV)

This table shows estimates of the remote share effect at the occupation-by-firm level in a specification of the form

$$100 \times \text{GenAIJobShare}('23-'34)_i = \beta \text{RemoteWorkShare}('21-'22)_i + \alpha_{firm} + \alpha_{occ} + \text{Controls}_i + \varepsilon_i$$

for the generative AI share for the YTD Sep. 2024 period and the remote share during 2021-2022 (excl. Q4 2022). The instrument consists of the interaction between firm-level exposure to MSA commuting costs through its hiring labor markets (measured in 2019), firm level teleworkability in 2019, and occupation-level remote work adoption as of 2021-2022. T-test statistics based on heteroskedasticity-robust standard errors double-clustered at the company and 6-digit occupation level in parentheses: * p<0.10, ** p<0.05, *** p<0.01.

<i>Dependent variable:</i>	100 × Generative AI Job Share (%)			
<i>Estimation:</i>	IV (1)	IV (2)	IV (3)	IV (4)
Remote Job Share (%)	4.34*** (5.71)	7.73*** (4.75)	8.69*** (4.03)	24.01** (2.37)
Observations	2,166,999	2,166,992	1,953,368	2,166,368
1st-stage KP F-stat.	551	99	58	12
Firm remote work & telework. (2019)	X	X	X	
Firm exposure to MSA remote work & telework. (2019)	X	X	X	
Firm exposure to MSA commuting time ('15-'19)	X	X	X	
Firm Gen. AI potential			X	
Firm Teleworkability			X	
Firm tech skill hiring: AI, DL, ML, Data Mgmt. (2019)			X	
Firm × Occ. adv. educ. requirements			X	X
2-dig. Industry FEs			X	
Occupation FEs		X	X	X
Firm FEs				X

B Data collection

Return-to-Office Data from Flex Index

I obtain information on firms' return-to-office (RTO) policies by using a webscraper to collect the crowdsourced public information from the website Flex Index (<https://www.flexindex.com>).¹¹

In total, I am able to collect the RTO data of 13,251 companies from the Flex Index site. This data includes each firm's basic characteristics, including headquarter state and company size, and a classification of its office attendance requirements for employees.

After removing duplicates, I match the Flex Index companies to Compustat data to identify the standard company code (GVKEY) associated with the Flex Index companies—which are listed only by name. The matching process proceeds in the following order:

1. Fuzzy match company names, as well as state and/or city between the Compustat HQ location and the Flex Index company location.
2. Include all exact 1:1 company name matches in the final data
3. Drop imprecise matches that do not have city and state information
4. For remaining imprecise matches, drop all with fewer than 1000 employees, as such small firms are unlikely to be included in Compustat.
5. Manually review all remaining imprecise matches to see if both company name and location are plausibly the same, i.e. the company name only has generic differences, e.g. omitting "Corp." or "Inc.", and the company location is in the same or a neighboring metro area in both data sets. Include in the final data if both location and name are plausible matches.

The final dataset includes data on 1,336 companies with office requirements (8 Flex Index classifications, see Table 1) for which it was possible to determine a Compustat GVKEY, and Compustat company name.

See Table B.3 for details on the different RTO policy classifications. I classify a company as having provided a "Return-to-Office" mandate, if its RTO policy requires either a full-time presence, or requires a non-zero minimum number of days or time in the office (categories 3-8).

¹¹The data was collected on March 3rd, 2025 by my research assistant Yi Li, who provided invaluable assistance in this task.

No.	Flex Index Category	Description
1	Fully Remote	Organization does not have offices; all employees work remotely.
2	Employee's Choice	Each employee can choose how often (or never) to come to the office.
3	Minimum Days a week	There is a minimum number of days employees must come to the office each week.
4	Specific Days a week	There are specific days of the week employees must come to the office.
5	Minimum & Specific Days a week	
6	Minimum % of Time per week	There is a minimum % of time that employees must be in the office.
7	Full Time in Office	Employees are expected to be in the office full time.
8	No Requirements	

Table B.3: Flex Index Return-to-Office Policy Classifications. Source: <https://www.flexindex.com/about>

C Derivations

A Optimal decision-making intensity

We will work backwards to derive the firm's optimal choices for technology investment. After a firm chooses whether to adopt a technology, it determines the efficient amount of time for its workers to spend on decision-making. There is a trade-off, as time spent on decisions D_f is time spent not producing task output, so spending more time on decisions will only be optimal for firms that have tasks are sensitive to good decision-making, i.e. where $\delta(x)$ is high on average.

Maximizing the expression in equation 4 with regard to D_f , the firm's first-order condition is

$$\frac{j - k_{fj}}{1 - D_f} = \frac{\eta(e^j - e^{k_{fj}})}{\eta D_f + \rho^{\mathbb{I}[\text{Remote}]} M_f},$$

where the LHS captures the cost in terms of lost output of spending more time on decisions, while the RHS represents the benefit in terms of additional quality of decisions for the firm's tasks. Some algebra then results in the optimal decision intensity as

$$D_f^* = \frac{(e^j - e^{k_{fj}}) - \frac{1}{\eta}(j - k_{fj})\rho^{\mathbb{I}[\text{Remote}]} M_f}{(e^j - e^{k_{fj}}) + (j - k_{fj})}$$

This expression directly implies that firms will choose greater decision-making intensity for

their workers in (1) firms where local worker decisions are more important (higher η_f); (2) firms where management decision support is worse (lower M_f); (3) firms that work remotely (as $\rho_f^{\mathbb{1}[\text{Remote}]} < 1$).

B Remote work decision quality.

Starting from the optimal decision intensity

$$D_f^* = \frac{(e^j - e^{k_{fj}}) - \frac{1}{\eta}(j - k_{fj})\rho^{\mathbb{1}[\text{Remote}]}M_f}{(e^j - e^{k_{fj}}) + (j - k_{fj})},$$

note that this can be written as

$$D_f^* = \frac{(e^j - e^{k_{fj}})}{(e^j - e^{k_{fj}}) + (j - k_{fj})} - \frac{(j - k_{fj})}{\eta((e^j - e^{k_{fj}}) + (j - k_{fj}))} \rho^{\mathbb{1}[\text{Remote}]}M_f.$$

As decision quality is defined as

$$Q_{fj} = \ln\left(\eta D_f + \rho^{\mathbb{1}[\text{Remote}]}M_f\right),$$

it follows that the sign of the change in overall decision quality with regard to changes in central decision support is the same as the sign of

$$1 - \frac{(j - k_{fj})}{(e^j - e^{k_{fj}}) + (j - k_{fj})},$$

such that

$$\frac{\partial Q_{fj}}{\partial(\rho^{\mathbb{1}[\text{Remote}]}M_f)} > 0 \Leftrightarrow e^j - e^{k_{fj}} > 0.$$

That is, as long as the decision quality has a positive effect on output, then an increase in central decision support is not going to be fully offset by a reduction in local decision intensity. This implies that *remote firms have lower overall decision quality even though they invest more worker effort in local decision-making*