NBER WORKING PAPER SERIES

TECHNOLOGY AND LABOR DISPLACEMENT: EVIDENCE FROM LINKING PATENTS WITH WORKER-LEVEL DATA

Leonid Kogan
Dimitris Papanikolaou
Lawrence D.W. Schmidt
Bryan Seegmiller

Working Paper 31846
http://www.nber.org/papers/w31846

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2023, Revised November 2025

We are grateful to Daron Acemoglu, Philippe Aghion, David Autor, Effi Benmelech, Nicholas Bloom, Carter Braxton, Julieta Caunedo, Martin Beraja, Carola Frydman, Tarek Hassan, David Hemous, Anders Humlum, Nir Jaimovich, David Lagakos, Joseba Martinez, Michael Peters, Pascual Restrepo, Jonathan Rothbaum, Miao Ben Zhang, along with seminar participants at University of Amsterdam, BI-SHoF Conference, Boston University, CIREQ Macroeconomics Conference, Columbia GSB, FIRS, Johns Hopkins, HKUST, Labor and Finance Group, NBER (EFG, PRMP, LS, PIE), Macro-Finance Society, MIT Sloan, Michigan State, the Econometric Society, Rice University, University of Rochester, the Society of Economic Dynamics, University College London, University of Illinois at Urbana Champaign, University of Toronto, UZH Workshop on Automation, Tsinghua PBC, WFA, and Wharton for valuable discussions and feedback. We thank Carter Braxton, Will Cong, and Jonathan Rothbaum for generously sharing code. Huben Liu provided outstanding research support. The paper has been previously circulated under the title "Technology, Vintage-Specific Human Capital, and Labor Displacement: Evidence from Linking Patents with Occupations". The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product (Data Management System (DMS) number: P-7503840, Disclosure Review Board (DRB) approval numbers: CBDRB-FY21-POP001-0176, CBDRB-FY22-SEHSD003-006, CBDRB-FY22-SEHSD003-023, CBDRB-FY22-SEHSD003-028, CBDRB-FY23-SEHSD003-0350, CBDRB-FY23-SEHSD003-064). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Technology and Labor Displacement: Evidence from Linking Patents with Worker-Level Data
Leonid Kogan, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller
NBER Working Paper No. 31846
November 2023, Revised November 2025
JEL No. E0, E01, J01, J23, J24, O3, O4

ABSTRACT

We develop measures of labor-saving and labor-augmenting technology exposure using textual analysis of patents and job tasks. Using US administrative data, we show that exposure to labor-saving technologies negatively affects the earnings of exposed workers. This negative effect is pervasive across both blue-and white-collar workers and across workers of different ages or earnings relative to their peers. In contrast, labor-augmenting technologies have a heterogeneous impact on exposed workers. While the wage bill paid to affected groups rises, this increase is driven primarily by an increase in employment, while earnings rise for new entrants but decline for incumbent workers. This decline is primarily present among white-collar, older, and higher-paid workers, highlighting the importance of vintage-specific human capital. Last, we find positive spillovers of both types of innovation at the industry level, benefiting other workers in the same industry who are not directly exposed to these innovations.

Leonid Kogan
Massachusetts Institute of Technology
MIT Sloan School of Management
and NBER
lkogan@mit.edu

Dimitris Papanikolaou
Northwestern University
Kellogg School of Management
Department of Finance
and NBER
d-papanikolaou@kellogg.northwestern.edu

Lawrence D.W. Schmidt
Massachusetts Institute of Technology
ldws@mit.edu

Bryan Seegmiller
Northwestern University
bryan.seegmiller@kellogg.northwestern.edu

Economists and workers alike have long worried about the prospect of technological displacement of labor. New technologies can displace incumbent workers, either because these technologies can perform certain tasks formerly done by workers or because these technologies require new skills that incumbent workers lack. Labor-saving technologies directly substitute for labor and can reduce wages as demand for labor falls. By contrast, labor-augmenting technologies can benefit workers as a whole, although individual workers who cannot adapt to the new technologies may be left behind. The distinction between these two types of technology exposure is relevant for understanding the effects of technology on the cross-section of workers. We provide evidence that is consistent with both of these channels by first developing direct measures of worker exposure to these two types of technologies and then linking them to individual worker outcomes using administrative data on worker earnings from the United States.

We have four key findings. First, the development of new labor-saving technologies reduces the earnings of directly affected workers. This negative effect is among both blue- and white-collar workers and is largely not predictable by worker characteristics such as age, education level, or prior level of earnings relative to their peers. Second, in sharp contrast, the development of new labor-augmenting technologies has a heterogeneous impact on directly exposed workers. While the total wage payments to affected occupation–industry cells rise, this increase is primarily due to higher employment levels. Average worker earnings remain largely unchanged, with new entrants seeing higher earnings, while incumbent workers experience modest declines. Importantly, this decline in earnings amongst incumbents is concentrated among older and higher-paid workers, which is consistent with the importance of vintage-specific human capital. Third, job destruction accounts for an economically significant share of these earnings declines in response to exposure to labor-saving and labor-augmenting technologies. Last, we find evidence of positive spillovers: when aggregated at the industry level, both measures of technology exposure are positively related to industry productivity and the earnings of incumbent workers in that industry—after controlling for these workers' direct task exposure to these technologies.

To motivate our empirical analysis, we start with a model that combines elements of Acemoglu and Restrepo (2018) and Caunedo, Jaume, and Keller (2023). Workers perform a combination of two types of tasks: we can regard these two types of task as either routine and non-routine or high- and low-skill, respectively. Our main assumption is that these two types of tasks differ in the elasticity of substitution between capital and labor inputs in performing each task. As in Caunedo et al. (2023), capital is specific to each task. A given technological innovation can improve the quality of capital used in multiple tasks, and hence can simultaneously complement labor inputs in some tasks and be a substitute for others. For a given occupation, a technology is labor-augmenting if it improves the quality of capital used in its non-routine (or high-skill) tasks, and labor-saving if it improves the quality of capital in its routine (or low-skill) tasks. Since the arrival rate of new technologies is neither constant nor uniform across tasks, different occupations are exposed to technology improvements at different points in time. Our model therefore implies that to identify periods during which new technologies substitute for or complement the tasks of a given occupation, we need to estimate how specific technologies are related to specific tasks performed by workers in a given occupation.

We use textual analysis of patents and occupation task descriptions to construct technology exposure measures that closely resemble their counterparts generated by the model. First, we classify each task performed by an occupation as a routine or non-routine (alternatively as high- or low-required experience) using generative AI (GPT4). The resulting task classifications aggregated at the occupation level correlate highly with the Acemoglu and Autor (2011) measure of routine-task-intensity (RTI) and ONET's classification of jobs into high and low-skill. Second, we use natural language processing methods to estimate the similarity between the textual description of the routine or non-routine tasks performed by an occupation and that of major technological breakthroughs. We identify these breakthroughs using the methodology of Kelly, Papanikolaou, Seru, and Taddy (2021) who define a breakthrough innovation as one that is both novel (i.e. distinct from prior patents) and impactful (i.e. related to subsequent patents). Overall, we see that highly paid occupations are more exposed to our measure of labor-augmenting technologies, whereas the exposure to labor-saving technologies is highest for occupations in the middle of the earnings distribution—consistent with the literature on job polarization (e.g. Autor and Dorn, 2013; Goos, Manning, and Salomons, 2014; Jaimovich and Siu, 2020).

Our methodology allows us to control for common shocks to labor demand and supply at the industry or occupation level, and thus complements existing work that focuses on measures of technology exposure that are invariant over time (see, for instance Autor, Levy, and Murnane, 2003; Autor and Dorn, 2013; Acemoglu and Autor, 2011; Webb, 2020). In particular, our empirical analysis leverages the granularity of our patent–occupation measures to generate time variation in worker exposure at the occupation–industry level: this variation is driven by differences in the rate at which firms in different industries develop and patent new technologies that are related to the routine and non-routine tasks of a given occupation at a point in time. Our empirical estimates are identified by comparing two workers in different occupation–industry cells that are otherwise similar in their observable characteristics and past earnings history. For most of the paper, we focus on the direct effect of technology exposure, by controlling for variation related to changes in industry productivity—by including industry-year and occupation-year fixed effects in our specification.

Naturally, the emergence of new labor-saving or labor-augmenting technologies can be influenced by the current state of the labor market for specific types of workers. To strengthen our interpretation of the findings as identifying the causal effect of technology on worker earnings, we develop a shift-share identification strategy inspired by Acemoglu, Akcigit, and Kerr (2016). In essence, we construct our instrument by interacting the share of patents in a given class from a specific industry related to a particular occupation (the share) with knowledge spillovers from past breakthrough innovations in different technology classes (the shift). Repeating our analysis with our shift-share instrument, we find not only qualitatively but also quantitatively similar estimates, which alleviates our endogeneity concerns.

We begin our analysis by examining aggregate outcomes for affected workers. Using US Census data at the occupation–industry level, we find that an increase in workers' exposure to labor-saving technologies is associated with a decline in total labor compensation, which is driven both by a decline in employment and a modest decline in average worker earnings. By contrast, an increase in exposure to technologies related to non-routine tasks is associated with a rise in the total wage payments to affected occupation–industry cells, although this increase is mostly driven by higher employment levels—average worker earnings are not significantly affected. Importantly, we obtain quantitatively similar estimates to our OLS estimates when we employ our shift-share IV, which helps strengthen our interpretation of these correlations as partly capturing causal effects.

The response of employment validates our interpretation of our two technology exposure measures as labor-saving and labor-augmenting. Nevertheless, the response of average earnings is somewhat at odds with this interpretation. However, we should emphasize that an important shortcoming of analyzing the effect of technology at an aggregate level is that the response of average earnings can be influenced by changes in worker composition. For example, if lower-paid workers exit an occupation or industry in response to labor-saving technologies, average earnings can rise even if individual workers experience declines.

To study the importance of technology exposure for individual worker earnings, we next use employer–employee matched administrative earnings records from the US Social Security Administration for the 1978 to 2016 period, which are linked with information on each worker's occupation and education from the Current Population Survey. Thus, relative to existing work which has mostly studied repeated cross-sections of workers, we are able to measure a worker's occupation prior to the introduction of related technologies and estimate how her earnings evolve in future years even if she switches employers, industries, and/or occupations.

We find that improvements in labor-saving technologies are negatively related to the wage earnings of incumbent workers employed in exposed occupation–industry cells. The decline in earnings implied by our worker-level regressions is significantly larger than those implied by the response of average earnings at the occupation–industry level, highlighting the importance of accounting for changes in worker composition. These effects are pervasive for both blue- and white-collar workers: they are present in both services and manufacturing, in occupations that emphasize manual or cognitive skills, and among workers with or without a college degree. Importantly, the magnitude of these estimates are essentially unrelated to worker age or relative worker earnings. Though these declines are not predictable from worker characteristics, they are concentrated on workers that are (likely involuntarily) separated from their employer. In particular, the increased likelihood of job destruction—defined as a worker switching employers while simultaneously experiencing significant earnings declines—in response to increased exposure to labor-saving technologies accounts for over one-half of the overall average effect.

Perhaps surprisingly, new labor-augmenting technologies also lead to a decline in earnings for exposed incumbent workers, although the average magnitudes are smaller than for labor-saving technologies. This decline in earnings for incumbent workers stands in sharp contrast to the increase in total labor compensation estimated using aggregate data. In particular, the fact that wage earnings for incumbent workers decline, whereas average earnings for the affected occupation–industry cells do not respond is consistent with the redistributive effects of skill displacement, in which new vintages of labor-augmenting technologies primarily benefit new workers. Consistent with this view, we document a positive effect of new labor-augmenting technologies on the earnings of new entrants. Importantly, the decline in earnings among incumbent workers is particularly pronounced for older workers and workers who are paid more relative to other workers with similar characteristics in the same industry and occupation. For these affected workers, the possibility of job destruction accounts for a significant fraction of the average effect. The fact that it is the most highly paid workers (relative to their peers) that are most adversely affected suggests that specific skills account for at least part of the differences in pay within jobs.

Do our results imply that technological improvements displace all workers? Not necessarily: in the model, there are also positive spillovers to worker earnings due to the increase in labor demand in response to productivity increases following new technological improvements. Thus far, we have focused on workers whose job tasks are directly related to new technologies: spillovers to all other workers in the same industry whose tasks are not directly related to these new technologies are absorbed by our controlling for granular industry-specific trends. By moving to a coarser industry definition, we are able to document positive spillover effects. We find that, consistent with our model, the average technology exposure of workers in a given industry is positively related to industry productivity. More importantly, when controlling for workers' direct technology exposure, we find that these improvements in technology are associated with increased earnings for all workers in the same industry.

Overall, our work contributes to the literature along two key dimensions. First, we propose a new direct measure of workers' technology exposure that differentiates between labor-saving and labor-augmenting technologies and varies over time within occupation and industry groups. Second, we provide a systematic analysis of the heterogeneous effects of both labor-substituting and labor-augmenting technologies on worker-level earnings outcomes in the United States since the early 1980s. Our new measures allow us to elucidate two distinct channels through which improvements in technology can displace incumbent workers: 1) new technologies substitute for labor and therefore reduce firms' demand for workers; 2) new technologies helps workers perform their jobs more efficiently, but using these technologies may require new skills that certain incumbent workers lack. Considering both of these channels paints a more nuanced picture of the effects of technology on the labor market.

Existing measures of workers' technology exposure have focused primarily on labor-saving technologies, with a particular emphasis studying the adoption of robots on employment.[1] While focusing on narrow measures of automation helps us understand the effects of particular technologies, they paint an incomplete picture of the impact of labor-saving technologies on workers: Benmelech and Zator (2022) argue that investment in robots is small and highly concentrated in a few industries, and accounts for less than 0.3% of aggregate expenditures on equipment. In comparison, our measure of exposure to labor-saving improvements is broader and identifies technologies that relate to both blue- and white-collar workers.

In recent work, Mann and Püttmann (2023); Dechezlepretre, Hemous, Olsen, and Zanella (2023) provide measures of automation using patents by analyzing the textual description of the patent technology classifications. Compared to these automation measures, our approach has two advantages. First, we have a broader view of labor-saving technologies than pure automation alone, and second, we can identify the impact of a specific technology on a relatively narrow group of workers—those in a specific 4-digit NAICS code and a specific 6-digit occupation code—rather than the broad groups considered in previous studies (for instance, all workers in a given region or all production workers). This level of granularity allows us not only to partial out industry- and occupation-specific trends in our empirical analysis, but also to differentiate between direct effects on worker productivity for workers whose tasks are complemented or substituted by a specific technology, and to account for indirect effects through shifts in labor demand for workers whose tasks are not directly affected. Our measure has some similarities to Webb (2020), who also analyzes the similarity between patents identified as being related to robots, AI, or software and occupation task descriptions. However, there are some key differences: our measure differentiates between technologies related to workers' routine or non-routine tasks, which allows us to distinguish between labor-saving and labor-augmenting technologies, and varies within occupation and over calendar time as different patents are issued to different industries at different points in time.

In addition to proposing a new measure of technology exposure, a key advantage of our approach is the use of panel data on individual worker earnings. Linking a granular measure of technology exposure to the earnings of individual workers allows us to analyze the heterogeneous impact

[1] An incomplete list includes Acemoglu and Restrepo (2020, 2022); Graetz and Michaels (2018); Humlum (2019); Dauth, Findeisen, Suedekum, and Woessner (2021); Koch, Manuylov, and Smolka (2021); Bonfiglioli, Crinò, Fadinger, and Gancia (2020); de Souza and Li (2023). Among these, the work closest to ours is de Souza and Li (2023), who studies the adoption of both labor-saving and labor-augmenting technologies—although the focus of their work is on employment and wages at the occupation level rather than the earnings of individual workers.

of technology on workers while avoiding some of the pitfalls of existing work that focuses on aggregate wage earnings, which can mask significant heterogeneity and is affected by shifts in worker composition. Nevertheless, there is some existing work that studies the impact of labor-saving technologies on individual workers. These studies have largely focused on European labor markets over short time periods.[2] The main focus thus far has been on the impact of firms' adoption of robots (Humlum, 2019), or automation technologies (Bessen, Goos, Salomons, and van den Berge, 2023). Other work exploits more aggregate sources of variation (Dauth et al., 2021). The conclusion of this literature is somewhat mixed: Humlum (2019) and Bessen et al. (2023) find negative effects of automation for production workers, especially older workers; Dauth et al. (2021) finds a positive effect on the earnings of incumbent workers but negative effects on new entrants. Examining a broad range of labor-saving technologies over a longer period in the United States, we find negative effects on workers directly affected by these technologies, across both blue- and white-collar workers, with both the youngest and oldest workers experiencing similarly adverse effects.[3] Moreover, in addition to these negative effects we also document a positive spillover effect of industry innovation on other workers in the same industry whose occupations are not directly affected.

In contrast to labor-saving technologies, the impact of labor-augmenting technologies on individual worker earnings has received significantly less attention. Akerman, Gaarder, and Mogstad (2015) investigates the adoption of broadband internet on the earnings of Norwegian workers and finds evidence in favor of skill complementarity. Autor, Chin, Salomons, and Seegmiller (2022) builds on the methods introduced in this paper to develop a measure of labor-augmenting technology exposure based on Census job description write-ins; however, their focus is on analyzing the factors driving demand for new tasks and their implications for aggregate employment rather than individual workers. Kalyani, Bloom, Carvalho, Hassan, Lerner, and Tahoun (2021) show that new technologies are associated with an increase in job postings for both high- and low-skill workers. We contribute to this literature by documenting the heterogeneous impact of labor-augmenting technologies—defined as technologies that complement labor in its existing tasks. While these technologies increase the earnings of all workers in the same industry and occupation, this increase is primarily concentrated among new entrants, with the earnings of incumbents declining, especially the oldest and most highly paid workers. These results support the view that human capital is partly specific to a

[2] A notable exception is Feigenbaum and Gross (2020), who finds that women were more likely to be in lower-paying occupations following the adoption of mechanical switching technology by AT&T (a labor-saving technology) in the early 20th century in the United States.

[3] To reconcile our estimates with existing studies based on European data, one possibility is that institutional differences between U.S. and European labor markets may lead to different effects on worker earnings. Studies of automation based on aggregate data also reveal sharp differences between these two regions. For example, Acemoglu and Restrepo (2020, 2022) argue that automation has resulted in large negative demand shifts in the United States. In contrast, Graetz and Michaels (2018) find a positive impact of automation technologies on productivity across a sample of 17 countries and little evidence of a negative effect on local labor demand, while Aghion, Antonin, Bunel, and Jaravel (2021) reports similar findings at the firm level in France.

particular technology vintage.[4]

Along these lines, our work complements the existing empirical evidence emphasizing the displacing effect of new technologies (Deming and Noray, 2020; Braxton and Taska, 2020; Hombert and Matray, 2021; Atack, Margo, and Rhode, 2019, 2022). Closest to our paper is Braxton and Taska (2020), who show that displaced workers in occupations facing changing skill requirements fare worse than other displaced workers; Deming and Noray (2020), who finds that the wage premium of jobs with frequent changes in skill demands (STEM) declines with worker age; and Hombert and Matray (2021) who show that French workers starting out in the ICT sector at the time of the 1990s internet boom had large declines in wages relative to other sectors. Additionally, our results relate to Dixon, Hong, and Wu (2021) who find that firms adopting robots also reduce the number of managers they employ.

More generally, our work is connected to the literature aiming to understand the complementarity between capital (broadly defined) and labor. Existing work emphasizes the complementarity between technology and certain types of worker tasks or skills (Goldin and Katz, 1998; Autor et al., 2003; Autor, Katz, and Kearney, 2006; Goos and Manning, 2007; Autor and Dorn, 2013; Autor et al., 2022) or the substitution between workers and capital (Krusell, Ohanian, Ríos-Rull, and Violante, 2000; Hornstein, Krusell, and Violante, 2005; Karabarbounis and Neiman, 2013; Hemous and Olsen, 2021; Caunedo et al., 2023). The main focus of these papers is on how technology affects differences in wages between groups with different ex ante skill levels (typically education). Among these papers, the closest to our work is by Caunedo et al. (2023), who allows for heterogeneity: some types of capital substitute for labor, while others are complements. Overall, our findings provide direct evidence consistent with the view in Autor et al. (2003) regarding the potential for labor-saving technologies to substitute for labor in routine cognitive tasks.

1 Model

We motivate our empirical analysis with a model that captures the key channels through which technology affects the earnings of individual workers. Technological progress is embodied in new vintages of capital. Our model builds on Acemoglu and Restrepo (2018) in that workers perform different tasks and on Caunedo et al. (2023) in that capital improvements are specific to a particular set of tasks and can be either complements or substitutes. Thus, our model allows for both labor-saving and labor-augmenting technology improvements, depending on whether these improvements reflect capital-specific technological changes to routine or non-routine tasks performed by workers.

[4]Models of vintage-specific human capital imply that workers who have accumulated greater levels of skill in existing technologies are more likely to be displaced when newer vintages of technology arrive (Chari and Hopenhayn, 1991; Jovanovic and Nyarko, 1996; Violante, 2002).

1.1 Setup

The model has two periods; in the second period, there is a stochastic improvement in technology, which we model as a decline in the (quality-adjusted) price of capital. To economize on notation, we omit time subscripts and denote log growth rates across periods by Δ. Unless otherwise noted, our discussion refers to pre-shock variables. Appendix A contains detailed derivations.

Output Aggregate output (the numeraire good) is produced by a continuum of competitive industries,

$$\bar{Y} = \left(\int_k Y(k)^{\frac{\chi-1}{\chi}} dk \right)^{\frac{\chi}{\chi-1}} . \tag{1}$$

There is free entry of all firms in the production sector. In equilibrium, the price of each industry's output is equal to its marginal cost and firms make zero profits.

The output of a single industry $Y(k)$ is a CES aggregate of a large number of intermediate tasks

$$Y(k) = \left(\sum_{j=1}^{J} y(j,k)^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}} . \tag{2}$$

The parameter $\psi > 0$ indexes the elasticity of substitution across tasks and the absolute value of the demand elasticity for each task output. To simplify exposition, we focus on a given industry and drop the k subscripts unless needed. Task j is produced using capital $n(j)$ and labor $l(j)$

$$y(j) = \left((1-\gamma_j)\, n(j)^{\frac{\nu_j-1}{\nu_j}} + \gamma_j\, l(j)^{\frac{\nu_j-1}{\nu_j}} \right)^{\frac{\nu_j}{\nu_j-1}} . \tag{3}$$

Our notion of capital $n(j)$ incorporates not only machines used in production but also process methods or software. The parameter $\nu_j > 0$ determines the elasticities of substitution between $l(j)$ and $n(j)$ when performing each task, and $\gamma_j \in (0,1)$ governs factor shares.

We allow for both labor-saving and labor-augmenting technology improvements. To capture this distinction, we partition the set of tasks (2) into two types, which we refer to as routine and non-routine, $J = J_R \cup J_N$, and assume that γ_j and ν_j is constant within each set. At this point, routine and non-routine are simply labels we attach to these two types of tasks that differ only in the elasticity of substitution between capital and labor. We expect that the degree of capital–labor complementarity is higher for non-routine than routine tasks: $\nu_N < \nu_R$. As a result, labor-augmenting technologies complement workers' non-routine tasks; for these tasks, improvements embodied in capital reflect improvements in the tools that workers use in their job, which in principle should enhance their productivity. By contrast, labor-saving technologies substitute for workers' routine tasks; for these tasks, improvements in the quality of capital will lower its effective price and therefore lead to greater degrees of substitution by capital.

Capital and Technology. Firms choose the demand for capital deployed in each task taking prices $q(j)$ as given. The supply of capital suitable for task j is perfectly elastic. Technological innovation consists of declines in the (quality-adjusted) price of capital,

$$\Delta \log q(j) = -\varepsilon(j). \tag{4}$$

Technological progress can simultaneously affect multiple tasks: $\varepsilon \equiv [\varepsilon_1 \dots \varepsilon_J]$ is a vector of weakly positive random variables jointly distributed according to $f(\varepsilon)$, where draws of ε are independent and identically distributed across industries. Under this assumption, the evolution of aggregate output (1) is deterministic.

Labor Supply. In the initial period, there is a continuum of incumbent workers of measure I who supply labor across the different tasks j. The aggregate supply of labor in task j is given by:

$$L(j) = \int_0^I l(i,j)di, \tag{5}$$

where $l(i,j)$ is the number of efficiency units of labor supplied by incumbent worker i in task j, and $L(j)$ is the aggregate supply of labor in task j from workers who are incumbents in the first period. For now, we regard these as constant.

Occupations. Each incumbent worker is associated with a single occupation $o(i)$ and industry $k(i)$. Specifically, workers in occupation o supply labor in only a small number of routine and non-routine tasks: we denote by $J_o \subset J$ the set of tasks performed by occupation o. As in Acemoglu and Restrepo (2022), each task can only be performed by a single occupation, so $\bigcup J_o = J$ and $\bigcap J_o = \emptyset$. For each incumbent worker, we assume that frictions to switching occupations or industries are sufficiently large that her occupation and industry is fixed in both periods. For simplicity, we assume that in the first period each worker has the same skill in each task she performs

$$l(i,j) = \bar{l}(i), \quad \forall j \in J_{o(i)} \tag{6}$$

that is, initial productivity differences across workers only reflect absolute advantage—and the distribution of $\bar{l}(i)$ is i.i.d. across occupations. We allow for the aggregate supply of labor for task j to also vary on the extensive margin. Rather than explicitly modeling the flow of workers, we assume that aggregate labor supply satisfies

$$\Delta \log L(j) = \bar{\zeta} + \zeta_j \Delta \log w(j)\,. \tag{7}$$

1.2 Worker earnings growth and technology exposure

The level of wage earnings for an individual worker is equal to the total compensation for the tasks she supplies. An individual worker i's earnings growth evolves according to

$$\begin{aligned}\Delta \log W(i) \approx & \left[\frac{\psi - \nu_R}{\nu_R + \zeta_R}\, \Gamma_R\right] \xi^R(i) + \left[\frac{\psi - \nu_N}{\nu_N + \zeta_N}\, \Gamma_N\right] \xi^N(i) \\ & + \Big[A_N + [A_R - A_N]\, \theta(i)\Big] \left[(\chi - \psi)\Delta \log X + \Delta \log \bar{Y}\right], \end{aligned} \tag{8}$$

where X is industry productivity (number of units produced per of dollar of input expenditure), Γ_j is the elasticity of the marginal cost of each task $p(j)$ with respect to improvements in capital specific to that task $q(j)$, and A_j is the elasticity of task-level wages $w(j)$ to aggregate productivity $\bar{Y}$. To obtain (8), we assume that labor supply elasticities and factor shares are constant within the set of routine and non-routine tasks. Appendix A contains additional details.

A worker's exposure to labor-saving (R) and labor-augmenting (N) technologies is given by

$$\xi^R(i) \equiv \theta(i) \sum_{j \in J_R} \tilde{s}^R(i,j)\, \varepsilon(j) \quad \text{and} \quad \xi^N(i) \equiv (1 - \theta(i)) \sum_{j \in J_N} \tilde{s}^N(i,j)\, \varepsilon(j). \tag{9}$$

Here, $\theta(i)$ denotes the share of labor compensation to worker i due to performing the routine tasks, and the weights $s(i,j)$ denote the share of task j in the earnings of worker i,

$$\theta(i) \equiv \sum_{j \in J_R} s(i,j), \qquad s(i,j) \equiv \frac{w(j)\, l(i,j)}{W(i)}\, I(j \in J_o). \tag{10}$$

Last, $\tilde{s}^R$ and $\tilde{s}^N$ are the task weights normalized to sum to one within the worker's set of routine or non-routine tasks, respectively.

Ideal Exposure Measure. Equation (9) states that a worker's exposure to labor-saving technologies is a function of the contribution of routine tasks to her compensation θ and the extent to which specific technological improvements are related to her routine tasks. Similarly, the worker's exposure to labor-augmenting innovation is related to the contribution of non-routine tasks to her compensation $1 - \theta$ and the degree to which technological improvements apply to her non-routine tasks. To fully implement (9), we would need data on compensation at the task level and the task-specific skills of individual workers. This data is not available, so we construct our exposure measures at the occupation, rather than at the individual worker, level.

Equation (9) highlights the distinction between ours and existing work that focuses on measuring workers' exposure to labor-saving technologies based on variation in the share of routine tasks $\theta(i)$ across occupations (Autor et al., 2003). An occupation's routine share captures the *potential* for these tasks to be automated. By contrast, the ideal exposure measure (9) captures the extent to which actual technological improvements substitute (or complement) for labor in these tasks. Thus, a key advantage of our approach is that we can focus on variation in technology exposure while holding constant an occupation's routine share $\theta(i)$.

Testable Predictions. Equations (8) and (9) form the foundation of our empirical analysis. The first term in (8) captures the direct effect of labor-saving improvements on worker earnings—technology improvements $\varepsilon(j)$ in capital specific to routine tasks. The sign of the elasticity of task compensation $w(j)$ to quality improvements $\varepsilon(j)$ depends on the sign of $\psi - \nu_j$, where ν_j is the elasticity of substitution between capital and labor and ψ, which is equal to the elasticity of demand for that task (Hicks, 1932). If $\nu_j > \psi$, the high substitutability between cheaper capital and labor more than offsets the increase in labor demand for task j, which leads to lower wages for task j. By contrast, if $\nu_j < \psi$, demand for the output of task j is sufficiently elastic to lead to higher labor demand and therefore wages. For routine tasks $j \in J_R$, a reasonable parametrization is that $\nu_R >> \psi$, so improvements in labor-saving technologies will lower worker earnings. The second term in (8) captures the effect of labor-augmenting technologies on the compensation of labor assigned to the same task. Its sign is ambiguous, as it depends on whether the complementarity across tasks ψ is greater or smaller than the complementarity between capital and labor in non-routine tasks ν_N. Since most estimates of ψ are close to one (see, e.g., Humlum, 2019), we would expect this term to be positive.

The last two terms in (8) capture the impact of technology spillovers on firm labor demand. Worker earnings grow with changes in industry $\Delta \log X$ and aggregate productivity $\Delta \log \bar{Y}$ due to increases in labor demand. Unlike the direct effects discussed above, the strength of this channel is uniform across occupations and is a function solely of the occupation's routine share $\theta(i)$—this heterogeneity disappears if $A_N = A_R$. In the first part of the paper, we include dummies for granular industry classifications interacted with calendar year in our estimation, and hence focus on the direct effect of technology on exposed workers. We revisit these spillover effects of innovation in Section 5.

2 Data and Measurement

Here, we discuss the data and methodology used in our empirical analysis.

2.1 Methodology

Our goal in this section is to construct a measure of technology exposure based on the distance between the description of a technology in a patent document and the tasks performed by an occupation. Occupations perform several tasks, and the labor input in each task may be either complemented or substituted by new technologies. Constructing an empirical analogue of Equation (9) requires us to overcome three challenges: 1) identifying which tasks are likely to be complemented or substituted by technology; 2) identifying the arrival of new technologies in the data; and 3) determining whether a specific technology complements or substitutes the tasks performed by a specific occupation depending on the textual similarity between the description of the technology in the patent document and the two disjoint sets of tasks identified in the first step. We next discuss our approach to addressing these challenges.

2.1.1 Classifying tasks based on their capital/labor complementarity

Most occupations perform a combination of tasks, and each task can be substituted or complemented by capital. To construct a measure of exposure to labor-saving and labor-augmenting technologies, we therefore need to classify each task a worker performs as either substituted or complemented by capital (what we refer to as R and N in the model, respectively). To separate tasks into these two types, we rely on two ideas in the literature: whether the task involves performing routine actions or whether the task requires a significant amount of prior experience. To partition the set of tasks performed by a set of occupations into these groups, we rely on recent advances in generative AI (GPT4) applied to the description of job tasks a given occupation performs from the Dictionary of Occupational Titles (DOT). To conserve space, we summarize our methodology here and relegate all details to Appendix B.2.

Our first approach to identifying which tasks are complemented or substituted by technology relies on the notion of routine tasks (Autor et al., 2003; Acemoglu and Autor, 2011; Autor and Dorn, 2013). The key idea is that routine tasks can be codified into repeatable instructions, which implies that technologies such as software or robotics could in principle be programmed to do the task instead of a worker. Thus, routine tasks are those that *could* be performed by machines. By contrast, non-routine tasks cannot be performed by (tangible or intangible) capital alone; instead, capital increases workers' productivity when performing these tasks. Autor et al. (2003) and Autor and Dorn (2013) argue that routine tasks were particularly exposed to labor-saving technological change in our sample (1980-2016).

To classify each task as routine or non-routine, we query GPT4 on whether a given task can be characterized as routine or non-routine. GPT4 tags approximately 62% of the tasks performed by all occupations as routine (we tag the remainder as non-routine). To validate this classification of tasks into routine and non-routine, we collapse the task scores at the occupation level and compare the fraction of routine tasks performed by a specific occupation to the RTI score constructed by Acemoglu and Autor (2011). As we see in Panel A of Figure 1, the two occupation-level metrics are highly correlated (81%).

As a second measure of the degree of substitution or complementarity between technology and tasks, we also construct an alternative task partition based on the degree of required prior experience that relies on the notion of skill-biased technical change (Goldin and Katz, 1998; Krusell et al., 2000). The assumption here is that tasks that require high levels of prior expertise are more likely to require workers to exercise critical thinking, creativity, or judgment calls—and therefore these tasks are less likely to be mundane and repeatable and thus less likely to be substituted by capital. To classify each task into a high- vs. low-skill task, we again rely on GPT4, which tags approximately 61 percent of all tasks as low-skill (low required experience). To validate this classification, we compare the fraction of high-skill tasks performed by a given occupation to ONET's job zone classification, which assigns occupations a score between 1 (lowest) and 5 (highest) depending on the extent of their required preparation, which includes education requirements, related experience, and on-the-job training. As we see in Panel B of Figure 1, there is a strong correlation (85%) between the share of tasks classified as high-skill by GPT4 and ONET's job classification at the occupation level.

Overall, this classification of tasks into routine/non-routine or high-/low-skill will allow us to construct two separate measures of technology exposure for each worker depending on which part of her tasks based on this classification are most related to new technologies at a given point in time. These two classifications are highly correlated: approximately 70 percent of the tasks we classify as non-routine are also classified as high skill, while over 81 percent of the tasks we classify as routine are also classified as low skill. Our operating assumption is that technologies that relate to routine (or low-skill) tasks are labor-saving, whereas technologies related to non-routine (high-skill) tasks are labor-augmenting. Whether this assumption is appropriate is ultimately an empirical question.

2.1.2 Breakthrough innovations

Next, we need to identify major improvements in technology. Technology improvements in the model are embodied in capital—they correspond to declines in the quality-adjusted price of capital goods. However, the definition of capital can be broad: it can relate not only to physical capital (for example, machines) but also intangibles (e.g., production methods or software). To identify these innovations, we rely on patent data and follow the methodology of Kelly et al. (2021), henceforth KPST. KPST identify breakthrough innovations as those that are both novel (whose descriptions are distinct from their predecessors) and impactful (they are similar to subsequent innovations). In particular, KPST first create a measure of importance for each patent that combines novelty and impact and then define a 'breakthrough' patent as one that falls in the top 10% of the distribution of importance. KPST show that these breakthrough technologies are associated with increases in measured productivity at both the aggregate and industry level. KPST provide several indices; given that our administrative data sample has a shorter time dimension, we use the breakthrough definition of KPST that relies on 5-year forward similarity.

2.1.3 Measuring similarity between occupation tasks and patents

The last step in constructing our measure requires us to estimate the distance between a given technology and the set of routine and non-routine tasks performed by a specific occupation. To calculate the semantic similarity between these two sets of documents, we use word embeddings provided by Pennington, Socher, and Manning (2014). Specifically, we represent each document (either a patent or a part of an occupation task description) as a weighted average of the set of word vectors x_k for the terms contained in the document:

$$\mathcal{X}_i = \sum_{x_k \in A_i} w_{i,k}\, x_k. \tag{11}$$

The weights $w_{i,k}$ are based on the 'term frequency inverse document frequency' (TF-IDF), which overweights word vectors for terms that occur relatively frequently within a given document and underweights terms that occur commonly across all documents. Next, we calculate the cosine similarity between a breakthrough patent b and the routine or non-routine component of the task description of occupation o,

$$\rho^j(b,o) = \frac{\mathcal{X}_b}{||\mathcal{X}_b||} \cdot \frac{\mathcal{X}_o^j}{||\mathcal{X}_o^j||}, \qquad j \in \{R, N\}. \tag{12}$$

Our measure that is based on the high- and low-skill classification of tasks is constructed analogously. We perform two adjustments to (12). First, we remove year fixed effects to account for language and structural differences in patent documents over time; patents have become substantially longer and use considerably more technical language over the sample period. Second, we impose sparsity to focus only on high degrees of similarity. Specifically, after removing the fixed effects we set all patent–occupation pairs that are below the 80th percentile to zero in this fixed-effect-adjusted similarity. Thus, only similarity scores sufficiently high in the distribution receive any weight. Last, we scale the remaining nonzero pairs such that a patent/occupation pair at the 80th percentile of yearly adjusted similarities has a score equal to zero and the maximum adjusted score equals one. We denote by $\tilde{\rho}^j(b,o)$ the adjusted similarity metric. We also repeat the exercise for the alternative classification of tasks into high- and low-experience.

In sum, we propose a method that uses a combination of word embeddings and TF-IDF weights in constructing a distance metric between a patent document (which includes the abstract, claims, and the detailed description of the patented invention) and the two subsets of the detailed description of the tasks performed by occupations that we believe a priori are likely to be complemented or substituted by technology. Appendix B.3 contains further details.

2.2 Examples

Here, we briefly discuss examples that illustrate the key features of our technology exposure measures—Appendix Table A.1 includes additional examples.

An example of a likely labor-saving technology is patent 6,044,352 for "Method and system for processing and recording the transactions in a medical savings fund account", which is most closely related to the routine part of the job description for insurance claims processing clerks. Other examples of labor-saving technologies include patents 5,737,710 and 6,796,499 for automated payment systems for parking facilities, which are most closely related to the routine tasks of parking lot attendants. As an example of a labor-augmenting technology, consider patent 5,189,606 for a system for estimating construction costs, which is the patent most closely related to the non-routine tasks performed by architects. Another example is radio operators. Among the technologies that are most closely related to their non-routine task descriptions are patents 5,123,112 and 5,212,804 for air-to-ground communication systems, which are likely complementary technologies.

Sometimes, new technologies are related to both the routine and the non-routine tasks of an occupation, so it is not obvious whether they are overall complements or substitutes. For instance, US Patent 5,870,721 for an automated loan processing system is related to both the routine and non-routine task components of loan officers. Last, some innovations can be simultaneously labor-saving for some occupations and labor-augmenting for others. For example, consider patent 4,683,202 for a "Process for amplifying nucleic acid sequences." The closest occupation in terms of its non-routine tasks is biochemists, and our method will therefore interpret this technology as complementary to these workers. By contrast, the closest occupations in terms of their routine tasks are medical and clinical laboratory technologists and technicians; thus, our procedure will treat this technology as labor-saving for these occupations.

It is important to emphasize that each individual patent likely has a small impact on worker earnings; given that we are aggregating across thousands of patents, the time-series behavior of our measure will be driven by a collection of related technologies. For instance, one labor-saving technology captured by our measure is the rise of e-commerce—specifically the automatic fulfillment of retail purchase orders. As we see in Figure 2, our labor-saving exposure measure indicates that the 1990s featured a significant surge in innovation related to the routine component of the tasks performed by order-fulfillment clerks.[5] Subsequent to this wave of innovations related to the routine component of order processing clerks, the hourly wages of these clerks declined by approximately 20 percent compared to all other clerical occupations.

[5] Examples include patents 5,696,906 and 5,884,284 for "Telecommunication user account management system and method"; patent 5,892,900 for "Systems and methods for secure transaction management and electronic rights protection"; patent 5,848,396 for "Method and apparatus for determining behavioral profile of a computer user"; and patent 5,627,973 for "Method and apparatus for facilitating evaluation of business opportunities for supplying goods and/or services to potential customers".

Figure 2 illustrates the advantages and disadvantages of our approach. The key advantage of our approach is that we are able to identify time variation in the exposure of specific occupations to labor-saving or labor-augmenting technologies. The key disadvantage, which is apparent in this figure, is that technology generally improves gradually, so there are no clear structural breaks in our exposure measures that can be identified with the arrival of a specific technology. A separate but compounding factor is that, even if the technology did arrive in a discrete fashion, we lack direct data on adoption. In particular, the actual implementation of the technology can occur earlier, if the firm developing the technology deploys it before the patent application is granted, or later, if there are lags in firms broadly adopting the technology. In our empirical analysis, we settle on the patent issuance year as our baseline, but we recognize that there is likely significant measurement error in the exact timing, which will smooth the time-series behavior of our exposure measures. The lack of observed clear structural breaks can raise concerns about identification, since the time-series behavior of the measures that we construct could simply reflect slow-moving trends in the underlying labor markets.

In our empirical analysis, we address this identification challenge in multiple ways. First, we exploit cross-industry variation in the rate of innovation to isolate our effects from industry- or occupation-specific time trends. Second, we employ a shift-share IV approach, using past breakthrough innovations to predict future labor-saving or labor-augmenting innovations, building on Acemoglu et al. (2016). Third, we focus on individual worker outcomes, which allows us to include a rich set of worker-specific controls and directly measure individual earnings. Focusing on individual workers also lets us avoid the bias in average earnings data (as in Figure 2) due to changes in worker composition—for example, average wages can rise if the lowest-paid workers exit an occupation following the arrival of a labor-saving technology.

2.3 Worker technology exposure

The final step is to create worker-level exposure metrics. Our data contain detailed information on both the industry of a particular worker and the industry of origination of each patent. This allows us to exploit additional sources of variation: we can not only compare workers across occupations in the same industry but also in the same occupation across industries. We denote by b breakthrough patents; $\mathcal{B}_{k,t}$ denote the set of breakthrough patents issued in industry k in year t; o denote occupations; and k denote industry at the NAICS4 level.

We measure the exposure of workers in industry k and occupation o to technology at time t as

$$\xi^j(k,o,t) = \theta^j(o)\,\log\left(1 + \sum_{b\in\mathcal{B}_{k,t}} \tilde{\rho}^j(o,b)\right), \qquad j \in \{R,N\}. \tag{13}$$

Our measure $\xi^j(k,o,t)$ is the empirical analogue of (9); it varies over time t and industry k due to the differential arrival of breakthrough technologies across industries and it varies across occupations o as these breakthrough technologies have different levels of similarity with the tasks performed by each occupation. Specifically, $\xi(k, o, t)$ aggregates our patent-occupation similarity scores $\tilde{\rho}^j(o, b)$ across all breakthrough patents $\mathcal{B}_{k,t}$ issued to a firm in industry k in period t. We match each patent to the industry of the patent assignee at the 4-digit NAICS level using restricted access Census data. Given that just under half of the occupation–industry pairs have zero breakthrough patents in a given year, we apply a log transform to smooth out the skewed distribution. As in Equation (9), $\xi^j(k, o, t)$ is proportional to the shares θ_o^R and θ_o^N of routine or non-routine tasks performed by occupation o, respectively. As we discussed previously, we date breakthrough patents by the year when the patent is issued to a firm in industry k.

Our measure $\xi^j(k, o, t)$ is the empirical analogue of (9); it varies over time t and industry k due to the differential arrival of breakthrough technologies across industries, and it varies across occupations o because these breakthrough technologies have different levels of similarity with the tasks performed by each occupation. Specifically, $\xi(k, o, t)$ aggregates our patent–occupation similarity scores $\tilde{\rho}^j(o, b)$ across all breakthrough patents $\mathcal{B}_{k,t}$ issued to a firm in industry k in period t. We match each patent to the industry of the patent assignee at the 4-digit NAICS level using the Census links from patents to establishments. Given that just under half of the industry–occupation pairs have zero breakthrough patents in a given year, we apply a log transform to smooth out the skewed distribution. As in Equation (9), $\xi^j(k, o, t)$ is proportional to the shares θ_o^R and θ_o^N of routine or non-routine tasks performed by occupation o, respectively. As discussed previously, we date breakthrough patents by the year when the patent is issued to a firm in industry k.

Overall, the key innovation in our approach is that it allows us to identify periods through which technology allows for specific tasks of a given occupation to be complemented or substituted by technology (which is a function of the distance between new technologies and different aspects of the job description) rather than whether the tasks of a given occupation *could* be automated (which is a function of whether the occupation mainly performs specific types of tasks).

Figure 3 shows that the exposure of different types of occupations (different rows) has varied over time across broad technology classes (different columns). Examining the figure reveals several key patterns. First, across most technology classes, blue-collar workers (workers in occupations emphasizing manual tasks) tend to be more exposed to labor-saving than labor-augmenting technologies, and the gap has increased over time in electronics and IT (columns two and four). By contrast, the opposite pattern obtains for workers in occupations that emphasize interpersonal skills. Third, and more importantly, the second row reveals that workers in occupations emphasizing cognitive tasks are, on average, roughly equally exposed to labor-saving and complementary technologies. Over time, these workers face increasing exposure to labor-saving technologies coming from ICT (electronics and IT). Among all white-collar jobs, office workers are among the occupations most exposed to labor-saving technologies. By contrast, workers in high-skill service jobs, especially business and STEM occupations, are relatively more exposed to complementary technologies than technologies that substitute for labor. Overall, this figure offers caution when interpreting some technologies as always labor-saving or labor-augmenting, given the significant heterogeneity across both time and types of jobs.

Figure 4 shows how our technology exposures vary across occupation wage levels. We see that the most exposed occupations to labor-saving technologies tend to be found in the middle of the income distribution, consistent with the prevailing view regarding job polarization in the United States and in Europe (Autor and Dorn, 2013; Jaimovich and Siu, 2020; Goos et al., 2014; Bárány and Siegel, 2018). By contrast, the exposure of occupations to labor-augmenting technologies is increasing in their average wages. The patterns are in line with the traditional view of skill-biased technical change, in the sense that labor-augmenting technologies are more related to higher-paying occupations.

2.4 Administrative employer–employee data

We combine employer-employee matched Social Security Administration (SSA) administrative data (the Detailed Earnings Records file) with survey information for a random sample of individual workers tracked by the Current Population Survey (CPS) Annual Social and Economic Supplement (ASEC). We limit the sample to individuals who are older than 25 and younger than 55 years old, to periods where the CPS interview date is within the past 3 years so that the occupation information is relatively recent, and we restrict the sample to workers with as high attachment to the labor force by imposing the restrictions in Braxton, Herkenhoff, Rothbaum, and Schmidt (2021). Appendix B.9 contains further details on sample construction.

Our main variable of interest in our empirical analysis is the total growth in the earnings of an individual worker. To smooth out the impact of transitory earnings spikes, we follow Autor, Dorn, Hanson, and Song (2014) and Guvenen, Ozkan, and Song (2014) and consider the growth in average wage earnings, adjusted for life-cycle effects,

$$\Delta w^i_{t+h} \equiv w^i_{t+1,t+h} - w^i_{t-2,t} \qquad \text{and} \qquad w^i_{t,t+h} \equiv \log\left(\frac{\sum_{j=0}^{h} \text{W2 earnings}_{i,t+j}}{\sum_{j=0}^{k} D(\text{age}_{i,t+j})}\right). \tag{14}$$

Table 1 summarizes the sample. Our sample contains approximately 2.8 million person–year observations spanning the period from 1981 to 2016. Approximately 54% of the sample is male, and 34% of the observations correspond to workers with a four-year college degree. The median worker in the sample is 41 years old and earns approximately \$50k per year in 2015 dollars. There is considerable variation in worker earnings growth, and much of it is persistent—the standard deviation of (14) increases from 0.53 at the 5-year horizon to 0.61 at the 10-year horizon. The distribution of earnings is rather skewed: the average is equal to \$66k, while the 5th and 95th percentiles are equal to \$16k and \$152k, respectively. Importantly, much of these differences in pay (approximately 58 percent of the cross-sectional dispersion) exist within industry–occupation groups. These within-industry–occupation differences in wage earnings are only partly driven by firm heterogeneity: using data from the Longitudinal Business Database (LBD), we find that the firms that employ the workers in the top 5 percent of the within-industry–occupation earnings distribution have approximately 45% higher labor productivity than the firms that employ the workers in the bottom 25 percent. In our empirical analysis, we will interpret this sizable heterogeneity in worker earnings within an industry–occupation cell as reflecting within-job differences in skill across workers.

2.5 A shift-share instrument for worker technology exposure

One potential concern with our empirical analysis is that the arrival of new labor-saving or labor-augmenting technologies may be endogenous to the current state of the labor market for a particular type of worker. For instance, a scarcity of workers in a particular industry or occupation could incentivize the development of labor-saving technologies targeting these workers. To address this concern, we develop a shift-share identification strategy in the spirit of Bartik (1991); Blanchard and Katz (1992). In essence, our instrument predicts the exposure of workers in specific occupation/industry cells to labor-saving or labor-augmenting technologies based on knowledge spillovers from breakthrough innovations in other technology classes in the past. This builds upon ideas in Acemoglu et al. (2016), Liu and Ma (2021), and Bloom, Schankerman, and Van Reenen (2013), who show that technological progress can propagate through an innovation network, whereby prior innovations in one technology class spur downstream innovations in separate connected technology classes in the future. To conserve space we briefly describe the construction here. Appendix B.4 contains all details.

The construction of the shift-share entails two steps. The first step (the shift) involves predicting the arrival of breakthrough patents in a given tech class c at time t based on innovation in other tech classes in the past,

$$\lambda_{c,t,\tau} = \sum_{c' \neq c} \Omega_{c' \to c,t,\tau} \times I_{c',t-\tau}. \tag{15}$$

Here, $\Omega_{c' \to c,t,\tau}$ is a technology diffusion matrix constructed based on the textual similarity of patents: its elements are the average similarity of patents in technology class c to patents in technology class c' for patents issued in tech class c at time t and tech class c' at time $t-\tau$; hence τ represents a diffusion lag period for innovation to propagate from tech class c' to c. When constructing Ω, we set its diagonal to zero so that we only use spillovers to class c from other technology classes c'. $I_{c,t}$ gives the intensity of breakthrough patents in class c and year t, as measured by the share of patents that are breakthroughs at time t and in technology class c. Using (15), we predict $\mathrm{N}^B_{c,t}$, the number of breakthrough patents in time period t for tech class c, using a Poisson regression where the independent variable is the average value of $\lambda_{c,t,\tau}$ across diffusion lags $\tau = 5, \ldots, 20$; this measure is a strong predictor of subsequent breakthrough patenting in tech class c (t-stat of 6.61 in the Poisson model).

To construct the second step (the shares) we estimate the likelihood that a patent j from technology class c coming from industry k is related to the routine $j = R$ or non-routine $j = N$ tasks of occupation o:

$$\alpha_{o,k,c,t} = E_t[\tilde{\rho}^j(p, o) \times \mathbf{1}(p, k) | p \in B_{c,t}]. \tag{16}$$

Our shift-share instrument for the innovation exposure of occupation o in industry k is

$$Z^j_{o,k,t} = \theta^j(o) \times \sum_c \overline{\alpha}^j_{o,k,c,t-5} \times \widehat{N}^B_{c,t}, \tag{17}$$

where $\overline{\alpha}^j_{o,k,c,t-5}$ is the average value of (16) from $t-5$ to $t-10$ and $\theta^j(o)$ gives the occupation's share of routine or non-routine tasks as before. We choose a shorter lag here because we only observe the industry of origin for patents going back to 1976. Since the instrument Z in (17) has a very right-skewed distribution, we winsorize the right tail at the 5% level each year to limit the impact of outliers. Alternatively, taking the log of one plus the IV predicted patenting exposure leads to qualitatively and quantitatively similar results.

Our shift-share instrument (17) is aimed to address the concern that certain shocks within an occupation–industry cell may cause firms to optimally shift production away from certain kinds of labor, and the direction of innovation may simply move in parallel with such a trend, which would lead us mistakenly attribute the change in worker earnings to our measure of technology exposure. Our controls for industry—year and occupation—year dummies may only partially alleviate this concern. For our shift-share instrument to be valid, we would need that the knowledge spillovers (15) from 'upstream' technologies in the past are unrelated to these labor market shocks in occupation–industries that are exposed to these 'downstream' technologies.

3 Technology Exposure and Worker Outcomes

Here, we examine the relation between our measures of technology exposure and worker outcomes.

3.1 Technology exposure and aggregate outcomes

As a prelude to our worker-level analysis, we first examine the relation between our technology exposure measures and labor demand at the occupation–industry level. In particular, our model implies that the growth in employment in occupation o is a function of workers' technology exposure to labor-saving (R) and labor-augmenting (N) technologies,

$$\Delta \log WB(o) \approx (1+\zeta_R)\left[\frac{\psi-\nu_R}{\nu_R+\zeta_R}\,\Gamma_R\right]\xi^R(o) + (1+\zeta_N)\left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\,\Gamma_N\right]\xi^N(o) + \text{Spillover Terms.} \tag{18}$$

Examining (18), we see that, as long as $\nu_R > \psi$, total labor compensation should decline in occupations exposed to labor-saving technologies, whereas it should increase in occupations affected by labor-augmenting technologies if $\nu_N < \psi$.

We estimate the empirical analogue of (18) using data from the Decennial Census and American Community Survey (see Appendix B.9 for details)

$$\log(X_{o,k,t+h}) - \log(X_{o,k,t}) = \gamma\,\xi^R_{o,k,t} + \delta\,\xi^N_{o,k,t} + c\,\mathbf{Z}_{o,k,t} + u_{o,k,t}. \tag{19}$$

The dependent variable in (19) is the cumulative growth rate in the wage bill, employment, and average worker earnings over the next decade $h = 10$ at the occupation–industry–year level. Our two independent variables of interest are the technology exposure measures $\xi^j_{o,k,t}$ that we assign to workers in industry k in occupation o at time t constructed in (13). Our vector of controls $\mathbf{Z}$ includes occupation and industry (NAICS4) fixed effects interacted with calendar year. These fixed effects absorb sources of time-series variation that affect specific groups of workers in a given occupation or industry; as before, they also absorb the omitted terms in Equation (18) capturing cross-occupation spillovers, which allows us to focus our attention on the direct effects of labor-saving and labor-augmenting technologies on employment in affected occupations. In addition we include controls for the share of female workers in the occupation–industry cell, and the logarithm of the average years of education, worker age, average wage, and employment at the start of the period.

Table 2 reports the estimated coefficients γ and δ from Equation (19). Examining Column (1), we see that the signs of the estimated coefficients are consistent with our model and our interpretation of our technology exposure measures ξ^R and ξ^N as labor-saving and labor-augmenting, respectively. The estimate of γ implies that an increase from the median to the 90th percentile in the exposure to technologies related to workers' routine tasks ξ^R implies a 4.6 percent cumulative decline in labor compensation over the next decade—consistent with our interpretation of these technologies as labor-saving. By contrast, moving from the median to the 90th percentile in exposure to technologies related to workers' non-routine tasks ξ^N implies a 9.1 percent increase in total labor compensation over the next decade, which is consistent with our interpretation that technologies related to workers' non-routine tasks are complementary to labor in these tasks.

Column (2) of Table 2 reports the estimated coefficients γ and δ from the same specification, where we now use our shift-share instruments Z^R and Z^N in (17) to instrument for ξ^R and ξ^N. We note that the estimated coefficients are not only qualitatively but also quantitatively similar between the IV and OLS specification—even though the OLS coefficients are more precisely estimated. The fact that the estimated coefficients are close in magnitude helps assuage somewhat concerns about the endogeneity of ξ to labor market conditions. Last, we note that our shift-share IV is strongly predictive of our exposure measures (the F-statistics are well over 600) which assuages concerns about weak instruments—see Column (1) of Appendix Table A.2.

The remaining four columns of Table 2 decompose the effect on total compensation into the effect on employment and average worker earnings. Focusing on exposure to labor-saving innovations ξ^R, we see that the decline in the total wage bill is driven by a decline in both employment (2.7 to 3.9 percent) and modest decline in average earnings (0.7 to 1.7 percent) over the next decade. By contrast, the increase in the total labor compensation in response to our exposure measure to labor-augmenting innovations is essentially driven by an increase in employment since average earnings do not show a statistically significant response.

The muted response of average worker earnings in columns (5) and (6) of Table 2 to our technology exposure measures can be surprising when viewed through the lens of our model. In particular, if labor-saving or labor-augmenting technologies significantly affect the marginal product of labor, we would expect to see a significant response of average worker earnings. However, focusing on average wage changes can be misleading. For instance, if workers transition across industry—occupation pairs in response to increases in technology exposure, then the response of average wages will be affected by shifts in worker composition. To see this, consider the following decomposition in the change in the log total wage bill in a given occupation–industry cell,

$$\Delta \log WB_{t+1} \approx s_t \, \Delta \log \bar{w}^s_{i,t+1} + \frac{E_{t+1}}{N_t} \left(\log \bar{w}^e_{i,t+1} - \log \bar{w}_{i,t} \right) - \frac{X_t}{N_t} \left(\log \bar{w}^x_{i,t} - \log \bar{w}_{i,t} \right) + \Delta \log N_{t+1}. \tag{20}$$

Here, s_t denotes the wage-bill share of stayers, N_t denotes employment, E_t denotes the number of entrants and X_t denotes the number of exiting workers. The first three terms correspond to the change in the average wage earnings at the occupation level, while the last term corresponds to the change in employment. The first term captures the earnings dynamics of continuing workers (stayers) multiplied by their share in the total wage bill. The second term is equal to the share of new entrants times the log difference between the wage earnings of new entrants relative to the average wage earnings in the occupation. The third term captures the impact of worker exit: average earnings in the occupation–industry cell can vary if workers with wage earnings above or below the average exit the cell.

3.2 Technology exposure and individual workers

Equation (20) clearly shows that focusing on aggregate wage earnings can obscure the impact of technology on individuals. For example, while some workers may see wage increases due to higher productivity, others may experience wage stagnation or even job loss. Therefore, to truly understand the impact of technology on worker earnings, we next focus on individual workers.

3.2.1 Incumbent workers

Using individual data on worker earnings, we estimate the following specification,

$$\Delta w^{i}_{t+h} = \gamma\, \xi^{R}_{i,t} + \delta\, \xi^{N}_{i,t} + c\,\mathbf{Z}_{i,t} + u_{i,t}. \tag{21}$$

The dependent variable in (21) is the growth in worker i's average W2 earnings defined in (14) over the next five years, relative to the prior three years.

Equation (21) maps to Equation (8) in the model that describes the earnings growth of individual workers. We are interested in the estimated coefficients γ and δ on our technology exposure measures $\xi^{j}_{i,t}$ that we can assign to an individual worker i at time t based on her occupation $o(i)$ and the (NAICS 4-digit) industry $k(i,t)$ of the worker in each year (13). These coefficients are identified by comparing future earnings for a highly exposed worker to less-exposed workers that are the same age, have the same level of earnings and prior earnings history, and either work in the same industry in a different occupation, or in the same occupation in other industries. Specifically, the vector of controls $\mathbf{Z}$ includes flexible nonparametric controls for worker age and the level of past worker earnings as well as recent earnings growth rates.[6] The first four columns of Panel A of Table 3 report the estimated coefficients γ and δ for different combinations of year, occupation, worker prior income, and industry fixed effects.

The contrast between Table 3 and Table 2 reveals two of the main findings of our paper. First, as before, the impact of labor-saving innovations on worker earnings is negative: an increase in ξ^{R} from the median to the 90th percentile is associated with approximately a 2.2 to 2.9 percent decline in worker cumulative earnings over the next five years. To put these estimates into perspective, note that the peak-to-trough change in average 5-year earnings growth during the Great Recession was 5.2 percentage points (see Global Repository of Income Dynamics). The fact that the estimate of γ is negative in both (20) and (21) is consistent with the idea that technologies that are related to routine tasks act as a substitute for workers ($\nu_R > \psi$). Note, however, that compared to the corresponding estimates in Table 2, the magnitudes are significantly larger (the corresponding OLS estimates in Table 2 imply a 0.4 percent cumulative decline in average worker earnings over

[6] We construct controls for worker age and lagged earnings by linearly interpolating between 3rd degree Chebyshev polynomials in workers' lagged income quantiles within an industry–age bin at 10-year age intervals. In addition, to remove some potential variation related to potential mean-reversion in earnings (which could be the case following large transitory shocks), we also include 3rd degree Chebyshev polynomials in workers' lagged income growth rate percentiles, and we allow these coefficients to vary by past income levels based on five earnings bins; we additionally include gender by 10-year age bin dummies.

the next five years).[7] This difference in magnitudes is likely driven by composition effects—recall Equation (20)—and highlights the value of individual worker data to study the effect of technology on workers.

Second, note that in Table 3, the exposure to labor-augmenting technologies ξ^N is now significantly *negatively* related to the subsequent earnings growth of individual incumbent workers: an increase in ξ^N from the median to the 90th percentile is associated with a 1 to 1.7 percent decline in worker cumulative earnings over the next five years. Recall that in Table 2, the response of average worker earnings to an increase in ξ^N is only marginally statistically different from zero for our OLS estimate, and essentially zero for our IV estimate. How can we then make sense of this negative estimate? In the model, the sign and magnitude of δ depends on the sign of $\nu_N - \psi$. Thus, one possibility is that the elasticity of substitution across tasks ψ is larger than the elasticity of substitution between labor and technology in non-routine tasks ν_N. However, this case would be at odds with our findings in the previous section: the increase in total labor compensation and employment in response to our measure of labor-augmenting technologies suggests that $\nu_N < \psi$. One potential resolution of this apparent inconsistency is the presence of vintage human capital effects (skill displacement) that would imply that some groups of workers would benefit from technological advancements while other groups are adversely affected. We revisit this issue in the remainder of this section and in Section 4.

Once again, we emphasize our continuing focus on the direct effects of technology—the first two terms of Equation (8) in the model. That is, our baseline specification interacts occupation, worker prior income, and industry fixed effects with calendar year to account for occupation- or industry-specific time trends. This saturated specification allows us to partial out sources of time-series variation that affect specific groups of workers in a given occupation, industry, or income level. The cost of saturating our specification is that these fixed effects largely absorb the last term in Equation (8), which captures productivity spillovers across workers. Comparing Columns (1) to (3) of Table 3 to our preferred specification in Column (4), we see that including industry times calendar year fixed effects has a modest negative impact on the estimated γ coefficient, whereas the estimate of δ becomes significantly more negative—almost double the magnitude in some cases. This finding is consistent with our conjecture that the industry-year fixed effects in our specification likely absorb the potential positive benefits of technology exposure on worker wages through increases in industry

[7] Comparing between the estimates in Table 2 and Table 3 is complicated by the fact that the first set of estimates refers to point-to-point growth rates across a decade, whereas the latter refers to the growth in cumulative earnings over the last five years—see Equation (14). If the effect of a new technology on earnings is immediate and permanent, the two sets of estimates are directly comparable. However, if this is not the case, two opposing effects come into play. If new technologies have an immediate transitory effect on earnings that subsequently dissipates, the decline in cumulative earnings will be larger than the point-to-point estimates. Conversely, if there are lags in the impact of the new technology, the growth in cumulative earnings will be understated relative to the point-to-point estimates. The fact that Appendix Table A.3 shows a somewhat larger impact on cumulative earnings over the next 10 years compared to the 5-year estimates in Table 3 suggests that the second effect is likely more important the first—implying that the growth in cumulative earnings is understated relative to the point-to-point comparison.

productivity—a topic we revisit in Section 5. By contrast, allowing for occupation-specific trends in earnings growth (occupation interacted with year fixed effects) has only a modest quantitative impact on our results, suggesting that this is not the dominant source of variation that identifies our estimates.

3.2.2 New entrants

Taken together, the results in Table 2 and Table 3 imply that an increase in workers' exposure to labor-augmenting technologies ξ^N leads to a decline in earnings for incumbent workers, but average earnings for workers in the same industry and occupation are unaffected (while total employment increases). Recalling Equation (20), we see one interpretation for this pattern: the earnings of new entrants can respond positively to ξ^N, even as earnings for incumbent workers decline, thus leaving the average earnings for the group unaffected.

To study this issue further, we now focus our attention on new entrants. Given the structure of our data, we need to make some assumptions to identify new entrants into industry k and occupation o. First, we require that new entrants into occupation o are those for which we observe their occupation reported in the CPS over the next three years. Second, we identify entrants into industry k as those who switch into k from a different industry k' in years $t+1$ to $t+3$ and remain in that industry for a minimum of two years. Third, motivated by the decomposition (20), we focus on the difference in average earnings that new entrants earn from $t+1$ to $t+h$ relative to the earnings that *similar incumbents* earned in the periods before t. Thus, we match our sample of occupation–industry entrants to incumbent workers based on age, gender, college education attainment, and yearly income rank within age/gender/college cells. These restrictions imply that the resulting matched sample of entrant workers is quite small (approximately 57,500 observations for our entrant–incumbent matched regression sample). Appendix B.5 contains further details.

We estimate a version of Equation (21), where now our focus is on the difference between the future earnings of new entrants to the prior earnings of the incumbent workers,

$$w^i_{e,t+1,t+h} - w^i_{I,t-2,t} = \gamma\, \xi^R_{i,t} + \delta\, \xi^N_{i,t} + c\, \mathbf{Z}_{i,t} + k\, \left(w^i_{e,t-2,t} - w^i_{I,t-2,t} \right) + u_{i,t}. \tag{22}$$

Here $w^i_{I,t-2,t}$ denotes an average of pre-period earnings taken across all incumbents matched to worker i in the entrant sample, while $w^i_{e,t+1,t+h}$ gives the actual subsequent earnings of entrant worker i. As before, $\mathbf{Z}$ includes flexible nonparametric controls for worker age and the level of past worker earnings as well as recent earnings growth rates. In addition, we control for the prior difference in income between incumbents and entrants to account for differential selection in terms of worker ability into the industry. Thus, the coefficients γ and δ are identified by comparing the earnings of new entrants to the prior earnings of similarly paid incumbent workers across jobs that have experienced increases in their exposure to labor-saving or labor-augmenting technologies.

If new labor-augmenting technologies require skills that incumbent workers lack, new entrants with these skills are likely to benefit at the expense of incumbent workers. Therefore, we expect $\delta > 0$ in equation (22). Conversely, workers who switch into jobs affected by labor-saving improvements would only do so if they are indifferent between switching and staying in their old jobs, suggesting that $\gamma = 0$. As we see in Panel B of Table 3, the data is consistent with this hypothesis. Focusing on the estimate of δ, we see that it is positive and statistically significant. Thus, an increase in the exposure to labor-augmenting technologies ξ^N from the median to the 90th percentile implies a 2.4 to 4.1 percent increase in the earnings of new entrants relative to incumbent workers. By contrast, the estimated coefficient γ is not statistically different from zero.

These results are suggestive of the importance of vintage-specific capital. However there are important reasons to be cautious with this interpretation. First, the structure of our data limits the precision with which we can identify new entrants—recall that we cannot observe occupational changes at the frequency at which we estimate (22); we can only identify workers who are observed in the affected occupations at some point in the next three years. Second, measurement issues aside, entry into the industry is an endogenous decision on the part of the worker and our matching procedure may be an imperfect control for this endogenous selection effect. This issue complicates the interpretation of these coefficients. To sidestep these concerns about selection, in Section 4 we will examine the role of vintage-specific human capital by exploiting heterogeneity among incumbent workers across different measures of skill.

3.3 Do our worker-level results reflect causal estimates?

Establishing a direct causal link between our technology exposure measure and worker earnings is challenging for several reasons. For instance, one possibility is that our negative estimates of γ reflect the endogenous development of labor-saving technologies targeted toward workers (occupation–industry cells) that are temporarily 'overpaid' relative to their marginal productivity, an effect that could lead to a negative bias in our estimates if worker earnings were to mean-revert even in the absence of new technologies. More generally, the arrival of labor-saving and labor-augmenting technologies may be endogenous to underlying trends in the labor markets.

We address these concerns in two ways. First, we explore whether past wage earnings growth of individual workers predicts our technology exposure measures. Second, we use the shift-share identification strategy we develop in Section 2.5 that exploits prior variation in the arrival of breakthrough technologies and their spillovers across different types of technologies.

3.3.1 Wage growth and future technology exposure

We examine whether past worker earnings predict the change in our measure of technology exposure using

$$\xi^j_{t+5+h} = \kappa\,\Delta w^i_{t+5} + \gamma\,\xi^R_{i,t} + \delta\,\xi^N_{i,t} + c\,\mathbf{Z}_{i,t} + u_{i,t}, \tag{23}$$

using our most saturated specification with industry–year and occupation–year fixed effects. To ensure that there is no overlap between the time the technology is implemented and the period over which we measure wage growth, we examine horizons h of two to five years.[8]

Examining Figure 6 and focusing on the left panel, we see that there is little evidence that wage growth positively predicts our labor-saving technology exposure measure ξ^R. The right panel similarly shows that past wages are essentially unrelated to future realizations of our labor-augmenting technology exposure measure ξ^N. We obtain very similar patterns when we use our alternative definitions of labor saving and labor augmenting technologies based on high/low skill task classifications (see Appendix Figure A.2). We conclude that even though innovation effort could in principle be directed to certain worker groups, the resulting innovation outcomes in our sample are unrelated to past worker wage earnings growth.[9]

3.3.2 Shift-share IV estimates

Next, we repeat our worker-level analysis using the shift-share IV described in Section 2.5. In particular, we re-estimate Equations (21) and (22) but now use the shift-share IV in (17) to instrument for ξ^R and ξ^N. To conserve space, we only focus on specifications with industry–year fixed effects. Columns (2) and (3) of Appendix Table A.2 reports the first stage: we see that the shift-share IVs for ξ^R and ξ^N strongly predict the respective technology exposure measures, and the F-statistic for the first-stage regression is over 1,000, suggesting that weak instruments is not a concern.

We report the IV results in Columns (5) and (6) of Table 3. We see that the IV estimates are somewhat noisier than but quantitatively similar to our OLS estimates. In Panel A, we see that exposure to labor-saving innovations is significantly negatively related to the wage earnings of incumbent workers. The point estimate of δ is comparable to the OLS estimate but is more noisily estimated (t-statistic of -1.42 in the most saturated specification). In Panel B, we again see the same pattern as before: ξ^R is unrelated to the level of wage earnings for entrants relative to incumbents, while ξ^N is still significantly positively related to the relative wage difference between entrants and incumbents.

[8] Recall that we are dating the technology exposure measure based on the year the patents are issued. There is a significant lag between when the patent application is filed and when the patent is issued: over 75 percent of the breakthrough patents in our sample have at least a 2-year lag between the time the application is filed to when the patent is approved, and the median lag is 3 years.

[9] This lack of a correlation between wage changes and subsequent development of automation technologies feels at odds with Dechezlepretre et al. (2023), who find evidence that the development of automation technologies responds to wage differentials. We suspect this difference is driven by two forces. First, Dechezlepretre et al. (2023) focus on a relatively narrow measure of mechanization of manual tasks by classifying a small subset of patent technology classes as automation-related, and they show that this measure of automation is predictable by a broad measure of worker earnings (wage earnings of all production workers). By contrast, we have a more general measure of labor-saving technologies that is linked to the earnings of specific workers in closely related occupations. Second, the sample is different: our results focus on the US while Dechezlepretre et al. (2023) focus on Europe. It is likely that wages in the US labor market are significantly more flexible than in Europe, which dampens the incentive to direct innovation to certain worker groups.

3.4 Alternative specifications and exposure measures

Here, we briefly discuss results using alternative specifications and measures of labor-saving and labor-augmenting technology exposure. Appendix Table A.3 estimates earnings growth for individual worker earnings over the next ten years (as opposed to five years in our baseline specification). The estimated magnitudes are comparable or somewhat larger in this case, suggesting these are persistent effects. Next, we report results using two alternative exposure measures. The first alternative exposure measure is based on the partition of worker tasks based on the degree of required experience—a proxy for low- vs high-skill—discussed in Section 2.1.1. For the second alternative exposure measure, we modify our methodology for assigning a patent to a specific industry. Rather than simply assigning the patent to the industry of the firm that files for the patent, we now allow patents to be applied in downstream industries based on the weights in the BEA Input-Output matrix. As we see in Appendix Tables A.4 and A.5, we obtain qualitatively and quantitatively similar results using these alternative exposure measures.

4 Heterogeneous impact of technological change

The difference in the estimated impact of labor-augmenting innovations on worker earnings between incumbent and new entrants suggests that new technologies come with new skill requirements that incumbent workers lack—that is, worker human capital has a component that is specific to a particular technology. Next, we provide additional evidence for the importance of vintage-specific human capital by exploiting variation in outcomes among incumbent workers in terms of their level of skill.

4.1 The importance of vintage-specific human capital

If human capital has a vintage-specific component, we expect to see that the negative impact of labor-augmenting technologies on worker earnings is concentrated among older and more highly skilled workers. In the absence of data on the full occupation history of the worker, we expect age to be a useful proxy for skill in the existing technology since older workers have likely accumulated more human capital than younger workers. The level of a worker's earnings relative to her peers—workers in the same occupation and industry—is also likely correlated with her level of skill in the existing technology. Panels A and B of Table 4 confirm our hypothesis: both the OLS and IV estimates of δ are significantly more negative for older and more highly paid workers relative to their peers.

Panel A of Table 4 reports the estimated coefficients γ and δ across workers in different age groups—25 to 35, 35 to 45, and 45 to 55. To conserve space, we report results for our most saturated specification with both occupation–year and industry–year fixed effects. We see that the estimated coefficient δ shows a strong age pattern: the estimate is essentially zero for the youngest workers but is negative and statistically significant for older workers (workers in either the 35–45 or in the 45–55 group). The difference in the estimated coefficients δ between the groups of oldest and youngest workers ranges between 1.4 and 1.7 percentage points across the OLS and IV specifications and is both statistically and economically significant. By contrast, the estimated coefficient γ capturing the impact of labor-saving technologies is consistently negative for workers in all age groups but shows no discernible pattern across groups.

In Panel B, we allow the coefficients γ and δ to vary with the worker's level of earnings (over the last year) relative to her peers—defined as workers in the same occupation and industry. As before, the estimated coefficient γ corresponding to the impact of labor-saving technologies on worker earnings is consistently negative for all income groups, but does not vary significantly across income levels. By comparison, the coefficient δ corresponding to labor-augmenting technologies is only significantly negative for the most highly paid workers—workers in the top half of the within-occupation–industry pay distribution, especially the top quartile. The magnitudes are sizable: an increase in our labor-augmenting exposure measure from the median to the 90th percentile has essentially no effect on workers whose relative earnings fall below the median relative to their peers but leads to a 4.6 percent decline in earnings for workers in the top 5 percent of the relative pay distribution.

Next, in Panel C, we explore the extent to which the response of worker earnings varies by the education level of the worker. Here, we do not have a strong a priori view: it depends on whether education conveys general or vintage-specific human capital. To the extent that a worker's education correlates with her skill at operating the current vintage of capital, but does not relate to the degree to which her skills are transferable across technology vintages, we expect that the estimated coefficient δ is more likely to be negative for college-educated workers than for workers without a college degree. By contrast, if college relates more to the acquisition of general skills, we would expect the opposite pattern. We see that college-educated workers are somewhat more negatively exposed to labor-augmenting technical change. This difference is statistically significant but rather small in magnitude (a 0.5 to 0.6 percentage-point difference). Again, we find no meaningful difference in our estimate of γ across these two groups implying that the effects of labor-saving technologies are pervasive across different worker groups.

Last, in Figure 5 we allow the estimated coefficients γ and δ to vary across different sectors or job types. Overall, we see that the displacement of workers in response to labor-saving technologies is not purely a blue-collar worker phenomenon; rather it is increasingly present in white-collar occupations, supporting the view in Autor et al. (2003) regarding the potential for labor-saving technologies to substitute for labor in routine cognitive tasks. By contrast, the negative effect of labor-augmenting technologies on the earnings of incumbents is primarily concentrated among white-collar workers—services (NAICS codes 42-92) and in occupations emphasizing cognitive skills. Last, the earnings of workers in occupations emphasizing interpersonal tasks are essentially unaffected by either technology exposure measure.

In sum, we find that the negative impact of labor-saving technologies is largely uniform in the cross-section of workers. By contrast, the response of worker earnings to labor-augmenting technologies displays significant heterogeneity, with patterns consistent with the importance of vintage-specific human capital. These patterns are robust to several changes in methodology. First, Appendix Table A.5 shows that the results are quantitatively very similar under alternative definitions of technology exposure. Second, Appendix Table A.6 shows that our baseline results are robust to alternative methods of ranking workers that remove heterogeneity in earnings that may be unrelated to skill. In particular, in Column (2), we report estimates when we drop workers hired within the last year, while in Column (3), we rank workers based on their average income over the last two years relative to their peers. In Column (4), we compute worker wages relative to the average wage offered by employers using data from the LBD. In Column (5), we rank workers based on their residual wage earnings net of occupation, industry, commuting zone, age, and gender fixed effects, all interacted with calendar year. Column (6) combines both adjustments from Columns (4) and (5). Last, in Column (7), we also compute wage earnings residuals net of the worker's unionization status interacted by calendar year.

4.2 The role of job destruction

Our results thus far have focused on the average response of worker earnings to our technology exposure measures. However, these average responses may mask considerable heterogeneity in ex post outcomes for workers: our estimates cannot distinguish between each worker experiencing a 2 percent decline in earnings or 10 percent of the workforce experiencing a 20 percent decline in earnings due to job loss. However, these two possibilities have very different implications for worker welfare costs in the presence of risk aversion and incomplete markets. Our goal here is to understand the extent to which these average losses are concentrated among a subset of workers by analyzing the importance of job destruction.

In Table 5, we re-estimate (21) but now replace the outcome variable with a dummy for involuntary exit from the firm (Panel A), involuntary exit from the industry (Panel B), or long-term unemployment (Panel C). Since we cannot directly distinguish between voluntary and involuntary job transitions, we construct a set of dummy variables for involuntary exit that take value one if the worker switched employer (EIN) or industry (NAICS4) at any point over the next five years and experienced significant declines in worker earnings (earnings growth below the 20th percentile in that year). In addition, we create a measure of long-term unemployment: a dummy that takes value one if the worker experiences a consecutive 3-year period with zero W2 earnings at any point over the next five years. In addition to the homogenous specification, we also allow the coefficients γ and δ to vary with the worker's relative income or age.

Job destruction. In Panel A, we see that our labor-saving technology exposure measure is both statistically and economically significantly related to our measure of job destruction. The magnitude of the estimated coefficient γ is sizable: an increase in ξ^R from the median to the 90th percentile is associated with a 1.1 percentage-point increase in the likelihood of the worker experiencing job loss, compared to a base probability of approximately 17 percent. Given our definition of job loss, the expected growth rate of cumulative earnings for these workers is -94.1 log points (Table 1). Thus, our estimates imply that the increased likelihood of job loss can account for approximately one-half of the response of mean worker earnings to labor-saving technologies in Panel A of Table 3. The estimated magnitudes decrease somewhat with the worker's prior income relative to her peers, but there are no significant differences in the estimates based on worker age.

Exposure to labor-augmenting technologies has a positive but smaller impact on the likelihood of job destruction. However, as before, these estimates vary considerably with worker age and prior income. An increase in ξ^N from the median to the 90th percentile is associated with 1.5 and 3.8 percentage-point increases in the likelihood of job loss for workers in the 75th to 95th and over the 95th percentile of prior earnings. By contrast, the estimated coefficients are essentially zero for workers below the 75th percentile in terms of prior earnings. Similarly, for a worker aged 45 to 55 years, an increase in labor-augmenting technology exposure from the median to the 90th percentile is associated with a 1 percentage-point increase in the likelihood of involuntary exit—compared to essentially zero for workers aged 25 to 35. These patterns are consistent with skill displacement: for the workers at the top of the earnings distribution, the increased likelihood of job loss can account for almost three-quarters of the overall decline in mean earnings.

Involuntary exit from the industry. Panel B focuses on (involuntary) exit from the industry as the worker outcome. We see similar patterns as in Panel A. An increase in exposure to labor-saving technologies is associated with a higher likelihood of industry exit, particularly for workers who are paid less than their peers. In contrast, increased exposure to labor-augmenting technologies raises the likelihood of industry exit for older and higher-paid workers, consistent with the patterns of skill displacement observed throughout this section.

Long-term unemployment. In Panel C, we focus on long-term unemployment spells as the outcome variable. Increased exposure to labor-augmenting technologies is associated with long-term unemployment, particularly for higher-paid workers. Workers at the top of the relative pay distribution experience a 0.8 percentage-point increase in the likelihood of long-term unemployment, compared to a 0.3 percentage point increase for workers at the bottom of the distribution. These magnitudes are economically significant, as the base likelihood of our long-term unemployment indicator is approximately 3.3 percent (see Table 1) and are consistent with skill displacement. The response of long-term unemployment to labor-saving technologies is somewhat smaller than the response to labor-augmenting technologies and is larger among older and lower-paid workers.

4.3 Allowing for skill displacement in the model

Thus far, we have documented a significant degree of heterogeneity in the response of worker earnings to our measure of labor-augmenting technologies: our estimated coefficient δ in our empirical specification (21) is larger in magnitude for older, more highly educated, and more highly paid workers than for their peers. Importantly, this heterogeneity in the estimated coefficient δ is significantly larger than the average effect: the response of earnings for top workers to the same improvement in labor-augmenting technologies is almost four times larger than the magnitude of the average coefficient.

Is the larger decline in earnings for more highly paid workers in response to labor-augmenting technologies inconsistent with capital-skill complementarity? Not necessarily: the existing evidence in favor of capital-skill complementarity is inferred by trends in average wage differentials *between groups* of workers with different skill levels. Our findings are consistent with this view: recall that workers in highly paid occupation–industry cells are, on average, more exposed to labor-augmenting innovations than lower-skilled workers (Figure 4). Furthermore, the first two columns of Table 2 show that, *as a group*, workers that are more exposed to labor-augmenting technologies experience higher total compensation. By contrast, our findings in this section refer to within-occupation heterogeneity in earnings responses to *the same technology improvement*—recall that our exposure measure (13) is constructed at the industry–occupation level. Our interpretation of this fact is that, although these workers as a group may benefit, individual incumbent workers may be left behind.

A simple modification to the model in Section 1 that allows for skill displacement at the individual level can jointly reconcile all our findings. Specifically, worker i's effective labor supply in task j at time t is now given by

$$\Delta \log l(i,j) = \begin{cases} -\beta I[j \in J_N]\, \varepsilon(j) + (u_{i,j} - \log l(i,j)) & \text{w/ prob. } \omega\, I[j \in J_N]\, \varepsilon(j) \\ -\beta I[j \in J_N]\, \varepsilon(j) & \text{otherwise} \end{cases}, \tag{24}$$

where ω is a positive constant indicating the strength of the skill displacement effect and $u_{i,j}$ is an i.i.d. draw from the distribution of $\log l(i,j)$ for incumbent workers in the same occupation in the initial period. Equation (24) implies that the adoption of a new vintage of technology leads to a change not only in wages but also in the quantity of skill available for performing a task. The first term implies that larger shifts in the technology frontier are likely to cause greater skill displacement among incumbent workers, consistent with Violante (2002). The second term in (24) captures the idea that incumbent workers who are highly skilled in the prior vintage of technology will be, on average, at a relative disadvantage when a new vintage is introduced.

In Appendix Table A.9 we show how the modified model can be calibrated to quantitatively match our empirical facts with reasonable estimates of elasticities of substitution ($\nu_R = 1.7, \nu_N = 0.9, \psi = 1.2$). In particular, equations (7) and (24) together imply that new entrants are, on average, more highly skilled in the new technology than incumbents—which is consistent with our evidence in Section 3.2. In addition, this extension implies that workers with greater levels of skill in the previous technology are more likely to experience earnings declines in response *to the same level of technology exposure* than their less-skilled coworkers in the same occupation. Specifically, Equation (8) that describes the earnings for incumbent workers now has an additional term that captures the redistributive effects of skill displacement,

$$\Delta \log W(i) = \cdots - \left(\beta + \omega \left[\log \bar{l}(i) - \int \log \bar{l}(i)\, dF(i)\right]\right) \xi^N(i), \tag{25}$$

where now $F(i)$ is the cumulative distribution function of $\log \bar{l}(i)$ across workers. Relative to the baseline model, there are two differences. First, the sign of the earnings exposure of the average incumbent worker to labor-augmenting technologies is ambiguous, as it depends on the relative strength of two forces: whether the complementarity across tasks ψ is greater or smaller than the complementarity between capital and labor in non-routine tasks ν_N and now also on the degree of skill displacement β among incumbent workers. Second, the same improvement in labor-augmenting technology leads to winners and losers among incumbent workers in the same occupation that are consistent with our evidence in Section 4.1. See Appendix A.7 for details.

5 Indirect Effects of Technology Exposure

Do our results thus far imply that, in general, technology improvements decrease earnings for incumbent workers? Not necessarily: recall that we have focused on estimating the direct effect of technology exposure on the earnings of incumbent workers—the first two terms in (8). In doing so, we omitted the indirect effect of technological advancements on worker earnings (the last two terms in the same equation). In particular, if new technologies enhance the productivity of certain tasks, they will increase overall labor demand and affect the wage earnings of all workers, regardless of whether these technologies are labor-augmenting or labor-saving for the directly affected tasks. Since our empirical design includes time fixed effects (interacted with industry and occupation), we can only identify the spillover effects of industry innovation (the term multiplying $\Delta \log X$). Focusing on this term, we see that the sign of these spillovers depends on the relative values of ψ and χ. If the complementarity among tasks ψ is greater than the complementarity among industries χ, a higher rate of innovation in a specific industry will tend to uniformly raise wage earnings for all workers in that industry, in addition to the direct effects of technology on their own labor productivity.

This section focuses on estimating these spillover effects by relaxing the granularity of our fixed-effect specification. We proceed in two steps. First, we verify that technological improvements at the industry level are indeed related to measured industry productivity. Second, we show that these technological improvements that raise productivity also affect the wage earnings of workers in the same industry, after accounting for their direct technology exposure.

Industry productivity and technology exposure. In the model, industry productivity growth (the decline in the unit cost of production) evolves according to

$$\Delta \log X(k) \approx \frac{\Gamma_R}{1+\epsilon_x} \frac{LS}{LS_R} \left[\sum_o s(o)\xi^R(o,k)\right] + \frac{\Gamma_N}{1+\epsilon_x} \frac{LS}{LS_N} \left[\sum_o s(o)\xi^N(o,k)\right] + \text{Constant}, \qquad (26)$$

where ϵ_x denotes the elasticity of the price of industry output to an improvement in productivity, and $s(j)$ denotes the expenditure share of task j. In short, productivity growth in the model is a function of weighted averages of the occupation-level exposures (9) weighted by the share $s(o)$ of occupation o in the industry wage bill.

To estimate Equation (26) in the data, we first construct direct empirical analogues of these industry-level technology exposures,

$$\bar{\xi}^R_{k,t} \equiv \left[\sum_o s_{o,k,t}\, \xi^R_{o,k,t}\right], \qquad \bar{\xi}^N_{k,t} \equiv \left[\sum_o s_{o,k,t}\, \xi^N_{o,k,t}\right], \qquad (27)$$

where $s_{o,k,t}$ is the share of occupation o in industry k in year t. Second, we estimate

$$\log(X_{k,t+h}) - \log(X_{k,t}) = \gamma\, \bar{\xi}^R_{k,t} + \delta\, \bar{\xi}^N_{k,t} + c\, \mathbf{Z}_{k,t} + u_{k,t} \qquad (28)$$

where the left-hand side is the cumulative decline in the unit cost of production from the BLS at the NAICS4 industry level. To ensure that our results are not affected by broad industry trends, we include the interaction of coarse industry (NAICS2) fixed effects interacted by calendar year. We also control for the log of industry employment and the lagged 3-year growth rates of the dependent variable. We weight observations by the industry employment share, and we focus on a horizon of five years. Appendix B.7 contains further details.

Columns (1) to (5) of Table 6 report the estimated coefficients γ and δ. Columns (1) and (2) show that both coefficients are positive and close in magnitude: a one-standard-deviation increase in either measure of technology exposure leads to a nearly 6 percentage-point increase in the rate of industry productivity over the next five years. However, in Column (3), we see that, although they are individually highly significant, these coefficients are rather noisy when they are jointly estimated, mainly because the two-industry technology exposure measures are highly correlated (approximately 80%). Therefore, in Column (4), we also estimate a version of Equation (28) where we combine the two exposure measures into one, $\bar{\xi}_{k,t} = \bar{\xi}^R_{k,t} + \bar{\xi}^N_{k,t}$. The estimated coefficient for the combined exposure measure is more precisely estimated and approximately equal to the sum of the two coefficients $\hat{\gamma}$ and $\hat{\delta}$ from Column (3), suggesting that we cannot reject the hypothesis that the combined exposure measure summarizes all relevant information for industry productivity. Column (5) reports the corresponding IV estimate, which is close in magnitude (5 vs. 6 percentage points) to the OLS estimate in Column (4).

Spillovers to worker-level earnings. Having established a link between our technology exposure measures and industry productivity, we next turn our attention to estimating the spillover term of industry innovation on worker earnings using a modified version of (21),

$$\Delta w^i_{t+h} = a\,\bar{\xi}_{i,t} + \gamma\,\xi^R_{i,t} + \delta\,\xi^N_{i,t} + c\,\mathbf{Z}_{i,t} + \varepsilon_{i,t}. \tag{29}$$

where now $\bar{\xi}_{k,t} = \bar{\xi}^R_{k,t} + \bar{\xi}^N_{k,t}$ is the combined technology exposure measure, which, given the discussion above, contains all the relevant information in predicting industry productivity in (28). To identify our coefficient of interest α, we now replace the granular industry-year fixed effects with a coarser version (we move from 4-digit NAICS to 2-digits) and we now control for total industry employment (in logs). All other controls remain as in (21).

We find evidence consistent with positive spillovers. Examining Columns (6) and (7) of Table 6, we now see that estimated coefficient a in (29) is both economically and statistically significant. Our OLS estimate of a implies that a one-standard-deviation increase in $\bar{\xi}$ is associated with a 2.5 percent increase in the growth of cumulative worker earnings over the next five years—the IV estimates imply a similar increase of 2.6 percent. In addition, the estimates of γ and δ from (29) are very close to our baseline estimates from (21).

Summary. In sum, the results in this section imply that the direct effects on which we have focused throughout the paper only tell part of the story; technological improvements that increase industry productivity can also raise earnings for all workers in an industry through aggregate increases in labor demand, regardless of whether these improvements are labor-augmenting or labor-saving. For a given worker, the overall effect on her earnings would therefore depend on the relative magnitude of these direct and indirect effects. That said, the results in this section come with some caveats: to identify the coefficient a in (29), we rely on a different source of variation than that used to identify the coefficients γ and δ on which we have focused throughout the paper. In particular, a is identified by comparing two workers in the same occupation and broad industry category (2-digit NAICS level) that are employed in different industries (defined at the 4-digit NAICS level). Nevertheless, it is somewhat reassuring that the estimated coefficients γ and δ under this alternative specification are comparable to our baseline.

6 Conclusion

We develop a methodology for identifying the arrival of labor-saving and labor-augmenting technologies that relies only on the textual description of the patent document and the tasks performed by workers in an occupation. Labor-saving technologies can perform routine tasks in place of human workers; our assumption is that these technologies are related to tasks performed by workers that can be classified as routine (or alternatively, require little related experience). Labor-augmenting technologies, by contrast, complement certain worker tasks; our assumption is that they are related to tasks that can be classified as non-routine (or alternatively, require high levels of related experience).

Combining our measures with administrative data on worker earnings reveals three new findings. First, the development of new labor-saving technologies has a negative effect on the earnings of directly affected workers. This negative effect estimated from worker-level data is over ten times larger than the corresponding estimate from occupation- and industry-average wage earnings, highlighting that composition effects are likely important. This negative effect is pervasive among both blue- and white-collar workers and is largely unrelated to workers' age, education level, and prior earnings relative to their peers. Second, the development of new labor-augmenting technologies has a heterogeneous impact on directly exposed workers. While the total earnings of these workers as a group increase, driven by the rise in the earnings of new entrants, the earnings of incumbent workers decline. This decline is primarily present among white-collar workers and is concentrated among older and higher-paid workers, consistent with skill displacement. Job destruction accounts for an economically significant fraction of these negative effects on earnings. Last, we find evidence for positive spillovers of both types of innovations: improvements in technology that increase industry productivity also increase earnings for all workers in the industry after controlling for their direct exposure to these innovations.

References

Acemoglu, D., U. Akcigit, and W. R. Kerr (2016). Innovation network. *Proceedings of the National Academy of Sciences 113*(41), 11483–11488.

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics, Volume 4. Amsterdam: Elsevier-North*, pp. 1043–1171.

Acemoglu, D. and P. Restrepo (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review 108*(6), 1488–1542.

Acemoglu, D. and P. Restrepo (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy 128*(6), 2188–2244.

Acemoglu, D. and P. Restrepo (2022). Tasks, automation, and the rise in u.s. wage inequality. *Econometrica 90*(5), 1973–2016.

Aghion, P., C. Antonin, S. Bunel, and X. Jaravel (2021). What are the labor and product market effects of automation? new evidence from france. Working paper.

Akerman, A., I. Gaarder, and M. Mogstad (2015). The skill complementarity of broadband internet. *The Quarterly Journal of Economics 130*(4), 1781–1824.

Arora, S., Y. Liang, and T. Ma (2017). A simple but tough-to-beat baseline for sentence embeddings. In *ICLR*.

Atack, J., R. A. Margo, and P. W. Rhode (2019). Automation of manufacturing in the late nineteenth century: The hand and machine labor study. *Journal of Economic Perspectives 33*(2), 51–70.

Atack, J., R. A. Margo, and P. W. Rhode (2022). “Mechanization takes command?": Powered machinery and production times in late nineteenth-century american manufacturing. *The Journal of Economic History 82*(3), 663–689.

Autor, D., C. Chin, A. Salomons, and B. Seegmiller (2022). New frontiers: The origins and content of new work, 1940–2018. Working paper, MIT.

Autor, D., D. Dorn, G. H. Hanson, and J. Song (2014). Trade adjustment: Worker-level evidence. *The Quarterly Journal of Economics 129*(4), 1799–1860.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review 103*(5), 1553–97.

Autor, D. H., L. F. Katz, and M. S. Kearney (2006). The Polarization of the U.S. Labor Market. *American Economic Review 96*(2), 189–194.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Bárány, Z. L. and C. Siegel (2018). Job polarization and structural change. *American Economic Journal: Macroeconomics 10*(1), 57–89.

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* Number wbsle in Books from Upjohn Press. W.E. Upjohn Institute for Employment Research.

Benmelech, E. and M. Zator (2022). Robots and firm investment. Working Paper 29676, National Bureau of Economic Research.

Bessen, J., M. Goos, A. Salomons, and W. van den Berge (2023). What Happens to Workers at Firms that Automate? *The Review of Economics and Statistics*, 1–45.

Blanchard, O. and L. Katz (1992). Regional evolutions. *Brookings Papers on Economic Activity 23*(1), 1–76.

Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying technology spillovers and product market rivalry. *Econometrica 81*(4), 1347–1393.

Bonfiglioli, A., R. Crinò, H. Fadinger, and G. Gancia (2020). Robot Imports and Firm-Level Outcomes. Crc tr 224 discussion paper series, University of Bonn and University of Mannheim, Germany.

Braxton, J. C., K. F. Herkenhoff, J. L. Rothbaum, and L. Schmidt (2021). Changing income risk across the US skill distribution: Evidence from a generalized kalman filter. Technical report, National Bureau of Economic Research.

Braxton, J. C. and B. Taska (2020). Technological change and the consequences of job loss.

Caunedo, J., D. Jaume, and E. Keller (2023). Occupational exposure to capital-embodied technical change. *American Economic Review 113*(6), 1642–85.

Chari, V. V. and H. Hopenhayn (1991). Vintage human capital, growth, and the diffusion of new technology. *Journal of Political Economy 99*(6), 1142–1165.

Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association 19*(6), 3104–3153.

de Souza, G. and H. Li (2023). Robots, Tools, and Jobs: Evidence from Brazilian Labor Markets. Working Paper Series WP 2023-42, Federal Reserve Bank of Chicago.

Dechezlepretre, A., D. Hemous, M. Olsen, and C. Zanella (2023). Induced automation innovation: Evidence from firm-level patent data. Working paper.

Deming, D. J. and K. Noray (2020). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics 135*(4), 1965–2005.

Dixon, J., B. Hong, and L. Wu (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science 67*(9), 5586–5605.

Feigenbaum, J. and D. P. Gross (2020). Automation and the fate of young workers: Evidence from telephone operation in the early 20th century. Working Paper 28061, National Bureau of Economic Research.

Goldin, C. and L. F. Katz (1998). The origins of technology-skill complementarity. *The Quarterly Journal of Economics 113*(3), 693–732.

Goos, M. and A. Manning (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics 89*(1), 118–133.

Goos, M., A. Manning, and A. Salomons (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review 104*(8), 2509–26.

Graetz, G. and G. Michaels (2018). Robots at Work. *The Review of Economics and Statistics 100*(5), 753–768.

Guvenen, F., S. Ozkan, and J. Song (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy 122*(3), 621–660.

Hemous, D. and M. Olsen (2021). The rise of the machines: Automation, horizontal innovation, and income inequality. *American Economic Journal: Macroeconomics*.

Hicks, J. (1932). *The Theory of Wages*. Macmillan.

Hombert, J. and A. Matray (2021). Technology boom, labor reallocation, and human capital depreciation. *Working Paper*.

Hornstein, A., P. Krusell, and G. L. Violante (2005). The Effects of Technical Change on Labor Market Inequalities. In P. Aghion and S. Durlauf (Eds.), *Handbook of Economic Growth*, Volume 1 of *Handbook of Economic Growth*, Chapter 20, pp. 1275–1370. Elsevier.

Humlum, A. (2019). Robot adoption and labor market dynamics. *Princeton University*.

Jaimovich, N. and H. E. Siu (2020). Job Polarization and Jobless Recoveries. *The Review of Economics and Statistics 102*(1), 129–147.

Jovanovic, B. and Y. Nyarko (1996). Learning by doing and the choice of technology. *Econometrica 64*(6), 1299–1310.

Kalyani, A., N. Bloom, M. Carvalho, T. A. Hassan, J. Lerner, and A. Tahoun (2021). The diffusion of new technologies. Working Paper 28999, National Bureau of Economic Research.

Karabarbounis, L. and B. Neiman (2013). The Global Decline of the Labor Share*. *The Quarterly Journal of Economics 129*(1), 61–103.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021). Measuring technological innovation over the long run. *American Economic Review: Insights 3*(3), 303–20.

Koch, M., I. Manuylov, and M. Smolka (2021). Robots and Firms. *The Economic Journal 131*(638), 2553–2584.

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull, and G. L. Violante (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica 68*(5), 1029–1053.

Liu, E. and S. Ma (2021). Innovation Networks and Innovation Policy. NBER Working Papers 29607, National Bureau of Economic Research, Inc.

Mann, K. and L. Püttmann (2023). Benign Effects of Automation: New Evidence from Patent Texts. *The Review of Economics and Statistics 105*(3), 562–579.

Pennington, J., R. Socher, and C. D. Manning (2014). Glove: Global vectors for word representation. In *EMNLP*.

Violante, G. L. (2002). Technological acceleration, skill transferability, and the rise in residual inequality. *The Quarterly Journal of Economics 117*(1), 297.

Webb, M. (2020). The impact of artificial intelligence on the labor market.

Figures and Tables

Figure 1: Validation of ChatGPT4 task classification

A. Validation of task split into Routine/Non-routine using RTI

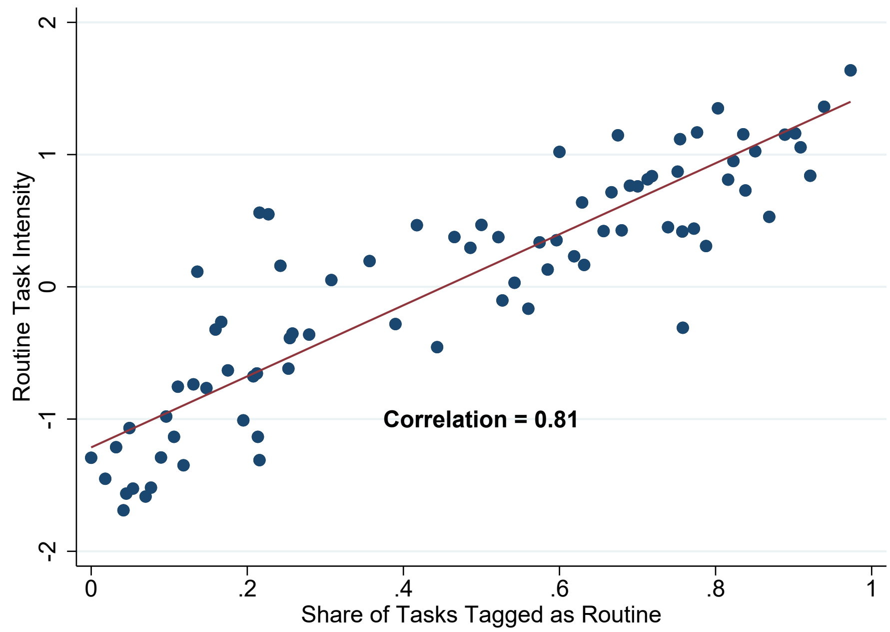

B. Validation of task split into low/high related experience using ONET job zones

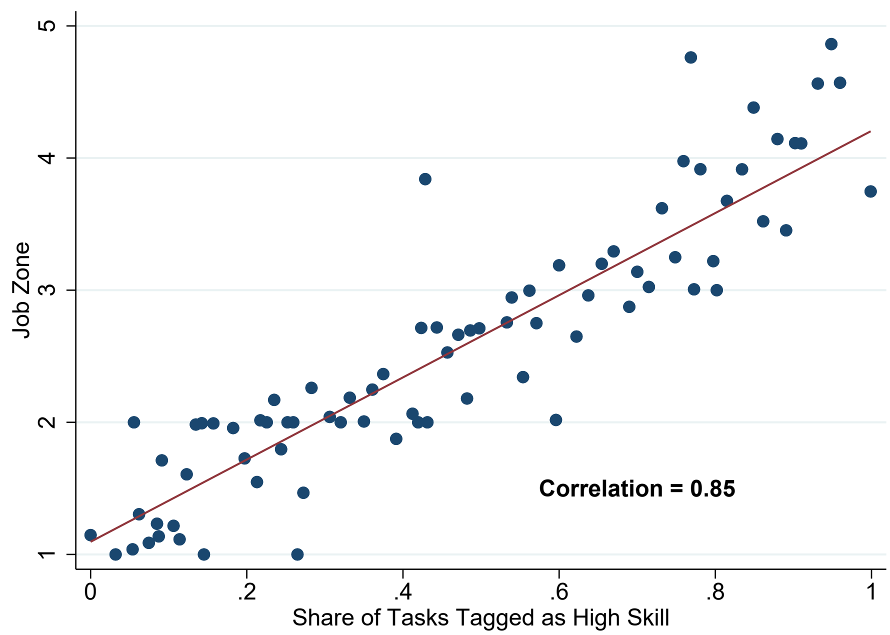

Note: This figure compares occupation routine task intensity scores constructed from Acemoglu and Autor (2011) task categories, with the share of occupation tasks from the Dictionary of Occupational Titles (DOT) tagged as routine by ChatGPT in Panel A; Panel B compares the share of tasks tagged as high experience required with ONET job zone category, which varies between 1 and 5, with 5 corresponding to high amounts of prior preparation and/or experience required. We weight 6-digit SOC occupations by employment shares from Acemoglu and Autor (2011). See Appendix B.2 for further details on the classification of tasks into routine/non-routine or low/high experience and the construction of the data in these figures.

Figure 2: Example: The rise of e-commerce technology and order clerks

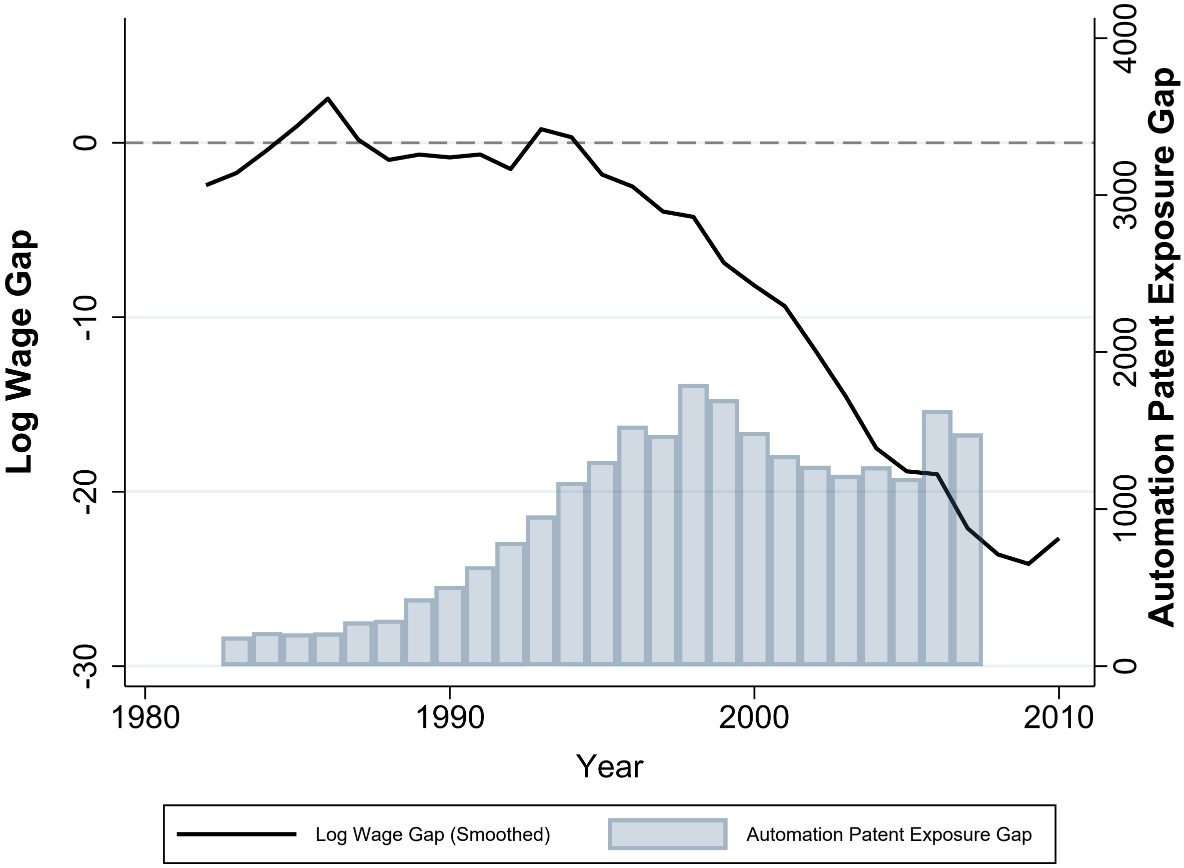

Note: This figure compares the total exposure of order clerks to routine-focused technological change relative to other clerk occupations (gray bars) and the relative log wages of order clerks compared to all other clerk occupations (black line). Specifically, we sum up the total order clerk occupation-patent similarity ($\tilde{\rho}^j(b,o)$) across all breakthrough patents issued in year t , and we do the same for all other clerk occupations (taking an employment-weighted average across the other clerk occupations). We then take the difference between these two innovation exposure series to arrive at the gray bars. We also plot the 3-year moving average of the log wage for order clerks minus all other clerk occupations. We center this series by subtracting off the average difference from 1985 to 1995. Wage and occupational employment counts come from the CPS merged outgoing rotation group sample.

Figure 3: Technology exposure of occupations emphasizing manual vs cognitive tasks vs interpersonal tasks

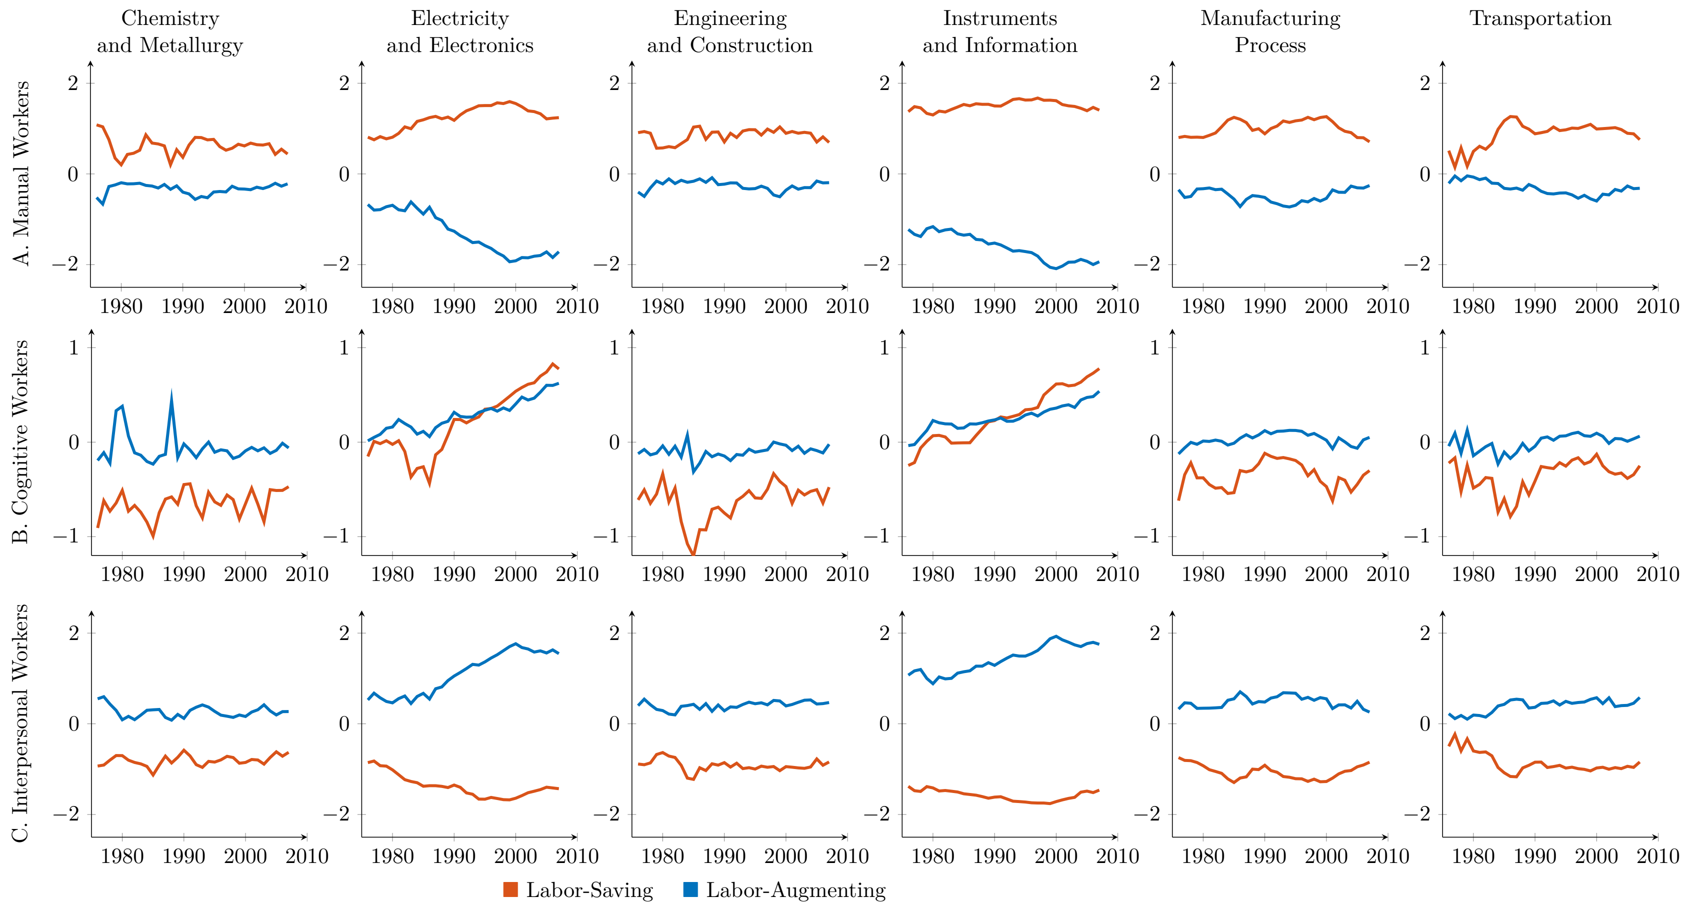

Note: This figure shows how the exposure of different types of occupations (those emphasizing manual or cognitive tasks) has varied over time across broad technology classes. We designate occupations as primarily focusing on either manual, cognitive, or interpersonal tasks using task scores from Acemoglu and Autor (2011). The Acemoglu and Autor (2011) non-routine manual (physical) and routine manual task category scores gives the manual score; average non-routine cognitive (analytical) and routine cognitive represents the cognitive score; and non-routine manual (interpersonal) and non-routine cognitive (interpersonal) scores yields the interpersonal score. We group patents into broad technology classes, with CPC codes in parentheses: Chemistry and Metallurgy (C); Electricity (H0); Engineering (E0, E2, F0, F1); Instruments and information (G, Y1); Manufacturing (B0, B2, B3, B4, B8, D0, D1, D2); and Transportation (B6). We plot residuals net of broad technology classes times calendar-year fixed effects to account for differences in patenting propensity across technologies. We assign occupations into one of these three disjoint categories depending on which of the three task scores corresponds to the highest percentile ranking across all occupations. We generate this figure using a version of our measure that is computed using publicly available sources based on patent-to-industry probabilistic links from Goldschlag, Lybbert, and Zolas (2016). See Appendix Section B.8 for further details.

Figure 4: Technology exposure by occupation skill (wage)

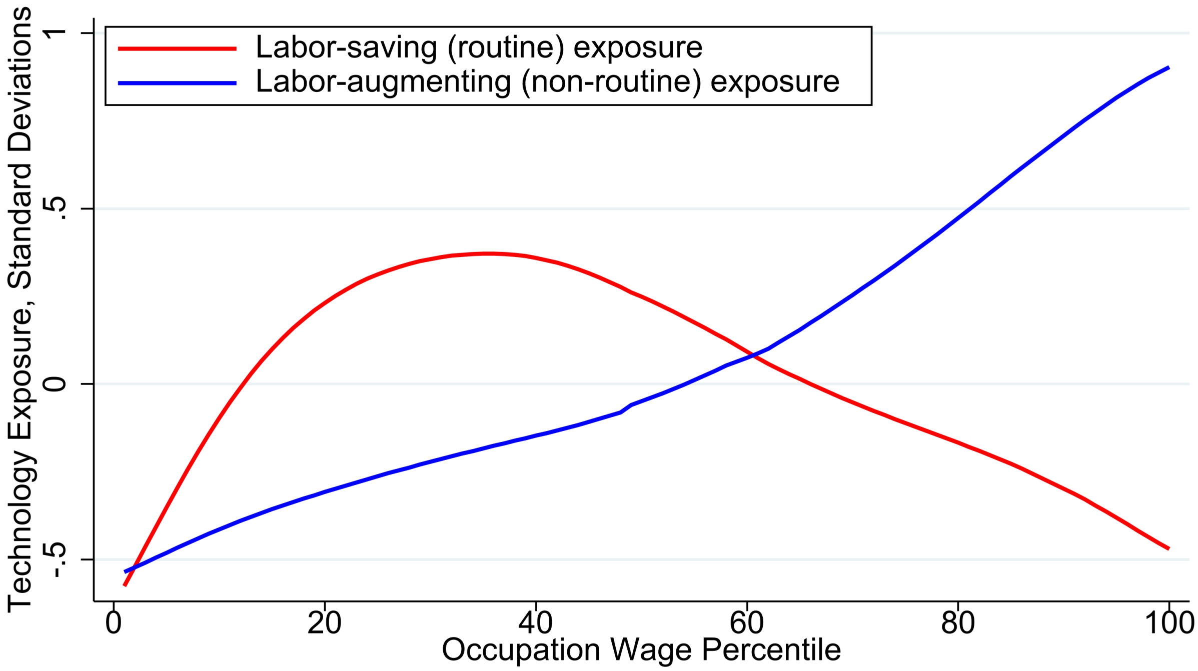

Note: This figure shows average routine or non-routine technological exposure (given by (13) in the main text) plotted against occupational cross-sectional wage percentile rank (lowess smoothed using a bandwidth of 5 bins). We compute averages of technological exposure in employment-weighted standard deviation units by averaging the occupation–industry level measure at the occupation level each Census yea from 1980 to 2000, and then we rank occupations into percentiles based on their average wages in a given year. We generate this figure using a version of our measure that is computed using publicly available sources based on patent-to-industry probabilistic links from Goldschlag et al. (2016), as well as Decennial Census data on occupational employment and wages for the 1980, 1990, and 2000 Census years. See Appendix Section B.8 for further details.

Figure 5: Technology exposure and worker earnings growth, by industry or occupation type

A. Manufacturing vs Non-Manufacturing

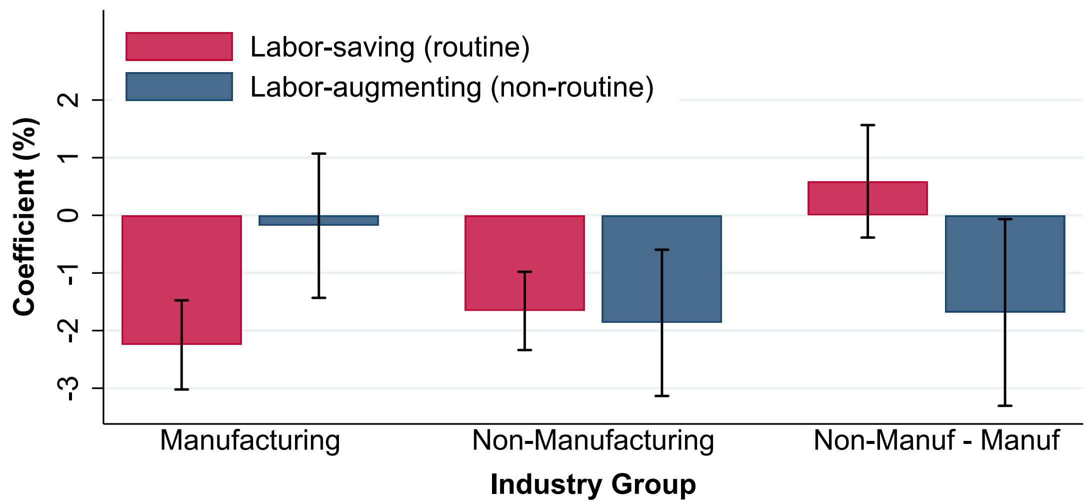

B. Occupation Task Type: Manual/Cognitive/Interpersonal

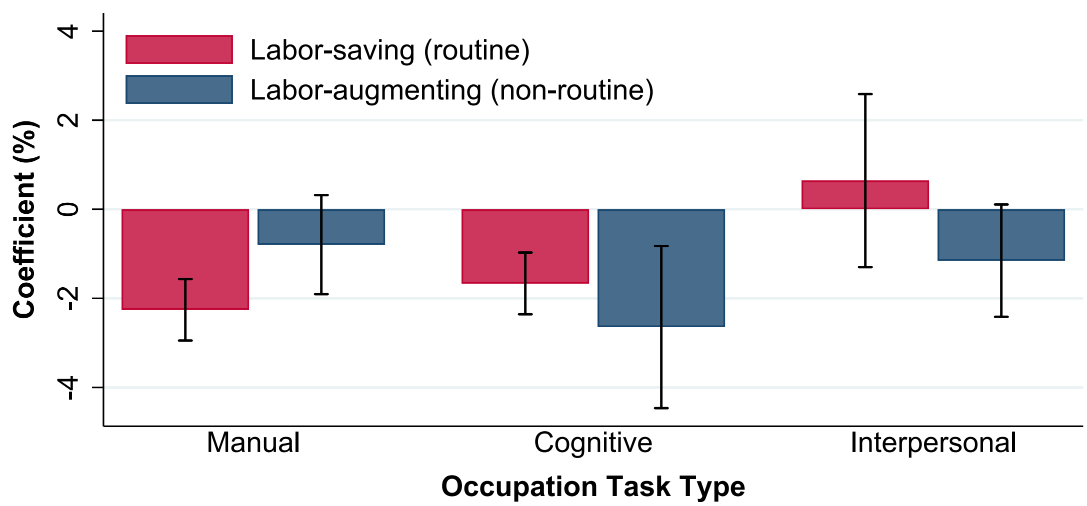

Note: This figure shows the estimated slope coefficients γ and δ (times 100) from Equation (21) in the main text; we allow these coefficients to vary by industry type (Panel A) or occupation task type (Panel B). In Panel A, we compare coefficients for individuals employed in or outside of manufacturing (broadly defined as 2-digit NAICS codes 11 through 33). Our broad definition of manufacturing also includes construction, mining/extraction, utilities, and agriculture, but a large majority of total employment from this wider set of industries still comes from the standard NAICS manufacturing industries 31 to 33. In Panel B, we designate occupations as primarily focusing on either manual, cognitive, or interpersonal tasks using task scores from Acemoglu and Autor (2011). The Acemoglu and Autor (2011) non-routine manual (physical) and routine manual task category scores gives the manual score; average non-routine cognitive (analytical) and routine cognitive represents the cognitive score; and non-routine manual (interpersonal) and non-routine cognitive (interpersonal) scores yields the interpersonal score. Each year we assign occupations into one of these three disjoint categories depending on which of the three task scores corresponds to the highest cross-sectional percentile ranking for that occupation. The dependent variable is workers' cumulative earnings growth (net of lifecycle effects) over the next 5 years. We plot 95% confidence intervals from standard errors clustered at the occupation–industry level beneath coefficient estimates, and normalize the coefficients to correspond to a shift from the median to the 90th percentile. All specifications include industry × year, occupation × year, within occupation–industry income bin × year fixed effects, dummies at the level of coefficient interaction, and controls listed under Table 3. See the main text and notes to Table 3 for further details.

Figure 6: Wage growth and future technology exposure

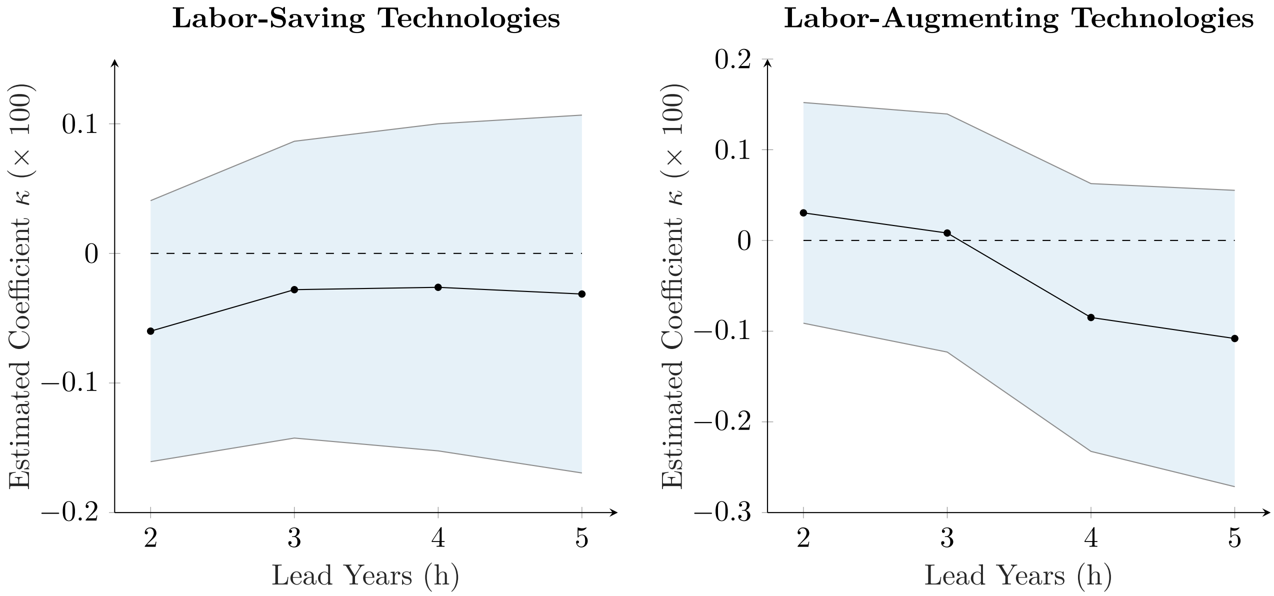

Note: The figure plots the estimated coefficients κ (multiplied by 100) of Equation (23) in the main document together with 95% error bands. The coefficients correspond to a regression of future technology exposure ξ^j_{t+5+h} for horizons $h = 2 \ldots 5$ on the wage growth of affected workers Δw^i_{t+5}. The left figure shows the response of future labor-saving technologies (defined based on the distance to routine tasks), whereas the figure on the right shows the response of future labor-augmenting (based on the distance to non-routine tasks).

Table 1: Summary Statistics: Census–CPS merged sample (worker-level data)

Variable	Mean	SD	5%	10%	25%	Median	75%	90%	95%	Observations
W2 Earnings	66,150	145,500	15,500	20,710	32,390	50,190	76,070	114,000	152,200	2,782,000
Age	40.9	7.4	29	31	35	41	47	51	53	2,782,000
Age, workers in bottom-25 income bin	40.0	7.6	29	30	33	40	46	51	53	632,000
Age, workers in top-5 income bin	43.2	6.8	31	33	38	44	49	52	53	109,000
occupation–industry technology exposure (ξ)	0.647	0.971	0	0	0	0.232	0.866	2.028	2.922	1,495,000
Lifecycle-adjusted earnings growth, 5-years	-0.095	0.526	-1.116	-0.622	-0.174	0.002	0.142	0.337	0.507	2,596,000
Lifecycle-adjusted earnings growth, 10-years	-0.145	0.609	-1.363	-0.825	-0.280	-0.023	0.160	0.388	0.576	1,697,000
Male	0.542	0.498	0	0	0	1	1	1	1	2,782,000
Has four-year college degree	0.344	0.475	0	0	0	0	1	1	1	2,782,000
Involuntary EIN exit proxy	0.1775	0.3821	0	0	0	0	0	1	1	2,147,000
5-year earnings growth | Inv. EIN exit = 1	-0.9447	0.6705	-2.55	-2.06	-1.229	-0.7033	-0.4407	-0.331	-0.2924	381,000
Involuntary NAICS exit proxy	0.1731	0.3783	0	0	0	0	0	1	1	2,147,000
No W2 income 3-yrs straight within next 5 yrs	0.0334	0.181	0	0	0	0	0	1	1	2,210,000
Exit firm within 1 year	0.175	0.380	0	0	0	0	0	1	1	2,676,000
Exit firm within 5 years	0.516	0.500	0	0	0	0	1	1	1	2,210,000
Union member	0.165	0.271	0	0	0	0	0	1	1	553,000
Industry unionization rate	0.163	0.163	0	0	0.031	0.111	0.250	0.436	0.475	2,398,000

Note: This table reports summary statistics for our wage earnings data from the Census Detailed Earnings Record (DER)-CPS merged sample, which covers the 1981 to 2016 period. The sample includes all workers whose unique identifiers (PIK codes) can be matched between the DER and CPS data for CPS years between 1981 and 2016 and who satisfy labor force attachment sampling criteria. W2 earnings are reported in terms of 2015 dollars. The occupation–industry technology exposure ξ is defined as in (13) from the main text. Patents are matched to industry of origination using information from the confidential Census SSL and LBD datasets. The variable “Has four-year college degree" denotes whether a given individual has completed a 4-year degree at the time they were observed in the CPS. The involuntary exit proxy equals 1 if a worker leaves the firm within the next 5 years and also has a 5-year earnings growth rate that is below the yearly 20th percentile. Workers are required to be between the ages of 25 and 55 to be included in the sample. Lifecycle-adjusted earnings growth rates follow Guvenen et al. (2014) and are constructed following (14) in the main text. We winsorize earnings growth rates, as well as the summation component of worker technology exposure (defined in Equation (13)), at the 1% level each year. Observation counts are rounded in accordance with Census disclosure rules. For further details on the construction of the CPS-DER matched sample and the linking of patents to industries, see Appendix Section B.9.

Table 2: Technology exposure and aggregate labor market outcomes
(occupation–industry level)

	Labor Compensation (Wage Bill)					
	Total Effect		Employment		Worker Earnings	
	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Exposure to labor-saving (ξ^R)	-4.60	-4.33	-3.87	-2.67	-0.73	-1.66
	(-3.91)	(-2.11)	(-3.38)	(-1.37)	(-2.79)	(-2.63)
Exposure to labor-augmenting (ξ^N)	9.10	6.66	8.29	7.09	0.81	-0.43
	(5.07)	(2.06)	(4.72)	(2.30)	(1.73)	(-0.46)
Observations	64,500	64,500	64,500	64,500	64,500	64,500
Controls						
Demographic controls	Y	Y	Y	Y	Y	Y
Lagged Employment	Y	Y	Y	Y	Y	Y
Lagged Wage	Y	Y	Y	Y	Y	Y
Occupation $\times$ Year FEs	Y	Y	Y	Y	Y	Y
Industry $\times$ Year FEs	Y	Y	Y	Y	Y	Y

Note: This table shows the estimated slope coefficients γ and δ (times 100) from Equation (19) in the main text. The dependent variables are the cumulative log changes in the wage bill, employment, and average worker earnings over the next 10 years for an aggregate occupation–industry cell. Columns 1 and 2 examine the impact on total labor compensation (the wage bill). The next four columns decompose the effect on labor compensation into the effect on employment (Columns 3 and 4) and the effect on average earnings (Columns 5 and 6). We report t-statistics based on standard errors clustered at the occupation–industry level beneath coefficient estimates, and normalize the coefficients to correspond to a shift from the median to the 90th percentile. We winsorize the total wage bill and employment growth at the 1% level each year, and take the difference between the two to get growth in average worker earnings. The data come from the 1980, 1990, and 2000 Decennial Census, as well as the 2008-2012 5-year ACS panel. We aggregate occupational employment, wages, wage bill, and demographic characteristics by occupation code and industry cells. For this analysis we create a modified version of 4-digit NAICS codes that allows us to construct a consistent crosswalk between the Census industry codes found in the Decennial/ACS and modified NAICS codes, and we compute all measures at this modified NAICS industry level. We use restricted-access versions of these surveys available on Census data servers; the main advantage of this version of the data compared ot publicly-available versions is that earnings survey responses are not top-coded. Observation counts are rounded in accordance with Census disclosure rules. See Appendix B.9 for details on the construction of the aggregated occupation–industry-level panel.

Table 3: Technology exposure and individual worker earnings

	OLS				IV	
	(1)	(2)	(3)	(4)	(5)	(6)
A. Incumbent Workers						
Labor-saving Exposure ξ^R	-1.68	-1.76	-1.86	-2.06	-1.94	-2.33
	(-7.90)	(-8.39)	(-6.95)	(-7.25)	(-3.54)	(-4.24)
Labor-augmenting Exposure ξ^N	-0.69	-0.62	-0.96	-1.31	-0.57	-1.21
	(-2.26)	(-2.11)	(-1.95)	(-2.53)	(-0.69)	(-1.42)
B. Entrants						
Labor-saving Exposure ξ^R	-0.20	-0.24	-0.27	0.34	2.23	1.93
	(-0.34)	(-0.38)	(-0.31)	(0.34)	(1.14)	(0.91)
Labor-augmenting Exposure ξ^N	2.38	2.54	3.12	4.10	7.37	6.77
	(3.00)	(3.06)	(2.29)	(2.57)	(2.04)	(1.73)
Worker Controls (age, gender, and past earnings)	Y	Y	Y	Y	Y	Y
Fixed Effects						
Industry	Y	Y				
Occupation	Y		Y		Y	
Industry × Year			Y	Y	Y	Y
Occupation × Year		Y		Y		Y
Prior Income Rank × Year	Y	Y	Y	Y	Y	Y

Note: We report the estimated slope coefficients γ and δ (times 100) from Equation (21) in the main text in Panel A, and Equation (22) in Panel B. The dependent variable is workers' cumulative earnings growth (net of lifecycle effects) over the next five years. We report t statistics corresponding to standard errors clustered at the occupation–industry level in parentheses. We scale the coefficients to correspond to a shift from the median to the 90th percentile of our technology exposure measures. The vector of controls $\mathbf{Z}$ includes flexible nonparametric controls using 3rd degree Chebyshev polynomials for worker age and the level of past worker earnings and recent earnings growth rates. We allow the prior growth rate controls to differ by past income levels based on 5 earnings bins, and we additionally include gender by 10-year age bin dummies. Prior income rank bins are based on workers' yearly earnings rank within their occupation–industry pair. We define occupations based on David Dorn's revised Census occ1990 occupation codes. We define industries based on the 4-digit NAICS code of a worker's primary employer, unless there are fewer than 10 workers in the occupation–industry–year cell in which case we move to the broader 2-digit NAICS industry classification when computing income ranks. Within these groups, we partition workers into the following earnings bins: between the bottom and 25th percentiles; between the 25th percentile and median; between the median and 75th percentile; between the 75th and 95th percentiles; and the 95th percentile and above. In panel B we perform a nonparametric matching procedure to link industry entrants with comparable incumbents, and we additionally control for the preperiod earnings gap between a given entrant and matched incumbents. See Appendix B.5 for details on this matching procedure. There are approximately 1,213,000 observations in our regression sample in Panel A and 57,500 in Panel B (both rounded to match Census disclosure rules).

Table 4: Worker technology exposure and incumbent worker earnings, by measures of worker skill

	Technology Exposure			
	OLS		IV	
	Labor-Saving (ξ^R)	Labor-Augmenting (ξ^N)	Labor-Saving (ξ^R)	Labor-Augmenting (ξ^N)
A. Worker Age				
25-35 y/o	-2.15	-0.04	-2.09	-0.48
	(-6.52)	(-0.06)	(-3.68)	(-0.49)
35-45 y/o	-1.54	-1.65	-1.94	-1.55
	(-5.05)	(-3.01)	(-3.48)	(-1.59)
45-55 y/o	-2.77	-1.75	-3.19	-1.91
	(-7.69)	(-3.23)	(-5.19)	(-1.95)
Older - Younger	-0.62	-1.71	-1.10	-1.43
	(-1.82)	(-4.25)	(-2.84)	(-3.38)
B. Income (relative to Ind × Occ peers)				
0–25th percentile	-2.14	-0.89	-2.65	-0.65
	(-6.14)	(-1.51)	(-4.31)	(-0.65)
25–50th percentile	-1.65	-0.89	-1.97	-0.90
	(-5.27)	(-1.56)	(-3.51)	(-0.91)
50–75th percentile	-2.06	-1.09	-2.16	-1.29
	(-6.67)	(-2.06)	(-3.84)	(-1.31)
75–95th percentile	-2.29	-1.99	-2.42	-2.31
	(-6.67)	(-3.55)	(-4.16)	(-2.37)
95–100th percentile	-2.58	-4.64	-2.61	-4.63
	(-4.39)	(-5.42)	(-3.38)	(-3.82)
Top - Bottom	-0.44	-3.75	0.04	-3.98
	(-0.78)	(-4.24)	(0.07)	(-4.11)
C. Education				
No College Education	-2.04	-1.03	-2.26	-1.02
	(-7.17)	(-2.03)	(-4.12)	(-1.07)
College Educated	-2.38	-1.63	-2.62	-1.56
	(-5.53)	(-3.01)	(-3.94)	(-1.60)
College - No College	-0.34	-0.60	-0.36	-0.54
	(-0.91)	(-2.24)	(-0.82)	(-1.76)

Note: This table shows the estimated slope coefficients γ and δ (times 100) from Equation (21) in the main text; we allow these coefficients to vary by age group (Panel A); within occupation–industry income rank (Panel B); or 4-year college graduate status (Panel C). The dependent variable is workers' cumulative earnings growth (net of lifecycle effects) over the next 5 years. We report t-statistics based on standard errors clustered at the occupation–industry level beneath coefficient estimates and normalize the coefficients to correspond to a shift from the median to the 90th percentile. All specifications include industry $\times$ year, occupation $\times$ year, within occupation–industry income bin $\times$ year fixed effects, dummies at the level of coefficient interaction, and controls listed under Table 3. See the main text and notes to Table 3 for further details.

Table 5: Drivers of worker earnings losses

	All	By Relative Income					By Age		
		Low	(2)	(3)	(4)	High	Younger	Mid	Older
A. Job Loss									
Labor-saving Exposure ξ^R	1.09	1.38	1.10	1.00	0.80	0.90	1.09	0.83	1.49
	(5.64)	(5.99)	(4.94)	(4.49)	(3.48)	(2.11)	(4.70)	(3.90)	(6.12)
Labor-augmenting Exposure ξ^N	0.71	0.28	0.18	0.54	1.46	3.80	0.02	0.88	0.99
	(2.00)	(0.72)	(0.46)	(1.47)	(3.90)	(7.60)	(0.05)	(2.31)	(2.63)
B. Industry Exit									
Labor-saving Exposure ξ^R	0.95	1.26	0.94	0.85	0.66	0.96	0.96	0.70	1.34
	(5.01)	(5.52)	(4.28)	(3.92)	(2.88)	(2.19)	(4.10)	(3.33)	(5.69)
Labor-augmenting Exposure ξ^N	0.50	0.03	-0.06	0.39	1.24	3.57	-0.14	0.65	0.75
	(1.44)	(0.08)	(-0.15)	(1.09)	(3.41)	(7.30)	(-0.37)	(1.77)	(2.01)
C. Long-term Unemployment									
Labor-saving Exposure ξ^R	0.19	0.39	0.13	0.17	0.09	-0.09	0.07	0.13	0.37
	(2.68)	(3.92)	(1.58)	(2.08)	(1.15)	(-0.65)	(0.80)	(1.63)	(3.88)
Labor-augmenting Exposure ξ^N	0.42	0.25	0.34	0.48	0.58	0.81	0.49	0.53	0.24
	(3.53)	(1.75)	(2.58)	(3.79)	(4.40)	(4.08)	(3.71)	(4.21)	(1.77)

Note: We report the estimated slope coefficients γ and δ (times 100) from Equation (21) in the main text, where we replace the income growth dependent variable with proxies for worker displacement. In Panel A, the dependent variable is instead an indicator for whether workers leave their firm in the next 5 years and experiences income growth below the yearly 20th percentile (proxy for involuntary exit from the firm); in Panel B, the dependent variable is an indicator for whether workers left their industry (NAICS4) and had income growth beneath the yearly 20th percentile (proxy for involuntary exit from the industry); in Panel C, the dependent variable is an indicator for whether a worker reports a period of 3 consecutive years with no W2 earnings sometime in the next 5 years. We report homogeneous coefficients in the first column and allow the coefficients to vary by income rank in Columns 2–6, and age in Columns 7–9. We report t-statistics corresponding to standard errors clustered at the occupation–industry level in parentheses. We scale the coefficients to correspond to a shift from the median to the 90th percentile of our technology exposure measures. The vector of controls $\mathbf{Z}$ includes flexible nonparametric controls using 3rd degree Chebyshev polynomials for worker age and the level of past worker earnings and recent earnings growth rates. We allow the prior growth rate controls to differ by past income levels based on 5 earnings bins, and we additionally include gender by 10-year age bin dummies. Prior income rank bins are based on workers' yearly earnings rank within their occupation–industry pair. We define occupations based on David Dorn's revised Census occ1990 occupation codes. We define industries based on the 4-digit NAICS code of a worker's primary employer, unless there are fewer than 10 workers in the occupation–industry–year cell in which case we move to the broader 2-digit NAICS industry classification when computing income ranks. Within these groups we partition workers into the following earnings bins: between the bottom and 25th percentiles; between the 25th percentile and median; between the median and 75th percentile; between the 75th and 95th percentiles; and the 95th percentile and above.

Table 6: Technology spillovers

	Industry Productivity					Worker Earnings	
	OLS				IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Industry Exposure to Labor-saving ($\bar{\xi}^R$)	5.73		2.82				
	(4.40)		(1.40)				
Industry Exposure to Labor-augmenting ($\bar{\xi}^N$)		5.57	3.49				
		(6.20)	(2.27)				
Overall Industry Technology Exposure ($\bar{\xi}$)				6.02	5.07	2.53	2.60
				(6.12)	(5.08)	(4.89)	(2.69)
Worker Exposure to Labor-saving (ξ^R)						-2.77	-3.14
						(-7.70)	(-4.37)
Worker Exposure to Labor-augmenting (ξ^N)						-2.01	-2.19
						(-2.74)	(-1.79)
Controls							
Log Industry Employment	Y	Y	Y	Y	Y		
Lagged 3-year Growth Rate	Y	Y	Y	Y	Y		
NAICS2 $\times$ Year FE	Y	Y	Y	Y	Y	Y	Y
Occupation $\times$ Year FE						Y	Y
Prior Income Rank $\times$ Year						Y	Y
Observations	2,600	2,600	2,600	2,600	2,600	1.2m	1.2m

Note: Columns 1–5: The first five columns of this table show the estimated slope coefficients γ and δ (times 100) from Equation (28) in the main text. The dependent variable is the cumulative log changes in industry productivity (decline in the unit cost of production) in these columns, where we report t-statistics based on standard errors clustered at the industry level beneath coefficient estimates; the coefficients correspond to a one-standard deviation increase in the given measure. These regressions additionally control for the log of industry employment and lagged 3-year growth rates in the dependent variable. Industry-level measures are the yearly employment-weighted averages of the worker-level occupation–industry exposures as in Equation 13. The overall technology exposure is the standardized average of the two individual technology exposures. Data on industry outcomes are from the Bureau of Labor Statistics and our sample period covers 1987-2012. See Appendix B.7 for details on the construction of the industry-level outcomes from the BLS data. The first-stage F-statistic in Column (5) is 103.2. Columns 6-7: Columns (6) and (7) of this table show the estimated slope coefficients α, γ and δ (times 100) from Equation (29) in the main text. The dependent variable is workers' cumulative earnings growth (net of lifecycle effects) over the next 5 years. We report t-statistics from standard errors clustered at the occupation–industry level beneath coefficient estimates, and we normalize the coefficients to correspond to a shift from the median to the 90th percentile for individual-level exposures ξ^R and ξ^N and a one-standard deviation increase for $\bar{\xi}$. Since we use 4-digit NAICS level average innovation exposure $\bar{\xi}$, we replace the usual 4-digit NAICS $\times$ year fixed effects with 2-digit NAICS $\times$ year fixed effects. Specifications also control for the log of industry employment and include occupation $\times$ year, within occupation–industry income bin $\times$ year fixed effects, and controls listed under Table 3. See the main text and notes to Table 3 for further details.

Figure A.1: Technology exposure and worker earnings growth, by industry or occupation type (alternative measure based on low-/high-skill tasks)

A. Manufacturing vs Non-Manufacturing

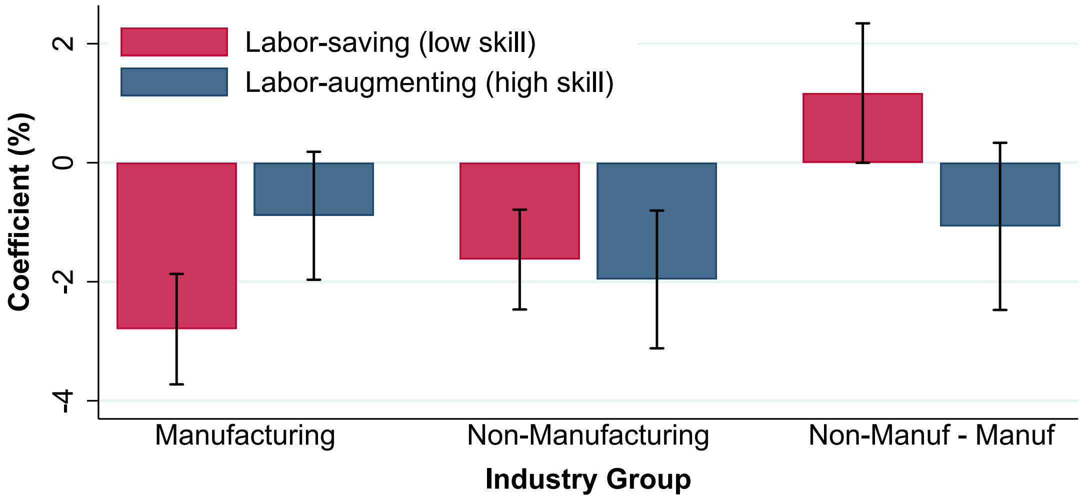

B. Occupation Task Type: Manual/Cognitive/Interpersonal

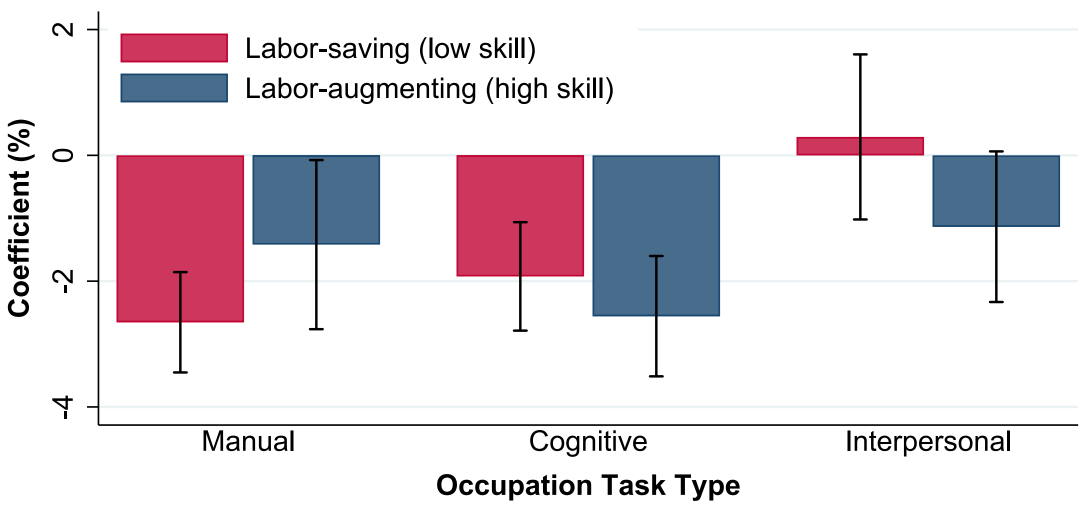

Note: This figure shows the estimated slope coefficients γ and δ (times 100) from Equation (21) in the main text, except that we now use ChatGPT to classify individual tasks as requiring either low or high specific preparation. We allow these coefficients to vary by industry type (Panel A) or occupation task type (Panel B). In Panel A, we compare coefficients for individuals employed in or out of manufacturing (broadly defined as 2-digit NAICS codes 11 through 33); in Panel B, we designate occupations as primarily focusing on either manual, cognitive, or interpersonal tasks using task scores from Acemoglu and Autor (2011). The dependent variable is workers' cumulative earnings growth (net of lifecycle effects) over the next 5 years. We plot 95% confidence intervals from standard errors clustered at the occupation–industry level beneath coefficient estimates, and normalize the coefficients to correspond to a shift from the median to the 90th percentile. All specifications include industry $\times$ year, occupation $\times$ year, within occupation–industry income bin $\times$ year fixed effects, dummies at the level of coefficient interaction, and controls listed under Table 3. See the main text and notes to Table 3 for further details.

Figure A.2: Wage growth and future technology exposure
Alternative definition (low-/high-skill)

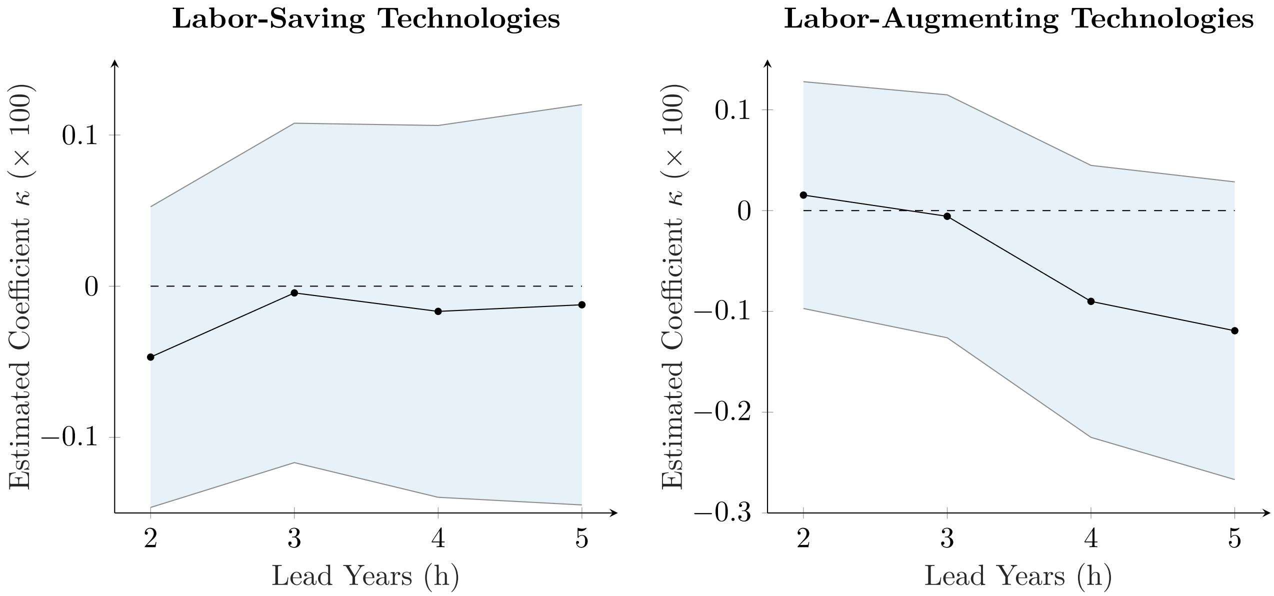

Note: This plot depicts the estimated coefficients κ (multiplied by 100) of Equation (23) in the main document together with 95% error bands. The coefficients correspond to a regression of future technology exposure ξ^j_{t+5+h} for horizons $h = 2 \ldots 5$ on the wage growth of affected workers Δw^i_{t+5}. The left figure shows the response of future labor-saving technologies (defined based on the distance to low-skill tasks), whereas the figure on the right shows the response of future labor-augmenting (based on the distance to high-skill tasks).

Table A.1: Specific examples of labor-saving and labor-augmenting Technologies

A. Labor-Saving			
Patent	Issue Number(s)	Occupation	SOC Code
System for processing transactions in a medical savings fund account	6,044,352	Insurance claims processing clerks	43-9041
Automated manufacture of semi conductors	3,933,538/5,071,776/5,135,608/5,353,498	Semiconductor assembler	51-9141
Automated payment systems for parking facilities	5,737,710/6,796,499	Parking lot attendants	53-6021
Automated securities trading system	4,674,044	Brokerage clerks	43-4011
Method for automatically installing software	6,948,168	computer programmers	15-1131
Automated disease management system	6,234,964	medical assistants	31-9092
Process for amplifying nucleic acid sequences	4,683,202	medical and clinical laboratory technicians	29-2011/29-2012
Telecommunication user account management system and method	5,696,906/5,884,284	Order-fulfillment clerks	43-4150
System for secure transaction management and electronic rights protection	5,892,900	Order-fulfillment clerks	43-4150
Method and apparatus for determining behavioral profile of a computer user	5,848,396	Order-fulfillment clerks	43-4150
Method and apparatus for facilitating evaluation of business opportunities for supplying goods and/or services to potential customers	5,627,973	Order-fulfillment clerks	43-4150

B. Labor-Augmenting			
Patent	Issue Number(s)	Occupation	SOC Code
System for estimating construction costs	5,189,606	Architects	11-9041
Air-to-ground communication systems	5,123,112/5,212,804	Radio operators	27-4013
Environmental screening system	5,687,093	Surveyors	17-1022
Health monitoring system	5,544,649	Radiologists	29-2034
System for managing customer orders	5,666,493	Production schedulers	43-5060
Movie recording system	3,934,268	Motion picture projectionists	39-3021
Process for amplifying nucleic acid sequences	4,683,202	Biochemists	19-1021

C. Labor-Saving and Labor-Augmenting			
Patent	Issue Number(s)	Occupation	SOC Code
Automated loan processing system	5,870,721	Loan officers	13-2072
Automatic theater ticket concierge	5,797,126	Ushers and ticket takers	39-3031
Waiting line management system	5,502,806	Ushers and ticket takers	39-3031
Sous-vide method of cooking	3,966,980	Line cooks	35-2012
Automated trading platform	4,799,156 and 5,794,207	securities brokers	41-3031
Computerized travel reservation system	5,644,727	Ticketing agents	43-4181
Automatic manage a client's budgetary and financial affairs	5,644,727	Financial managers	11-3031
Spelling and grammar checking system	6,424,983	Word processors and typists	43-9022

Note: This plot lists select examples of patented technologies that are primarily labor-saving for a particular occupation (Panel A); primarily labor-augmenting for a particular occupation (panel B); or both labor-saving and labor-augmenting (panel C). See the main text for further details.

Table A.2: First stages for shift-share IV regressions

	Industry–Occupation		Worker			
	(1)		(2)		(3)	
First-stage for:	ξ^R	ξ^N	ξ^R	ξ^N	ξ^R	ξ^N
Z^R	0.62	-0.20	0.61	-0.16	0.64	-0.17
	(36.21)	(-20.30)	(46.97)	(-20.02)	(54.83)	(-19.98)
Z^N	-0.12	0.39	-0.20	0.47	-0.20	0.47
	(-6.04)	(25.61)	(-12.78)	(28.39)	(-13.52)	(29.14)
F-stat	667.9	704.4	1,348	1,381	1,725	1,427
Observations	64,500	64,500	1,213,000	1,213,000	1,213,000	1,213,000
Controls	Y		Y		Y	
Fixed Effects						
Industry (NAICS 4-digit)						
Occupation (occ1990dd)			Y			
Industry × Year	Y		Y		Y	
Occupation × Year	Y				Y	
Prior Income Rank × Year			Y		Y	

Note: Columns (1) to (3) of this table report the first-stage estimates corresponding to Column (2) of Table 2 and Columns (5) and (6) of Table 3, respectively. We report t-statistics clustered by occupation–industry in parentheses and resulting F-statistics for each first-stage regression. See the notes to Tables 2 and 3 for further details. The shift-share IVs Z^R and Z^N are constructed following Equation (17) in the main text.

Table A.3: Technology exposure and individual worker earnings, longer horizons (10 years)

	OLS				IV	
	(1)	(2)	(3)	(4)	(5)	(6)
Labor-saving Exposure ξ^R	-2.00	-2.00	-2.16	-2.35	-1.95	-2.24
	(-7.95)	(-8.08)	(-6.56)	(-6.87)	(-2.92)	(-3.33)
Labor-augmenting Exposure ξ^N	-0.83	-0.77	-1.27	-1.66	-0.70	-1.30
	(-2.40)	(-2.28)	(-2.06)	(-2.70)	(-0.62)	(-1.11)
Worker Controls	Y	Y	Y	Y	Y	Y
Fixed Effects						
Industry (NAICS 4-digit)	Y	Y				
Occupation (occ1990dd)	Y		Y		Y	
Industry × Year			Y	Y	Y	Y
Occupation × Year		Y		Y		Y
Prior Income Rank × Year	Y	Y	Y	Y	Y	Y

Note: This table shows the estimated slope coefficients γ and δ (times 100) from Equation (19) in the main text, where we extend the forward-looking horizon h to 10 years instead of 5 years. See the notes to Table 3 for further details.

Table A.4: Technology exposure and aggregate labor market outcomes
(alternative measure based on low-/high-skill tasks or assigning patents to industries using the BEA IO table)

	Labor Compensation (Wage Bill)					
	Total Effect		Employment		Worker Earnings	
	L/H Skill	BEA IO	L/H Skill	BEA IO	L/H Skill	BEA IO
	(1)	(2)	(3)	(4)	(5)	(6)
Exposure to labor-saving (ξ^L/ξ^R)	-5.72	-5.36	-5.18	-4.41	-0.54	-0.95
	(-4.28)	(-3.48)	(-4.05)	(-2.87)	(-1.75)	(-2.08)
Exposure to labor-augmenting (ξ^H/ξ^N)	7.60	9.19	7.20	8.74	0.40	0.45
	(4.79)	(3.81)	(4.66)	(3.69)	(0.95)	(0.79)
Observations	64,500	64,500	64,500	64,500	64,500	64,500
Fixed Effects						
Occupation × Calendar Year	Y	Y	Y	Y	Y	Y
NAICS4 × Calendar Year	Y	Y	Y	Y	Y	Y
Demographic controls	Y	Y	Y	Y	Y	Y
Lagged Employment	Y	Y	Y	Y	Y	Y
Lagged Wage	Y	Y	Y	Y	Y	Y

Note: This table shows the estimated slope coefficients γ and δ (times 100) from Equation (19) in the main text, except that we now use ChatGPT to classify individual tasks as requiring either low or high specific preparation (columns labeled "L/H skill"), or assign patents to downstream industries using BEA input–output tables (columns labeled "BEA IO"). The dependent variables are the cumulative log changes in the wage bill, employment, and average worker earnings over the next 10 years for an aggregate occupation–industry cell. Columns 1 and 2 examine the impact on total labor compensation (the wage bill). The next four columns decompose the effect on labor compensation into the effect on employment (Columns 3 and 4) and the effect on average earnings (Columns 5 and 6). We report t-statistics based on standard errors clustered at the occupation–industry level beneath coefficient estimates and normalize the coefficients to correspond to a shift from the median to the 90th percentile. We winsorize the total wage bill and employment growth at the 1% level each year; the difference between the two gives the growth rate average worker earnings. The data come from the 1980, 1990, and 2000 Decennial Census, as well as the 2008-2012 5-year ACS panel. We aggregate occupational employment, wages, wage bill, and demographic characteristics by occupation code and industry cells. For this analysis we create a modified version of 4-digit NAICS codes that allows us to construct a consistent crosswalk between the Census industry codes found in the Decennial/ACS and modified NAICS codes and we recompute all measures analogously at the modified NAICS industry level. We use restricted-access versions of these surveys available on Census data servers; the main advantage of this version of the data is that earnings survey responses are not top-coded, as is the case with the publicly available versions. Observation counts are rounded in accordance with Census disclosure rules. See Appendix B.9 for details on the construction of the aggregated occupation–industry–level panel and Appendix B.6 for information on the BEA IO assignment of patents to downstream industries.

Table A.5: Worker technology exposure: homogeneous and interacted with age and prior income (alternative measures based on low-/high-skill tasks or assigning patents to industry using BEA IO table)

	Alternative Technology Exposure			
	Low-/High-Skill		BEA IO Assignment	
	Labor-Saving (ξ^L)	Labor-Augmenting (ξ^H)	Labor-Saving (ξ^R)	Labor-Augmenting (ξ^N)
A. Homogeneous	-1.93	-1.54	-2.09	-1.14
	(-5.82)	(-3.22)	(-5.82)	(-1.64)
B. Worker Age				
25-35 y/o	-1.89	-0.40	-1.99	0.38
	(-4.77)	(-0.76)	(-4.83)	(0.49)
35-45 y/o	-1.34	-1.95	-1.48	-1.46
	(-3.83)	(-3.84)	(-3.78)	(-1.98)
45-55 y/o	-2.87	-1.76	-3.06	-1.71
	(-6.99)	(-3.46)	(-7.01)	(-2.41)
Older-Younger	-0.98	-1.36	-1.07	-2.09
	(-2.57)	(-3.50)	(-2.65)	(-4.19)
C. Income (relative to Ind × Occ peers)				
0–25th percentile	-2.30	-0.78	-1.90	-0.73
	(-5.91)	(-1.41)	(-4.31)	(-0.93)
25–50th percentile	-1.75	-0.90	-1.63	-0.83
	(-4.82)	(-1.69)	(-4.13)	(-1.06)
50–75th percentile	-1.84	-1.41	-2.21	-0.96
	(-5.12)	(-2.89)	(-5.76)	(-1.36)
75–95th percentile	-1.69	-2.65	-2.57	-1.69
	(-4.10)	(-5.11)	(-5.82)	(-2.30)
95–100th percentile	-2.03	-5.27	-2.72	-4.30
	(-3.29)	(-6.37)	(-3.59)	(-4.21)
Top-Bottom	0.27	-4.49	-0.82	-3.57
	(0.45)	(-5.25)	(-1.09)	(-3.32)

Note: This table shows the estimated slope coefficients γ and δ (times 100) from Equation (21) in the main text, except that we now use ChatGPT to classify individual tasks as requiring either low or high specific preparation required (first 2 columns), or assign patents to downstream industries using BEA input–output tables (last 2 columns). We allow these coefficients to vary by age group (Panel A); within occupation–industry income rank (Panel B); or 4-year college graduate status (Panel C). The dependent variable is workers' cumulative earnings growth (net of lifecycle effects) over the next 5 years. We report t-statistics based on standard errors clustered at the occupation–industry level beneath coefficient estimates, and normalize the coefficients to correspond to a shift from the median to the 90th percentile. All specifications include industry × year, occupation × year, within occupation–industry income bin × year fixed effects, dummies at the level of coefficient interaction, and controls listed under Table 3. See the main text and notes to Table 3 for further details on the specification and Appendix B.6 for information on the BEA IO assignment of patents to downstream industries.

Table A.6: Technology exposure and worker earnings growth by income rank, alternative income rankings

Income Rank	Baseline		Drop Recent Hires		2-yr Avg Income		Firm Adjusted		Residual Earnings					
									Age/Sex/CZ		Age/Sex/CZ/F		Age/CZ/Union	
	(1)		(2)		(3)		(4)		(5)		(6)		(7)	
	ξ^R	ξ^N	ξ^R	ξ^N	ξ^R	ξ^N	ξ^R	ξ^N	ξ^R	ξ^N	ξ^R	ξ^N	ξ^R	ξ^N
0–25th percentile	-2.14	-0.89	-2.43	-0.79	-2.48	-0.85	-2.49	-0.82	-2.76	-0.58	-2.82	-0.17	-2.53	-0.80
	(-6.14)	(-1.51)	(-6.41)	(-1.29)	(-7.55)	(-1.58)	(-6.84)	(-1.38)	(-7.34)	(-0.94)	(-7.31)	(-0.28)	(-3.66)	(-0.79)
25–50th percentile	-1.65	-0.89	-1.40	-0.68	-1.41	-1.07	-1.68	-0.83	-1.80	-0.24	-1.70	-0.49	-1.80	-0.90
	(-5.27)	(-1.56)	(-4.24)	(-1.14)	(-4.75)	(-2.08)	(-5.00)	(-1.44)	(-5.61)	(-0.41)	(-5.28)	(-0.87)	(-3.07)	(-0.95)
50–75th percentile	-2.06	-1.09	-1.85	-0.85	-1.69	-1.46	-1.93	-1.02	-1.86	-1.07	-1.61	-1.02	-1.98	-0.78
	(-6.67)	(-2.06)	(-5.80)	(-1.53)	(-5.73)	(-2.92)	(-5.83)	(-1.82)	(-6.01)	(-1.92)	(-5.16)	(-1.79)	(-3.44)	(-0.83)
75–95th percentile	-2.29	-1.99	-2.09	-1.50	-2.02	-1.95	-1.85	-1.90	-2.17	-1.37	-2.02	-1.27	-2.53	-1.22
	(-6.67)	(-3.55)	(-5.92)	(-2.57)	(-6.30)	(-3.68)	(-5.33)	(-3.27)	(-6.38)	(-2.42)	(-5.79)	(-2.25)	(-4.03)	(-1.22)
95–100th percentile	-2.58	-4.64	-2.40	-4.00	-2.05	-4.36	-2.84	-3.56	-2.79	-4.50	-2.73	-4.13	-1.31	-5.02
	(-4.39)	(-5.42)	(-4.07)	(-4.63)	(-9.63)	(-5.02)	(-8.48)	(-3.91)	(-4.94)	(-5.49)	(-4.40)	(-4.82)	(0.34)	(-3.40)
Top-Bottom (p-val)	0.43	0.00	0.96	0.00	0.40	0.00	0.54	0.00	0.96	0.00	0.89	0.00	0.40	0.00

Note: This table shows the estimated slope coefficients γ and δ (times 100) from Equation (21) in the main text, where these coefficients vary with worker earnings rank. The dependent variable is workers' cumulative earnings growth (net of lifecycle effects) over the next 5 years. We report standard errors clustered at the occupation–industry (NAICS 4-digit) level beneath coefficient estimates, and normalize the coefficients to correspond to a shift from the median to the 90th percentile. Columns (1) through (7) use different methods for ranking workers based on earnings. Column (1) represents our baseline method of sorting workers each year within occupation–industry. Column (2) uses our baseline income sort but drops workers who were hired within the last year. Column (3) ranks workers on average earnings over the past two years. Column (4) adjusts workers' earnings for average firm wages by subtracting off the log average wage of a worker's employer in the Longitudinal Business Database from the worker's log wage. Column (5) residualizes log earnings with respect to yearly fixed effects for occupation–industry, commuting zone, and 10-year age bin (25-35, 35-45, 45-55) interacted with gender. Column (6) residualizes income with respect to the same characteristics as column (5), plus year-specific Longitudinal Business Database firm wage decile bins; Column (7) restricts attention to the subsample of workers who are in-universe for the CPS union membership question (approximately 20% of the sample), and residualizes log earnings with respect to yearly fixed effects for occupation, industry, worker union status, and commuting zone. The bottom panel of the table reports the p-value corresponding to a two-sided test of coefficient equality between the top and bottom income bin coefficients in each column. All specifications include industry $\times$ year, occupation $\times$ year, within occupation–industry income bin $\times$ year fixed effects, dummies at the level of coefficient interaction, and controls listed under Table 3. See the main text and notes to Table 3 for further details.

Table A.7: Technology spillovers (alternative measures based on low- and high-skill tasks or using BEA IO table)

	Industry Productivity		
	Baseline	L/H Skill	BEA IO
Overall Industry Technology Exposure ($\bar{\xi}$)	6.02	6.02	6.09
	(6.12)	(6.14)	(5.96)
Controls			
Log Industry Employment	Y	Y	Y
Lagged 3-year growth rate	Y	Y	Y
NAICS2 $\times$ year FE	Y	Y	Y
Observations	2,600	2,600	2,600

Note: This table shows the estimated slope coefficients γ and δ (times 100) from Equation (28) in the main text, except thaat we now use ChatGPT to classify individual tasks as requiring either low or high specific preparation required (second column), or assign patents to downstream industries using BEA input–output tables (last column). The dependent variable is the annualized log changes in the industry productivity index. We report t-statistics based on standard errors clustered at the industry level beneath coefficient estimates; the coefficients correspond to one-standard deviation increases in the given measure. Regressions additionally control for the log of industry employment and lagged 3-year growth rates in the dependent variable. Industry-level measures are the yearly averages of the worker-level occupation–industry exposures as in Equation 13. The overall technology exposure is the standardized average of the two individual technology exposures. Data on industry outcomes are from the Bureau of Labor Statistics and our sample period covers 1987–2012. Observation counts are rounded in accordance with Census disclosure rules. See Appendix B.7 for details on the construction of the industry-level outcomes from the BLS data, and Appendix B.6 for information on the BEA IO assignment of patents to downstream industries.

Table A.8: Technology spillovers (industry labor share)

	Labor Share Growth	
	OLS	IV
Industry Exposure to Labor-saving ($\bar{\xi}^R$)	-2.32	-2.57
	(-2.46)	(-2.20)
Industry Exposure to Labor-Augmenting ($\bar{\xi}^N$)	0.92	1.56
	(1.32)	(1.64)
Controls		
Log Industry Employment	Y	Y
Lagged 3-year Growth Rate	Y	Y
NAICS2 $\times$ Year FE	Y	Y
Observations	2,600	2,600

Note: This table shows the estimated slope coefficients γ and δ (times 100) from Equation (28) in the main text, except the dependent variable is now the 5-year cumulative log change in industry labor share. We report t-statistics based on standard errors clustered at the industry level beneath coefficient estimates; the coefficients correspond to one-standard deviation increases in the given measure. Regressions additionally control for the log of industry employment and lagged 3-year growth rates in the dependent variable. Industry-level measures are the yearly employment-weighted averages of the worker-level occupation–industry exposures as in Equation 13. The overall technology exposure is the standardized average of the two individual technology exposures. Data on industry outcomes are from the Bureau of Labor Statistics and our sample period covers 1987–2012. Observation counts are rounded in accordance with Census disclosure rules. See Appendix B.7 for details on the construction of the industry-level outcomes from the BLS data. The first-stage F-stat for industry exposure to labor-substituting is 136.7; for labor-augmenting the F-stat is 44.01.

Table A.9: Model estimates

A. Target Moments		
Moment or regression coefficient in (%)	Model	Data
Response of worker earnings to ξ^R, homogeneous coefficient	-2.38	-2.06
Response of worker earnings to ξ^N, homogeneous coefficient	-1.30	-1.31
Response of worker earnings to ξ^N, homogeneous coefficient, new entrants	4.50	4.10
Response of worker earnings to ξ^N, 0–25th percentile in relative income	-0.25	-0.89
Response of worker earnings to ξ^N, 25–50th percentile in relative income	-1.03	-0.89
Response of worker earnings to ξ^N, 50–75th percentile in relative income	-1.57	-1.09
Response of worker earnings to ξ^N, 75–95th percentile in relative income	-2.19	-1.99
Response of worker earnings to ξ^N, 95–100th percentile in relative income	-3.01	-4.64
Response of worker earnings to ξ^R (regression with $\bar{\xi}$), homogeneous coefficient	-2.23	-2.77
Response of worker earnings to ξ^N (regression with $\bar{\xi}$), homogeneous coefficient	-1.42	-2.01
Response of worker earnings to $\bar{\xi}$, homogeneous coefficient	2.67	2.53
Response of industry productivity to composite exposure measure $\bar{\xi}$	6.98	6.02
Response of industry labor share to labor-saving tech exposure $\bar{\xi}^R$	-1.84	-2.32
Response of industry labor share to labor-augmenting tech exposure $\bar{\xi}^N$	0.88	0.92
Response of employment to labor-saving tech exposure ξ^R	-2.24	-2.10
Response of employment to labor-augmenting tech exposure ξ^N	4.23	4.14
Labor share (%), mean value	82.85	80.00

B. Parameter Estimates		
Parameter	Symbol	Estimate
Elasticity of substitution across tasks	ψ	1.212
Elasticity of substitution across industries	χ	1.715
Elasticity of substitution between capital and labor, routine tasks	ν_R	1.699
Elasticity of substitution between capital and labor, non-routine tasks	ν_N	0.863
Elasticity of labor supply, routine tasks	ζ_R	0.819
Elasticity of labor supply, non-routine tasks	ζ_N	0.951
Mean skill loss for incumbent workers	β	0.040
Capital-labor expenditure ratio, routine tasks	κ_R	0.157
Capital-labor expenditure ratio, non-routine tasks	κ_N	0.233
Skill loss across technology vintages	ω	0.011
Ratio of noise volatility to signal volatility for industry $\bar{\xi}^R$, $\bar{\xi}^N$ and $\bar{\xi}$	γ	1.076

Note: This table presents GMM estimates of our model from Section 1 in the main text taken by minimizing the objective function (A.78) across 11 parameters (plus the two cross-equation restrictions $\Gamma_R = \Gamma_N$ and $A_R = A_N$) to match 16 target moments with their model-implied counterparts. We calibrate coefficients to a 5-year horizon, and we scale the individual-level worker earnings responses to technology exposure to correspond to an increase from the median to the 90th percentile of their empirical distributions, and we scale responses to industry-level aggregates $\bar{\xi}$, $\bar{\xi}^R$, and $\bar{\xi}^N$ to correspond to an empirical standard-deviation increase. The labor share target comes from Koh, Santaeulalia-Llopis, and Zheng (2020), which we further adjust by excluding the payments to structures in total capital income by using an estimate of structures' income share from Ohanian, Orak, and Shen (2023). See Appendix Section A for model details, including subsection A.7 for an in-depth discussion of model estimation. Additionally, see Table A.10 for the list of closed-form model-implied coefficients in terms of model parameters.

Online Appendix

A Model Appendix

In this section, we provide additional details about the model derivation and calibration.

A.1 Setup

Here, we provide the complete solution of the model with skill displacement. For convenience, we repeat some of the assumptions of the model from section 1 of the main text here.

Aggregate output is a CES aggregate over output produced across a continuum of industries

$$\bar{Y} = \left(\int_k Y(k)^{\frac{\chi-1}{\chi}} dk \right)^{\frac{\chi}{\chi-1}}, \tag{A.1}$$

where $\bar{Y}$ is a constant and $\chi > 0$ is a parameter which captures the elasticity of demand for industry output. Industries are competitive and there is free entry of all firms in the production sector. In equilibrium, prices of industry output are equal to its marginal cost, and given the constant returns to production firms make zero profits. As stated in the main text, we then focus on a given industry and drop the k subscripts unless needed.

The industry-level output Y is a CES aggregate of a large number of intermediate tasks, indexed by $j \in \{1, \ldots, J\}$:

$$Y = \left(\sum_{i=1}^{J} y(j)^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}}, \tag{A.2}$$

where the parameter $\psi > 0$ indexes the elasticities of substitution across tasks j, as well as the absolute value of the demand elasticity for each task output. Each task j is produced using $n(j)$ and labor $l(j)$ according to

$$y(j) = \left((1-\gamma_j) n(j)^{\frac{\nu_j - 1}{\nu_j}} + \gamma_j l(j)^{\frac{\nu_j - 1}{\nu_j}} \right)^{\frac{\nu_j}{\nu_j - 1}}, \tag{A.3}$$

where the parameter $\nu_j > 0$ governs the elasticities of substitution between $l(j)$ and $n(j)$ when producing each task and the parameter $\gamma_j \in (0, 1)$. As stated in the main text, we partition the set of tasks into routine and non-routine tasks, so $J = J_R \cup J_N$, and $\nu_j \in \{\nu_R, \nu_N\}$ with $\nu_N < \nu_R$.

We model the cumulative impact of an industry's breakthrough innovations between the first and second period as lowering the price of capital specific to task j,

$$\Delta \log q(j) = -\varepsilon(j) \tag{A.4}$$

where $\varepsilon \equiv [\varepsilon_1 \ldots \varepsilon_J]$ is a vector of weakly positive random variables jointly distributed according to $f(\varepsilon)$.

There is a continuum of measure I workers supplying labor across the different tasks j, so that the aggregate supply of labor in task j, $L(j)$, is given by:

$$L(j) = \int_0^I l(i,j)di, \tag{A.5}$$

where $l(i,j)$ is the number of efficiency units of labor supplied by worker i in task j.

Each worker i is associated with a single occupation $o(i)$. Workers in occupation o supply positive output in for a small number of routine tasks and non-routine tasks, we denote by $J_o \in J$ the set of tasks performed by occupation o, and we have $\cup J_o = J$ and $\cap J_o = \emptyset$.

There are changes in aggregate and individual labor supply after technology shocks. We allow for skill displacement at the individual level: the adoption of a new vintage of technology leads to a change not only in wages but also in the number of efficiency units of a worker's human capital. Specifically, changes in worker i's productivity in task j satisfies

$$\Delta \log l(i,j) = \begin{cases} -\beta I[i \in J_N]\varepsilon(j) + u_{i,j} - \log l(i,j) & \text{w/ prob. } I[i \in J_N]\,\omega\,\varepsilon(j) \\ -\beta I[i \in J_N]\varepsilon(j) & \text{otherwise} \end{cases}, \tag{A.6}$$

where ω is a positive constant indicating the strength of the skill displacement effect, and $u(i,j)$ a new, i.i.d. draw from the ergodic distribution of $\log z(i,j)$ for incumbent workers in the same occupation. The first term says that larger shifts in the technology frontier are likely to generate greater amounts of skill displacement among incumbent workers, while the second term is a force of redistribution: there is some mean reversion in workers' productivity levels, so the most productive workers in the current vintage are most likely to experience declines in their productivity if technology changes.

The aggregate labor supply for task j varies on the extensive margin. We assume that aggregate labor supply satisfies

$$\Delta \log L(j) = \bar{\zeta} + \zeta_j \Delta \log w(j), \tag{A.7}$$

where ζ_j indexes the elasticity of aggregate labor supply in performing task j to labor compensation and $\bar{\zeta}$ is a constant which plays no material role in our analysis. Equation (A.7) makes implicit assumptions about the labor supply of new workers. In particular, the quantity of labor supplied by incumbent workers is pinned down by equation (A.6). Changes in aggregate labor supply are also affected by the entry of new workers. Equation (A.7) and equation (A.6) imply that new entrants are, on average, more highly skilled in the new technology than incumbents – exactly so that the aggregate loss of skill for incumbent workers is offset by the higher skill of new entrants, implying the absence of skill displacement in the aggregate.

For further notational convenience, according to (A.7), we capture the labor supply function in each period

$$L(j) = L_0(j)w(j)^{\zeta(j)} \tag{A.8}$$

where $L_0(j)$ is a constant which captures impacts of $\bar{\zeta}$ and also potentially other shifters of the aggregate supply labor in task j.

A.2 Equilibrium conditions

A competitive equilibrium in this economy is a set of prices and allocations which satisfies the following conditions:

- The representative firm chooses $l^*(j), n^*(j)$ to maximize profits, taking factor prices $w^*(j), q(j)$ and the output price c_y^* as given,
- Labor supply equals labor demand, $L(j) = l^*(j)$, and
- Output markets clear: the quantity demanded equals the quantity supplied given output price c_y^* (which equals the marginal cost of production), $l^*(j)$, and $n^*(j)$.

Consider the firm's cost minimization problem across tasks j:

$$\min_{y(j)} \sum_{j \in J} p(j)y(j) \qquad s.t.\;\; Y = \left(\sum_{j \in J} y(j)^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}} \tag{A.9}$$

where $y(j) = \left[(1-\gamma_j)n(j)^{\frac{\nu_j-1}{\nu_j}} + \gamma_j\, l(j)^{\frac{\nu_j-1}{\nu_j}} \right]^{\frac{\nu_j}{\nu_j-1}}$. Here $p(j)$ is the per-unit cost index for task j after optimal input choices have been made within task j. Under perfect competition, $p(j)$ equals the firm's per-unit marginal cost of producing $y(j)$. Our CES structure admits the following Hicksian demand for $y(j)$ from the first order condition for problem (A.9):

$$y(j) = \frac{1}{p(j)^{\psi}} \left[J^{-\psi} \sum_{j \in J} p(j)^{1-\psi} \right]^{\frac{\psi}{1-\psi}} Y = p(j)^{-\psi} X^{\chi-\psi} \bar{Y}, \tag{A.10}$$

where the second equality in (A.10) holds from defining the level of industry productivity (output per dollar of input expediture) X as

$$X \equiv \left(\sum_{j=1}^{J} p(j)^{1-\psi} \right)^{-\frac{1}{1-\psi}}. \tag{A.11}$$

and applying the market clearing condition in the output market:

$$\left(\frac{Y}{\bar{Y}} \right)^{-\frac{1}{\chi}} = c_y = X^{-1}. \tag{A.12}$$

Next, we solve the cost minimization problem within task j:

$$\min_{l(j), n(j)} q(j)n(j) + w(j)l(j) \qquad s.t.\;\; y(j) = \left[(1-\gamma_j)n(j)^{\frac{\nu_j-1}{\nu_j}} + \gamma_j\, l(j)^{\frac{\nu_j-1}{\nu_j}} \right]^{\frac{\nu_j}{\nu_j-1}}. \tag{A.13}$$

Defining the constants $a_i \equiv \gamma_j^{\nu_j}$, $b_j \equiv (1-\gamma_j)^{\nu_j}$, solving (A.13) yields the per-unit cost of input j, $p(j)$:

$$p(j) = \left(a_j w(j)^{1-\nu_j} + b_j q(j)^{1-\nu_j}\right)^{\frac{1}{1-\nu_j}}. \tag{A.14}$$

Plugging (A.10) and (A.14) in the CES Hicksian demand for $n(j)$ and $l(j)$ and imposing labor market clearing gives

$$L(j) = L_0(j) w(j)^{\zeta(j)} = \frac{a_j}{w(j)^{\nu_j}} \left(a_j w(j)^{1-\nu_j} + b_j q(j)^{1-\nu_j}\right)^{\frac{\nu_j-\psi}{1-\nu_j}} X^{\chi-\psi}\bar{Y} = l(j). \tag{A.15}$$

Since $q(j)$ is exogenous, (A.15) provides a condition which implicitly defines $w(j)$. $n(j)$ is then uniquely determined by $w(j)$ and the exogenous $q(j)$:

$$n(j) = \frac{b_j}{q(j)^{\nu_j}} \left(a_j w(j)^{1-\nu_j} + b_j q(j)^{1-\nu_j}\right)^{\frac{\nu_j-\psi}{1-\nu_j}} X^{\chi-\psi}\bar{Y}. \tag{A.16}$$

Taking the ratio of (A.15) and (A.16), we immediately obtain the relative expenditure shares of capital and labor in producing task j:

$$\kappa(j) \equiv \frac{q(j)n(j)}{w(j)l(j)} = \frac{b_j q(j)^{1-\nu_j}}{a_j w(j)^{1-\nu_j}}. \tag{A.17}$$

Implicitly differentiating the above equations thus allows us to perform comparative statics to advances in technology.

A.3 Comparative statics at the task level

We then analyze the impact of technology in producing capital specific to task j on the labor compensation for performing task j. The total effect of innovation on equilibrium wages is the sum of a direct effect on wages holding constant the aggregate productivity X, and an indirect effect which captures the effect of aggregate productivity X on equilibrium wages. All cross-task dependencies operate only through indirect effects.

The direct effect can be obtained by applying the implicit function theorem to (A.15) that

$$-\frac{\partial \log w(j)}{\partial \log q(j)} = \frac{\partial \log w(j)}{\partial \varepsilon(j)} = \frac{(\psi-\nu_j)\kappa(j)}{(\psi+\zeta_j)+(\nu_j+\zeta_j)\kappa(j)} = \frac{\psi-\nu_j}{\nu_j+\zeta_j}\Gamma_j. \tag{A.18}$$

The last equality in (A.18) holds because we define Γ_j as the partial elasticity of the marginal cost of producing task j to changes in the capital price $q(j)$ holding aggregate productivity constant:

$$\Gamma_j \equiv \frac{\partial \log p(j)}{\partial \log q(j)} = \left.\frac{\partial \log p(j)}{\partial \log w(j)}\right|_{q(j)} \left.\frac{\partial \log w(j)}{\partial \log q(j)}\right|_{X} + \left.\frac{\partial \log p(j)}{\partial \log q(j)}\right|_{w(j)} = \frac{(\nu_j+\zeta_j)\kappa(j)}{(\psi+\zeta_j)+(\nu_j+\zeta_j)\kappa(j)}. \tag{A.19}$$

Next, we examine the indirect effect through changes of overall level of industry productivity. Analogous to (A.18), the partial elasticity of the wage level to industry productivity holding $q(j)$

fixed is:

$$A_j \equiv \frac{1}{(\chi-\psi)}\frac{\partial \log w(j)}{\partial \log X} = \frac{(1+\kappa_j)}{\psi+\zeta_j+(\nu_j+\zeta_j)\kappa_j}. \tag{A.20}$$

Using (A.18) and (A.20), we can now approximate wage changes due to innovation as

$$\Delta \log w(j) \approx \frac{\partial \log w(j)}{\partial \varepsilon(j)}\varepsilon(j) + \frac{\partial \log w(j)}{\partial \log X}\Delta \log X = \frac{\psi-\nu_j}{\nu_j+\zeta_j}\Gamma_j \varepsilon(j) + (\chi-\psi)A_j \Delta \log X. \tag{A.21}$$

We derive comparative statics for productivity X and other industry aggregate outcomes in section A.5 below.

A.4 Worker earnings growth and technology exposure

Next, we aggregate across tasks to derive predictions for expected changes in incumbent worker earnings in response to innovation. As stated in the main text, we assume that parameters such as Γ_j, A_j do not vary within task type, and denote the corresponding values as Γ_R and A_R for routine tasks and Γ_N and A_N for non-routine tasks.

The level of wage earnings for an individual worker is equal to the total compensation for the tasks she supplies,

$$W(i) = \sum_{j\in J} w(j) l(i,j) I(j \in J_{o(i)}). \tag{A.22}$$

Taking a first order approximation, earnings growth for a worker i is approximately equal to a weighted average of changes in task prices and skills

$$\Delta \log W(i) \approx \sum_{j\in J} s(i,j)\Big(\Delta \log w(j) + \Delta \log l(i,j)\Big), \tag{A.23}$$

where the weights depend on the contribution of each task her occupation performs to her wage

$$s(i,j) = \frac{w(j) l(i,j) I(j \in J_{o(i)})}{W(i)}. \tag{A.24}$$

Note that we assume that workers had similar levels of initial skill in each task they are performing $l(i,j) = \bar{l}(i), \forall j \in J_{o(i)}$. Applying (A.6) and (A.21) we can rewrite (A.23) as

$$\begin{aligned}\Delta \log W(i) &\approx \sum_{j\in J} s(i,j)\left(\frac{\psi-\nu_j}{\nu_j+\zeta_j}\Gamma_j\varepsilon(j) + (\chi-\psi)A_j\Delta\log X - \left(\beta + \omega\Big[\log \bar{l}(i) - \int \log \bar{l}(i)\, dF(i)\Big]\right) I(j\in J_N)\varepsilon(j)\right) \\ &= \left[\frac{\psi-\nu_R}{\nu_R+\zeta_R}\Gamma_R\right]\xi^R(i) + \left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\Gamma_N - \beta - \omega\Big[\log\bar{l}(i) - \int \log \bar{l}(i)\, dF(i)\Big]\right]\xi^N(i) \\ &\quad + (\chi-\psi)\Big[(A_R - A_N)\theta(i) + A_N\Big]\Delta \log X\end{aligned} \tag{A.25}$$

where $F(i)$ denotes the cross-sectional distribution of $\log \bar{l}(i)$ across workers. $\theta(i)$ measures the share of labor compensation to worker i due to performing the routine tasks in her occupation. $\xi^R(i)$ and

$\xi^N(i)$ denote the worker's exposure to labor-saving and labor-augmenting technologies, respectively. Also, we define $\tilde{s}^R(i,j)$ and $\tilde{s}^N(i,j)$ as shares normalized to one within each task type:

$$\theta(i) \equiv \sum_{j\in J_R} s(i,j) \tag{A.26}$$

$$\begin{aligned} \xi^R(i) &\equiv \sum_{j\in J_R} s(i,j)\varepsilon(j) = \theta(i) \sum_{j\in J_R} \frac{s(i,j)}{\sum_{i\in J_R} s(i,j)}\varepsilon(j) = \theta(i) \sum_{j\in J_R} \tilde{s}^R(i,j)\varepsilon(j) \\ \xi^N(i) &\equiv \sum_{j\in J_N} s(i,j)\varepsilon(j) = (1-\theta(i)) \sum_{i\in J_N} \frac{s(i,j)}{\sum_{i\in J_N} s(i,j)}\varepsilon(j) = (1-\theta(i)) \sum_{j\in J_N} \tilde{s}^N(i,j)\varepsilon(j). \end{aligned} \tag{A.27}$$

A.5 Aggregate Effects of Technology Exposure

In this section we examine the changes of aggregate quantiles, including relation between technology exposure and occupation employment, industry productivity, and industry labor share.

A.5.1 Notation at task, occupation, and industry levels

Before discussing comparative statics, we define a number of task- and occupation-level objects which appear in various aggregations below.

We begin at the task level. For task j, the wage bill share $s(j)$, the expenditure share o$s^p(j)$, satisfy

$$s(j) \equiv \frac{w(j)l(j)}{\sum_{k\in J} w(k)l(k)}, \qquad s^p(j) \equiv \frac{p(j)y(j)}{\sum_{k\in J} p(k)y(k)}, \qquad \theta \equiv \sum_{j\in J_R} s(j). \tag{A.28}$$

Given our assumption that $\kappa_j = \kappa_R$ for $j \in J_R$ and $\kappa_j = \kappa_N$ otherwise, we can convert between wage bill share and expenditure shares using our knowledge of the total and task-level labor shares. Specifically,

$$\frac{s(j)}{s^p(j)} = \frac{w(j)l(j)}{\sum_{k\in J} w(k)l(k)} \frac{\sum_{k\in J} p(k)y(k)}{p(j)y(j)} = \frac{LS(j)}{LS}, \tag{A.29}$$

where $LS \equiv \frac{\sum_{k\in J} w(k)l(k)}{\sum_{k\in J} p(k)y(k)}$ is the total industry labor share, and $LS(j) \equiv \frac{w(j)l(j)}{w(j)l(j)+n(j)q(j)} = \frac{1}{1+\kappa_j}$ is the task level labor share. Note there are only two possible value for $LS(j)$, thus we define routine and non-routine labor shares via $LS_R \equiv \frac{1}{1+\kappa_R}$ and $LS_N \equiv \frac{1}{1+\kappa_N}$, respectively.

With these definitions in hand, we can properly aggregate across tasks to define analogous occupation-level variables. We define $s(o)$ the aggregate wage bill share of the occupation, and $s_o(j)$ the relative wage bill share of task j within that occupation. The share of the wage bill in routine tasks for a given occupation $\theta(o)$ is defined by

$$s(o) \equiv \sum_{j\in J_o} s(j), \qquad s_o(j) \equiv \frac{s(j)}{s(o)}, \qquad \theta(o) \equiv \sum_{j\in J_o\cap J_R} s_o(j). \tag{A.30}$$

Next, we denote exposure of labor-saving and labor-augmenting technologies to occupation o $\xi^R(o)$

and $\xi^N(o)$, respectively, in a manner analogous to (A.27) at the worker level:

$$\xi^R(o) \equiv \sum_{j \in J_o \cap J_R} s_o(j)\varepsilon(j), \qquad \xi^N(o) \equiv \sum_{j \in J_o \cap J_N} s_o(j)\varepsilon(j). \tag{A.31}$$

$\xi^R(o)$ and $\xi^N(o)$ are weighted average of task-level innovation measures, where the weight on each task is the task-level wage share across workers within the occupation.

Next, we define wage bill and aggregate employment at the occupation level. It is straightforward to define the wage bill for occupation o $WB(o)$, which satisfies

$$WB(o) = \sum_{j \in J_o} w(j)l(j). \tag{A.32}$$

In contrast, because worker productivity is potentially heterogeneous between workers and across different tasks within a given worker, the definition of the number of efficiency units of labor employed in occupation o is not unique. We define the growth rate in the quantity of labor utilized in occupation o using a Divisia index –i.e., as a wage-bill weighted-average of task level employment growth

$$\frac{\Delta L(o)}{L(o)} = \left[\sum_{j \in o} s_o(j)\frac{\Delta l(j)}{l(j)}\right]. \tag{A.33}$$

Finally, we turn to industry-level measures. Further, we define industry-level measures of aggregate exposure of workers to labor-saving and labor-augmenting technologies – $\bar{\xi}^R$ and $\bar{\xi}^N$, respectively – as wage bill-weighted averages of task level technology exposure:

$$\bar{\xi}^R = \sum_{j \in J_R} s(j)\varepsilon(j) = \sum_o s(o)\xi^R(o), \qquad \bar{\xi}^N = \sum_{j \in J_N} s(j)\varepsilon(j) = \sum_o s(o)\xi^N(o) \tag{A.34}$$

Finally, θ is the total wage bill share of all labor-saving tasks, which is defined according to

$$\theta \equiv \sum_{j \in J_R} s(j). \tag{A.35}$$

A.5.2 Comparative statics for occupation-level labor compensation and employment

Using our definitions and comparative statics above, the occupation level wage bill evolves according to

$$\begin{aligned}
\Delta \log WB(o) &\approx \sum_{j \in J_o} s_o(j)\left[\Delta \log w(j) + \Delta \log l(j)\right] \\
&= \sum_{j \in J_o} s_o(j)(1+\zeta_j)[\frac{\psi - \nu(j)}{\nu(j)+\zeta(j)}\Gamma_j \varepsilon(j) + (\chi - \psi)A_j \Delta \log X] \\
&= \left((\psi - \nu_R)\frac{1+\zeta_R}{\nu_R+\zeta_R}\Gamma_R\right)\xi^R(o) + \left((\psi - \nu_N)\frac{1+\zeta_N}{\nu_N+\zeta_N}\Gamma_N\right)\xi^N(o) \\
&+ (\chi - \psi)[A_R(1+\zeta_R)\theta_o + A_N(1+\zeta_N)(1-\theta_o)]\Delta \log X.
\end{aligned}$$

The response of the wage bill across all workers in the same occupation is equal to the change in wages plus the change in employment. Given our definition of employment in (A.33) above, we obtain a closely related comparative static for employment:

$$\begin{aligned}
\Delta \log L(o) \approx& \sum_{j\in o} s_o(j)\Delta \log l(j) \\
=& \sum_{j\in o} s_o(j)\zeta(j)\Delta \log w(j) \\
\approx& \zeta_R\left[\frac{\psi-\nu_R}{\nu_R+\zeta_R}\Gamma_R\right]\xi^R(o)+\zeta_N\left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\Gamma_N\right]\xi^N(o)+(\chi-\psi)\left[\zeta_R\theta_o A_R+\zeta_N(1-\theta_o)A_N\right]\Delta \log X.
\end{aligned} \tag{A.36}$$

The first two terms in equation (A.36) capture the direct effects of labor-saving and labor-augmenting technologies on employment in affected occupations. The last term capture the employment changes due to changes in aggregate industry productivity; we take a closer look at innovation, industry productivity, and other industry outcomes next.

A.5.3 Comparative statics for industry productivity

Here, we explore how industry productivity changes in response to the set of task specific innovation shocks $\varepsilon(j)$ experienced. We apply the implicit function theorem to the second equal sign of the industry market clearing condition (A.12). Specifically, we take the log of both sides of (A.12) and apply the implicit function theorem to take the derivative with respect to $\varepsilon(j)$, which yields the expression

$$\frac{\partial \log X}{\partial \varepsilon(j)} = -\frac{\frac{\partial \log c_y}{\partial \varepsilon(j)}}{1+\frac{\partial \log c_y}{\partial \log X}}. \tag{A.37}$$

Since $\log X = -\log c_y$, we obtain the numerator of (A.37) directly by computing the elasticity of $\log X$ with respect to $\varepsilon(j)$.

$$\frac{\partial \log c_y}{\partial \varepsilon(j)} = -\left(\frac{p(j)}{c_y}\right)^{1-\psi}\frac{\partial \log p_j}{\partial \log q_j} = -s^p(j)\Gamma_j, \tag{A.38}$$

where in (A.38) we apply the fact that

$$\left(\frac{p(j)}{c_y}\right)^{1-\psi} = \frac{p(j)y(j)}{c_y Y} = s^p(j). \tag{A.39}$$

Recall that $s^p(j)$ is the expenditure share of task j.

Next, we implicitly differentiate (A.12), using the definition of aggregate productivity (A.11), with respect to X in order to find the term in the denominator of (A.37),

$$\frac{\partial \log c_y}{\partial \log X} = \sum_{j\in J}\left(\frac{p(j)}{c_y}\right)^{1-\psi}\left[\frac{\partial \log p(j)}{\partial \log w(j)}\frac{\partial \log w(j)}{\partial \log X}\right] = (\chi-\psi)\sum_{j\in J} s^p(j)\frac{A_j}{1+\kappa_j} = (\chi-\psi)LS\left[(A_R-A_N)\theta+A_N\right], \tag{A.40}$$

where ϵ_c, the elasticity of marginal cost of output c_y to an increase in output, satisfies

$$\epsilon_c \equiv \frac{\partial \log c_y}{\partial \log Y} = \frac{1}{\chi}\frac{\partial \log c_y}{\partial \log X} = LS\frac{\chi - \psi}{\chi}\left[(A_R - A_N)\theta + A_N)\right]. \tag{A.41}$$

Plugging (A.38) and (A.41) into (A.37) implies that, up to a first order approximation, productivity X changes according to

$$\begin{aligned} \Delta \log X &\approx \sum_{j\in J} - \frac{\frac{\partial \log c_y}{\partial \varepsilon(j)}}{1 + \frac{\partial \log c_y}{\partial \log X}}\varepsilon(j) = \frac{1}{1+\chi\epsilon_c}\sum_{j\in J} s^p(j)\Gamma_j\varepsilon(j) \\ &= \frac{\Gamma_R}{1+\chi\epsilon_c}\frac{LS}{LS_R}\bar{\xi}^R + \frac{\Gamma_N}{1+\chi\epsilon_c}\frac{LS}{LS_N}\bar{\xi}^N. \end{aligned} \tag{A.42}$$

A.5.4 Comparative statics for industry labor share

In this section, we analyze the evolution of the labor share of output in response to innovation. While the model makes clear predictions about the evolution of the labor share at the task and occupational levels, data are not typically collected at these levels of aggregation. As a result, we focus attention on LS, the aggregate labor share at the industry level.

Differentiating the definition of LS yields the following decomposition:

$$\begin{aligned} \Delta \log LS &\approx \frac{1}{LS}\sum_{j\in J} s^p(j)LS(j)\left(\Delta \log LS(j) + \Delta \log \frac{p(j)y(j)}{\sum_{k\in J}p(k)y(k)}\right) \\ &= \underbrace{\sum_{j\in J} s(j)\Delta \log LS(j)}_{\text{average within-task LS change}} + \underbrace{\sum_{j\in J} s(j)\left(\Delta \log[p(j)y(j)] - \sum_{k\in J} s^p(k)\Delta \log[p(k)y(k)]\right)}_{\text{between task redistribution}}, \end{aligned} \tag{A.43}$$

where we make use of the identities in (A.29) to simplify the expression. The change in LS is captured by two terms. The first term captures the weighted average change in the labor share within each task, where tasks with higher task shares $s(j)$ get higher weight. The second term captures a reallocation effect, capturing the fact that the aggregate labor share can change if innovation reallocates input expenditures across tasks with different labor share. This term captures a covariance between initial labor shares at the task level and change in shares of input expenditures across different tasks.

First, we consider how the within task term changes with $\varepsilon(j)$, holding industry productivity fixed:

$$\begin{aligned} \frac{\partial \log LS(j)}{\partial \varepsilon(j)} &= \partial \log\left(\frac{\frac{w(j)l(j)}{q(j)n(j)}}{1+\frac{w(j)l(j)}{q(j)n(j)}}\right)/\partial\varepsilon(j) = [1 - LS(j)]\frac{\partial\left(\log[w(j)l(j)] - \log[q(j)n(j)]\right)}{\partial \varepsilon(j)} \\ &= (1-\nu_j)\frac{\psi + \zeta_j}{\nu_j + \zeta_j}\Gamma_j. \end{aligned} \tag{A.44}$$

where in (A.44) we applied the fact that

$$\frac{\partial(\log[w(j)l(j)] - \log[q(j)n(j)])}{\partial\varepsilon(j)} = (1-\nu_j)\frac{\partial(\log w(j) - \log q(j))}{\partial\varepsilon(j)} = (1-\nu_j)(\frac{\partial \log w(j)}{\partial\varepsilon(j)} + 1) \tag{A.45}$$

by taking log of (A.17) and then taking the derivative with respect to innovation. Following the same approach and differentiating with respect to X, we obtain that

$$\frac{\partial \log LS(j)}{\partial \log X} = (1-\nu_j)(1-LS(j))(\chi-\psi)A_j. \tag{A.46}$$

Combining (A.44) and (A.46), we obtain an expression how the within-task labor share changes

$$\Delta \log LS(j) \approx \frac{\partial \log LS(j)}{\partial\varepsilon(j)}\varepsilon(j) + \frac{\partial \log LS(j)}{\partial \log X}\Delta \log X = (1-\nu_j)\left[\frac{\psi+\zeta_j}{\nu_j+\zeta_j}\Gamma_j\varepsilon(j) + (1-LS(j))(\chi-\psi)A_j\Delta\log X\right] \tag{A.47}$$

the direction of whom depends on the elasticity of substitution between capital and labor.

Next, we turn to the between-task re-distribution term, which requires us to obtain changes in task expenditure. As before, we take separate derivatives with respect to $\varepsilon(j)$ and X. By taking the log of (A.10) and applying a simple transformation, we obtain that

$$\log p(j)y(j) = (1-\psi)\log p_i + (\chi-\psi)\log X + \log \bar{Y}. \tag{A.48}$$

Differentiating (A.48) with respect to $\varepsilon(j)$, we find that

$$\frac{\partial \log p(j)y(j)}{\partial\varepsilon(j)} = -(1-\psi)\Gamma_i \tag{A.49}$$

and

$$\frac{\partial \log p(j)y(j)}{\partial \log X} = (1-\psi)(\chi-\psi)LS(j)A_j + \chi - \psi. \tag{A.50}$$

Adding up the direct ($\varepsilon(j)$) and indirect ($\Delta \log X$) effects gives us

$$\begin{aligned}\Delta \log p(j)y(j) &\approx \frac{\partial \log p(j)y(j)}{\partial\varepsilon(j)}\varepsilon(j) + \frac{\partial \log p(j)y(j)}{\partial \log X}\Delta \log X \\ &= (1-\psi)\Bigg[-\Gamma_j\varepsilon(j) + (\chi-\psi)A_jLS(j)\Delta\log X\Bigg] + (\chi-\psi)\Delta\log X.\end{aligned} \tag{A.51}$$

where the first term with the brackets depends on elasticity of substitution between tasks. Notice that terms outside the bracket will not contribute to the covariance term since this term is uniform across tasks.

Combining (A.43), (A.47) and (A.51), we can rewrite the change in the total labor share as

$$
\begin{aligned}
\Delta \log LS \approx \quad & \Gamma_R\Bigg[(1-\nu_R)\frac{\psi+\zeta_R}{\nu_R+\zeta_R}+(\psi-1)\left(1-\frac{LS}{LS_R}\right)+\frac{LS}{LS_R}\frac{\vartheta}{1+\chi\epsilon_c}\Bigg]\bar{\xi}^R \\
& +\Gamma_N\Bigg[\underbrace{(1-\nu_N)\frac{\psi+\zeta_N}{\nu_N+\zeta_N}}_{\substack{\text{within-task substitution}\\ \text{(aggregated direct effect)}}}+\underbrace{(\psi-1)\left(1-\frac{LS}{LS_N}\right)}_{\substack{\text{between-task substitution}\\ \text{(aggregated direct effect)}}}+\underbrace{\frac{LS}{LS_N}\frac{\vartheta}{1+\chi\epsilon_c}}_{\substack{\text{aggregated}\\ \text{indirect effect}}}\Bigg]\bar{\xi}^N,
\end{aligned}
\tag{A.52}
$$

where

$$
\begin{aligned}
\vartheta = \frac{\partial \log LS}{\partial \log X} = \quad & (\chi-\psi)A_R\theta(1-\nu_R)(1-LS_R)+(\chi-\psi)A_N(1-\theta)(1-\nu_N)(1-LS_N) \\
& +(1-\psi)(\chi-\psi)\Bigg[A_R\theta(LS_R-LS)+A_N(1-\theta)(LS_N-LS)\Bigg].
\end{aligned}
\tag{A.53}
$$

Examining (A.52), we see that the response of the industry labor share and overall labor-saving $\bar{\xi}_R$ and labor-augmenting $\bar{\xi}_N$ innovation depends on three separate terms. The first two terms inside each one of the brackets multiplying our exposure measures aggregate up the direct effects of productivity improvements on the labor share. The first captures the average change in the labor share (holding X constant) induced, holding the original share of output produced by each task constant. The sign of this first term depends on whether capital and labor are complements or substitutes in performing that task. Given our assumption that $\nu_R > 1$, the within-task reallocation force lowers the labor share for routine tasks, whereas the relationship between ν_N and unity is less obvious.

The second terms within each set of brackets illustrate how technological innovation alters the allocation of expenditures across tasks, holding productivity fixed. The signs of these terms partly depend on the interaction between elasticity of substitution across tasks ψ and the relative labor intensity of routine ($\frac{LS_R}{LS}$) and non-routine ($\frac{LS_N}{LS}$) tasks. The interaction gives that, if tasks are complements ($\psi < 1$), then labor-saving technological improvements—increases in $\bar{\xi}_R$—would shift the share of expenditures towards nonroutine tasks. If it is further the case that routine tasks are more capital-intensive than non-routine tasks, this force would tend to increase the labor share since now high-labor share tasks constitute a larger component of the industry's total expenditures. By contrast, if tasks are sufficient substitutes ($\psi > 1$), then increases in $\bar{\xi}_R$ would induce firms to increase their expenditure towards routine tasks. The labor intensity term also captures how much scope for reallocation there is between two type of tasks.

The last term inside each set of brackets captures indirect effects – the extent to which aggregate productivity changes increase or decrease the labor share. Specifically, ϑ captures the impact of aggregate productivity improvements on the labor share, holding current marginal costs at the task level constant. The sensitivity of the labor share to X ϑ also has within and between effects somewhat analogous to the first two terms in (A.52). The terms outside of the bracket in (A.53) capture within-task substitution between capital and labor and the terms in the bracket capture between-task reallocation associated with X changes, which depends on ψ. If $\psi > 1$, resources get diverted to other tasks and the input share of an unaffected task will shrink. On the contrary, if $\psi < 1$, the input share of an unaffected task will rise, pulling wages up.

A.6 Testing the model's implications for the labor share

Here, we examine our model's prediction about the labor share. The industry labor share LS, defined as the ratio of the industry wage bill to the value of industry output, evolves according to

$$\begin{aligned}\Delta \log LS \approx \quad & \Gamma_R\left[(1-\nu_R)\frac{\psi+\zeta_R}{\nu_R+\zeta_R} + (\psi-1)\left(1-\frac{LS}{LS_R}\right) + \frac{LS}{LS_R}\frac{\vartheta}{1+\chi\epsilon_c}\right]\bar{\xi}^R \\ & + \Gamma_N\left[(1-\nu_N)\frac{\psi+\zeta_N}{\nu_N+\zeta_N} + (\psi-1)\left(1-\frac{LS}{LS_N}\right) + \frac{LS}{LS_N}\frac{\vartheta}{1+\chi\epsilon_c}\right]\bar{\xi}^N,\end{aligned} \tag{A.54}$$

where ϑ is the elasticity of the industry labor share to changes in productivity. In the model, the relation between the industry labor share and the overall level of overall labor-saving $\bar{\xi}_R$ and labor-augmenting $\bar{\xi}_N$ technologies is subtle, as it depends on three separate terms. The first term in each bracket captures a within-task reallocation effect: holding the task share of output constant, labor-saving technologies lower the industry labor share (assuming $\nu_R > 1$), while the effect of labor-augmenting technologies is positive if $\nu_N < 1$. The second term in brackets illustrate how technological innovation alters the allocation of expenditures across tasks, holding productivity fixed; their sign depends on the elasticity of substitution across tasks ψ and the relative labor intensity of routine and non-routine tasks. The last term in each bracket accounts for cross-task spillovers through changes in industry productivity and depends on the sign of ϑ.

We estimate the relation between our technology exposure measures and the industry labor share using equation (28). As we see in column (1) of Table A.8, an increase in exposure to labor-saving technologies leads to a 2.32 percent cumulative decline in the labor share over the next five years; by contrast, an increase in exposure to labor-augmenting technologies is followed by approximately a 0.92 percent increase, though the estimate is not statistically different from zero. IV estimates in column (2) paint a similar picture.

A.7 Model Calibration

Are the different regression estimates from the aggregate and micro data quantitatively consistent with the model? To answer this question, we next use these coefficient estimates to estimate the underlying parameters of the model using GMM. To do so, however, we need to make some additional assumptions on how the empirical measures relate to their model equivalents: we assume that our empirical measure of technology exposure (13) maps directly into its model counterpart, up to the presence of industry-level measurement error—capturing, for example, differences in patenting rates across industries. In brief, we target the coefficient estimates corresponding to our worker level regressions (21), conditioning on income; the response of employment at the occupation level to our exposure measures (19); and the response of productivity, labor share, and individual worker wages to our industry-level exposure measure (equations (28) and (29)). Since the IV estimates are both statistically indistinguishable and of very similar magnitude to the OLS estimates throughout, we favor using the far more precisely-estimated OLS coefficients as targets for our calibration. The model is over-identified: we have 17 moments plus two cross-equation restrictions to estimate 11 parameters. Table A.9 reports the moments we target in our estimation and the resulting parameter estimates. We also report the model expressions for all targeted moments together in Table A.10.

In brief, the top panel of Table A.9 shows that the overall fit of the model is rather good. Examining the parameters in the bottom panel, we see that they are largely consistent with the literature. Our estimate of ψ is close to one, in line with the estimates in Caunedo et al. (2023). The estimate of χ is a little under 2, which is consistent within the range of estimates reported by Broda and Weinstein (2006). The estimates of $\nu_R = 1.70$ and $\nu_N = 0.86$ are consistent with technology substituting routine and complementing non-routine tasks. The estimates of labor supply elasticities ζ_R and ζ_N are in line with existing estimates of the elasticity of labor supply based on micro data (Chetty, Guren, Manoli, and Weber, 2011). Last, our estimates of skill loss are modest: the decline in incumbent worker earnings in response to ν_N identifies the mean skill loss among incumbent workers to equal $\beta = 0.04$, while worker skill is reasonably persistent across vintages ($\omega = 0.011$).

In what follows, we provide a detailed discussion of model expressions for targeted moments and their mappings to empirical counterparts.

A.7.1 Individual responses without industry-level productivity spillover terms

The model-implied worker-level responses without industry-level productivity spillover terms (i.e. when including industry × year fixed effects) come from equation (A.25). Empirically, we have the technology exposures $\xi^R(i)$ and $\xi^N(i)$ on the right-hand side for the earnings' response regression. The standard deviations for the labor-saving and labor-augmenting technology shock are denoted by $\tilde{\sigma}_R$ and $\tilde{\sigma}_N$ and the CDF of $\log \bar{l}(i)$ is $F(i)$.

$$\begin{aligned}\Delta \log W(i) \approx & \underbrace{\left[\frac{\psi-\nu_R}{\nu_R+\zeta_R}\Gamma_R\right]\tilde{\sigma}_R}_{\text{worker earning to } \xi^R\text{, Homogenous}} \frac{\xi^R(i)}{\tilde{\sigma}_R} \\ & + \Bigg[\underbrace{\left(\frac{\psi-\nu_N}{\nu_N+\zeta_N}\Gamma_N-\beta\right)\tilde{\sigma}_N}_{\text{worker earning to } \xi^N\text{, Homogenous}} - \omega \underbrace{\left[\log \bar{l}(i) - \int [\log \bar{l}(i)]\, dF(i)\right]\tilde{\sigma}_N}_{\text{0 at homogenous level}} \Bigg] \frac{\xi^N(i)}{\tilde{\sigma}_N} + \text{Industry-Time Fixed Effe}\end{aligned} \tag{A.55}$$

When skill displacement is not included (in the data, this is for our matched entrant regression estimates) the income response to $\xi^N(i)$ is then $\frac{\psi-\nu_N}{\nu_N+\zeta_N}\Gamma_N\tilde{\sigma}_N$.

Heterogeneous responses of workers' earnings to technology exposure can be derived using the expression above conditional on empirical income sort bins. Since the model implies no heterogeneity in earnings response to labor-saving technologies $\xi^R(i)$ within different income bins, we only examine heterogeneous responses to $\xi^N(i)$, corresponding to the second term in the coefficient on $\xi^N(i)$ in (A.55).

We assume that the CDF of log worker-specific productivity $F(i)$ follows a normal distribution, so that $\log \bar{l}(i) - \int [\log \bar{l}(i)]\, dF(i) \sim N(0, \sigma)$, where we denote the dispersion to be σ; we map the parameter σ to the within-occupation-industry volatility of log earnings in the data. For the heterogeneous coefficients, now consider workers whose earnings are between percentiles a and b in the distribution. The expectation of $\tilde{Z} \sim N(0, \sigma)$ conditional on being in the percentile $[\alpha_a, \alpha_b]$ can be derived as

$$E[\tilde{Z}|\tilde{Z}_a \leq \tilde{Z} \leq \tilde{Z}_b] = \sigma E[Z|Q_a \leq Z \leq Q_b] = \sigma \frac{\int_{Q_a}^{Q_b} z f(z) dz}{b-a} = \sigma \frac{-\int_{Q_a}^{Q_b} f'(z) dz}{b-a} = -\sigma \frac{f(Q_b) - f(Q_a)}{b-a} \tag{A.56}$$

where Z follows a standard normal distribution, $\tilde{Z}_x$ is the level of $\tilde{Z}$ at x percentile, Φ is the standard normal CDF while f is the standard normal PDF. $Q_x = \Phi^{-1}(x)$. Using (A.56), we have that for the $[a, b]$ income bin, the income response to labor-augmenting technology is

$$\tilde{\sigma}_N \left[\left(\frac{\psi - \nu_N}{\nu_N + \zeta_N} \Gamma_N - \beta \right) + \omega\sigma \frac{f(Q_b) - f(Q_a)}{b-a} \right] \tag{A.57}$$

We compare the model-implied coefficient on $\xi^R(i)$ and both homogenous and heterogeneous coefficients in responses to technological exposure to their empirical equivalents. For the homogenous coefficient on $\xi^N(i)$ we set $a = 0$ and $b = 1$ in (A.57), which sends the second term to zero and recovers the unconditional coefficient; for the heterogeneous coefficients we examine each of the 5 income bins found in Table 4. We use the model to target the coefficient estimates in column (4), panel A of Table 3 for the homogeneous coefficients; for the no displacement coefficient we target the earnings response of entrants to $\xi^N(i)$ in Table 3, panel B and column (4). For the heterogeneous coefficients on $\xi^N(i)$ we target the second column of panel B in Table 4. To align with our empirical analysis, we re-scale coefficients to correspond to a shift from the empirical median to 90th percentile of $\xi^R(i)$ or $\xi^N(i)$ by multiplying the standardized coefficients by the difference between the 90th percentiles and medians in the data of the respective standardized empirical measures.

A.7.2 Occupational employment responses

The response of overall occupational employment within the industry to technology exposure can be obtained directly from (A.36). As in the previous section, we also include industry × year fixed effects in order to net out productivity spillover terms:

$$\Delta \log L(o) \approx \underbrace{\zeta_R \left[\frac{\psi - \nu_R}{\nu_R + \zeta_R}\Gamma_R\right]\tilde{\sigma}_{R_o}}_{\text{employment to } \xi^R(o)} \frac{\xi^R(o)}{\tilde{\sigma}_{R_o}} + \underbrace{\zeta_N \left[\frac{\psi - \nu_N}{\nu_N + \zeta_N}\Gamma_N\right]\tilde{\sigma}_{N_o}}_{\text{employment to } \xi^N(o)} \frac{\xi^N(o)}{\tilde{\sigma}_{N_o}} + \text{Industry-Time Fixed Effect} \tag{A.58}$$

Here $\tilde{\sigma}_{R_o}$ and $\tilde{\sigma}_{N_o}$ represent the empirical standard deviations for $\xi^R(o)$ and $\xi^N(o)$ within the regression sample for our estimates of (19). To capture the total employment effect including hours worked, we map (A.58) to the growth in worker total employment hours (sum of person weight times hours worked per week times weeks worked per year within an occupation–industry cell) to technology exposure. This is a slightly different object from the response of the raw employment count (sum of person weights) reported in Table 2. The relevant responses are -4.20% (t-stat of -3.59) for ξ^R and 8.28% (t-stat of 4.67) for ξ^N at the 10-year horizon. Since these employment response coefficients are cumulative over a 10-year horizon, we convert to the 5-year horizon by dividing them by 2, arriving at employment-growth targets of 2.10% and 4.14% for ξ^R and ξ^N, respectively. We also again re-scale closed coefficient expressions to correspond to a shift from the median to 90th percentile of exposure within the regression sample to match our empirical scaling of $\xi^R(o)$ and $\xi^N(o)$.

A.7.3 Individual homogenous responses including industry-level productivity spillover terms

In this section we relax the industry × year fixed effects and examine specifications which instead directly control for industry-level innovative outcomes, as in specification (29). Going forward we assume that specifications involving industry-level aggregates are subject to a random observation error at the industry level–for example, due to potential mismeasurement of overall industry innovation by proxying for it using patent grants. We introduce the following variables: $\bar{\xi}^R_{jt}$ and $\bar{\xi}^N_{jt}$ are the true labor-saving and labor-augmenting technology shocks to industry j at time t, respectively. $\xi^R_{ijt} + \bar{\xi}^R_{jt}$ and $\xi^N_{ijt} + \bar{\xi}^N_{jt}$ are the true labor-saving and labor-augmenting technology shocks to individual j at time t, respectively. Instead of observing $\bar{\xi}^R_{jt}$ and $\bar{\xi}^N_{jt}$, we observe $\bar{\xi}^R_{jt} + e^R_{jt}$ and $\bar{\xi}^N_{jt} + e^N_{jt}$, where e^R_{jt} and e^N_{jt} are random observation errors. The covariance matrix for shock pairs $\begin{pmatrix} \xi^R_{ijt} \\ \xi^N_{ijt} \end{pmatrix}$, $\begin{pmatrix} \bar{\xi}^R_{jt} \\ \bar{\xi}^N_{jt} \end{pmatrix}$ and $\begin{pmatrix} e^R_{jt} \\ e^N_{jt} \end{pmatrix}$ are $\begin{pmatrix} \sigma^2_R & 0 \\ 0 & \sigma^2_N \end{pmatrix}$, $\begin{pmatrix} \sigma^2_{\bar{R}} & \sigma_{\bar{R}\bar{N}} \\ \sigma_{\bar{R}\bar{N}} & \sigma^2_{\bar{N}} \end{pmatrix}$, $\begin{pmatrix} \sigma^2_{e^R} & \sigma_{e^R e^N} \\ \sigma_{e^R e^N} & \sigma^2_{e^N} \end{pmatrix}$, respectively. We denote the correlations between the second and third variable pairs $\rho_{\bar{R}\bar{N}}$ and ρ_e, respectively, such that $\sigma_{\bar{R}\bar{N}} = \rho_{\bar{R}\bar{N}}\sigma_{\bar{R}}\sigma_{\bar{N}}$ and $\sigma_{e^R e^N} = \rho_e \sigma_{e^R}\sigma_{e^N}$. All other interactions of random variables have 0 correlation.

We assume that the noise-to-signal ratio is γ for $\bar{\xi}^R_{jt}$ and $\bar{\xi}^N_{jt}$:

$$\sigma^2_{e^S} = \gamma^2 \sigma^2_{\bar{S}}, \qquad S \in \{R, N\} \tag{A.59}$$

We also impose that the industry-level composite $\bar{\xi}^R_{jt} + \bar{\xi}^N_{jt}$ has the same noise-to-signal ratio, which implies[10]

$$\gamma^2 = \frac{Var(e^R_{jt} + e^N_{jt})}{Var(\bar{\xi}^R_{jt} + \bar{\xi}^N_{jt})} = \gamma^2 \frac{\sigma^2_{\bar{R}} + \sigma^2_{\bar{N}} + 2\rho_e \sigma_{\bar{R}}\sigma_{\bar{N}}}{\sigma^2_{\bar{R}} + \sigma^2_{\bar{N}} + 2\rho_{\bar{R}\bar{N}}\sigma_{\bar{R}}\sigma_{\bar{N}}} \tag{A.60}$$

which requires $\rho_e = \rho_{\bar{R}\bar{N}}$, and thus

$$\sigma_{e^R e^N} = \rho_e \sigma_{e^R}\sigma_{e^N} = \gamma^2 \rho_{\bar{R}\bar{N}}\sigma_{\bar{R}}\sigma_{\bar{N}} = \gamma^2 \sigma_{\bar{R}\bar{N}} \tag{A.61}$$

For GMM estimation, we pin down the true volatilities using empirical volatilities and γ. The mapping between empirical volatilities with measurement error and true volatilities is as follows.

$$\begin{aligned} \tilde{\sigma}^2_{\bar{S}} &\equiv Var(\bar{\xi}^S_{jt} + e^S_j) = (1+\gamma^2)\sigma^2_{\bar{S}}, \qquad S \in \{R, N\} \\ \tilde{\sigma}_{\bar{R}\bar{N}} &\equiv Cov(\bar{\xi}^R_{jt} + e^R_{j,t}, \bar{\xi}^N_{jt} + e^N_{j,t}) = \sigma_{\bar{R}\bar{N}} + \rho_e \sigma_{e^R}\sigma_{e^N} = (1+\gamma^2)\sigma_{\bar{R}\bar{N}} \\ \tilde{\sigma}^2_S &\equiv Var(\xi^S_{ijt} + \bar{\xi}^S_{jt} + e^S_j) = \sigma^2_S + (1+\gamma^2)\sigma^2_{\bar{S}}, \qquad S \in \{R, N\} \end{aligned} \tag{A.62}$$

Using the expressions in (A.62), we can represent the covariance matrices for $\begin{pmatrix} \xi^R_{ijt} \\ \xi^N_{ijt} \end{pmatrix}$, $\begin{pmatrix} \bar{\xi}^R_{jt} \\ \bar{\xi}^N_{jt} \end{pmatrix}$ and $\begin{pmatrix} e^R_{jt} \\ e^N_{jt} \end{pmatrix}$ using empirical values and γ as

$$\begin{pmatrix} \tilde{\sigma}^2_R - \tilde{\sigma}^2_{\bar{R}} & 0 \\ 0 & \tilde{\sigma}^2_N - \tilde{\sigma}^2_{\bar{N}} \end{pmatrix}, \quad \frac{1}{1+\gamma^2}\begin{pmatrix} \tilde{\sigma}^2_{\bar{R}} & \tilde{\sigma}_{\bar{R}\bar{N}} \\ \tilde{\sigma}_{\bar{R}\bar{N}} & \tilde{\sigma}^2_{\bar{N}} \end{pmatrix}, \quad \frac{\gamma^2}{1+\gamma^2}\begin{pmatrix} \tilde{\sigma}^2_{\bar{R}} & \tilde{\sigma}_{\bar{R}\bar{N}} \\ \tilde{\sigma}_{\bar{R}\bar{N}} & \tilde{\sigma}^2_{\bar{N}} \end{pmatrix} \tag{A.63}$$

For the individual regression including productivity spillover terms, the true model implies that

$$y_{ijt} = \beta_1(\xi^R_{ijt} + \bar{\xi}^R_{jt}) + \beta_2(\xi^N_{ijt} + \bar{\xi}^N_{jt}) + \delta(\bar{\xi}^R_{jt} + \bar{\xi}^N_{jt}) + \epsilon_{ijt} \tag{A.64}$$

where

$$\beta_1 = \frac{\psi - \nu_R}{\nu_R + \zeta_R}\Gamma_R, \qquad \beta_2 = \frac{\psi - \nu_N}{\nu_N + \zeta_N}\Gamma_N - \beta, \qquad \delta = \frac{1}{1+\chi\epsilon_c}A_R\Gamma_R \tag{A.65}$$

The last equation imposes $\Gamma_R = \Gamma_N$, $A_R = A_N$, as we do in the model estimation. Responses are normalized by the regressors' empirical volatility.

With the observation error, our empirical estimates yield the potentially biased coefficients $\tilde{\beta}_1$, $\tilde{\beta}_2$ and $\tilde{\delta}$ given by the following regression expression

$$y_{ijt} = \tilde{\beta}_1(\xi^R_{ijt} + \bar{\xi}^R_{jt} + e^R_{jt}) + \tilde{\beta}_2(\xi^N_{ijt} + \bar{\xi}^N_{jt} + e^N_{jt}) + \tilde{\delta}(\bar{\xi}^R_{jt} + \bar{\xi}^N_{jt} + e^R_{jt} + e^N_{jt}) + \tilde{\epsilon}_{ijt} \tag{A.66}$$

As Abel (2019) shows, in the case of multivariate measurement error the coefficient estimates

[10] The model technically implies we should measure the industry-level composite as $\bar{\xi} = \frac{LS}{LS_R}\bar{\xi}^R + \frac{LS}{LS_N}\bar{\xi}^N$; since $\frac{LS}{LS_R}$ and $\frac{LS}{LS_N}$ are both close to one in practice we ignore this nuance and simply take $\bar{\xi} = \bar{\xi}^R + \bar{\xi}^N$.

have the following probability limit:

$$\begin{pmatrix} \frac{1}{\tilde{\sigma}_R}\tilde{\beta}_1 \\ \frac{1}{\tilde{\sigma}_N}\tilde{\beta}_2 \\ \frac{1}{\tilde{\sigma}_{\text{comp}}}\tilde{\delta} \end{pmatrix} = \begin{pmatrix} \beta_1 \\ \beta_2 \\ \delta \end{pmatrix} - (V+\Sigma)^{-1}\Sigma \begin{pmatrix} \beta_1 \\ \beta_2 \\ \delta \end{pmatrix} \tag{A.67}$$

where

$$V = \begin{pmatrix} \sigma_R^2 + \sigma_{\bar{R}}^2 & \sigma_{\bar{R}\bar{N}} & \sigma_{\bar{R}}^2 + \sigma_{\bar{R}\bar{N}} \\ \sigma_{\bar{R}\bar{N}} & \sigma_N^2 + \sigma_{\bar{N}}^2 & \sigma_{\bar{N}}^2 + \sigma_{\bar{R}\bar{N}} \\ \sigma_{\bar{R}}^2 + \sigma_{\bar{R}\bar{N}} & \sigma_{\bar{N}}^2 + \sigma_{\bar{R}\bar{N}} & \sigma_{\bar{R}}^2 + \sigma_{\bar{N}}^2 + 2\sigma_{\bar{R}\bar{N}} \end{pmatrix}, \quad \Sigma = \begin{pmatrix} \sigma_{e^R}^2 & \sigma_{e^R e^N} & \sigma_{e^R}^2 + \sigma_{e^R e^N} \\ \sigma_{e^R e^N} & \sigma_{e^N}^2 & \sigma_{e^N}^2 + \sigma_{e^R e^N} \\ \sigma_{e^R}^2 + \sigma_{e^R e^N} & \sigma_{e^N}^2 + \sigma_{e^R e^N} & \sigma_{e^R}^2 + \sigma_{e^N}^2 + 2\sigma_{e^R e^N} \end{pmatrix} \tag{A.68}$$

and $\tilde{\sigma}_{\text{comp}}$ gives the empirical standard deviation of the industry composite $\bar{\xi}_{jt}^R + \bar{\xi}_{jt}^N$ within the regression sample.

By substituting β_1, β_2 and δ using equation (A.65), and substituting the entries in (A.68) using the empirical mapping shown in (A.63), we can represent the targeted moment $\tilde{\beta}_1$, $\tilde{\beta}_2$ and $\tilde{\delta}$ in terms of the underlying parameters of interest. We target the coefficients in column (6) of Table 6; finally, to match the scaling of our empirical estimates we further re-scale coefficients β_1 and β_2 so that marginal effects correspond to a shift from the median to 90th percentile of the distributions of exposure in the data.

A.7.4 Industry productivity and labor share responses

For the responses of industry productivity to composite exposure measure, we assume the true expression is

$$y_{jt} = \delta_1(\bar{\xi}_{jt}^R + \bar{\xi}_{jt}^N) + \epsilon_{jt} \tag{A.69}$$

where the model implies

$$\delta_1 = \frac{1}{1+\chi\epsilon_c}\Gamma_R\tilde{\sigma}_{\text{comp}} \tag{A.70}$$

according to (A.42) after imposing $\Gamma_R = \Gamma_N, A_R = A_N$. The scale factor $\tilde{\sigma}_{\text{comp}}$ now gives the standard deviation of the industry-level composite proxy $\bar{\xi}_{jt}^R + \bar{\xi}_{jt}^N + e_{jt}^R + e_{jt}^N$ within this regression sample. As in the previous section, this empirical industry-level aggregates are subject to measurement error, so the empirical specification is instead

$$y_{jt} = \tilde{\delta}_1(\bar{\xi}_{jt}^R + \bar{\xi}_{jt}^N + e_{jt}^R + e_{jt}^N) + \tilde{\epsilon}_{jt} \tag{A.71}$$

With the noise-to-signal ratio γ, the standard expressions for the OLS coefficient under measurement error gives

$$\tilde{\delta}_1 = \frac{Cov(\delta_1(\bar{\xi}_{jt}^R + \bar{\xi}_{jt}^N) + \epsilon_{ij}, \bar{\xi}_{jt}^R + \bar{\xi}_{jt}^N + e_{j,t}^R + e_{j,t}^N)}{Var(\bar{\xi}_{jt}^R + \bar{\xi}_{jt}^N + e_{j,t}^R + e_{j,t}^N)} = \frac{\delta_1}{1+\gamma^2} \tag{A.72}$$

For the industry labor share response to the industry technology exposure, the true expression is

$$y_{jt} = \eta_1 \bar{\xi}^R_{jt} + \eta_2 \bar{\xi}^N_{jt} + \epsilon_{jt} \tag{A.73}$$

According to (A.52), we have that

$$\begin{aligned}
\eta_1 &= \Gamma_R \left[(1-\nu_R)\frac{\psi+\zeta_R}{\nu_R+\zeta_R} + (\psi-1)\left(1-\frac{LS}{LS_R}\right) + \frac{LS}{LS_R}\frac{\vartheta}{1+\chi\epsilon_c}\right]\tilde{\sigma}_{\bar{R}} \\
\eta_2 &= \Gamma_N \left[(1-\nu_N)\frac{\psi+\zeta_N}{\nu_N+\zeta_N} + (\psi-1)\left(1-\frac{LS}{LS_N}\right) + \frac{LS}{LS_N}\frac{\vartheta}{1+\chi\epsilon_c}\right]\tilde{\sigma}_{\bar{N}}
\end{aligned} \tag{A.74}$$

Under observation error, we measure $\tilde{\eta}_1$ and $\tilde{\eta}_2$:

$$y_{jt} = \tilde{\eta}_1(\bar{\xi}^R_{jt} + e^R_{j,t}) + \tilde{\eta}_2(\bar{\xi}^N_{jt} + e^N_{j,t}) + \tilde{\epsilon}_{jt} \tag{A.75}$$

Again using the expression from Abel (2019) for coefficient estimates under multivariate measurement error, we have that

$$\begin{pmatrix}\tilde{\eta}_1 \\ \tilde{\eta}_2\end{pmatrix} = \begin{pmatrix}\eta_1 \\ \eta_2\end{pmatrix} - \begin{pmatrix}\sigma^2_{\bar{R}} + \sigma^2_{e^R} & \sigma_{\bar{R}\bar{N}} + \sigma_{e^R e^N} \\ \sigma_{\bar{R}\bar{N}} + \sigma_{e^R e^N} & \sigma^2_{\bar{N}} + \sigma^2_{e^N}\end{pmatrix}^{-1} \begin{pmatrix}\sigma^2_{e^R} & \sigma_{e^R e^N} \\ \sigma_{e^R e^N} & \sigma^2_{e^N}\end{pmatrix} \begin{pmatrix}\eta_1 \\ \eta_2\end{pmatrix} = \frac{1}{1+\gamma^2}\begin{pmatrix}\eta_1 \\ \eta_2\end{pmatrix} \tag{A.76}$$

which follows immediately from applying the fact that $\rho_e = \rho_{\bar{R}\bar{N}}$ (as shown in equation (A.60)), to get that $\sigma_{\bar{R}\bar{N}} = \gamma^2 \sigma_{e^R e^N}$. Thus the assumption that the noise-to-signal ratios for $\bar{\xi}^R_{jt} + e^R_{j,t}$, $\bar{\xi}^N_{jt} + e^R_{j,t}$, and $\bar{\xi}^R_{jt} + e^R_{j,t} + \bar{\xi}^N_{jt} + e^N_{j,t}$ are all equal to γ implies that the coefficients collapse to the classical univariate measurement error expression here.

We map the model-implied productivity response to the point estimate on composite exposure in column (4) of Table 6; for the labor share responses we compare the model with the two point estimates in column (1) of Table A.8, where we replace growth in the industry-level labor productivity index with industry-level labor share growth as the dependent variable.

The last moment we target is the average labor share LS. We have that

$$LS = \frac{\sum_{k\in J} w(k)l(k)}{(1+\kappa_R)\sum_{k\in J_R} w(k)l(k) + (1+\kappa_N)\sum_{k\in J_N} w(k)l(k)} = [(1+\kappa_R)\theta + (1+\kappa_N)(1-\theta)]^{-1} \tag{A.77}$$

To obtain our labor share target, we first take the labor share estimate of 70.56% from Koh et al. (2020). Since we do not have a notion of capital payments to buildings/structures in our model, we further adjust for the share of income accruing to structures based on the estimate of 11.7% from Ohanian et al. (2023). This gives us a labor share target of $0.7056/(1-0.117) \approx 80\%$.

A.7.5 Estimation Procedure

We form a parameter vector Θ using the 11 parameters of interest: $\Theta = [\psi, \chi, \nu_R, \nu_N, \zeta_R, \zeta_N, \beta, \kappa_R, \kappa_N, \gamma, \omega]'$. A few other model parameters needed for estimating model moments are pre-calibrated directly from data. The parameter σ controls the standard deviation of earnings within occupation-industry; this value is 0.531 in the data. We calibrate a routine share of total wage bill (model parameter θ) of 0.340 by taking the total routine wagebill share in the data $\frac{\sum_o \text{Routine Share}_o \times \text{Total Wagebill}_o}{\sum_o \text{Total Wagebill}_o}$ for each of the 1980, 1990, and 2000 IPUMS Samples and 2010 ACS samples, and then averaging across years. The worker-level average routine share is $\theta(i)$ is 0.373, but since we impose $A_R = A_N$, this parameter does not affect our calibration. Finally, the relevant empirical volatilities/covariances are: $\tilde{\sigma}_R = 0.537$, $\tilde{\sigma}_N = 0.904$, $\tilde{\sigma}_{R_o} = 0.585$, $\tilde{\sigma}_{N_o} = 0.871$, $\tilde{\sigma}_{\bar{R}} = 0.436$, $\tilde{\sigma}_{\bar{N}} = 0.749$, $\tilde{\sigma}_{\bar{R}\bar{N}} = 0.279$, $\tilde{\sigma}_{\text{comp}} = 1.107$.

Let $\hat{F}(\Theta)$ be a function mapping model-implied parameters to the 17 target moments. We choose the parameter vector $\hat{\Theta}$ by minimizing the distance between the model implied $\hat{F}(\Theta)$ and the actual empirical estiamtes F, with the 2 constraints,

$$
\begin{aligned}
\hat{\Theta} &= \underset{\Theta}{\operatorname{argmin}} \left(F - \hat{F}(\Theta)\right)' W \left(F - \hat{F}(\Theta)\right) \\
&s.t. \quad A_R = A_N \quad \text{and} \quad \Gamma_R = \Gamma_N
\end{aligned}
\tag{A.78}
$$

where the weight matrix W is a diagonal matrix with W_{ii} equal to the absolute value of the empirical moment.[11]

[11] We make a couple slight modifications to this weighting. In order to de-emphasize the bottom income bin moment for ξ_N and highlight the top income bin moment, we further divide the weight on the bottom moment by 3 and multiply weight the top bin moment by 3. We also multiply the weight on the aggregate average labor share moment by 2.

Table A.10: Summary of Targeted Moments

#	Parameter Combination	Description
1	$\frac{\psi-\nu_R}{\nu_R+\zeta_R}\,\frac{(\nu_R+\zeta_R)\,\kappa_R}{\psi+\zeta_R+(\nu_R+\zeta_R)\,\kappa_R}\tilde{\sigma}_R$	worker earnings to ξ^R, homogeneous
2	$\left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\,\frac{(\nu_N+\zeta_N)\,\kappa_N}{\psi+\zeta_N+(\nu_N+\zeta_N)\,\kappa_N}-\beta\right]\tilde{\sigma}_N$	worker earnings to ξ^N, homogeneous
3	$\left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\,\frac{(\nu_N+\zeta_N)\,\kappa_N}{\psi+\zeta_N+(\nu_N+\zeta_N)\,\kappa_N}\right]\tilde{\sigma}_N$	worker earnings to ξ^N, homogeneous, new entrants
4	$\left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\,\frac{(\nu_N+\zeta_N)\,\kappa_N}{\psi+\zeta_N+(\nu_N+\zeta_N)\,\kappa_N}-\beta+\sigma\omega\frac{f(Q_{0.25})-f(Q_0)}{\Phi(Q_{0.25})-\Phi(Q_0)}\right]\tilde{\sigma}_N$	worker earnings to ξ^N, 0%-25%
5	$\left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\,\frac{(\nu_N+\zeta_N)\,\kappa_N}{\psi+\zeta_N+(\nu_N+\zeta_N)\,\kappa_N}-\beta+\sigma\omega\frac{f(Q_{0.5})-f(Q_{0.25})}{\Phi(Q_{0.5})-\Phi(Q_{0.25})}\right]\tilde{\sigma}_N$	worker earnings to ξ^N, 25%-50%
6	$\left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\,\frac{(\nu_N+\zeta_N)\,\kappa_N}{\psi+\zeta_N+(\nu_N+\zeta_N)\,\kappa_N}-\beta+\sigma\omega\frac{f(Q_{0.75})-f(Q_{0.5})}{\Phi(Q_{0.75})-\Phi(Q_{0.5})}\right]\tilde{\sigma}_N$	worker earnings to ξ^N, 50%-75%
7	$\left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\,\frac{(\nu_N+\zeta_N)\,\kappa_N}{\psi+\zeta_N+(\nu_N+\zeta_N)\,\kappa_N}-\beta+\sigma\omega\frac{f(Q_{0.95})-f(Q_{0.75})}{\Phi(Q_{0.95})-\Phi(Q_{0.75})}\right]\tilde{\sigma}_N$	worker earnings to ξ^N, 75%-95%
8	$\left[\frac{\psi-\nu_N}{\nu_N+\zeta_N}\,\frac{(\nu_N+\zeta_N)\,\kappa_N}{\psi+\zeta_N+(\nu_N+\zeta_N)\,\kappa_N}-\beta+\sigma\omega\frac{f(Q_1)-f(Q_{0.95})}{\Phi(Q_1)-\Phi(Q_{0.95})}\right]\tilde{\sigma}_N$	worker earnings to ξ^N, 95%-100%
9	$\tilde{\beta}_1$ Given by (A.67)	worker earnings to ξ^R, regression with $\bar{\xi}$, homogeneous
10	$\tilde{\beta}_2$ Given by (A.67)	worker earnings to ξ^N, regression with $\bar{\xi}$, homogeneous
11	$\tilde{\delta}$ Given by (A.67)	worker earnings to $\bar{\xi}$, homogeneous
12	$\frac{1}{1+\gamma^2}\,\frac{1}{1+\chi\,\epsilon_c}\,\frac{(\nu_R+\zeta_R)\,\kappa_R}{\psi+\zeta_R+(\nu_R+\zeta_R)\,\kappa_R}\tilde{\sigma}_{\text{comp}}$ (imposing $\Gamma_R=\Gamma_N$ here)	industry prod to $\bar{\xi}$
13	$\frac{1}{1+\gamma^2}\,\frac{(\nu_R+\zeta_R)\,\kappa_R}{\psi+\zeta_R+(\nu_R+\zeta_R)\,\kappa_R}\left[(1-\nu_R)\frac{\psi+\zeta_R}{\nu_R+\zeta_R}+(\psi-1)(1-\frac{LS}{LS_R})+\frac{LS}{LS_R}\,\frac{\vartheta}{1+\chi\epsilon_c}\right]\tilde{\sigma}_{\bar{R}}$	industry labor share to $\bar{\xi}^R$
14	$\frac{1}{1+\gamma^2}\,\frac{(\nu_N+\zeta_N)\,\kappa_N}{\psi+\zeta_N+(\nu_N+\zeta_N)\,\kappa_N}\left[(1-\nu_N)\frac{\psi+\zeta_N}{\nu_N+\zeta_N}+(\psi-1)(1-\frac{LS}{LS_N})+\frac{LS}{LS_N}\,\frac{\vartheta}{1+\chi\epsilon_c}\right]\tilde{\sigma}_{\bar{N}}$	industry labor share to $\bar{\xi}^N$
15	$(\psi-\nu_R)\frac{\zeta_R}{\nu_R+\zeta_R}\,\frac{(\nu_R+\zeta_R)\,\kappa_R}{\psi+\zeta_R+(\nu_R+\zeta_R)\,\kappa_R}\tilde{\sigma}_{R_o}$	occupation employment to ξ^R
16	$(\psi-\nu_N)\frac{\zeta_N}{\nu_N+\zeta_N}\,\frac{(\nu_N+\zeta_N)\,\kappa_N}{\psi+\zeta_N+(\nu_N+\zeta_N)\,\kappa_N}\tilde{\sigma}_{N_o}$	occupation employment to ξ^N
17	$[\theta(\kappa_R+1)+(1-\theta)(\kappa_N+1)]^{-1}$	labor share

Note: This table displays close-form model-implied expressions for all moments we target in our GMM estimation. See appendix section A for model details, including subsection A.7 for an in-depth discussion of model estimation.

B Data Appendix

B.1 Converting Patent and Occupation Task Texts for Numerical Analysis

Here, we briefly overview our conversion of unstructured patent text data into a numerical format suitable for statistical analysis. We obtain text data for measuring patent/job task similarity from two sources. Job task descriptions come from the revised 4th edition of the Dictionary of Occupation Titles (DOT) database. We use the patent text data parsed from the USPTO patent search website in Kelly et al. (2021), which includes all US patents beginning in 1976, comprising patent numbers 3,930,271 through 9,113,586. Our analysis of the patent text combines the claims, abstract, and description section into one patent-level corpus for each patent. Since the DOT has a very wide range of occupations (with over 13,000 specific occupation descriptions) we first crosswalk the DOT occupations to the considerably coarser and yet still detailed set of 6-digit occupations in the 2010 edition of O*NET.[12] We then combine all tasks for a given occupation at the 2010 O*NET 6-digit level into one occupation-level corpus. The process for cleaning and preparing the text files for numerical representation follows the steps outlined below.

We first clean out all non-alphabetic characters from the patent and task text, including removing all punctuation and numerical characters. We then convert all text to lowercase. At this stage each patent and occupation-level task text are represented by a single string of words separated by spaces. To convert each patent/occupation into a list of associated words we apply a word tokenizer that separates the text into lists of word tokens which are identified by whitespace in between alphabetic characters. Since many commonly used words carry little semantic information, we filter the set of tokens by first removing all "stop words"– which include prepositions, pronouns, and other common words carrying little content–from the union of several frequently used stop words lists.

Stop words come from the following sources:

- https://pypi.python.org/pypi/stop-words
- https://dev.mysql.com/doc/refman/5.1/en/fulltext-stopwords.html
- http://www.lextek.com/manuals/onix/stopwords1.html
- http://www.lextek.com/manuals/onix/stopwords2.html
- https://msdn.microsoft.com/zh-cn/library/bb164590
- http://www.ranks.nl/stopwords
- http://www.text-analytics101.com/2014/10/all-about-stop-words-for-text-mining.html

[12] The DOT to SOC crosswalk is available at https://www.onetcenter.org/crosswalks.html. Since the time we originally obtained the crosswalk, O*NET has subsequently replaced the 2010 SOC code–based version we use in this paper with a crosswalk derived from the 2019 SOC code scheme.

- http://www.webconfs.com/stop-words.php
- http://www.nltk.org/book/ch02.html (NLTK stop words list)

We also add to the list of stop words the following terms that are ubiquitous in the patent text but don't provide information regarding the content and purpose of the patent: abstract, claim, claims, claimed, claiming, present, invention, united, states, patent, description, and background. The final stop word list contains 1337 unique terms that are filtered out.

Even after removing stop words, we expect much of the remaining text to offer little information regarding the purpose and use of a given patent or the core job functions expected to be performed by workers in a given occupation. In order to focus on the parts of the document most likely to contain relevant information, we retain descriptive and action words–i.e. nouns and verbs–and remove all other tokens. We do this using the part-of-speech tagger from the NLTK Python library. Finally, we lemmatize all remaining nouns and verbs, which is to convert them to a common root form. This converts all nouns to their singular form and verbs to their present tense. We use the NLTK WordNet Lemmatizer to accomplish this task. After these steps are completed, we have a set of cleaned lists of tokens for each patent and each occupation's tasks that we can then use to compute pairwise similarity scores.

B.2 Classifying tasks based on their capital/labor complementarity

Our model emphasizes the distinction between exposure to labor-saving technologies (technologies that substitute for labor in performing certain tasks) and labor-augmenting technologies (technologies that augment labor in performing certain tasks). A key issue then is to how empirically identify these two distinct sets of tasks. Conceptually, tasks in which capital can substitute for labor are tasks that can be performed by capital (automation). To that end, we rely on two complementary ideas in the literature. The first is the notion of routine tasks (Acemoglu and Autor, 2011); the key idea here is that routine tasks are the tasks that can *potentially* be performed by machines (or software or some other form of capital). By contrast, non-routine tasks cannot be performed by machines; instead workers use these machines to increase their productivity when performing these tasks. Second, we rely on the notion of skill-biased technical change, in which there are a certain type of tasks that require low levels of related skill that can be performed by machines, while capital is a complement to labor when performing high-skill tasks (Krusell et al., 2000; Goldin and Katz, 2008).

B.2.1 Classification into routine and non-routine tasks

To implement this distinction in the data we need to first partition the set of tasks for each occupation into routine and non-routine tasks. Our text data on job task descriptions come from the revised 4th edition of the Dictionary of Occupation Titles (DOT) database. Given that the main empirical analysis in this paper focuses on worker-level outcomes in the post-1980 sample, we use the task descriptions from the 1991 DOT—which is mostly identical to the 1977 DOT version beyond the addition of some IT-related occupations—rather than more recent versions from O*NET. Since the DOT has a very wide range of occupations (with over 13,000 specific occupation descriptions) we use a crosswalk from DOT occupations to the considerably coarser and yet still detailed set of 6-digit Standard Occupation Classification (SOC) codes from O*NET. We then combine all tasks for a given occupation at the 2010 SOC 6-digit level into one occupation-level corpus.

We then use GPT4 to classify individual tasks (defined by distinct sentences in DOT task descriptions) into routine or non-routine:

> A routine task can be defined as follows: A routine task involves carrying out a limited and well-defined set of work activities, those that can be accomplished by following explicit rules. These tasks require methodical repetition of an unwavering procedure, and they can be exhaustively specified with programmed instructions and performed by machines. Tell me whether the following task is primarily routine, primarily non-routine, or involves a mix of both routine and non-routine tasks; and, explain your reasoning in one sentence. [insert sentence describing task].

We apply this query for each one of the sentences of each occupation's task description in the DOT. Out of a total of approximately 90,000 sentences, GPT4 characterizes 62% of these as referring to routine tasks, 15% as non-routine and 22% as a mixture of routine and non-routine. We group the latter two categories into a non-routine category.

We validate the output produced by GPT4 with a routine task intensity (RTI) index, which we construct from the six standardized occupational task types for the 2010 ONET using code provided by Acemoglu and Autor (2011). Specifically, we take the sum of the standardized routine manual and routine cognitive scores and subtract off the standardized non-routine manual (interpersonal), non-routine manual (physical), non-routine cogntive (analytical), and non-routine cognitive (interpersonal) scores generated by the code from Acemoglu and Autor (2011).[13] To compare the two, we calculate the average share of tasks at the occupation level that are classified as routine by GPT4. Panel A of Figure 1 plots the RTI index versus the average share of routine tasks at the occupation level. We see that the two are highly correlated (81%) and the relation is approximately linear.

At this point, it is worth emphasizing that we view the 'routineness' of a task as an inherent property of a task that is invariant to the current state of technology. In this regard, we should think of routine tasks as those that can be potentially automated. The rate of arrival of breakthrough innovations determines which tasks are actually automated. Thus, for example, booking airline tickets is a routine task, but it required advances in communication before it could be (mostly) automated. That said, it is possible that the answers provided by GPT4 are influenced by the current state of technology.

[13]This is similar to what Autor et al. (2003) do to create a routine-task intensity index using the prior Dictionary of Occupational Titles occupation task scores. They take the log of routine score and subtract off the log of abstract and log of manual scores; we do not take logs since there are zeroes in the raw ONET scores, and we also follow Acemoglu and Autor (2011) in standardizing and centering the measures around zero.

B.2.2 Classification into high- and low-skill tasks

As a robustness check, we also explore an alternative classification of tasks to build our technology exposure measures that relies on the idea that low-skill tasks are more easily automated than high-skill tasks (Krusell et al., 2000). Given that the definition of skill can be somewhat vague, we use a definition that relies on related experience provided by ONET, termed Specific Vocational Preparation (SVP),

> Specific Vocational Preparation is the amount of lapsed time required by a typical worker to learn the techniques, acquire the information, and develop the facility needed for average performance in a specific job-worker situation. Tell me whether attaining proficiency in the below occupation task requires A) an extensive amount (more than 5 years); B) a fair amount (1 to 5 years); C) a moderate amount (3 months to 1 year); or D) very little (less than 3 months) of specific vocational preparation; and, explain your reasoning in one sentence. [insert sentence describing task].

As before, we apply this query for each one of the sentences in an occupation's description of tasks. ChatGPT classifies 61% of these task sentences in the D category, with the remainder classified as C(25%), B (13%), and A (0.01%). We label tasks in category D as low experience while tasks in categories A/B/C as high experience.

We validate our alternative measure using ONET's job zone classification, which assigns occupations a score between 1 (lowest) and 5 (highest) depending on the extent of their required preparation, which includes education requirements, related experience, and on-the-job training. As we see in Panel B of Figure 1, there is a strong correlation (85%) between the share of tasks classified as high-experience by GPT4 and ONET's job classification.

We conclude that our classification of occupation tasks into routine/non-routine or low/high skill captures economically meaningful variation that can be used to construct measures of technology exposure in line with our model. The correlation between these two sets of classification is relatively high, but not perfect: out of the tasks that are classified as routine, approximately 80 percent are classified as low experience. By contrast, out of the tasks that are classified as non-routine, approximately 70 percent are classified as high required experience. We primarily rely on the routine/non-routine classification as our baseline measure, and relegate the results using the low/high skill classification to the appendix.

B.3 Measuring the distance between patents and worker tasks using word embedding vectors

We identify technologies that are relevant to specific worker groups as those that are similar to the descriptions of the tasks performed by a given occupation. We do so by analyzing the textual similarity between the description of the innovation in the patent document and the worker's job description.

To identify the similarity between a breakthrough innovation and an occupation, we need to identify meaningful connections between two sets of documents that account for differences in the language used. The most common approach for computing document similarity is to create a matrix representation of each document, with columns representing document counts for each term (or some weighting of term counts) in the dictionary of all terms contained in the set of documents, and with rows representing each document. Similarity scores could then be computed simply as the cosine similarity between each vector of weighted or unweighted term counts:

$$Sim_{i,j} = \frac{V_i}{||V_i||} \cdot \frac{V_j}{||V_j||} \tag{B.1}$$

Here V_i and V_j denote the vector of potentially weighted terms counts for documents i and j.

This approach is often referred to as the 'the bag-of-words' approach, and has been used successfully in many settings. For example, Kelly et al. (2021) use a variant of this approach to construct measures of patent novelty and impact based on pairwise distance measures between patent documents. Since patent documents have a structure and a legalistic vocabulary that is reasonably uniform, this approach works quite well for patent-by-patent comparisons. However, this approach is less suited for comparing patent documents to occupation task descriptions. These two sets of documents come from different sources and often use different vocabulary. If we were to use the bag-of-words approach, the resulting vectors V_i and V_j would be highly sparse with most elements equal to zero, which would bias the distance measure (B.1) to zero.

The root cause of the problem is that the distance measure in (B.1) has no way of accounting for words with similar meanings. For example, consider a set of two documents, with the first document containing the words 'dog' and 'cat' and the other containing the words 'puppy' and 'kitten'. Even though the two documents carry essentially the same meaning, the bag of words approach will conclude that they are distinct: the representation of the two documents is $V_1 = [1, 1, 0, 0]$ and $V_2 = [0, 0, 1, 1]$, which implies that the two documents are orthogonal, $\rho_{1,2} = 0$.

To overcome this challenge, we leverage recent advances in natural language processing that allow for synonyms. The main idea behind this approach is to represent each word as a dense vector. The distance between two word vectors is then related to the likelihood these words capture a similar meaning. In our approach, we use the word vectors provided by Pennington et al. (2014), which contains a vocabulary of 1.9 million word meanings (embeddings) represented as (300-dimensional) vectors. The two most popular approaches are the "word2vec" method of Mikolov, Sutskever, Chen, Corrado, and Dean (2013) and the global vectors for word representation introduced by Pennington et al. (2014). These papers construct mappings from extremely sparse and high-dimensional word co-occurence counts to dense and comparatively low-dimensional vector representations of word meanings called word embeddings. Their word vectors are highly successful at capturing synonyms and word analogies (vec(king) − vec(queen) ≈ vec(man) − vec(woman) or vec(Lisbon) − vec(Portugal) ≈ vec(Madrid) − vec(Spain), for example). Thus they are well-suited for numerical representations of the "distance" between words. The word vectors provided by Pennington et al. (2014) are trained on 42 billion word tokens of web data from Common Crawl and are available at `https://nlp.stanford.edu/projects/glove/`.

To appreciate how our metric differs from the standard bag-of words approach it is useful to briefly examine how word embeddings are computed in Pennington et al. (2014). Denote the matrix X as a $V \times V$ matrix of word co-occurence counts obtained over a set of training documents, where V is the number of words in the vocabulary. Then $X_{i,j}$ tabulates the number of times word j appears in the context of the word i.[14] Denote $X_i = \sum_k X_{i,k}$ as the number of times any word appears in the context of word i, and the probability of word j occuring in the context of word i is $P_{i,j} \equiv X_{i,j}/X_i$. The goal of the word embedding approach is to construct a mapping $F(\cdot)$ from some d-dimensional vectors x_i, x_j, and $\tilde{x}_k$ such that

$$F(x_i, x_j, \tilde{x}_k) = \frac{P_{i,k}}{P_{j,k}} \tag{B.2}$$

Imposing some conditions on the mapping $F(\cdot)$, they show that a natural choice for modeling $P_{i,k}$ in (B.2) is

$$x_i^T \tilde{x}_k = \log(X_{i,k}) - \log(X_i) \tag{B.3}$$

Since the mapping should be symmetric for i and k they add "bias terms" (essentially i and k fixed effects) which gives

$$x_i^T \tilde{x}_k + b_i + b_k = \log(X_{i,k}) \tag{B.4}$$

Summing over squared errors for all pairwise combinations of terms yields the weighted least squares objective

$$\text{Min}_{x_i, \tilde{x}_k, b_i, b_k} \sum_{i=1}^{V} \sum_{j=1}^{V} f(X_{i,j}) \left(x_i^T \tilde{x}_k + b_i + b_k - \log(X_{i,j}) \right)^2 \tag{B.5}$$

Here the observation-specific weighting function $f(X_{i,j})$ equals zero for $X_{i,j} = 0$ so that the log is well defined, and is constructed to avoid overweighing rare occurrences or extremely frequent occurrences. The objective (B.5) is a highly-overidentified least squares minimization problem. Since the solution is not unique, the model is trained by randomly instantiating x_i and $\tilde{x}_k$ and performing gradient descent for a pre-specified number of iterations, yielding d-dimensional vector representations of a given word. Here d is a hyper-parameter; Pennington et al. (2014) find that $d = 300$ works well on word analogy tasks.

Since (B.5) is symmetric it yields two vectors for word i, x_i and $\tilde{x}_i$, so the final word vector is taken as the average of the two. The ultimate output is a dense 300-dimensional vector for each word i that has been estimated from co-occurence probabilities and occupies a position in a word vector space such that the pairwise distances between words (i.e. using a metric like the cosine similarity) are related to the probability that the words occur within the context of one another and within the context of other similar words. Note that the basis for this word vector space is arbitrary and has no meaning; distances between word embeddings are only well-defined in relation to one another and a different training instance of the same data would yield different word vectors

[14] Pennington et al. (2014) use a symmetric 10 word window to determine "context" and weight down occurences that occur further away from the word (one word away receives weight 1, two words away receives weight 1/2, etc.).

but very similar pairwise distances between word vectors.

The next step consists of using these word vectors to construct measures of document similarity. To begin, we first construct a weighted average of the word embeddings with a document (a patent or occupation description). Specifically, we represent each document as a (dense) vector X_i, constructed as a weighted average of the set of word vectors $x_k \in A_i$ contained in the document,

$$X_i = \sum_{x_k \in A_i} w_{i,k} x_k. \tag{B.6}$$

A key part of the procedure consists of choosing appropriate weights $w_{i,k}$ in order to emphasize important words in the document. We also note that there is not a Pennington et al. (2014) word embedding estimate for every word in every patent. For example, occasionally there are OCR text recognition errors in patent documents, which can lead to misspelled words that do not have an embedding estimate as a result; there may also be patents that use extremely specific and technical terms (i.e. chemical patents introducing an new and very specific molecule) which do not have a word embedding available. In such cases we simply calculate (B.6) for the set of words which do have a Pennington et al. (2014) embedding estimate—which is the vast majority of terms in practice.

In natural language processing, a common approach to emphasize terms that are most diagnostic of a document's topical content is the 'term-frequency-inverse-document-frequency' (TF-IDF). In brief, $TFIDF_{i,k}$ overweighs word vectors for terms that occur relatively frequently within a given document and underweighs terms that occur commonly across all documents. We follow the same approach: in constructing (B.6), we weigh each word vector by

$$w_{i,k} \equiv TF_{i,k} \times IDF_k. \tag{B.7}$$

The first component of the weight, term frequency (TF), is defined as

$$TF_{i,k} = \frac{c_{i,k}}{\sum_j c_{i,j}}, \tag{B.8}$$

where $c_{i,k}$ denotes the count of the k-th word in document i—a measure of its relative importance within the document. The inverse-document frequency is

$$IDF_k = \log\left(\frac{\text{\# of documents in sample}}{\text{\# of documents that include term } k}\right). \tag{B.9}$$

Thus, IDF_k measures the informativeness of term k by under-weighting common words that appear in many documents, as these are less diagnostic of the content of any individual document. We compute the inverse-document-frequency for the set of patents and occupation tasks separately, so that patent document vectors underweight word embeddings for terms appearing in many patents and occupation vectors underweight word embeddings for job task terms that appear in the task descriptions of many other occupations.

Armed with a vector representation of the document that accounts for synonyms, we next use the cosine similarity to measure the similarity between patent i and occupation j:

$$\mathrm{Sim}_{i,j} = \frac{X_i}{||X_i||} \cdot \frac{X_j}{||X_j||} \tag{B.10}$$

This is the same distance metric as the bag of words approach, except now X_i and X_j are dense vectors carrying a geometric interpretation akin to a weighted average of the semantic meaning of all nouns and verbs in the respective documents.

To illustrate the difference between our approach and the standard bag of words, consider the following example of two documents, with the first document containing the words 'dog' and 'cat' and the other containing the words 'puppy' and 'kitten'. Even though the two documents carry essentially the same meaning, the bag of words approach will conclude that they are distinct: the representation of the two documents is

$$V_1 = [1,1,0,0], \quad \text{and} \quad V_2 = [0,0,1,1] \tag{B.11}$$

which implies that the two documents are orthogonal, $\rho_{1,2} = 0$. Here, the TF-IDF weights in our simple example satisfy $TF_{1,dog} = 1/2$ and $IDF_{dog} = \log(2)$, with similar logic applying to "cat"; this proceeds analogously for document 2 containing "puppy" and "kitten".

By contrast, in the word embeddings approach, these two documents are now represented as

$$X_1 = (1/2) \times \log(2) x_{dog} + (1/2) \times \log(2) x_{cat} \tag{B.12}$$

and similarly for X_2. Here x_{dog}, x_{cat} would have been trained using the Pennington et al. (2014) method described above on a very large outside set of documents. Hence, in this case since word vectors are estimated such that $x_{dog} \approx x_{puppy}$ and $x_{cat} \approx x_{kitten}$, we now have $\mathrm{Sim}_{1,2} \approx 0.81$ using the word vectors estimated by Pennington et al. (2014). A weighted average word embedding approach has been shown in the natural language processing literature to achieve good performance on standard benchmark tests for evaluating document similarity metrics relative to alternative methods that are much more costly to compute (see, e.g. Arora, Liang, and Ma, 2017). A relative disadvantage is that it ignores word ordering—which also applies to the more standard 'bag of words' approach for representing documents as vectors. However, since we have dropped all stop words and words that are not either a noun or a verb, retaining word ordering in our setting is far less relevant.

In sum, we use a combination of word embeddings and TF-IDF weights in constructing a distance metric between a patent document (which includes the abstract, claims, and the detailed description of the patented invention) and the detailed description of the tasks performed by occupations. Our methodology is conceptually related, though distinct, to the method proposed by Webb (2020), who also analyzes the similarity between a patent and O*NET job tasks. Webb (2020) focuses on similarity in verb-object pairs in the title and the abstract of patents with verb-object pairs in the job task descriptions and restricts his attention to patents identified as being related to robots, AI, or software. He uses word hierarchies obtained from WordNet to determine similarity in verb-object pairings. By contrast, we infer document similarity by using geometric representations of word meanings (GloVe) that have been estimated directly from word co-occurence counts. Furthermore, we use not only the abstract but the entirety of the patent document—which includes the abstract, claims, and the detailed description of the patented invention. In addition to employing a different methodology, we also have a broader focus: we are interested in constructing time-series indices of technology exposures. As such, we compute occupation-patent distance measures for all occupations and the entire set of USPTO patents since 1976.

Last, we perform several adjustments to the raw measure of similarity (12). First, we remove yearly fixed effects. We do so in order to account for language and structural differences in patent documents over time.[15] Second, we impose sparsity: after removing the fixed effects we set all patent $\times$ occupation pairs to zero that are below the 80th percentile in this fixed-effect adjusted similarity. This imposes that the vast majority of patent–occupation pairs are considered unrelated to one another, and only similarity scores sufficiently high in the distribution receive any weight. Last, we scale the remaining non-zero pairs such that a patent/occupation pair at the 80th percentile of yearly adjusted similarities has a score equal to zero and the maximum adjusted score equals one. We denote by $\rho_{i,j}$ the adjusted similarity metric between patent j and occupation i.

While we compute $\rho_{i,j}$ for all patent–occupation pairs, we restrict to the set of patents identified as breakthroughs by the KPST procedure when we construct our time-varying measures occupation-industry exposure to technology—as in the main text equation (13). Additionally, though we initially calculate $\rho_{i,j}$ for DOT documents with occupations defined at the 2010 SOC code level, we compute η and ξ using the modified Census occ1990 scheme (commonly referred to as "occ1990dd" codes) from Autor and Dorn (2013). The occ1990dd codes are comparable over time and can be readily linked to the CPS or Census, which proves useful for our analysis of occupational- and worker-level outcomes across various datasets. After crosswalking SOC occupations to the occ1990dd level, we take the average of the original 2010 SOC version of $\rho_{i,j}$ within each linked occ1990dd occupation code; we then use the resulting occ1990dd-level measure of $\rho_{i,j}$ in equation (13).

B.4 Constructing the shift-share IV

Let $s_{p,p'}$ give the raw cosine textual similarity between patents p and p', and $\mathbf{1}(p,p')$ is an indicator taking a value of one if $s_{p,p'} \geq 0.50$. Denote a patent CPC class by c; patent issue year by t; the set of patents issued in year t to tech class c by $P_{t,c}$. The text-based knowledge spillover from patent tech class c' to c at diffusion lag τ and year t is given by:

$$\omega_{c' \to c,t,\tau} = \sum_{p \in P_{t,c}} \sum_{p' \in P_{t-\tau,c'}} \mathbf{1}(p,p') \tag{B.13}$$

[15] Patents have become much longer and use much more technical language over the sample period, and the OCR text recognition of very early patents is far from perfect.

And the normalized spillovers

$$\Omega_{c'\to c,t,\tau} = \frac{\omega_{c'\to c,t,\tau}}{\sum_{c'\neq c}\omega_{c'\to c,t,\tau}} \tag{B.14}$$

The breakthrough intensity of tech class c in period t $I_{c,t}$ is given by the share of patents in that tech class which were breakthroughs:

$$I_{c,t} = \frac{\sum_{p\in B_{c,t}} 1}{\sum_{p\in P_{c,t}} 1} \tag{B.15}$$

where $B_{c,t}$ is the set of breakthrough patents in tech class c. We then construct a predictor for technological spillovers to tech class c at time t and diffusion lag τ by

$$\lambda_{c,t,\tau} = \sum_{c'\neq c} \Omega_{c'\to c,t,\tau} \times I_{c',t-\tau} \tag{B.16}$$

Next, let $\overline{\lambda}_{c,t}^{5\to 20}$ denote the average of $\lambda_{c,t,\tau}$ across diffusion lags $\tau = 5,\ldots,20$.

We then predict the number of breakthrough patents in time period t for tech class c, $\mathrm{N}_{c,t}^{B}$ by running the Poisson regression:

$$\log\left(E[\mathrm{N}_{c,t}^{B}|\lambda_{c,t}]\right) = \beta\overline{\lambda}_{c,t}^{5\to 20} + \alpha_t \tag{B.17}$$

Thus we predict breakthrough patenting in tech class c at time period t using textual knowledge spillovers into tech class c from patents issued in other tech classes from $t-5$ to $t-20$. By estimating (B.17) we obtain a prediction of $\widehat{N}_{c,t}^{B} = E[\mathrm{N}_{c,t}^{B}|\lambda_{c,t}]$. The measure $\overline{\lambda}_{c,t}^{5\to 20}$ is a strong predictor of $\mathrm{N}_{c,t}^{B}$: we obtain an estimate of $\widehat{\beta} = 26.5$ and a t-stat of 6.61 when we estimate (B.17).

We next use tech class predicted patenting to estimate occupational exposure as follows. Index occupations by o, industry by k, and let j index either routine or non-routine tasks. We then compute the expected value of the occupation-patent routine/non-routine textual similarity times an indicator for whether the breakthrough came from industry k, given that the breakthrough patent came from tech class c:

$$\alpha_{o,k,c,t} = E_t[\tilde{\rho}^{j}(p,o) \times \mathbf{1}(p,k)|p\in B_{c,t}] \tag{B.18}$$

Here $\mathbf{1}(p,k)$ is an indicator which takes the value of one if patent p was issued in industry k.

Let $\overline{\alpha}_{o,k,c,t}^{j}$ give the average value of $\alpha_{o,k,c,t}^{j}$ from t to $t-5$. Our IV for the innovation exposure of occupation o in industry k is then given by:

$$Z_{o,k,t}^{j} = \theta^{j}(o) \times \sum_{c} \overline{\alpha}_{o,k,c,t-5}^{j} \times \widehat{N}_{c,t}^{B} \tag{B.19}$$

Our instrument therefore leverages occupation-patent textual similarities occuring in a given industry over a window from 5 to 10 years prior, and textual spillovers coming from innovations in other tech classes occuring 5 to 20 years prior. We choose a shorter lag for computing occupation-industry-patent similarity because we can only observe the industry of origin for patents going back to 1976, while we can estimate the cross-patent textual network for as long as we have patent data.

We do not take the log of one plus the IV predicted patenting exposure in our main specifications in order to preserve the typical linear structure of a shift-share instrument. This instrument for exposure has a very right-skewed distribution, and so we also winsorize the right tail of (B.19) at the 5% level each year. We have also explored specifications where we construct an IV that mimics our main exposure measure by taking the log of one plus predicted patent exposure:

$$Z_{o,k,t}^{j,Alternate} = \theta^j(o) \times \log\left(1 + \sum_c \overline{\alpha}_{o,k,c,t-5}^j \times \widehat{N}_{c,t}^B\right) \tag{B.20}$$

We winsorize this version at cross-sectionally the 1% level. Our key findings are robust to using either version of the IV.

B.5 Matching occupation-industry entrants and incumbents

We match workers currently in a given occupation-industry with future entrants as follows. For each 4-digit NAICS industry and year t, we identify those workers who are not currently employed in that industry but will be employed in the industry within the next 3 years, and will continue to do so for at least 2 years. We further require entrants' CPS occupation observation year to occur within the next 3 years. Next, we find incumbent workers in the same occupation and industry at time t who also match with a candidate entrant worker on the following characteristics: age (within plus or minus 2 years); college graduation status ; gender; and prior earnings quintile for their age, gender, and education cell. We then take the average pre-period earnings of all incumbents matched to a given entrant to get $w_{I,t-2,t}^i$ in regression equation (22) in the data. We drop all entrant workers without at least one matched incumbent, leaving us with 57,500 worker-years (rounded according to Census rules) in our regression sample for equation (22).

B.6 Assigning patents to downstream industries through Input-Output tables

Information and communication technology (ICT) patents such as software may sometimes be used outside of the origin industry (Mann and Püttmann, 2023), which may be problematic for our baseline assignment of patents to the industry of origination. As an alternative to our main measure, we utilize BEA input/output tables (specifically, the 2012 use tables from `https://www.bea.gov/industry/input-output-accounts-data`) to allow breakthrough ICT patents to receive non-zero weight outside of their industry of origination. We define ICT patents as those whose CPC codes start with G0, G1, Y, or H0, which comprises over 80% of breakthrough patents in our sample period; for non-ICT patents we continue to assign their origin industry.

To do this, we first aggregate the quantities in the BEA use tables to the 4-digit NAICS level (or highest level of aggregation if 4-digit NAICS is not available). For industry k that is downstream from j, we then assign the following fraction of an ICT patent originating in industry l to industry

k:

$$\text{share}_{l\to k} = \frac{\$\text{ amount of inputs domestic industry } k \text{ used from } l}{\text{industry } l \text{ net output}} \tag{B.21}$$

Where net output of industry l is its total gross outputs minus total foreign imports. Since the BEA occasionally does not allow complete 4-digit NAICS resolution, in such a case we assign a fraction of the weight $\text{share}_{l\to k}$ according to the employment share of a 4-digit industry within the NAICS industry category k in our main DER-CPS merged sample (which in practice may be at the 2- or 3-digit NAICS level). This results in many pairs of very small weights which have little impact in practice, so we set weights below 0.01 to 0 in order to reduce the dimensionality. We then re-compute occupation-industry technology exposures in equation (13) from the main text, where we now include in the sum the downstream ICT patents multiplied by their assignment weights:

$$\xi^{j,\text{BEA IO}}(k,o,t) = \theta^j(o)\,\log\left(1 + \sum_{b\in\mathcal{B}_{k,t}} \tilde{\rho}^j(o,b) + \sum_{l\neq k}\sum_{b'\in B_{l,t}^{ICT}} \text{share}_{l\to k} \times \tilde{\rho}^j(o,b')\right), \qquad j \in \{R,N\}. \tag{B.22}$$

Where $B_{l,t}^{ICT}$ gives the ICT patents of upstream industry l. In appendix tables A.4 and A.5 we see that this broader patent assignment procedure performs quite similarly to our baseline measure, except usually with some loss of precision in the estimates—most likely because there is some noise involved in assigning patents to downstream industries along the production network.

B.7 BLS Industry Productivity Data

We obtain data on 4-digit NAICS industry outcomes from the Bureau of labor Statistics (found here: https://www.bls.gov/productivity/tables/labor-productivity-detailed-industries.xlsx). The data cover 1987-2022, though we are constrained to the 1987-2012 period based off availability of our industry innovation measures. BLS total factor productivity estimates are only available for manufacturing industries, so we take the industry per-unit labor costs index as our analogue of industry productivity in the model (this corresponds to the "Unit Labor Costs" series from the BLS data). This also has the benefit of matching well with our model-based notion of industry productivity (A.11), which is the per-unit input cost of final industry output. The cost-based productivity series is an index normalized to 100 within industry in the year 2012, and growth in costs are computed relative to this within-industry baseline. We also take changes in the log levels of the BLS labor share and labor compensation series to measure growth in industry labor shares and wagebills, respectively. Finally, since we control for log employment and weight by employment shares in the regressions, we use the BLS "hours worked" series to measure industry employment.

B.8 Publicly Available IPUMS Census Survey Data

While we make use of restricted access data sources (described further in the next subsection) for our primary analysis in sections 3.2 and 5 of the main text, we also construct a version of the measure that can be created using publicly available sources, which we utilize to summarize variation in our measure in Figure 3 and Figure 4. Specifically, Figure 4 we use the 5% samples from the 1980, 1990, and 2000 Censuses obtained from the Integrated Public Use Microdata Series (IPUMS) database.[16] We construct wages and individual labor supply weights following code from Acemoglu and Autor (2011). For this data we cannot use our direct links of patents to NAICS industries, and the best time series consistent industry code is the Census ind1990 designation. Hence we use probabilistic weights from 3-digit cooperative patent technology classes to NAICS industries created by Goldschlag et al. (2016). We obtain the extended version which includes service industries and can be found here. Goldschlag et al. (2016) provide links at the 2007 NAICS level, so we exploit the fact that IPUMS provides both 2007 NAICS designations (variable “indnaics") and time-series consistent ind1990 codes in the 2008-2012 ACS, which we use to create a correspondence between ind1990 and NAICS codes; we then use this correspondence to create probabilistic assignments of patents to ind1990 codes for each of the Census survey years; finally, weighting patents by industry assignment probabilities allows us to construct a version of our main measure (13) that varies at the occ1990dd-by-ind1990-by-year level. We then merge this onto our IPUMS sample and compute employment-weighted averages of exposure at the occupation-by-year level to generate Figure 4. Since we don’t use Census wage information in Figure 3, we take a simpler approach by simply aggregating the Goldschlag et al. (2016) probabilistic weights to 2-digit NAICS industries, and them summing the weights within different technology classes and for occupations of the designated task type, as indicated in the figure.

B.9 Census administrative data

B.9.1 DER-CPS Sample

We use a random sample of individual workers tracked by the Current Population Survey (CPS) and their associated Detailed Earnings Records (DER) from the Census—which contains their W2 tax income. We limit the sample to individuals who are older than 25 and younger than 55 years old. Our sample includes individuals coming from the 1981-1991, 1994, and 1996-2016 ASEC waves for whom valid individual identifiers (a Census Protected Individual Key) can be assigned. In selecting which years to include in our sample, we exactly follow the labor force attachment restrictions imposed by Braxton et al. (2021).

The CPS includes information on demographic information such as age and gender, but more importantly occupation at the time of the interview. We assign workers to occupations based on their response to the CPS survey (CPS “occ” variable). We construct a crosswalk between the yearly CPS occupations codes and the occ1990dd classification scheme and assign all CPS occupations their corresponding occ1990dd code. We assign this occupation to the worker for the next 3 years, thus effectively dropping observations where the CPS interview date is older than 3 years—so that the occupation information is relatively recent. Workers’ employers are differentiated by the

[16]While in principle we could supplement with ACS survey data for later years, our measure cuts off in 2007 and so we use the 1980, 1990, and 2000 Censuses to preserve a consistent 10-year sample interval.

federal Employer Identification Number (EIN); if workers report earnings from more than one EIN in a calendar year then we assign them the EIN of highest W2 earnings in that year. Our final worker-year earnings panel spans the years 1981-2016.

B.9.2 Assigning Industry of Origination to Patents

We merge the individual worker records from the CPS-DER matched sample to patent data at the industry (4-digit NAICS) level. Specifically, we identify the industry of patent origination by relying on the Census SSEL patent–assignee database, which provides a corresponding SSEL firm identifier ("firmid") for each patent, which we then use to obtain the firm's 4-digit NAICS code. In particular, we use two SSEL patent–assignee crosswalks: the newer Business Dynamics Statistics of Patenting Firms database (BDS-PF) and an older patent-SSL crosswalk created by Kerr and Fu (2008). The BDS-PF links are available starting with the 2000 SSEL. We use the BDS-PF firmid-patent links for any patents for which it is available. Otherwise we use the union of links created by Kerr and Fu (2008) from the 1976–1999 SSL data. We obtain NAICS codes by using the "firmid" identifier to join with the Longitudinal Business Database (LBD), and we use the 2012 version of Fort and Klimek (2018) NAICS codes, which are adjusted to improve industry comparability over time. In cases where a firmid matches to multiple NAICS codes we apply the 4-digit NAICS code of highest employment based off the LBD. Finally, we drop from our analysis patents that are assigned to the 2-digit NAICS code 55, "Management of Companies and Enterprises", since this code consists of firms that manage and hold controlling interests in other companies, making the NAICS code uninformative about the industry where the actual production is taking place.

B.9.3 Decennial and ACS Surveys

For the analysis in Table 2, we use restricted access versions of Census Decennial surveys in 1980, 1990, and 2000; because ACS surveys have much smaller sample size, we average across ACS years 2008-2012 for observations in the year 2010. These surveys have the advantage of using the actual earnings responses rather than top-coded earnings.

In order to construct a crosswalk between the Census industry codes available in each of these survey years in the restricted-access survey data, we first obtain the mapping between the time-consistent Census 1990 industry scheme found in publicly available IPUMS surveys and the 2012 NAICS coding scheme. We use the 2013-2017 5-year ACS to do this because it is the first 5-year survey period containing 2012 NAICS codes. Because the 1990 Census industry scheme is available for all years, we can then construct a mapping between Census industry codes and 4-digit NAICS codes, which is the level of industry that we use throughout the main analysis in the paper. Because some NAICS codes are aggregated in the mapping to the 1990 Census industry scheme, we have to consolidate both Census and NAICS codes to create consistent mapping that is as close as possible to the original 4-digit NAICS codes. After consolidating we are left with 188 modified 4-digit NAICS codes from the original 288.

Following the age restrictions in our worker-level analysis, we include individuals between the ages of 25 to 55 in our sample. We compute employment by taking the total survey sample weight across individuals in an occupation–industry cell. We weight regressions by total labor hours in the occupation–industry cell, as given by survey sample weight times individuals' survey-reported hours worked per week, times weeks worked per year. If hours or weeks worked are reported in intervals for a given survey year, we take the median of that interval. As in our CPS-based analysis, we crosswalk occupations to the time-consistent occ1990dd coding scheme throughout. We then take occupation-industry aggregates of each outcome and control in each survey year. Specifically, we take the total labor employment weights, and compute the total labor hour-weighted averages of hourly wage, age, female population share, and years of education for regression controls.

B.9.4 Longitudinal Business Database

We merge the DER-CPS sample with the Longitudinal Business Database (hereafter LBD, Jarmin and Miranda (2002)) to get worker industry information and data on employers' payrolls and employment. We assign workers to a consistently-defined industry code based on their employer's industry codes in the LBD. To do this we use a deduplicated crosswalk from LBD firm identifiers ("firmid") to firm EINs, and we assign each EIN with the Fort and Klimek (2018) 2012 4-digit NAICS code based off LBD industry information. We compute firm-level wages by aggregating yearly LBD payroll and employment information to the EIN level, and we compute firm average wages for a given year by taking the ratio of payroll to total employment. Since LBD annual payroll information is reported for the last 12 months as of March of that year, while W2 earnings are reported over the calendar year, we use the average of the LBD-implied wage for years t and year $t+1$ as a measure of employers' average wages in year t. We use the 2016 of edition of the LBD to compute firm average wages and to obtain firm industry information. We have LBD information on firm industries and wages from the start of our sample in 1981 through 2015.

B.9.5 Restrictions for Assigning Worker Earnings Bins

To allow the effects to vary with prior income, we assign workers into five groups based on their income in the previous year compared to workers in the same occupation and NAICS4 industry. These groups are defined based on the following percentiles of prior income [0%, 25%), [25%, 50%), [50%, 75%), [75%, 95%), [95%, 100%] calculated within industry–occupation cells. In the (uncommon) case when NAICS4 industry-occupation cells have fewer than 10 individuals, we broaden the industry definition from 4-digit NAICS to 2-digit NAICS. Any Industry–Occupation cells that still have fewer than 10 individuals after moving to 2-digit NAICS codes are dropped.

Appendix References

Abel, A. B. (2019). Classical measurement error with several regressors. Working paper, University of Pennsylvania.

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics, Volume 4. Amsterdam: Elsevier-North*, pp. 1043–1171.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review 103*(5), 1553–97.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Braxton, J. C., K. F. Herkenhoff, J. L. Rothbaum, and L. Schmidt (2021). Changing income risk across the US skill distribution: Evidence from a generalized kalman filter. Technical report, National Bureau of Economic Research.

Broda, C. and D. E. Weinstein (2006). Globalization and the Gains From Variety*. *The Quarterly Journal of Economics 121*(2), 541–585.

Caunedo, J., D. Jaume, and E. Keller (2023). Occupational exposure to capital-embodied technical change. *American Economic Review 113*(6), 1642–85.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review 101*(3), 471–75.

Fort, T. C. and S. Klimek (2018). The effects of industry classification changes on us employment composition. Working Paper 18-28.

Goldin, C. and L. Katz (2008). *The Race Between Education and Technology.* Harvard University Press.

Goldschlag, N., T. J. Lybbert, and N. J. Zolas (2016). An 'algorithmic links with probabilities' crosswalk for uspc and cpc patent classifications with an application towards industrial technology composition. CES Discussion Paper 16-15, U.S. Census Bureau.

Jarmin, R. S. and J. Miranda (2002). The longitudinal business database. Technical report.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021). Measuring technological innovation over the long run. *American Economic Review: Insights 3*(3), 303–20.

Kerr, W. and S. Fu (2008). The survey of industrial r&d—patent database link project. *The Journal of Technology Transfer 33*, 173–186.

Koh, D., R. Santaeulalia-Llopis, and Y. Zheng (2020). Labor share decline and intellectual property products capital. *Econometrica 88*(6), 2609–2628.

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull, and G. L. Violante (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica 68*(5), 1029–1053.

Mikolov, T., I. Sutskever, K. Chen, G. Corrado, and J. Dean (2013). Distributed representations of words and phrases and their compositionality. *CoRR abs/1310.4546*.

Ohanian, L. E., M. Orak, and S. Shen (2023). Revisiting capital-skill complementarity, inequality, and labor share. *Review of Economic Dynamics 51*, 479–505.

Pennington, J., R. Socher, and C. D. Manning (2014). Glove: Global vectors for word representation. In *EMNLP*.

Webb, M. (2020). The impact of artificial intelligence on the labor market.