Labor, Despite AI *

Matteo Chequer Kyle Herkenhoff Dimitris Papanikolaou
Jonathan Rothbaum Lawrence Schmidt Bryan Seegmiller

March 26, 2026

Do not circulate.

Abstract

The hypothesis of this paper is that AI models must be trained on data generated by workers (both in- and out-of-house), and that this complementarity between labor and data fundamentally alters AI's effects on the labor market. We develop this hypothesis in three steps. We link the Annual Business Survey (ABS) to individual 1040 tax records and document (1) the importance of in-house R&D for AI adoption, (2) the importance of data and talent for AI adoption, and (3) the heterogeneous effects of AI adoption on wage outcomes. We find that exposed workers gain relative to non-exposed workers on average, particularly so in occupations where firms conduct significant in-house R&D on AI, while a significant share of exposed workers endure relative losses at downsizing firms. Second, we build a theory of AI adoption as an experience good (i.e., AI combines external and internal data on worker tasks) and we integrate this learning process into a directed search model with skill-weights. We allow firms to adapt their skill-weights as they partially or fully automate workers with AI, endogenously giving rise to new work, and we allow workers to retrain. Lastly, we use our empirics to estimate the model and exploit the block recursive nature of our model to compute transition paths of output, wages, and job flows in response to both sudden and gradual AI technology improvements. With few exceptions, we consistently find that the job creation margin wins the long run, even under large reductions in the cost of AI.

JEL codes: E24, O30

Keywords: Automation, Labor Market, AI, Future of Work, Directed Search, Data, Learning

*Chequer: University of Minnesota, Herkenhoff: University of Minnesota, Papanikolaou: Northwestern University, Rothbaum: U.S. Census Bureau, Schmidt: Massachusetts Institute of Technology, Seegmiller: Northwestern University. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily reflect the position of the Federal Reserve Bank of Minneapolis or the Federal Reserve System or the U.S. Census Bureau. The U.S. 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 (P-7518691); Disclosure Review Board (DRB) approval number: CBDRB-FY25-SEHSD003-067, CBDRB-FY26-SEHSD003-049, CBDRB-FY26-SEHSD003-052).

1 Introduction

Existing macro-labor models of AI assume that AI advances independently from the data generated by labor (i.e. AI "comes from the sky") and that workers cannot retrain in response to the arrival of AI. In this paper, we argue that these omissions are first-order when assessing the impact of AI on earnings and employment. We use new administrative firm surveys conducted by the Census Bureau to document the importance of in-house R&D, worker talent, and data for the adoption of AI. We then empirically assess the wage implications of AI adoption by providing the first linkage, to our knowledge, between worker-level tax records and firm-reported AI adoption. Motivated by our empirical findings, we develop a dynamic model of the labor market in which workers generate data by doing their jobs and may retrain in response to the arrival of AI. In-house and out-of-house data are then combined to adopt AI and make it more efficient.

We use our model address the pertinent questions of *How will AI affect the labor market? Who wins and loses? How does the nature of occupations change?* and we move beyond existing studies to answer *How should we insure workers displaced by AI?* We argue that across a broad range of potential paths for AI adoption costs, job creation wins and that high skill (who are also high-exposure) workers reap the majority of wage gains from AI. Moreover, under standard utilitarian welfare weights, an expansion of job re-training subsidies would yield an [X]%-[Y]% CEV welfare gain. Our finding that labor is ultimately better off with AI critically hinges on labor's role in the generation of data and advancement of AI.

We make three contributions. Our first contribution is empirical. We build the first link between between worker-level tax records, firm-reported adoption in the Annual Business Survey (ABS), and AI task exposure from resume data, (e.g. Hampole et al. (2025)). We document several facts: (1) approximately 6% of firms in the 2023 ABS report shedding workers over the prior 3 years due to AI adoption, whereas 10% report expanding; (2) 36% of AI adopters say their intent is to automate tasks performed by workers; (3) among the 1% firms who use AI *intensively* in the 2022 ABS, 50% report conducting in-house R&D on AI; (4) among firms who adopt AI and shed workers, highly exposed workers have 5% to 7% earnings losses 2 years after adoption relative to similar exposed workers at non-shedders; and (5) exposed workers at non-shedders experience earnings gains of 2% to 4% relative to non-exposed workers at those same firms. Our results suggest disparate worker-level outcomes: some workers gain, some workers lose, and the process of adopting AI intensively requires seom degree of in-house R&D. These empirical features motivate our model, which is necessary to conduct counterfactuals and simulate potential paths for the U.S. labor market.

Our second contribution is theoretical. To date, the rapidly expanding literature on the labor

market effects of AI and automation (Acemoglu and Restrepo (2018), Korinek and Suh (2024), Freund and Mann (2025)) has developed largely in isolation from the large existing literature on worker skill and job dynamics (Ljungqvist and Sargent (1998), Menzio and Shi (2011), Lise and Postel-Vinay (2020)) and has modeled AI in the same way as physical capital and investment, despite the importance of data in the development of AI (e.g., Beraja et al. (2023), Veldkamp and Chung (2024) among others). The few exceptions to this rule (e.g. Althoff and Reichardt (2026)) assume that workers may learn on-the-job but that AI technologies exogenously advance, and the few macro-labor models that incorporate some degree of AI's learning by doing (e.g. Wang and Wong (2026)) ignore the role of retraining and external data. Instead, we model AI investment as an experience good – AI requires firms to accumulate internal data on worker tasks – and we integrate this learning process into a directed search model (e.g., Menzio and Shi (2011)) in which jobs are differentiated by skill-weights (e.g. Lazear (2009) and Lise and Postel-Vinay (2020)). However, unlike Lazear (2009) which models skill-weights as a fixed intrinsic characteristics of a firm, we allow firms to adapt their skill-weights as they partially or fully automate workers with AI. The shifting skill-weights endogenously give rise to "new work" (Autor et al. (2024)). Moreover, by modeling AI as experience goods that require data, we marry elements of the data literature (e.g. Jones and Tonetti (2020) and Farboodi and Veldkamp (2021)) with job dynamics.

Our theory of innovation itself is novel and builds on Acemoglu and Restrepo (2018), Hémous and Olsen (2022), and Beraja and Zorzi (2025) in several ways: (1) AI automation (henceforth *automation*) is partly an experience good (firms collect in-house data on workers in order to adopt AI technologies capable of replicating the workers' tasks) and partly an inspection good (some part of AI can be adopted without in-house data), (2) we allow for both full and partial automation, (3) we let the nature and task content of work adapt and vary over time at the firm-level, (4) the speed of automation depends on worker characteristics, (5) we allow AI to continue to learn from workers after adoption, and lastly (6) we allow for non-stationary bursts of innovation. These features are motivated by our empirical findings, especially the importance of in-house R&D, as well as the staggered nature of AI adoption and its dependence on worker traits. What ultimately differentiates our predictions from those of other models is that AI is more productive when it is trained on in-house data, limiting the willingness of firms to substitute into AI and completely automate workers.

Our model is tractable and yields several theoretic predictions. Time to AI adoption is faster for higher wage workers, conditional on human capital (Dechezleprêtre et al. (2019), Acemoglu and Restrepo (2024)). Some workers enjoy wage gains and become "pace setters" (like a pace setting rabbit for a greyhound race) in which they have short stints at firms, generating enough

new data to enable AI adoption. Partial automation is preferred to full automation for workers with higher non-automatable human capital. If data and non-automatable human capital are complements, partial automation occurs faster. If data and non-automatable human capital are substitutes, full automation occurs faster.

Our third contribution is quantitative. We use our resume-Census linked data to estimate the model's parameters, and we exploit the block recursivity of our problem to compute paths of output, wages, and job flows in response to both sudden and gradual AI technology improvements. We find that even when the cost of automation falls sharply, workers are not broadly displaced. Instead, AI changes what workers do on the job: time shifts away from automatable tasks and toward tasks that AI cannot perform. Because partially automated matches become more productive, average wages rise by about 9 percent as automation gets cheaper. The largest gains accrue to workers with high overall human capital, who also tend to be the most exposed to AI, since partial automation raises their marginal product. By contrast, workers at the bottom of the exposure distribution are affected little, because their jobs are neither attractive targets for automation nor productive enough to justify adoption. The workers most at risk are those with high automatable but low non-automatable human capital, especially when external data is cheap. For these workers, retraining becomes the main adjustment margin, with the retraining rate rising from 14 to 22 percent in the cheap-data scenario.

The only force that generates wage losses in the model is a decline in the cost of external data, and even then the drop is only slightly above 1 percent. Cheaper automation itself instead produces sizable wage gains without increasing displacement. This distinction reflects our treatment of AI as an experience good. If firms could adopt AI without relying on internally generated data, job destruction would be substantially larger. Because workers help produce the data that makes AI valuable, labor remains complementary to AI even as adoption costs fall. In that sense, labor is an input into AI itself, and that complementarity limits the displacement effects that would arise in models where automation is purely an inspection good.

We are working on optimal UI exercises and further honing our quantitative results. We plan to build on earlier work (Jaimovich et al. (2021), Braxton and Taska (2025), Moll et al. (2022), Beraja and Zorzi (2025)) and reassess the gains from optimizing the safety and retraining policies in settings in which AI adoption is partial and requires data on worker tasks.

Literature. A rapidly growing body of work documents the pace and scope of AI adoption. Early work by McElheran et al. (2024) and Bonney et al. (2024) use the Annual Business Survey (ABS) and Business Trends and Outlook Survey (BTOS) to provide benchmark adoption

metrics, while Bick et al. (2026) find substantially higher adoption rates in worker-level data.[1] Babina et al. (2024) and Kogan et al. (2023) track diffusion through resume data; Acemoglu et al. (2022) and Hosseini Maasoum and Lichtinger (2025) measure demand for AI competencies in job postings. We build on this work by linking firm-reported adoption from the ABS directly to longitudinal worker outcomes in administrative tax data.

Several recent papers examine the worker-level effects of AI exposure. Humlum and Vestergaard (2025) use worker-reported adoption in key occupations in Denmark and find negligible average wage effects. Brynjolfsson et al. (2025) document relative earnings losses among junior workers in exposed occupations in private payroll data, and Hosseini Maasoum and Lichtinger (2025) corroborate this finding in Revelio data. Randomized controlled trials highlight the importance of in-house algorithm training and the difficulty that general-purpose models face with infrequent tasks (Brynjolfsson et al., 2025).[2] There is also a large literature on the effects of robot adoption on labor markets (Acemoglu and Restrepo, 2020), including evidence that higher-wage low-skill workers are more likely to be automated (Dechezleprêtre et al., 2019). Our administrative panel and rich occupation-level exposure metrics allow us to study these dynamics in a nationally representative U.S. data.

Our second contribution is theoretical. By modeling AI as an experience good that requires data, we marry elements of the data literature (Jones and Tonetti, 2020; Farboodi and Veldkamp, 2021) with job dynamics. Moreover, the expanding literature on the labor market effects of AI and automation (Acemoglu and Restrepo, 2018; Korinek and Suh, 2024; Freund and Mann, 2025) has developed largely in isolation from the existing literature on worker skill and job dynamics (Ljungqvist and Sargent, 1998; Menzio and Shi, 2011; Lise and Postel-Vinay, 2020) and data dynamics, with notable exceptions by (Althoff and Reichardt, 2026) and (Wang and Wong, 2026). Relative to every one of these papers, including (Althoff and Reichardt, 2026) and (Wang and Wong, 2026), we provide a novel model of AI investment as an experience good: firms must accumulate internal data on worker tasks before automation becomes productive. We allow firms to adapt their skill-weights as they partially or fully automate worker tasks, endogenously giving rise to "new work" (Autor et al., 2024),[3] and we allow workers to retrain (e.g. Jacobson et al. (2005) and in the context of AI, Hyman et al. (2025)). By allowing for retraining, our model moves us closer to making credible medium-run predictions.

On the normative side, several papers study optimal policy responses to automation in heterogeneous-agent models (Jaimovich et al., 2021; Braxton and Taska, 2025; Moll et al., 2022;

[1] See also the Atlanta Fed survey Baslandze et al. (2026).

[2] See Dell'Acqua et al. (2023) and Gathmann et al. (2024) for additional high-quality RCTs in specific occupations.

[3] See also Dillon et al. (2025) for discussion of the changing nature of occupations in the face of generative AI.

Beraja and Zorzi, 2025). We contribute to this literature by developing a tractable model of risk-averse workers in an environment with data-driven AI advancement. These advances allow us to discuss both the efficiency and redistributive aspects of optimal policy in response to AI in a meaningful way.

2 Empirics

Section 2.1 describes the firm-level effects of AI adoption based on the Annual Business Survey (ABS). Section 2.2 tracks worker outcomes after AI adoption among ABS firms. Section 2.4 measures the AI adoption timeline, i.e. the time between initially hiring an occupation and its (partial) AI automation.

2.1 AI Adoption at the Firm-level

Our primary source for firm-level effects of AI adoption is the ABS. The definition of artificial intelligence and the sets of questions fielded vary from year to year. We primarily use the 2023 ABS, which asked specific questions about the timing of AI adoption. We augment this with unique questions on AI R&D fielded in the 2022 ABS and barrier-to-adoption questions from the 2019 ABS. The 2023 ABS is the latest year of the ABS available to researchers in the Census data warehouse. We document seven facts:

I. Few firms use AI in production (2.9%–4.6%), but it affects many workers (14%–22% employment-weighted);

II. Among adopters, 37% of firms intend to automate tasks performed by labor (63% employment-weighted);

III. Among adopters, 83% of firms report no workforce change due to AI, 11% report employment expansions, and 6% report employment contractions;

IV. Among intensive AI users, 50% of firms conduct in-house R&D (77% employment-weighted); even among all adopters, 18% of firms (50% employment-weighted) perform in-house R&D;

V. AI adopters that conduct R&D are more likely to report employment expansions attributable to AI; contraction rates do not differ significantly by R&D status;

VI. AI adoption is associated with increased workforce skill requirements, and R&D-active adopters exhibit particularly strong upskilling; and

VII. Conditional on testing AI, firms cite data access (9%), data quality (8%), and talent constraints (11%) as barriers to adoption.

Facts I–III. We begin with adoption patterns, adoption intent, and self-reported workforce changes in the 2023 ABS.[4] The 2023 ABS defines artificial intelligence as "*a branch of computer science and engineering devoted to making machines intelligent*," where intelligence refers to the ability to perceive, analyze, determine a response, and act appropriately. We code firms as AI users if AI is "*used as part of the processes or methods*" of production.

Figure 1[5] shows that only 2.89 percent of firms report AI as part of their production, as opposed to testing AI. Using employment weights, however, 14 percent of workers are at firms using AI in production. AI is not yet pervasive across firms, but it is disproportionately present at larger employers. As shown in Figure 2, 37 percent of adopters (63 percent employment-weighted) report that their intent is to automate tasks performed by labor. Despite this stated intent, most adopters report no employment change: 83 percent of firms (78 percent employment-weighted) indicate no workforce adjustment attributable to AI (Figure 3). Still, non-trivial fractions report expansions (11 percent firm-weighted, 19 percent employment-weighted) and contractions (6 percent firm-weighted, 3 percent employment-weighted). AI adoption is not mechanically labor-reducing; short-run employment effects are heterogeneous and skewed toward larger establishments.

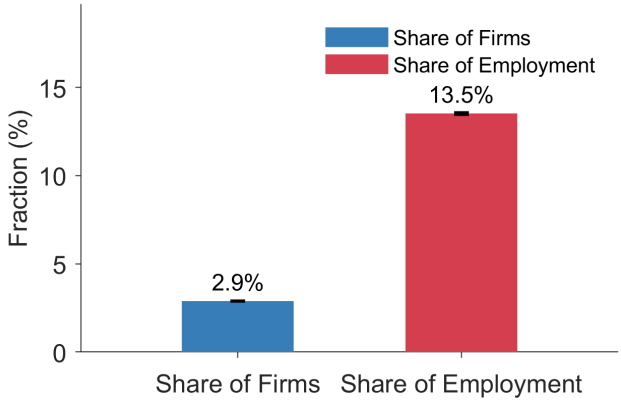


Figure 1: **AI adoption rate, 2023 ABS (reference years 2020–2022).**

Notes. Blue bar: share of firms using AI (survey-weighted); red bar: share of employment at AI-using firms (employment-weighted); $N \approx 460{,}000$. Whiskers show ± 1 SE. *Survey question:* "During the 3 years from 2020 to 2022, did this business adopt/use the following technologies? Artificial Intelligence." Figure shows the share responding "Used as part of processes or methods."

[4]The 2023 ABS was fielded in 2022, contemporaneous with the introduction of generative, pre-trained transformer models.

[5]See Table B1.

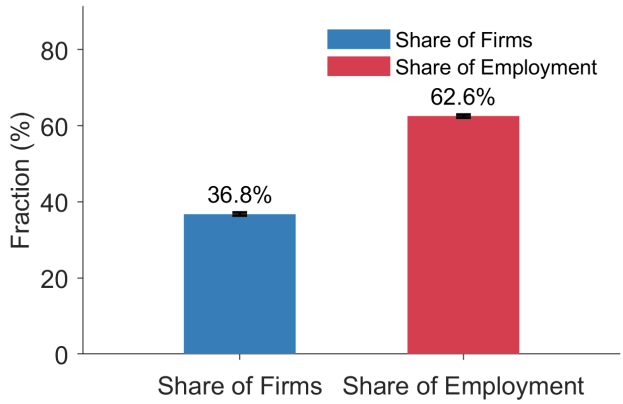


Figure 2: **Labor-saving intent among AI adopters, 2023 ABS (reference years 2020–2022).**

Notes. Sample: AI-adopting firms; $N \approx 16{,}000$. Blue bar: share of firms (survey-weighted); red bar: share of employment (employment-weighted). Whiskers show ± 1 SE. *Survey question:* "Among the following, what is this business's intention for adopting or using artificial intelligence? To automate tasks performed by labor."



Figure 3: **Employment effects of AI adoption, 2023 ABS (reference years 2020–2022).**

Notes. Sample: AI-adopting firms; $N \approx 16{,}000$. Panel (a): survey-weighted (share of firms); panel (b): employment-weighted (share of employment). Bars show the fraction reporting employment increased (left) and decreased (right). Whiskers show ± 1 SE. *Survey question:* "During the three years 2020 to 2022, what were the effects of adopting or using Artificial Intelligence on the following? The number of workers employed by this business."

Fact IV. Figure 4[6] examines adoption intensity and in-house R&D using the 2022 ABS, where AI is defined as "*Artificial intelligence (e.g., machine learning, planning, reasoning, and decision making).*" Overall adoption is 4.60 percent of firms (22 percent employment-weighted). Among all firms, only 1 percent report performing or funding in-house AI R&D (11 percent employment-weighted). Conditional on high-intensity AI use (self-reporting AI "*a lot*"), however, 50 percent of firms (77 percent employment-weighted) conduct in-house AI R&D.[7] Even among all

[6] See Table B2.

[7] In the 2019 ABS, using any technology "a lot" corresponded to using the technology for 25% of production.

adopters, 18 percent (50 percent employment-weighted) perform in-house R&D. Intensive AI use is therefore tightly linked to internal knowledge production: AI appears not merely as an off-the-shelf capital input but as a technology whose deployment is associated with firm-specific data accumulation and R&D investment.

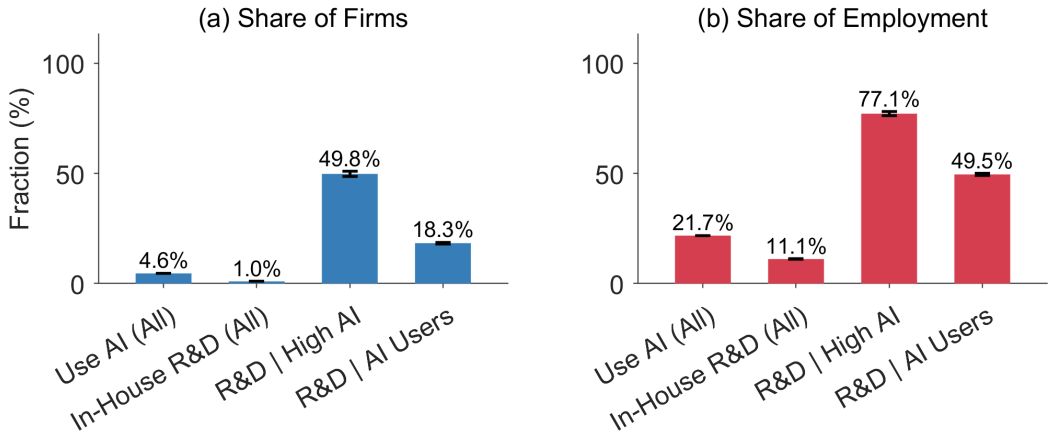


Figure 4: **AI use and in-house R&D rates, 2022 ABS (reference year 2021).**

Notes. Panel (a): survey-weighted (share of firms); panel (b): employment-weighted (share of employment). Four bars per panel: (i) "Use AI (All)": share of all firms using AI ($N \approx 179{,}000$); (ii) "In-House R&D (All)": share of all firms performing in-house R&D on AI ($N \approx 179{,}000$); (iii) "R&D | High AI": share with R&D among high-intensity AI users ($N \approx 2{,}000$); (iv) "R&D | AI Users": share with R&D among all AI-using firms ($N \approx 13{,}000$). Whiskers show ± 1 SE. *AI use question:* "To what extent did this business use Artificial Intelligence?" (response "A lot" = high intensity.) *R&D question:* "During 2021, did this business perform or fund R&D on the following technologies?"

Fact V. Figure 5 and Appendix Figure 20[8] compare employment outcomes between all AI adopters and those that also perform R&D, for the 2022 and 2023 ABS respectively. In the 2022 ABS, 15.9 percent of all AI adopters ($N \approx 5{,}000$) report employment increases, compared to 24.7 percent among AI adopters with in-house R&D ($N \approx 2{,}000$), a gap of 8.9 percentage points. In the 2023 ABS, 10.7 percent of all adopters report expansions versus 17.5 percent among R&D-active adopters ($N \approx 6{,}000$), a gap of 6.8 percentage points. By contrast, contraction rates are nearly identical across groups in both years: 7.9 versus 8.4 percent in 2022, and 6.1 versus 6.2 percent in 2023. While these patterns suggest that firms conducting in-house AI R&D enjoy greater employment gains conditional on adoption, the obvious caveat is that the choice to conduct R&D is endogenous. In what follows, we attempt to address this in our worker-level analysis.

Fact VI. Beyond headcount effects, the ABS asks adopters about changes in workforce skill composition and STEM content attributable to AI. Figure 6 and Appendix Figure 21[9] show skill-level outcomes by R&D status. In both the 2022 and 2023 ABS, reported skill-level increases

This output requirement was dropped from the 2022 ABS, so "a lot" has no quantitative meaning.

[8] See Tables B4 and B5.

[9] See Tables B6 and B7.

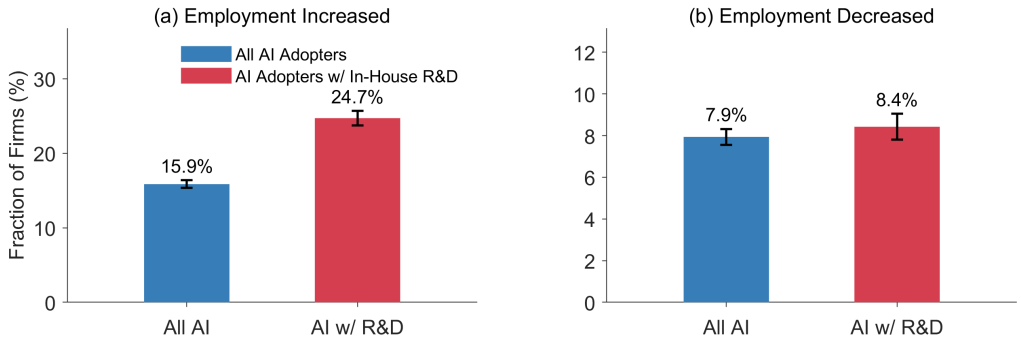


Figure 5: **Employment effects of AI adoption by in-house R&D status, 2022 ABS (reference year 2021).**

Notes. Blue bar: all AI adopters ($N \approx 5{,}000$); red bar: AI adopters with in-house R&D ($N \approx 2{,}000$). Panel (a): fraction reporting employment increased; panel (b): fraction reporting employment decreased. Whiskers show ± 1 SE. *Employment question:* "To what extent did this business's use of AI technology impact your workforce? The number of workers employed at this business." *R&D question:* "During 2021, did the business perform or fund R&D on the following technologies?"

substantially exceed reported skill-level decreases, and R&D-active adopters report higher rates of skill upgrading than the all-adopter average. Appendix Figures 22 and 23 document the same pattern for STEM composition: AI adoption is associated with increased STEM-intensive work, more pronounced among R&D-conducting firms. Together, these patterns suggest that AI deployment and internal knowledge production are complements—firms that invest in R&D restructure their workforces toward higher-skill and more STEM-intensive occupations.

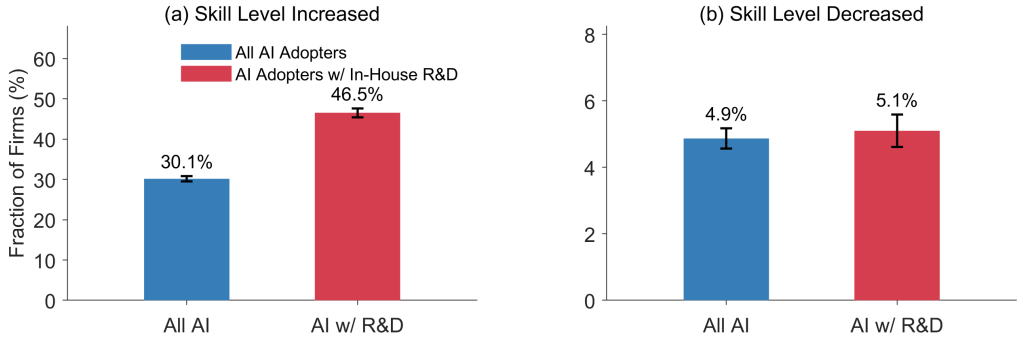


Figure 6: **Workforce skill-level effects of AI adoption by in-house R&D status, 2022 ABS (reference year 2021).**

Notes. Sample and R&D definition as in Figure 5. Panel (a): fraction reporting skill level increased; panel (b): fraction reporting skill level decreased. Whiskers show ± 1 SE. *Survey question:* "To what extent did this business's use of AI technology impact your workforce? The skill level of workers employed at this business."

Fact VII. Figure 7[10] reports barriers to AI adoption using the 2019 ABS, restricted to firms that tested AI between 2016 and 2018. Conditional on AI being applicable to the business, the most

[10]See Table B3.

common response for non-adoption is cost. However, a substantial fraction cite talent and data-related constraints: 9 percent of firms report lack of data access, 8 percent report poor data quality, and 11 percent report lack of talent. Under employment weights, these constraints are even more pronounced—36 percent of larger firms report poor data quality and lack of talent. Only 27 percent of firms (16 percent employment-weighted) report no barriers. These patterns indicate that AI adoption is data and human capital intensive, and that successful deployment depends critically on high-quality data and specialized talent rather than solely on access to off-the-shelf technologies.



Figure 7: **Data and talent barriers to AI adoption, 2019 ABS (reference years 2016–2018).**

Notes. Sample: firms that tested AI; $N \approx 2{,}000$. Panel (a): survey-weighted (share of firms); panel (b): employment-weighted (share of employment). Three bars: "No Data" (insufficient data availability), "Bad Data" (poor data quality), and "No Talent" (insufficient human capital). Cost and no-barrier responses are reported in Table B3. Whiskers show ± 1 SE. *Survey question:* "During the three years 2016 to 2018, indicate which factors adversely affected the adoption or utilization of the following technologies to produce goods or services."

Caveats. Our analysis uses the 2019, 2022, and 2023 ABS surveys available through the Census data warehouse, so it captures only the early phase of modern generative pre-trained transformer (GPT) models.[11] In the subsequent BTOS data, both reported adoption rates rise rapidly, suggesting that the trends we document in the ABS may shift sharply in later years. That said, the ABS and BTOS use different adoption thresholds: the ABS asks whether AI is used in production, whereas the newer BTOS asks whether AI is used in any business function. Even with this difference, the BTOS points to substantially higher adoption in recent years. It is also plausible that firms changed their in-house R&D behavior with the arrival of modern GPTs, and that several of the patterns we document may shift discretely in more recent data. At present, however, we do not have newer U.S. evidence on this margin.

[11] Earlier ABS surveys asked about the use of large language models (LLMs) and machine learning (ML) algorithms. LLMs were already being used in economically meaningful ways before the introduction of GPT. Appendix Figure 24 shows that a nontrivial number of firms were already using LLMs in the 2019 ABS.

2.2 AI Adoption at the Worker Level

Next, we move beyond self-reported firm outcomes and provide, to our knowledge, the first link of tax records to administrative firm adoption of AI. In this section, we study worker-level W-2 wage growth around firm-level AI adoption by linking four data sources: (i) Form 1040 individual tax returns (2015–2023), (ii) the 2023 Annual Business Survey (ABS) adoption metrics, (iii) occupation codes from the American Community Survey (ACS), and (iv) occupation-level AI exposure measures derived from Hampole et al. (2025). The resulting panel allows us to observe workers' earnings trajectories before and after AI adoption at their firms, while holding fixed pre-adoption occupation-based exposure.

Tax Data. The spine of the panel consists of 1040 tax returns from 2015 to 2023. These records provide annual W-2 wage and salary income for each taxpayer. Each return contains a taxpayer identification number (TIN) and employer identification number (EIN). We use an EIN-to-firmid crosswalk to assign workers to Census firm identifiers in each year, enabling us to track workers across firms and to merge firm-level AI adoption information.[12]

AI Adoption. We link firms to the Annual Business Survey (ABS) using the aforementioned firm identifiers. The 2023 ABS includes a retrospective question on the timing of AI adoption (including early generative AI in the 2023 ABS). We restrict attention to firms reporting adoption in the 2021–2022 window, and treat these firms as adopting in 2021. While this bucket may include some 2022 adopters, 2021 is the earliest clean post period observed in the tax data. We further require that each firm contain at least two distinct exposure categories in the year prior to adoption, ensuring within-firm variation in exposure.

Exposure. To measure worker-level exposure to AI, we link individuals to the ACS and recover occupation codes. For each worker, we retain the most recent occupation observed strictly before 2020 and fill this occupation forward and backward. Exposure from Hampole et al. (2025) is then assigned at the occupation level using 1990 Census occupation codes, which are mapped over time in a consistent way based on David Dorn's work (Autor and Dorn, 2013). Because exposure depends only on pre-adoption occupation, it is time-invariant and predetermined relative to AI adoption at the firm.

We define a high-exposure indicator

$$T_i = \mathbf{1}(e_i > e^{90}),$$

[12] "firmid" is the Census' identifier for firms.

where e_i denotes occupation-level AI exposure and e^{90} is the 90th percentile of the exposure distribution. We rank workers in the year before the event and then fill this forward and backwards. Workers in occupations above this threshold – including, for example, mathematicians and statisticians, management analysts, economists, market and survey researchers, operations and systems researchers, etc. – are classified as highly exposed.

Outcome. The outcome is annual W-2 wage growth. Let $w_{2,it}$ denote W-2 earnings for worker i in year t.[13] We residualize log W-2 earnings on a full set of age dummies:

$$\ln w_{2,it} = \sum_a \delta_a \cdot \mathbf{1}(\text{age}_{it} = a) + \epsilon_{it},$$

and define $y_{it} = \Delta\hat{\epsilon}_{it}$, the first difference of the residual. This removes lifecycle earnings dynamics while preserving idiosyncratic wage growth.

Let τ denote event time relative to AI adoption, with $\tau = 0$ corresponding to 2021. We restrict attention to an event window $\tau \in \{-3, \ldots, +2\}$, where $\tau = -3$ is the omitted time period.

Summary Statistics. Table 1 reports summary statistics for the estimation sample at $\tau = -1$. The sample consists of approximately 95,000 worker-year observations at 900 unique firms. Average W-2 earnings are $108,000. Workers are 44 years old on average, 55 percent male, and 46 percent hold a college degree. Twelve percent of workers are classified as highly exposed to AI ($T_i = 1$).[14]

[13] We deflate earnings to 2024 using the CPI.

[14] The fact that 10% are not above the 90th percentile is simply due to the coarseness of occupation codes.

Table 1: Summary Statistics

Variable	Mean	SD	N
W-2 Earnings ($)	108,100	94,030	95,000
Age	44.0	11.7	95,000
Male	0.545	0.498	95,000
White	0.707	0.455	95,000
College Degree	0.455	0.498	95,000
Some College	0.315	0.465	95,000
High School	0.204	0.403	95,000
Household Size	1.571	0.495	95,000
$\Delta \ln w_2$	-0.005	0.467	95,000
AI Exposure Score	0.070	0.043	95,000
High Exposure (T_i)	0.120	0.325	95,000
Unique firms: 900			

Notes. Sample restricted to workers at firms reporting AI adoption in 2021–2022 (ABS). W-2 earnings in 2023 dollars. $T_i = \mathbf{1}(e_i > e^{90})$: above 90th pctile of AI exposure. Statistics at $\tau = -1$ (2020). N rounded.

Earnings Dynamics by Exposure Group. Before turning to the event-study estimates, we document raw earnings dynamics for high- and low-exposure workers. Figure 8 plots real W-2 earnings and residualized log W-2 earnings for treatment ($T_i = 1$) and control ($T_i = 0$) workers over the event window $\tau \in \{-3, \ldots, +2\}$.

Panel (A) shows that treated workers earn substantially more in levels, approximately $130k-145k versus $104k-113k for controls, reflecting the higher-paying occupations that define the treatment group. Both groups experience a dip in real earnings around $\tau = 0$ (2021), consistent with the broader post-COVID earnings disruption. Treatment-group earnings fall by more and recover faster. Panel (B) shows the same pattern in residualized log earnings, with the treatment–control gap roughly 0.35 log points throughout and treatment-group earnings rebounding more quickly post-covid. Our baseline specification include firm by time fixed effects and income-bin by time fixed effects to address differential post-COVID dynamics. Several of our more aggressive specifications include occupation by time fixed effects in order to completely remove COVID dynamics, but occupation by time fixed effects are only identified in the triple interaction specifications.

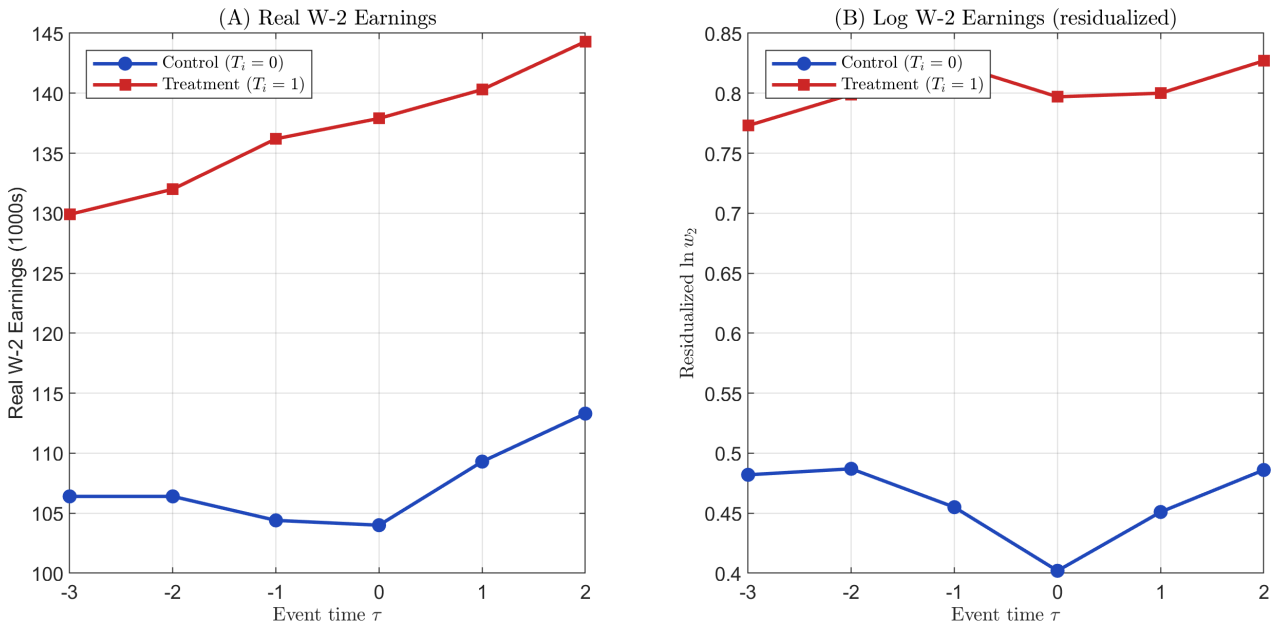


Figure 8: **Earnings dynamics by AI exposure group.**

Notes. Real W-2 earnings (Panel A, in $1000s) and residualized log W-2 earnings (Panel B) for treatment ($T_i = 1$, above 90th percentile AI exposure) and control ($T_i = 0$) workers. Treatment-group values sum the base period coefficient and the treatment interaction; no standard errors shown for the summed treatment coefficients. ABS survey-weighted. $\tau = 0$ corresponds to 2021.

Baseline Specification. Our primary specification exploits within-firm, within-year variation in exposure:

$$y_{it} = \alpha_i + \alpha_{jt} + \sum_{k=-2}^{2} \beta_k \cdot \mathbf{1}(\tau_{it} = k) \cdot T_i + \epsilon_{it}, \tag{1}$$

where α_i are worker fixed effects and α_{jt} are firm-by-year fixed effects. Identification comes from differential wage growth between high- and low-exposure workers within the same firm-year around adoption. The coefficients β_k trace out the event-time profile of high-exposure workers relative to low-exposure workers. We progressively augment this specification with income-bin-by-year fixed effects and a lagged dependent variable. To construct income bins, we rank individuals on their $\tau = -1$ income and then place them into deciles. We fill those deciles forward and backwards over the event window.

Triple-Difference Specification. Our second specification examines heterogeneity by firm-level employment responses (i.e., whether in the ABS they reported increasing or decreasing employment in response to AI). Let J_j^+ denote an indicator that firm j expands employment following AI adoption, and let J_j^- denote an indicator that firm j contracts employment. We

estimate:

$$y_{it} = \alpha_i + \alpha_{jt} + \sum_{k=-2}^{2} \beta_k \cdot \mathbf{1}(\tau_{it} = k) \cdot T_i + \sum_{k=-2}^{2} \gamma_k^+ \cdot \mathbf{1}(\tau_{it} = k) \cdot T_i \cdot J_j^+ + \sum_{k=-2}^{2} \gamma_k^- \cdot \mathbf{1}(\tau_{it} = k) \cdot T_i \cdot J_j^- + \epsilon_{it}. \tag{2}$$

The coefficients γ_k^+ and γ_k^- capture the additional dynamic effect of AI adoption on high-exposure workers at firms that expand or contract employment, respectively, relative to firms with stable employment. Because we include firm-by-year fixed effects, identification relies on within-firm-year comparisons across exposure groups. We further add occupation-by-year fixed effects to absorb any time-varying shocks common to workers in the same occupation nationwide. However, with occupation-by-year fixed effects, we lose identification of β_k.

Results. Figure 9 plots the baseline event-study coefficients $\hat{\beta}_k$ from Equation (1) with firm-by-year, worker, income-bin-by-year fixed effects, and a lagged dependent variable (Table 2, columns 4 and 5). Both the weighted and unweighted specifications show a positive earnings gradient that emerges after adoption ($\tau \geq 0$) and is statistically significant at $\tau = +2$. Pre-trends at $\tau = -2$ and $\tau = -1$ are close to zero, supporting the parallel-trends assumption. The shaded band marks the AI adoption window $\tau \in [0, 1]$.

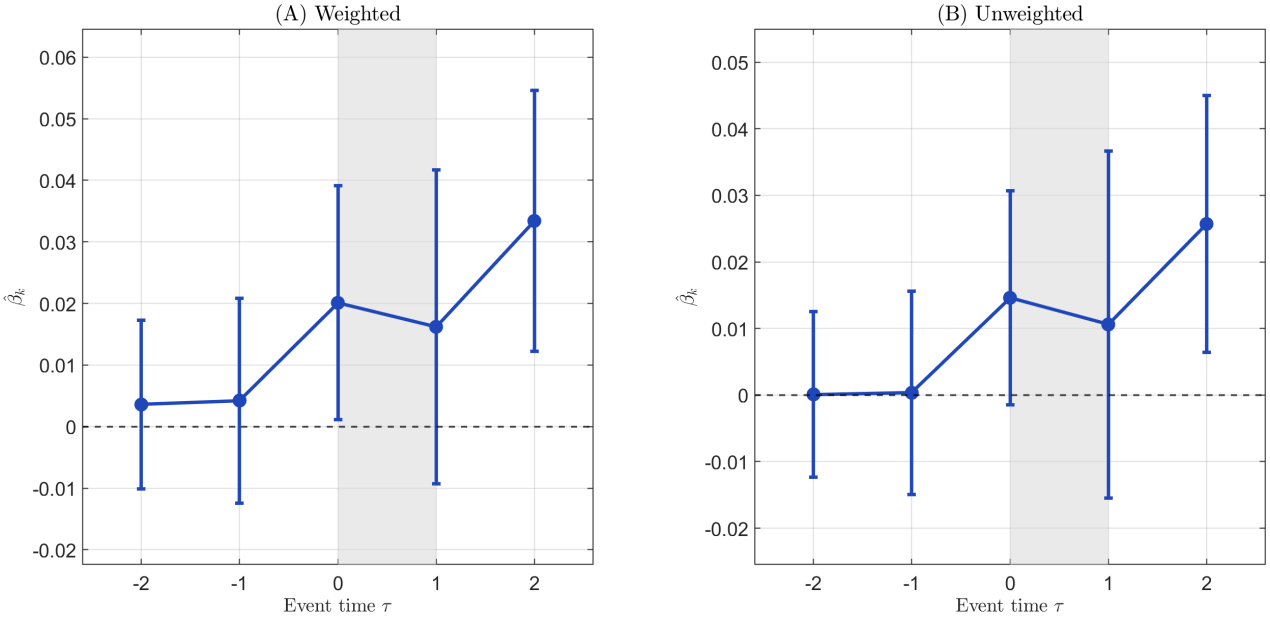


Figure 9: **Event study: $\Delta \ln w_2$ around AI adoption.**

Notes. Coefficients $\hat{\beta}_k$ from Equation (1) with firm×year, worker, and income-bin×year FE plus lagged outcome. Blue circles show point estimates; capped lines show ±1.96 SE (95% CI). Gray shaded band = $\tau \in [0, 1]$, the AI adoption window. Panel (A): ABS survey-weighted. Panel (B): unweighted. $\tau = 0$ corresponds to 2021.

Table 2 reports the full baseline results. Across all five specifications, the $\tau = +2$ coeffi-

cient is positive and statistically significant at the 1% level, ranging from 2.6 to 4.2 log points. These estimates imply that workers in high-exposure occupations experience meaningfully faster wage growth than their low-exposure coworkers within the same firm in the two years following AI adoption.

Table 2: Outcome: $\Delta \ln w_2$ — Baseline Specifications

	(1)	(2)	(3)	(4)	(5)
Outcome	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$
$\tau_{-2} \times T_i$	0.00530	0.00292	0.00616	0.00358	9.78e-05
	(0.00775)	(0.00795)	(0.00691)	(0.00700)	(0.00635)
$\tau_{-1} \times T_i$	0.0186*	0.00294	0.0209**	0.00423	0.000335
	(0.00974)	(0.00853)	(0.00987)	(0.00849)	(0.00779)
$\tau_0 \times T_i$	0.0100	0.0200**	0.0160*	0.0201**	0.0146*
	(0.00999)	(0.0102)	(0.00963)	(0.00969)	(0.00821)
$\tau_{+1} \times T_i$	0.0119	0.0130	0.00970	0.0162	0.0106
	(0.0121)	(0.0128)	(0.0126)	(0.0130)	(0.0133)
$\tau_{+2} \times T_i$	0.0423***	0.0306***	0.0419***	0.0334***	0.0257***
	(0.00956)	(0.0103)	(0.0102)	(0.0108)	(0.00984)
R^2	0.275	0.286	0.336	0.345	0.340
N	455K	455K	455K	455K	455K
Worker FE	Y	Y	Y	Y	Y
Firm × Year FE	Y	Y	Y	Y	Y
Inc. Bin × Year FE	N	Y	N	Y	Y
Occ. × Year FE	N	N	N	N	N
Lagged outcome	N	N	Y	Y	Y
ABS weights	Y	Y	Y	Y	N

Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. $T_i = \mathbf{1}(e_i > e^{90})$ is the high-AI-exposure indicator (above 90th percentile). $\tau = 0$ = 2021 (AI adoption year).

Table 3 presents the triple-difference results. The J_j^- interaction terms show that at firms contracting employment, high-exposure workers experience significantly *lower* wage growth. The $\tau_{+2} \times T_i \times J_j^-$ coefficient is negative and significant across all specifications, ranging from -5.7 to -7.3 log points. By contrast, the J_j^+ interaction terms are small and statistically insignificant throughout. These results suggest two findings: (1) that self-reported employment cuts translate into slower wage growth for exposed workers at those firms, and (2) significant heterogeneity, even within exposed workers. While our baseline results suggest that exposed workers gain on average from AI, these results suggest that many may lose when firms shed workers.

Figure 10 plots the $\hat{\gamma}_k^-$ coefficients for two specifications. Pre-trends are close to zero, and a pronounced negative effect emerges at $\tau = +2$ in both panels. This pattern indicates that the positive baseline effect is concentrated at firms that maintain or expand employment; at contracting firms, high-exposure workers face slower wage growth as well, although the terms

are not generally significant at conventional levels. These results cast some doubt on the validity of the self-reported employment expansion results (or at least their implications for the differential effect on exposed workers).

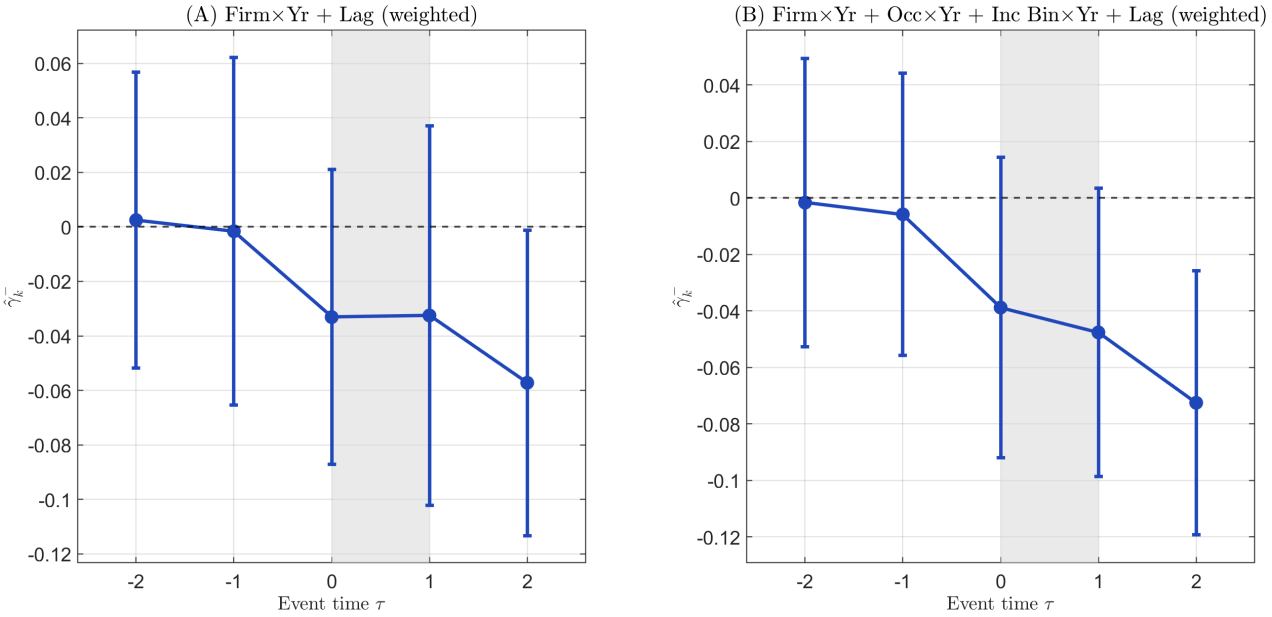


Figure 10: **Triple-difference J_j^- interaction: $\Delta \ln w_2$ around AI adoption.**

Notes. Coefficients $\hat{\gamma}_k^-$ from Equation (2). Blue circles show point estimates; capped lines show ±1.96 SE (95% CI). Gray shaded band = $\tau \in [0, 1]$, the AI adoption window. Panel (A): firm×year + worker FE + lagged outcome (weighted). Panel (B): adds occupation×year and income-bin×year FE (weighted). $\tau = 0$ corresponds to 2021.

Employment Change and Low Earnings. We next examine whether AI adoption affects worker mobility and the incidence of low earnings. Figure 11 plots event-study coefficients from Equation (1) for an employment-change indicator (equal to one if the worker switches firms) and a low-earnings indicator (equal to one if W-2 earnings fall below a low-earnings threshold). The employment-change coefficients are small and statistically insignificant, suggesting limited differential mobility on average. The low-earnings coefficients are also close to zero throughout the event window. Full regression results are reported in Appendix Tables C1 and C2.

2.3 R&D Occupation Exposure

We complement the AI exposure analysis with an alternative treatment based on occupation-level exposure to in-house research and development (R&D). Using the 2022 ABS linked to our W2 panel, we identify firms that report conducting in-house R&D and compute the share of each occupation's employment at R&D-performing firms. This yields an occupation-level

Table 3: Outcome: $\Delta \ln w_2$ — Triple-Difference Specifications

	(1)	(2)	(3)	(4)	(5)
Outcome	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$
$\tau_{-2} \times T_i$	0.00671	0.00362			
	(0.00793)	(0.00895)			
$\tau_{-1} \times T_i$	0.0200*	0.00154			
	(0.0112)	(0.00992)			
$\tau_0 \times T_i$	0.0218*	0.0256**			
	(0.0114)	(0.0120)			
$\tau_{+1} \times T_i$	0.0131	0.0157			
	(0.0148)	(0.0149)			
$\tau_{+2} \times T_i$	0.0489***	0.0375***			
	(0.0112)	(0.0114)			
$\tau_{-2} \times T_i \times J_j^+$	-0.00393	-0.00401	-0.00591	-0.00620	-0.00587
	(0.0165)	(0.0178)	(0.0165)	(0.0164)	(0.0160)
$\tau_{-1} \times T_i \times J_j^+$	0.0100	0.0144	0.00732	0.0103	0.0130
	(0.0281)	(0.0259)	(0.0266)	(0.0253)	(0.0255)
$\tau_0 \times T_i \times J_j^+$	-0.0286	-0.0299	-0.0303	-0.0344*	-0.0275
	(0.0196)	(0.0196)	(0.0192)	(0.0193)	(0.0179)
$\tau_{+1} \times T_i \times J_j^+$	-0.0112	-0.0111	-0.00748	-0.0117	-0.0124
	(0.0280)	(0.0266)	(0.0260)	(0.0258)	(0.0255)
$\tau_{+2} \times T_i \times J_j^+$	-0.0274	-0.0264	-0.0325	-0.0324	-0.0301
	(0.0295)	(0.0265)	(0.0272)	(0.0270)	(0.0268)
$\tau_{-2} \times T_i \times J_j^-$	0.00248	-9.15e-05	-0.00261	-0.00166	0.0129
	(0.0277)	(0.0355)	(0.0259)	(0.0260)	(0.0234)
$\tau_{-1} \times T_i \times J_j^-$	-0.00161	-0.00415	-0.00326	-0.00582	-0.0132
	(0.0325)	(0.0287)	(0.0281)	(0.0255)	(0.0244)
$\tau_0 \times T_i \times J_j^-$	-0.0331	-0.0267	-0.0397	-0.0388	-0.0401
	(0.0276)	(0.0266)	(0.0268)	(0.0271)	(0.0252)
$\tau_{+1} \times T_i \times J_j^-$	-0.0325	-0.0210	-0.0480*	-0.0476*	-0.0312
	(0.0355)	(0.0323)	(0.0267)	(0.0260)	(0.0233)
$\tau_{+2} \times T_i \times J_j^-$	-0.0572**	-0.0579**	-0.0693***	-0.0725***	-0.0627**
	(0.0286)	(0.0266)	(0.0237)	(0.0238)	(0.0252)
R^2	0.336	0.286	0.342	0.352	0.347
N	455K	455K	455K	455K	455K
Worker FE	Y	Y	Y	Y	Y
Firm×Year FE	Y	Y	Y	Y	Y
Inc. Bin×Year FE	N	Y	N	Y	Y
Occ.×Year FE	N	N	Y	Y	Y
Lagged outcome	Y	N	Y	Y	Y
ABS weights	Y	Y	Y	Y	N

Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. $T_i = \mathbf{1}(e_i > e^{90})$ is the high-AI-exposure indicator (above 90th percentile). $\tau = 0$ = 2021 (AI adoption year).

R&D exposure score analogous to the AI exposure measure. High R&D exposure occupations include aerospace engineers, electrical engineers, mechanical engineers, and atmospheric and space scientists.

We define the R&D exposure treatment indicator

$$T_i^{RD} = \mathbf{1}(e_i^{RD} > e_{RD}^{90}),$$

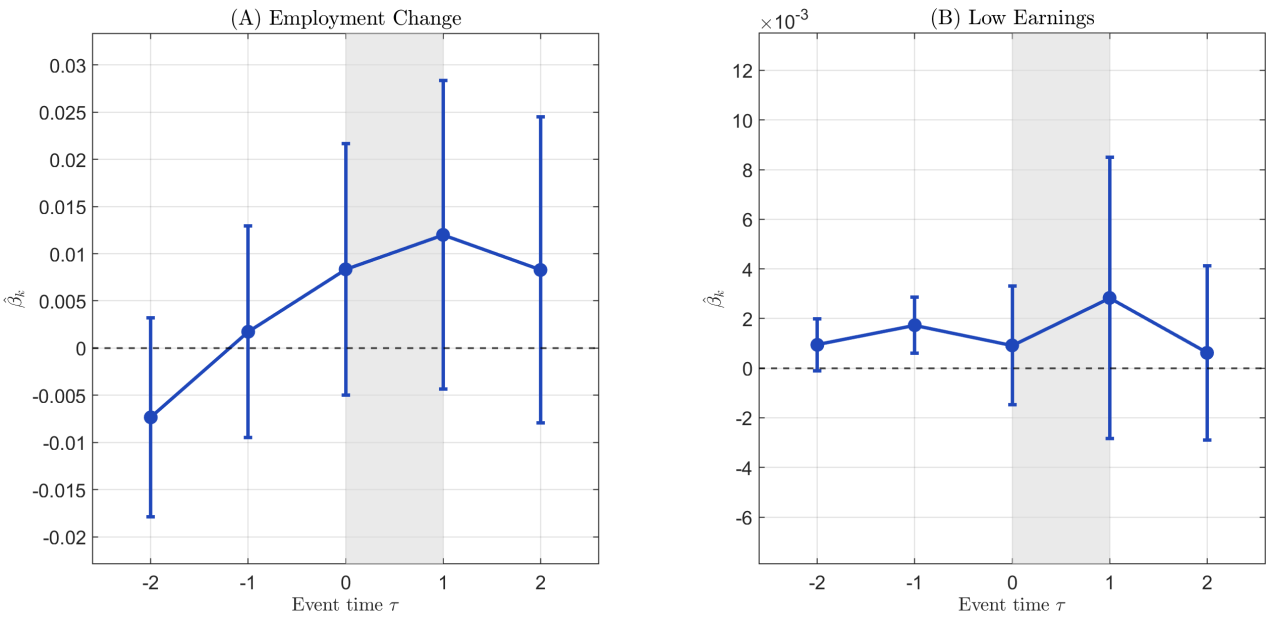


Figure 11: **Event studies: employment change and low earnings.**

Notes. Coefficients $\hat{\beta}_k$ from Equation (1) with firm×year, worker, and income-bin×year FE plus lagged outcome (unweighted). Blue circles show point estimates; capped lines show ±1.96 SE (95% CI). Gray shaded band = $\tau \in [0, 1]$. Panel (A): employment-change indicator. Panel (B): low-earnings indicator. $\tau = 0$ corresponds to 2021.

where e_i^{RD} is the occupation-level R&D exposure score and e_{RD}^{90} is its 90th percentile. We then estimate the same event-study specifications as in Section 2.2, using the 2023 ABS to identify AI adoption timing and the R&D exposure measure as the treatment. Importantly, all specifications control for the AI exposure measure T_i^{AI} to isolate the R&D channel from any direct AI exposure effect.

R&D Baseline Results. Figure 12 plots the baseline event-study coefficients for R&D exposure, controlling for AI exposure. The pattern is qualitatively similar to the AI exposure results. Those with higher R&D exposure have slightly faster wage growth. Across specifications, the $\tau = +2$ coefficient is significant and ranges from 1.9 to 2.2 log points after controlling for AI exposure. Full results are reported in Appendix Table C3.

R&D Triple-Difference Results. Our final exercise is to use the R&D exposure in the triple difference specifications. Figure 13 plots the $\hat{\gamma}_k^-$ coefficients for R&D exposure, again controlling for AI exposure. As with AI exposure, the J_j^- interaction terms are negative and significant at $\tau = +2$, ranging from -3.8 to -5.4 log points (Appendix Table C4). These results suggest that even among those in high R&D exposed occupation, AI adoption has heterogeneous effects. Those who have high R&D exposure at firms that report shedding employment in response to

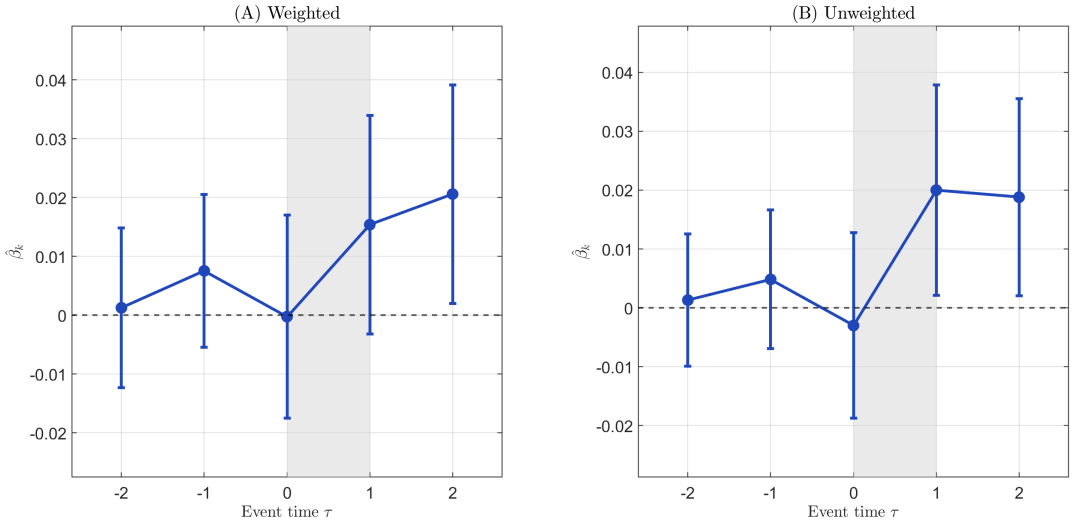


Figure 12: **Event study: $\Delta \ln w_2$ around AI adoption (R&D exposure, controlling for AI exposure).**

Notes. Coefficients $\hat{\beta}_k$ for R&D occupation exposure (T_i^{RD}) from the baseline DID specification, controlling for AI exposure (T_i^{AI}). Firm×year, worker, and income-bin×year FE plus lagged outcome. Panel (A): ABS survey-weighted. Panel (B): unweighted. $\tau = 0$ corresponds to 2021.

AI adoption also endure wage losses.

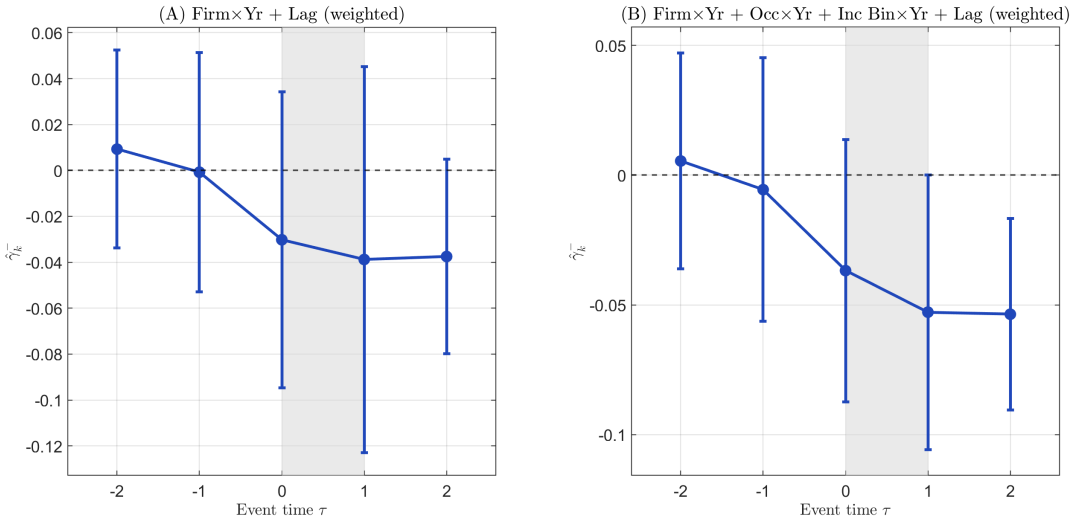


Figure 13: **Triple-difference J_j^- interaction: $\Delta \ln w_2$ (R&D exposure, controlling for AI exposure).**

Notes. Coefficients $\hat{\gamma}_k^-$ for R&D occupation exposure (T_i^{RD}) from the triple-difference specification, controlling for AI exposure (T_i^{AI}). Panel (A): firm×year + worker FE + lagged outcome (weighted). Panel (B): adds occupation×year and income-bin×year FE (weighted). $\tau = 0$ corresponds to 2021.

Caveats. The caveats in this section are more substantial than in the earlier ABS analysis. First, AI adoption as measured in the 2023 ABS may not fully capture the reach or effects of modern GPTs. More importantly, we observe only a single cohort of adopters, concentrated in 2021–2022, which overlaps heavily with the COVID period. Our use of firm-by-time and occupation-by-time fixed effects helps mitigate some of these concerns, but additional adoption cohorts will be necessary to place greater confidence in the results. In addition, our analysis ends in 2023 because of the timing with which W-2 data are received and processed at the Census. As a result, estimates at the end of the sample may partly reflect other contemporaneous trends that a longer stacked panel could better absorb. Finally, we observe some pre-trends in job switching, which may reflect anticipation effects or other differential trends not fully captured by our fixed effects. We view these estimates as preliminary and expect to refine them as additional data become available.

2.4 AI Adoption Timelines

Our final empirical exercise is to attempt to measure "learning time" for AI adoption. These timelines measure how long it takes for resumes of workers in a given sector to start reporting sufficient usage of AI (above a certain threshold that we discuss and measure below) directly related to a given occupation. This yields occupation-sector-level AI adoption timelines. We train this algorithm by combining resume-based measures of AI use with sector-level adoption targets from the 2023 Business Trends and Outlook Survey (BTOS). We first map AI-related text in resumes into task-level usage and then aggregate those task-level exposure rates up to occupation-sector-level exposure rates. Sector-specific adoption thresholds are estimated so that the resume-based firm adoption rates line up with BTOS sector adoption rates. We must train adoption thresholds at the sector level, since we do not have occupation-sector based adoption rates in administrative data.

Table 4 highlights the scale of the data used to estimate adoption timelines. The sample contains more than 700,000 AI resumes, over 1 million AI applications, and more than 85,000 firms. These large samples are important because the empirical strategy requires enough within-firm and within-sector variation to construct stable task-level usage measures and to train sector-specific thresholds that match BTOS adoption rates.

Resume-based AI usage at the task level. We use Revelio Labs resume data from 2000–2023 and follow the same text-processing pipeline used elsewhere in the paper. AI-related text is extracted from resumes, cleaned, and embedded in a vector space. We then compare that AI text to O*NET task descriptions. As Table 4 shows, the resume sample is large enough to support a

Table 4: **Summary statistics for the adoption-timing sample**

Statistic	Value
AI resumes	708,329
AI applications	1,056,747
Firms	85,472
Immediate adoption rate from CPS	0.895
Mean time to adoption from CPS (including immediate adoption)	0.20
Mean time to adoption below median exposure (including immediate adoption)	0.26
Mean time to adoption above median exposure (including immediate adoption)	0.14
Mean time to adoption from CPS (excluding immediate adoption)	1.89
Mean time to adoption below median exposure (excluding immediate adoption)	2.11
Mean time to adoption above median exposure (excluding immediate adoption)	1.65

Notes. The Revelio sample spans 2010–2023. CPS timing moments are computed in the Revelio-merged CPS sample using employment weights. The table reports both moments including immediate adoption and moments excluding immediate adoption. Including immediate adoption assigns zero waiting time to immediate adopters, while excluding immediate adoption conditions on strictly positive time to adoption.

high-dimensional mapping from AI text into detailed tasks and occupations. Let j index tasks, o occupations, s sectors (or industries), i AI applications extracted from resumes, and $N_{s,t}$ the number of AI applications observed in sector s in year t. For a fixed similarity cutoff x, we define

$$AI_{j,s,t}(x) = \frac{1}{N_{s,t}} \sum_{i=1}^{N_{s,t}} \mathbb{1}(\text{similarity}(\text{AI text } i,\ \text{task } j) \geq x)\,. \tag{3}$$

This object measures the share of AI applications in sector s and time t that map into task j.

Aggregation from tasks to occupation-by-sector measures. We first combine task-level usage into an occupation-sector-level measure using O*NET task-importance weights:

$$\mu_{o,s,t}(x) = \sum_{j} \omega_{o,j} \cdot AI_{j,s,t}(x), \tag{4}$$

where $\omega_{o,j}$ is the importance weight of task j within occupation o.

Sector-specific adoption thresholds trained on BTOS. Let τ_s denote the adoption threshold for sector s. We build a dummy $D_{o,s,t}$ that classifies whether an occupation-sector is an AI adopter in year t when its usage index exceeds the threshold for its sector:

$$D_{o,s,t} = \mathbb{1}(\mu_{o,s,t}(x) \geq \tau_s)\,. \tag{5}$$

We then train τ_s to match the BTOS adoption rates by sector. We then aggregate these indicators to the sector level by taking an unweighted mean across occupations (where N_s is the number of occupations in sector s):

$$\widehat{M}_s(\boldsymbol{\tau}) = \frac{1}{N_s} \sum_o D_{o,s,t}. \tag{6}$$

Let M_s^{BTOS} denote the 2023 BTOS share of firms in sector s that report AI use. We estimate the threshold vector by minimizing the distance between the resume-implied adoption rate and the BTOS target:

$$\boldsymbol{\tau} = \arg\min_{\boldsymbol{\tau}} \sum_{s=1}^{S} \left(\widehat{M}_s(\boldsymbol{\tau}) - M_s^{BTOS} \right)^2. \tag{7}$$

Model fit. Figure 14 compares the fitted sector adoption rates implied by the estimated threshold vector to the 2023 BTOS sector targets. The fit is close across the major 2-digit sectors in the sample, indicating that the trained sector thresholds reproduce the cross-sector distribution of AI adoption in the BTOS well. This is the key validation exercise for the adoption-timing measure, since the firm-year timelines are built by applying these fitted thresholds back to the full resume panel.

Adoption timelines. Once the sector-specific thresholds are estimated, we classify each occupation-sector-year in the full panel as adopted or non-adopted. We then measure the time from when the occupation-sector first hires an AI worker until it crosses the threshold.[15]

To summarize the timing moments in a worker-based sample, we merge the Revelio-based adoption measures to the CPS and compute employment-weighted statistics on time to automation. Table 4 reports both the moments including immediate adoption and the moments excluding immediate adoption. Including immediate adoption, the mean time to adoption is 0.20 years, reflecting the fact that roughly 89.5 percent of adopters are recorded with no measured delay. Excluding immediate adoption, the mean rises to 1.89 years. The same comparison also reveals a stable exposure gradient: below-median exposure jobs take 0.26 years to automate in the full sample and 2.11 years conditional on positive delay, while above-median exposure jobs take 0.14 years in the full sample and 1.65 years conditional on positive delay.

Figure 15 puts the timing evidence including immediate adoption and excluding immediate adoption on equal footing. Panel (a) shows the full adopter distribution: the overall employment-weighted mean time to adoption is 0.20 years, with means of 0.26 years below the exposure median and 0.14 years above it. Panel (b) isolates the delayed adopters by remov-

[15] If the occupation-sector already has AI workers in 2010, we code the first year as 2010, since the data are truncated at that date.

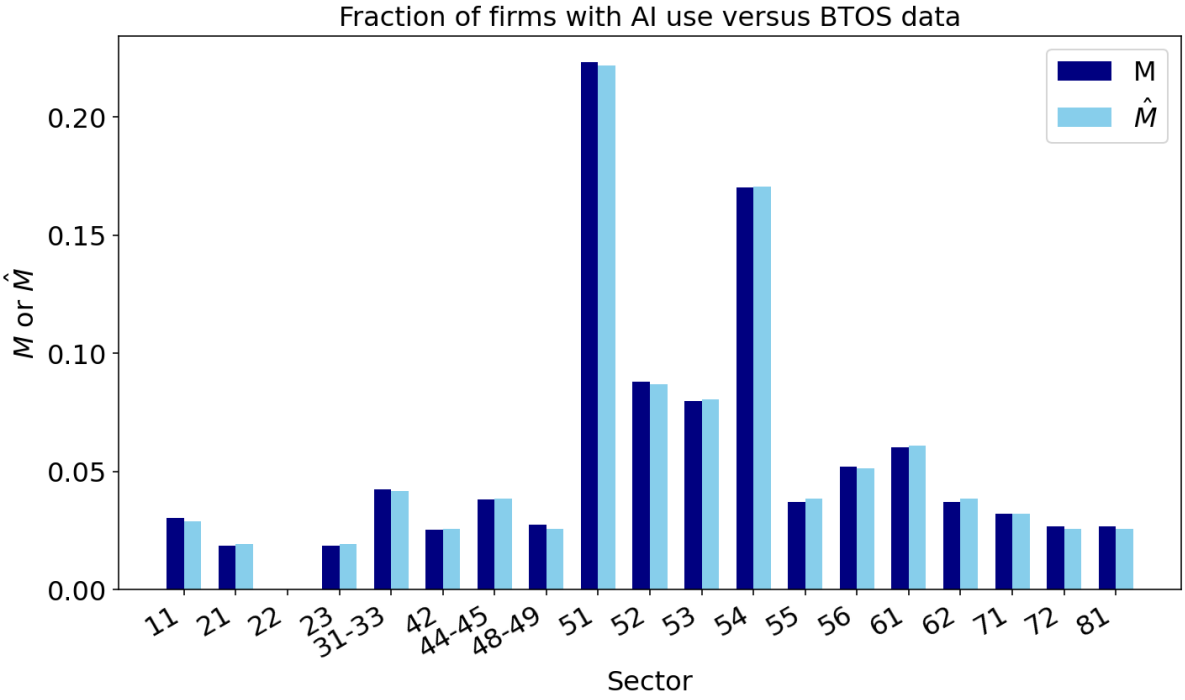


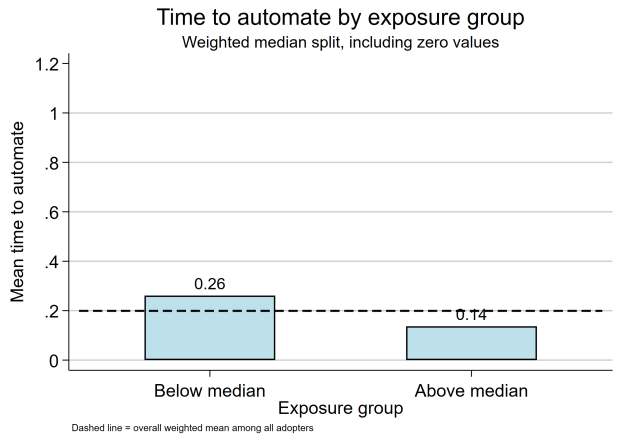
Figure 14: **Sector AI adoption rates in the fitted model and the 2023 BTOS**

Notes. Dark bars plot the target sector adoption rates from the 2023 BTOS, and light bars plot the corresponding fitted adoption rates implied by the estimated sector-specific thresholds $\{\tau_s\}_s$. Sectors are 2-digit NAICS groupings. Source for target moments: U.S. Census Bureau, 2023 Business Trends and Outlook Survey, sector AI-use tabulations. The close correspondence indicates that the trained thresholds recover the cross-sector pattern of AI use in the BTOS.

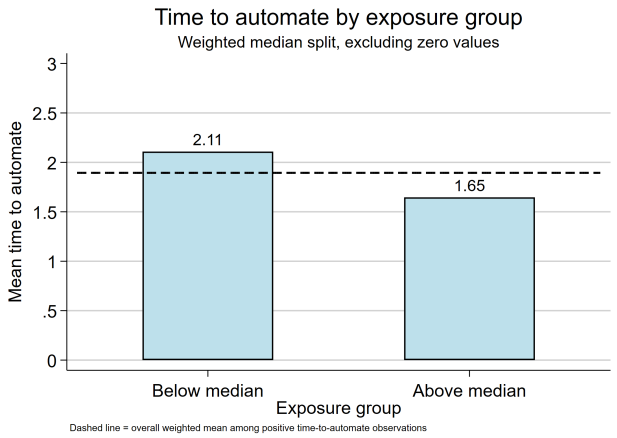
ing the immediate-adoption mass, which raises the overall mean to 1.89 years and yields means of 2.11 years below the median and 1.65 years above it. Read together, the two panels show that immediate adoption is common, but they also show that the faster adoption of higher-exposure jobs is present both in the full sample and in the positive-delay subsample. The corresponding numeric moments are reported in Table 4.

The main advantage of this approach is that it delivers a yearly firm-level adoption indicator that is disciplined by external sector evidence while remaining grounded in occupation-level task content. The BTOS targets anchor the level differences in adoption across industries, and the resume panel supplies the within-sector timing variation needed for event-study analysis. In that sense, the procedure combines survey discipline with high-frequency firm-level variation: BTOS determines where adoption should be prevalent, and the resume data determine when individual firms cross the estimated threshold.

Caveats. The resume adoption timelines are noisy for a number of reasons, including lags in when workers report AI-related developments. Moreover, the most recent GPT-based tools



(a) Including immediate adoption



(b) Excluding immediate adoption

Figure 15: **Time to automate by exposure group**

Notes. Panel (a) reports the employment-weighted mean time to automate in the full adopter sample, including immediate adoption, split into below-median and above-median exposure groups in the Revelio-merged CPS sample. Panel (b) reports the same moments after excluding immediate adoption. In both panels, the dashed line marks the corresponding overall weighted mean and the vertical axis begins at zero.

may not appear immediately on resumes. We therefore calibrate and consider a wide range of adoption times, ranging from near-immediate adoption to slower timelines of one to two years.

3 Theory

Our model economy is comprised of heterogeneous, risk-averse workers who face incomplete markets against job loss. We assume perpetual youth OLG. Workers are endowed with a vector of non-automatable human capital and automatable human capital; their skills evolve as they accumulate both types of human capital on the job and when unemployed they can invest on their non-automatable skills. Workers direct their search for jobs, and jobs are characterized by skill weights chosen ex-post by the firm (i.e., how intensively the job uses AI automatable human capital and non-AI automatable human capital). We allow firms to adjust skill weights as they automate tasks (i.e., the nature of jobs changes as AI is rolled out). Firms learn and accumulate data on the tasks their workers complete and choose to automate a job by comparing the stock of data accumulated with the automatable human capital level of the worker. Automation is partitioned into two groups: partial and full automation. Both incur a piece-rate cost on production, with the first keeping the worker executing non-automatable tasks, while

the second completely substitutes him. Therefore, there's no learning after full automation, as a worker is required for that. Finally, firms can purchase data at entry to expedite the learning or skip it altogether, however, under the requirement of at least a period of work before any automation.

3.1 Firms

Workers direct their search for vacancies as in Menzio and Shi (2011) and are differentiated by their skill vector $\mathbf{h} = (h_a, h_n)$, where h_a represents tasks that can eventually be conducted by AI with sufficient data, while h_n represents skills that cannot be replaced by AI. Firms choose which production piece-rate (ϕ) and skill $\mathbf{h}$ submarket to enter. We assume firms commit to ϕ. This assumption, instead of a commitment to a fixed wage, allows for on the job wage growth through human capital and/or technological changes.

Firms match with workers according to a CRS matching function. Following standard conventions, $q(\cdot)$ is the contact rate of firms, $p(\cdot)$ is the contact rate of workers, $u(\cdot)$ is the submarket unemployment rate, $v(\cdot)$ is the submarket vacancy rate, and $\theta(\cdot)$ is the submarket tightness:

$$q\left(\theta(\mathbf{h},\phi)\right) = \frac{M\left(u(\mathbf{h},\phi), v(\mathbf{h},\phi)\right)}{v(\mathbf{h},\phi)}\ ,\ p\left(\theta(\mathbf{h},\phi)\right) = \frac{M\left(u(\mathbf{h},\phi), v(\mathbf{h},\phi)\right)}{u(\mathbf{h},\phi)},\ \theta(\mathbf{h},\phi) = \frac{v(\mathbf{h},\phi)}{u(\mathbf{h},\phi)}. \quad (8)$$

Firms require data to replace workers with AI. When firms match with workers, they make an initial data purchase d subject to a convex cost $c(d)$. They then accumulate data on the worker's automatable human capital, h_a, at rate Δ, with probability p_d. If they exogenously separate, at rate δ, or the worker dies, at rate λ, they sell their data for the market price of $x(d)$. Thus our model nests pure experience models ($c(d) = \infty,\ \forall d$) in which all data must be accumulated through firm-worker experience, and pure inspection models ($\Delta = 0$) in which all data is bought up-front independent of firm-worker experience. The existing literature on AI and the labor market assumes that AI is an inspection good. While some component of AI can be pre-trained, a significant portion cannot. Banks, hospitals, universities, and firms have highly sensitive client and employee information that can only be used to train AI models in-house. Unmatched firms are allowed produce with fully automated production, but need a setup period where they do not produce. We believe our model of data accumulation provides an important alternative to existing theories.

Once matched, firms allocate workers' time across automatable and non-automatable tasks. The subtitutability across the two types of tasks is governed by ρ, where $\rho > 1$, $\rho < 1$ and $\rho = 1$ imply substitutes, complements and linearity, respectively. Workers supply 1 unit of labor inelastically. Let s denote the share of time spent on non-automatable tasks.

The cost of automating or partially automating is assumed to be a per-period operating cost κ.[16] Let $\{J, J^P, J^F\}$ denote the value functions of a firm that hasn't automated, that has partially automated and that has fully automated, respectively. Therefore, the value of a firm that hasn't automated is given by:

$$\begin{aligned} J(\mathbf{h}, \phi, d) = &\max_{s\in[0,1]} \left((sh_n)^\rho + ((1-s)h_a)^\rho\right)^{1/\rho} \cdot (1-\phi) \\ &+ \beta E_{\mathbf{h}', d'|\mathbf{h}, d}\left[(1-\lambda)(1-\delta)\widehat{J}(\mathbf{h}', \phi, d') + (\delta(1-\lambda)+\lambda)x(d')\right] \end{aligned} \tag{9}$$

where each period, assuming the match remains, firms choose between partial and full automation,

$$\widehat{J}(\mathbf{h}, \phi, d) = \max\left\{J(\mathbf{h}, \phi, d),\ J^P(\mathbf{h}, \phi, d),\ J^F(d)\right\} \tag{10}$$

Similarly, the value of a firm that has partially automated is:

$$\begin{aligned} J^P(\mathbf{h}, \phi, d) = &\left(h_n^\rho + d^\rho\right)^{1/\rho} (1-\kappa)(1-\phi) \\ &+ \beta E_{\mathbf{h}', d'|\mathbf{h}, d}\left[(1-\lambda)(1-\delta)\widehat{J}(\mathbf{h}', \phi, d') + (\delta(1-\lambda)+\lambda)x(d')\right] \end{aligned} \tag{11}$$

where the piece-rate falls on the firm's revenue post automation cost. The stock of data for these two types of firms evolve according to

$$d' = \begin{cases} d + \Delta, \text{ with probability } p_d \\ d, \text{ with probability } 1 - p_d \end{cases} \tag{12}$$

Notice that as the automatable tasks are not perfored by the worker, the weight assigned to the non-automatable tasks is 1. Finally, the value of a firm that has fully automated is:

$$J^F(d) = d(1-\kappa) + \beta[(1-\delta)J^F(d) + \delta x(d)] \Rightarrow J^F(d) = \frac{d(1-\kappa) + \beta\delta x(d)}{1-\beta(1-\delta)} \tag{13}$$

Therefore, the firm's automation and task policies f^P, f^F, s^* are given by

$$f^F(\mathbf{h}, \phi, d) = \mathbb{1}\left(\widehat{J}(\mathbf{h}, \phi, d) = J^F(\mathbf{h}, \phi, d)\right) \tag{14}$$

$$f^P(\mathbf{h}, \phi, d) = \mathbb{1}\left(\widehat{J}(\mathbf{h}, \phi, d) = J^P(\mathbf{h}, \phi, d)\right) \tag{15}$$

[16] Note that fixed costs are just as easy to incorporate, but it appears that services are being rented at low fixed costs.

$$\widehat{s}(\mathbf{h}) = arg\max_{s\in[0,1]} \left((sh_n)^\rho + ((1-s)h_a)^\rho\right)^{1/\rho} \tag{16}$$

At entry, firms face a submarket specific cost $\psi(\mathbf{h})$ of posting a vacancy to have access to the matching mechanism, which requires at least one period of non-automatable work, for matched firms. After the matching process is done, firms can purchase data up to the automatable human capital level of the submarket, h_a, to expedite the automation process (if matched) or to boost their fully automated production (if unmatched). Free-entry condition states that in every submarket $(\mathbf{h}, w)$, it must hold that:

$$\psi(\mathbf{h}) \geq q(\theta(\mathbf{h},\phi)) \left(\max_{d\leq h_a} J(\mathbf{h},\phi,d) - c(d)\right) + (1 - q(\theta(\mathbf{h},\phi))) \left(\max_{d\leq h_a} \beta J^F(d) - c(d)\right) \tag{17}$$

with equality for submarkets with vacancy posting, $v(\mathbf{h},\phi) > 0$. We denote the initial data investment policy of firms as $d(\mathbf{h},\phi)$, for matched firms, and $d^F(\mathbf{h})$, for unmatched. We can rewrite this condition as

$$\widetilde{\psi}(\mathbf{h}) \geq q(\theta(\mathbf{h},\phi)) \left(J(\mathbf{h},\phi,d(\mathbf{h},\phi)) - c(d(\mathbf{h},\phi)) - \left[\beta J^F(d^F(\mathbf{h})) - c(d^F(\mathbf{h}))\right]\right) \tag{18}$$

where $\widetilde{\psi}(\mathbf{h}) = \psi(\mathbf{h}) - \left[\beta J^F(d^F(\mathbf{h})) - c(d^F(\mathbf{h}))\right]$. Therefore, we need $\widetilde{\psi}(\mathbf{h}) \geq 0$. Otherwise, there would exist at least one submarket where the value of a fully automated dominates, and entering this submarket and choosing $v(\mathbf{h},\phi) = 0$ would be strictly preferred to participating in any matching process.

3.2 Workers

Workers are infinitely-lived. Each worker is endowed with two types of human capital – automatable, h_a, and non-automatable, h_n – summarized in a vector $\mathbf{h} = (h_n, h_a)$. Assume these are drawn from a joint distribution $G(\mathbf{h})$ at birth. Unemployed workers receive home production $z(\mathbf{h})$ and search for jobs indexed by its piece-rate and human capital. Workers have concave and strictly increasing utility $u(\cdot)$ over consumption.

We assume that each worker's skills evolve stochastically on-the-job $\mathbf{h}' \sim F(\mathbf{h}'|\mathbf{h})$ (e.g. Ljungqvist and Sargent (1998) and Lise and Postel-Vinay (2020)), but during unemployment workers can invest, R, in h_n at a unit cost of χ. Our baseline specification assumes that task intensity s does not enter skill evolution, but our framework is tractable enough to incorporate such features.

A worker can be unemployed, U, or working, W. The Bellman equation of an unemployed

worker is given by,

$$\begin{aligned} U(\mathbf{h}) = & \max_{R \geq 0} u(z(\mathbf{h}) - R \cdot \chi) \\ & + \beta(1-\lambda) E_{\mathbf{h}'|\mathbf{h},R,U} \left[\max_{\phi} p(\theta(\mathbf{h}', \phi)) W(\mathbf{h}', \phi, d(\mathbf{h}', \phi)) + \left(1 - p(\theta(\mathbf{h}', \phi))\right) U(\mathbf{h}') \right] \end{aligned} \tag{19}$$

where the initial data investment policy of firms $d(\mathbf{h}', \phi)$ and the market tightness $\theta(\mathbf{h}', \phi)$ are taken as given and $\mathbf{h}' = (h_n + R, h_a')$. Furthermore, denote the human capital investment policy $\widehat{R}(\mathbf{h})$. The Bellman equation for an employed worker is given by

$$\begin{aligned} W(\mathbf{h}, \phi, d) = & \, u\left(w(\mathbf{h}, \phi, d)\right) + \beta(1-\lambda) E_{\mathbf{h}', d'|\mathbf{h}, d, W} \Big[(1-\delta)(1 - f^F(\mathbf{h}', \phi, d')) \cdot W(\mathbf{h}', \phi, d') \\ & + \left\{ \delta + (1-\delta) f^F(\mathbf{h}', \phi, d') \right\} \cdot U(\mathbf{h}') \Big] \end{aligned} \tag{20}$$

where workers take as given the firm's task and automation policies, d' evolves according to (12) and the worker's labor income is a piece-rate of non-automated or partially automated firm's revenue,

$$w(\mathbf{h}, \phi, d) = \phi \cdot \begin{cases} \left((\widehat{s}(\mathbf{h}) h_n)^\rho + ((1 - \widehat{s}(\mathbf{h})) h_a)^\rho \right)^{1/\rho}, & \text{if } f^F(\mathbf{h}, \phi, d) + f^P(\mathbf{h}, \phi, d) = 0 \\ \left(h_n^\rho + h_a^\rho \right)^{1/\rho}, & \text{if } f^P(\mathbf{h}, \phi, d) = 1 \end{cases} \tag{21}$$

Workers are hand-to-mouth and are thus unable to insure against job loss, which occurs with probability δ, or full automation, which occurs with probability $f^F(\mathbf{h}, \phi, d)$. This gives rise to a role for government insurance to complete markets.

3.3 Competitive Equilibrium

A Competitive Equilibrium consists of firm policy functions $\{f^P(\mathbf{h}, \phi, d), f^F(\mathbf{h}, \phi, d), \widehat{s}(\mathbf{h}), d(\mathbf{h}, \phi)\}$, a search and human capital investment policy functions for workers $\{\widehat{\phi}(\mathbf{h}), \widehat{R}(\mathbf{h})\}$, and a market tightness function, $\{\theta(\mathbf{h}, \phi)\}$ such that:

1. Given $\{\theta(\mathbf{h}, \phi)\}$ and the firm policies, $\{\widehat{w}(\mathbf{h}), \widehat{R}(\mathbf{h})\}$ are the optimal search and human capital investment decision of workers;

2. Given $\{\theta(\mathbf{h}, \phi)\}$, $\{f^P(\mathbf{h}, \phi, d), f^F(\mathbf{h}, \phi, d), \widehat{s}(\mathbf{h}), d(\mathbf{h}, \phi)\}$ are the optimal automation, task assignment and data investment choices of firms.

3. Market tightness is such that free-entry condition (equation (17)) holds.

3.4 Characterization

In this section we present some theoretical results of the model, to make the key mechanisms clearer. We assume that $\mathbf{h} \in \mathbb{H}^2 \subseteq \mathbb{R}^2_{++}$, $\phi \in \mathbb{W} \subseteq \mathbb{R}_+$, $d \in \mathbb{H}$, where $\mathbb{H}, \mathbb{W}$ are compact sets. Moreover, let M be a constant returns to scale (CRS) matching function. All the proofs are in the appendix.

The first theoretical result sheds lights into the firm's task policy. **Lemma 1** states that if the relation between the automatable and non-automatable skills of the worker is of complementarity in production, then firms will optimally allocate the worker's time for the use of both skills, giving more weight to the most productive. At the same time, if the relation is of substitutability, then only the most productive skill will be used. Therefore, ρ is a parameter of utmost importance for the understanding of the firm's automation policy and which workers are more exposed to being replaced.

Lemma 1. *Consider a non-automated firm. If $\rho < 1$ (complements), then both automatable and non-automatable tasks will be used in production. If $\rho \geq 1$ (substitutes), then only the most productive task will be used in production* [17]*. Precisely, the optimal non-automatable skill weight is given by*

$$s^* = \begin{cases} h_n^{\frac{\rho}{1-\rho}} / (h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}}), & \text{if } \rho < 1 \\ 1, & \text{if } h_n > h_a \text{ and } \rho \geq 1 \\ 0, & \text{if } h_n < h_a \text{ and } \rho \geq 1 \\ \{0,1\}, & \text{if } h_n = h_a \text{ and } \rho > 1 \\ [0,1], & \text{if } h_n = h_a \text{ and } \rho = 1 \end{cases} \tag{22}$$

with production given by $(h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}}$ for $\rho < 1$ and $\max\{h_n, h_a\}$ for $\rho \geq 1$.

An equilibrium is said to have a block recursive property if the system of equations characterizing the equilibrium can be partitioned into an ordered sequence of blocks such that the variables in each block are uniquely determined given the solutions of earlier blocks, and do not depend on variables in later blocks. As our model is endowed with a CRS matching function, workers direct their submarket search, and our equilibrium concept includes a free-entry condition, then as **Proposition 1** states, this property is satisfied in our framework. More specifically, the first block is composed of the equations (9), (11), (13), that determines $\{J, J^P, J^F\}$, the

[17] Except in the linear case when $h_n = h_a$, where any weight is optimal.

second block is composed of (17), that determines θ, and finally, the last block composed of (19), (20) determines $\{U, W\}$. In light of the numerical exercises in later sections, it is important to state that the result still holds under the inclusion of aggregate shocks.

Consider Ω an aggregate exogenous state, such as the cost of automation, κ, or the rate of data accumulation, Δ, and assume they evolve according to some exogenous law of motion. Then, we can modify the equations that compose the blocks just mentioned, to include the additional aggregate state variable and the block recursive property will not be violated. This is due to the fact that the free-entry condition by itself determines the distributions of entrant firms and unemployed workers across submarkets by making entrants indifferent through the matching mechanism. For example, more productive markets are more crowded than less productive ones. The same logic applies for more productive economies when thinking about aggregate shocks. Therefore, no agent in the economy needs to keep track of other agent's policies, only of the aggregate distributions.

Proposition 1. *The competitive equilibrium has the block recursive property. Moreover, the equilibrium exists and is unique.*

3.4.1 Deterministic Analysis

Consider a deterministic version of the model presented, where human capital and the data stock don't grow ($\Delta = 0$). Moreover, workers cannot die and can only lose the job through full automation. First, firms post vacancies under free entry, deciding to enter or not and how much data to acquire. Observing the job-finding rate for all submarkets, unemployed workers, with $\mathbf{h}$ drawn from G, choose first how much to invest in their human capital (retraining) and then, in which submarket to search. Finally, payoffs are realized. We assume here, for simplicity that workers are risk-neutral, that the cost of data function is linear, $c(d) = c \cdot d$, that retraining is infinitely costly, $\chi = \infty$, the matching function is Leontief, $M(v, u) = \min\{v, u\}$, and that automatable and non-automatable human capitals are complements, i.e., $\rho < 1$. Additionally, we assume that when a firm is indifferent between fully automating and any other choice, it chooses not to fully automate. This exercise allows us to analyze thoroughly two key mechanisms for structural change: the nature of job openings (non, partially or fully automated) and their wages. Due to the block recursive property of the model, we can start by solving the entry problem of the firms (it doesn't depend on future worker choices), and recover the worker's search and wage policy.

At the entry stage, firms solve for each submarket $(\mathbf{h}, \phi)$

$$J^e(\mathbf{h},\phi) = \max_{d \le h_a} \left\{ \max_{s\in[0,1]} \left((sh_n)^\rho + ((1-s)h_a)^\rho \right)^{1/\rho} (1-\phi) + \beta \widehat{J}(\mathbf{h},\phi,d) - c\cdot d \right\} - \max_{d\le h_a} \left\{ \beta J^F(d) - c(d) \right\} \tag{23}$$

which is the static equivalent of (17). Notice that we can solve for the entry value of a firm, $J^e(\mathbf{h},\phi)$, in closed form under the stated assumptions, along with $\widetilde{c} := c(1-\beta)/\beta < 1$ [18],

$$\begin{aligned} J^e(\boldsymbol{h},\phi) = {} & (h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}} (1-\phi) \\ & + \frac{\beta}{1-\beta} \cdot \max\left\{ (h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}} (1-\phi) - h_a(1-\tilde{c}-\kappa), \left\{ (h_n^\rho + h_a^\rho)^{1/\rho}(1-\phi) - h_a \right\} (1-\kappa), 0 \right\} \end{aligned} \tag{24}$$

which comes from the fact that, as the states don't evolve in this deterministic setting, the automation policy in the second period forward will not change. As, under the assumptions, both partially and fully automated firms value is strictly increasing in the data stock [19], we have that their data purchase will be equal to h_a. Finally, it is optimal for non-automated firms to not buy data at entry. We can solve for the firm's automation policy in terms of the piece-rate of the submarket,

$$f^P(\mathbf{h},\phi) = \mathbb{1}\left(\phi \le \min\{\phi^{NP}(\mathbf{h}), \phi^{PF}(\mathbf{h})\} \right) \tag{25}$$

$$f^F(\mathbf{h},\phi) = \mathbb{1}\left(\phi > \max\{\phi^{NF}(\mathbf{h}), \phi^{PF}(\mathbf{h})\} \right) \tag{26}$$

where ϕ^{PF} and ϕ^{NF} represent the piece-rates at which the firm is indifferent between fully automating and not automating, and between partially automating and not automating, respectively. Meanwhile, ϕ^{NP} is the piece-rate that makes the firm indifferent between partial and no automation. Additionally, a firm enters a submarket only if it is profitable enough to pay the setup cost $\widetilde{\psi}(\mathbf{h})$.

Therefore, if ϕ^N and ϕ^P exhaust the non and partially automated firm's expected profit, respectively, a firm will post a non-automated job if

$$\phi^{NP}(\mathbf{h}) < \phi \le \min\{\phi^N(\mathbf{h}), \phi^{NF}(\mathbf{h})\} := \widehat{\phi}^N(\mathbf{h}) \tag{27}$$

and a partially automated job if

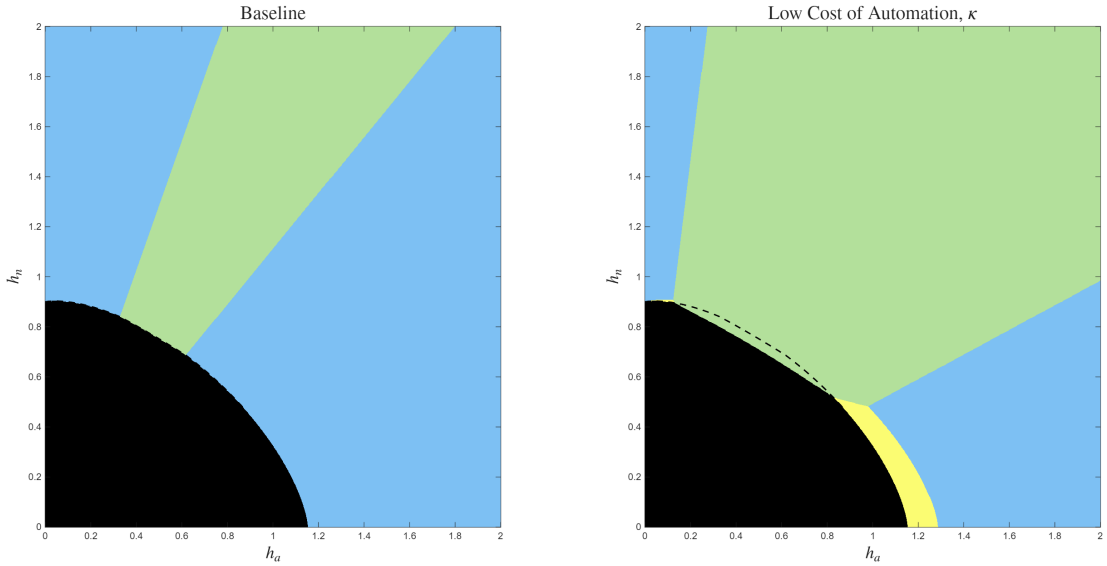
$$\phi \le \min\{\phi^P(\mathbf{h}), \phi^{PF}(\mathbf{h}), \phi^{NP}(\mathbf{h})\} := \widehat{\phi}^P(\mathbf{h}) \tag{28}$$

[18] otherwise, there wouldn't be any full automation.

[19] evident for the fully automated. For partially automated, assume not. Then, $\left(h_n^\rho + d^\rho\right)^{\frac{1-\rho}{\rho}} d^{\rho-1} - \widetilde{c} \le 0 \Rightarrow h_n^\rho \le d^\rho \left(\widetilde{c}^{\frac{\rho}{1-\rho}} - 1\right) < 0$, which is a contradiction.

both of these have job security (tenure). However, if ϕ is big enough that it satisfies (26), the firm's policy is to post a **pace-setter** job, i.e., a job with one initial period of human labor, followed by full automation.

Figure 16: Automation Policy



This figure illustrates the firm's entry policy for a combination of **h**. Blue = no automation, Green = partial automation, Black = no entry, Yellow = full automation. The baseline consists of a case with high wages and high cost of automation. The dashed lines separate the areas in the baseline and are kept in the other figures for comparative purposes.

Figure 16 presents the entry policy and automation status, for a given wage, along with how it is affected by changes in the cost of automation. The colors indicate the policy, with black: no entry - the cost of hiring is higher than the production of a worker; blue: no automation - workers are productive enough to compensate for wages, but have not enough non-automatable human capital to make it worth partially automating them; green: partial automation - worker's non-automatable skills are high enough to justify the payment of wages and the automation cost; finally, yellow: full automation - worker's non-automatable skills are low and automatable skills are high enough. The economy with a high cost of automation, κ, is reported as the baseline. As we can see, from a firm's point of view, most workers are productive enough to be worth hiring, but only the ones specially skilled in non-automatable tasks become partially automated.

The case of a reduction in the cost of automation (right graph), κ, is of special importance as we interpret it as a channel for technological development, which we explore further in later sections. Lowering the cost of automation leads to a rise of both automation states areas. The fact that some of the no entry zone becomes partially automated, while some of the non-

automated zone becomes fully automated, is an indicative that the effect of a lower cost of automation on unemployment is ambiguous.

Using the free-entry condition (17), we can back out the market tightness, θ, for each submarket. Given that the matching probability for the firm is given by $q(\theta) = \theta^{-1}$, under $v > u$,

$$\theta(\mathbf{h}, \phi) = \begin{cases} \dfrac{J^e(\mathbf{h}, \phi)}{\widetilde{\psi}(\mathbf{h})}, & \text{if } J^e(\mathbf{h}, \phi) \geq \max\{\widetilde{\psi}(\mathbf{h}), 0\} \\ 0, & \text{if otherwise} \end{cases} \tag{29}$$

The conditional $J^e(\mathbf{h}, \phi) \geq \max\{\widetilde{\psi}(\mathbf{h}), 0\}$ for firms to post vacancies in a submarket comes from the requirement that firms must have positive profits and that the probability of matching cannot exceed 1[20].

Knowing the distribution of unemployed and vacancies posted among the submarkets, workers can decide where to search, which, given that they are endowed with their human capital vector and cannot alter it while searching, is reduced to choosing their piece-rate. At this point, workers are aware that deciding their wage also affects their risk of being automated. Thus, given the firm's automation policy (26), (27) and (28), the problem of a worker searching for job security, i.e., a job that is either non-automated or partially automated, is given by

$$V^T = \max_{0 \leq \phi \leq 1} \phi \cdot \begin{cases} \dfrac{(h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}}}{1-\beta}, & \text{if } \phi^{NP}(\mathbf{h}) < \phi \leq \widehat{\phi}^N(\mathbf{h}) \\ (h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}} + \dfrac{\beta}{1-\beta}(h_n^\rho + h_a^\rho)^{\frac{1}{\rho}}(1-\kappa), & \text{if } \phi \leq \widehat{\phi}^P(\mathbf{h}) \end{cases} \tag{30}$$

which captures the trade-off between a higher piece-rate at a less productive match (non-automated), and a smaller piece-rate at a more productive match (partially automated). Finally, the problem of an unemployed worker is given by

$$U(\boldsymbol{h}) = z(\boldsymbol{h}) + \beta \max_{0 \leq \phi \leq 1} \begin{cases} V^T, & \text{if } \phi \leq \max\{\widehat{\phi}^N(\mathbf{h}), \widehat{\phi}^P(\mathbf{h})\} \\ \phi \cdot f(\mathbf{h}) + \beta U(\boldsymbol{h}), & \text{if } \max\{\phi^{NF}(\mathbf{h}), \phi^{PF}(\mathbf{h})\} < \phi \leq \phi^F(\mathbf{h}) \\ \dfrac{z(\boldsymbol{h})}{1-\beta}, & \text{otherwise.} \end{cases} \tag{31}$$

The solution to this equation, as shown in the appendix, can be expressed in terms of the

[20] In other words, if the profits weren't enough to pay the entry cost, $q(\theta)$ would have to exceed 1 to make (17) hold.

present value of wages,

$$U(\boldsymbol{h}) = z(\boldsymbol{h}) + \beta \max\left\{ w^N(\mathbf{h}), w^P(\mathbf{h}), \frac{w^F(\mathbf{h}) + \beta z(\mathbf{h})}{1-\beta^2}, \frac{z(\mathbf{h})}{1-\beta} \right\} \tag{32}$$

where the first options are comprised by the present value of the wage of a non-automated job, w^N, and a partially automated job, w^P, with the third option, being the welfare of a pace-setter, is given by the present value of a weighted average between the unemployment benefit[21], z, and the pace-setter's single period wage, w^F.

To understand how unexpected bursts of technological progress can affect the search policy and wages of the economy, we need to separate the workers whose wage policy is constrained by the firm's outside option of full automation, from those whose policy is not. Consider a worker who searches for a non-automated job, i.e., $w^N(\mathbf{h})$. If the present value of the outside option, $h_a(1-\widetilde{c}-\kappa)/(1-\beta)$, is smaller than the setup cost, $\psi(\mathbf{h})$, then the worker's piece rate choice, ϕ^N, is such that the present value of his wages matches the firms value. Alternatively, if this assumption didn't hold, he would receive his marginal product.

$$w^N(\mathbf{h}) = \frac{1}{1-\beta} \cdot \begin{cases} (h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}} - \psi(\mathbf{h})(1-\beta), & \text{if } \psi(\mathbf{h}) \leq \frac{h_a(1-\widetilde{c}-\kappa)}{1-\beta} \\ (h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}} - h_a(1-\widetilde{c}-\kappa), & \text{if otherwise} \end{cases} \tag{33}$$

We refer to this latter group as **marginally non-automated**, as a higher wage would lead him to being fully automated. Similarly, for partially automated workers,

$$w^P(\mathbf{h}) = \frac{1}{1-\beta} \cdot \begin{cases} (1-\beta)(f(\mathbf{h}) - \psi(\mathbf{h})) + \beta(f^P(\mathbf{h})(1-\kappa) - h_a\widetilde{c}), & \text{if } \frac{\psi(\mathbf{h})(1-\beta)+\widetilde{c}h_a\beta}{f(\mathbf{h})(1-\beta)+f^P(\mathbf{h})(1-\kappa)\beta} \leq \frac{h_a}{f^P(\mathbf{h})} \\ \left\{ f^P(\mathbf{h}) - h_a \right\} \left[\frac{f(\mathbf{h})}{f^P(\mathbf{h})}(1-\beta) + (1-\kappa)\beta \right], & \text{if otherwise} \end{cases} \tag{34}$$

where we refer to the group constrained by the fully automation option of the firm **marginally automated**. Finally, it is important to define the exposure measure in the model that matches the content of the empirical one:

$$e(\mathbf{h}) = \frac{h_a}{h_a + h_n} \tag{35}$$

in words, the share of a worker's total human capital that is automatable.

[21] This can be seen by writing $\frac{w^F(\mathbf{h})+\beta z(\mathbf{h})}{1-\beta^2} = \frac{1}{1-\beta}\left[w^F \frac{1}{1+\beta} + z(\mathbf{h}) \frac{\beta}{1+\beta} \right]$

3.4.2 Unexpected Permanent Shock in κ

Consider now a negative aggregate unexpected permanent shock in the cost of automation, κ. We interpret this experiment as analogous to the event study discussed in section 2.2 and, therefore, representative of the technological advancement that led to the creation of modern automation tools. Now, the implicit piece-rates chosen in the previous section come into light, as firms commit to it at the time of hiring, but not to total wage, which might be affected by the shock. The following results separate its impact on the wages and job stability of currently employed workers, and on the search policy of currently unemployed workers.

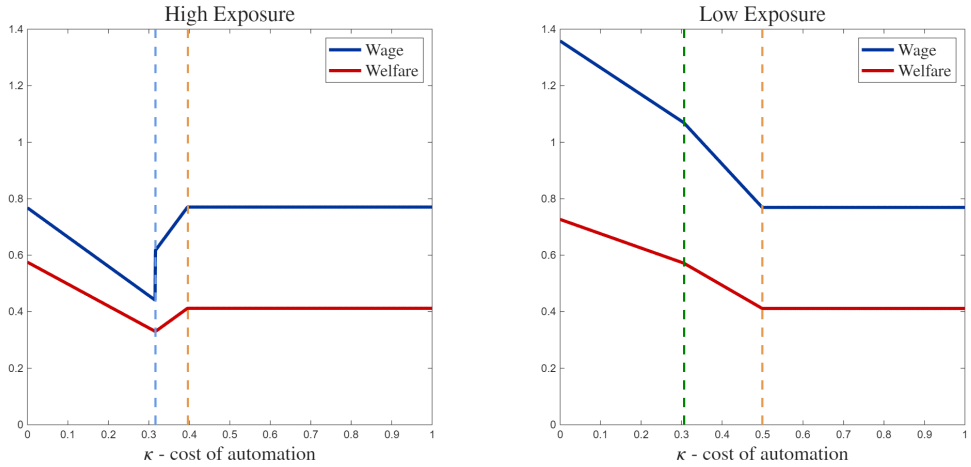
Proposition 2. *Consider a decrease in the cost of automation from κ to κ'. Then every unemployed worker who keeps the same search policy (i.e., chooses a job with the same automation state) experiences the following wage effects: (i) workers targeting partially automated, marginally automated, or pace-setter jobs experience wage gains; (ii) workers targeting marginally non-automated jobs experience wage losses; (iii) workers targeting non-automated jobs experience no change in wages. Workers who switch their search policy from non-automated to marginally non-automated experience wage losses, whereas workers who switch from marginally automated to partially automated experience wage gains.*

This proposition makes clear what unemployed worker profiles would observe gains or losses after a drop in the cost of automation. The first part of it is intuitive, as any worker that becomes or remains partially automated in the low κ world, receives an equal (or bigger) share of a growing match surplus. On the other end, non-automated workers whose marginal product has fallen and considerably restricted their choices, at the risk of full automation, will invariably observe losses. This certainty of loss doesn't hold true for the workers who transition into marginal automation, as the match surplus also changes, making a marginal product in a more productive match potentially higher than a bigger share of a less productive match. Finally, in exchange for his job tenure, a pace-setter reaps the rewards of helping a firm achieve full automation. As any displaced worker in the fully dynamic model must face this re-entry mechanism before and after an aggregate shock, the results presented in this section should make clear the patterns observed in the quantitative section for workers with same total human capital, but differing in its composition.

Figure 17 illustrates the changes in the search policy, along with a visual representation of the winners and losers discussed in **Proposition 2**. The left figure, presents the welfare and wages of a highly exposed worker to automation, with his automatable human capital composing over 90% of his total human capital, while the right figure presents a worker with less than 60% of automatable human capital, with both of them having the same total human capital[22].

[22] the exposures where chosen so that the workers have the same non-automated wages/welfare.

Figure 17: Search Policy



The figure on the left presents wages and welfare as a function of κ for a highly exposed worker. The figure on the right presents the same but for a low exposure worker. In the left figure, the areas from left to right separated by the vertical lines denote the search policy of the worker: pace-setter, marginal non-automation and non-automation. For the figure in the right, from left to right: marginal automation, partial automation and non-automation.

The exposed worker faces no wage losses while the outside option of the firm, to fully automate him, is not profitable enough. As κ decreases, this situation changes: after crossing the dashed orange line, the worker can retain employment only by accepting a wage no higher than his marginal product, which itself declines with further reductions in κ. In this region workers are marginally non-automated. Further to the left, after the dashed blue line, it becomes unfeasible to search for a tenured job, as the value of full automation rises and the worker has not enough non-automatable human capital to be partially automated. In this region the worker searches for a pace-setter job.

Under high automation costs, the low exposure worker receives the same as his counterpart. However, the firm's preference for partial automation becomes evident as κ falls, as they are able to extract the productivity gains of automation. Here it is clear that the composition of the human capital is key for determining the wage selection differences among workers, with the lesser exposed worker being able to reap the rewards of the technological advancements. For low values of κ, the outside option of full automation looms even for these workers, and they must accept becoming marginally automated, and earn exactly their marginal product, on which the efficiency gains are still reflected, albeit at a lower intensity that as partially automated workers. Finally, the welfare of the lesser exposed workers, compared to the more exposed, reflects the benefits of having the skills that are demanded in an economy with access to automation. The next proposition states that only highly exposed unemployed workers will switch their search policy to pace-setting.

Proposition 3. *Only highly exposed unemployed workers switch to a pace-setting search policy when the cost of automation κ falls. More specifically, if κ falls and $e(\mathbf{h}) \geq \bar{e}(\mathbf{h})$, the optimal search policy for unemployed workers with endowed with $\mathbf{h}$ is to search for a pace-setter job.*[23].

For a worker to be a pace-setter, he must offer little in terms of non-automatable human capital and the cost of automation must be small enough. Proposition 3 states exactly that. With a drop in the cost of automation only the workers that are exposed enough will prefer high earnings followed by a period of unemployment, instead of job security with lower wages. The following result concerns employed workers at the time of the shock.

Proposition 4. *Consider a decrease in the cost of automation from κ to κ'. Then every partially automated worker (marginal or not) experiences wage gains. If the decrease in κ is sufficiently large, the non-automated workers (marginal or not) ath the top or bottom of the exposure distribution, i.e., below some e^L or above some e^H, become fully automated, while the ones in the middle become partially automated and experience wage gains. If the decrease in κ is not sufficiently large, non-automated workers experience neither wage gains nor losses.*

The result that automated workers at the time of the shock have wage gains is expected, as they will be getting the same fixed share of the firm's product, which has now grown. For the non-automated employed workers, due to the complementarity between automatable and non-automatable human capital, the ones that have too little of either type are in risk of being fully automated. The ones outside of this area observe wage gains, as they become partially automated.

It is important to remember that the results presented in this section are under the assumption of no retraining. In an economy where agents are able to adjust the composition of their human capital to adapt to technological changes, many of the patterns observed here might not hold true. However, if the costs of retraining are considerable, there must be room for economic policy to alleviate the losses of the workers less prepared for the new demands of the job market, possibly through retraining.

4 Estimation

4.1 Calibrated Parameters

The unit of time is one quarter. Workers have CRRA utility, $u(c) = c^{1-\sigma}/(1-\sigma)$, with risk aversion $\sigma = 2$ and discount factor $\beta = 0.99$. Each period a worker dies with probability

[23] Under assumptions stated in the appendix

$\lambda = 1/(46 \times 4)$, matching 46 years of working life, and is replaced by a newborn who draws initial human capital. The matching function follows Den Haan et al. (2000),

$$M(v,u) = \frac{u \cdot v}{(u^\gamma + v^\gamma)^{1/\gamma}},$$

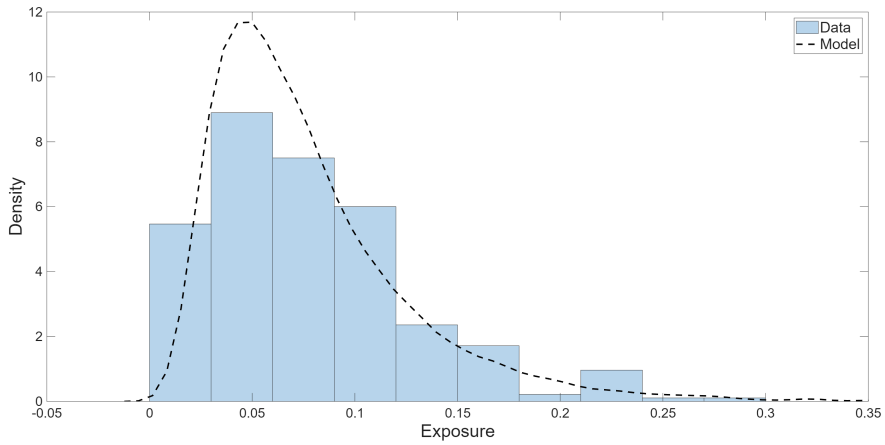
with $\gamma = 1.6$ as in Schaal (2017). Matches are destroyed exogenously at rate $\delta = 0.034$, calibrated to the average separation rate from JOLTS in 2010.

Home production is linear, $z(\mathbf{h}) = z_0$, and the cost of acquiring data is $c(d) = c_0 \cdot \bar{c} \cdot d$, where $\bar{c}$ is a scalar equal to the maximum production of a non-automated firm.[24] The setup cost in submarket $\mathbf{h}$ is proportional to the outside option of full automation: $\psi(\mathbf{h}) = \psi \cdot \left(\max_{d \leq h_a} \beta J^F(d) - c(d)\right)$.

4.2 Initial Human Capital Distribution

Newborn workers draw their initial exposure level, $e(\mathbf{h})$, from a log-normal distribution (Figure 18). Total human capital is monotonically related to exposure through $\bar{h} = h_a + h_n = \sigma_{h,e} \cdot \exp(e(\mathbf{h}))$, so that the automatable and non-automatable components are $h_a = e(\mathbf{h}) \cdot \bar{h}$ and $h_n = (1 - e(\mathbf{h})) \cdot \bar{h}$, respectively. We assume no human capital depreciation and that both types grow on the job with probability p_h.

Figure 18: Exposure Distribution



Notes. Empirical (bars) and parametrized (line) distribution of the AI exposure measure. The model assumes a log-normal distribution with mean 0.08 and standard deviation 0.05.

[24] The selling price of data upon separation equals the purchasing price, i.e., $c(d) = x(d)$.

4.3 Estimated Parameters

We estimate eleven parameters by targeting the moments in Table 5. The estimation proceeds in two stages: nine parameters are identified from steady-state moments characterizing the pre-2021 economy, and two additional parameters are identified from the natural experiment described in Section 2.2, in which the cost of automation drops unexpectedly and permanently by a fraction δ_κ (i.e., $\kappa' = (1 - \delta_\kappa)\kappa$).

Steady-state identification. The setup cost ψ governs vacancy creation: higher costs reduce the number of posted vacancies and raise the unemployment rate. We target the BLS unemployment rate of 6.2 percent. The per-period cost of automation κ discourages firms from adopting AI, so the ABS partial automation rate of 21.7 percent identifies this parameter. The home production level z_0 determines workers' outside option and is identified by the replacement rate of 88.0 percent from Braxton and Taska (2025).

The cost of retraining χ is pinned down by the retraining rate of 16.6 percent, taken from Braxton and Taska (2025), who in turn use the estimates of Jacobson et al. (2005). Recent estimates by Hyman et al. (2025) suggest positive and significant retraining gains even among AI-exposed workers. The learning probability p_d controls how quickly firms accumulate data on worker tasks; it is identified by the average time to automate (conditional on non-immediate adoption) of 19.9 percent, computed from Revelio and BTOS data. The price of data c_0 determines whether firms bypass the learning process and automate immediately; it is identified by the immediate partial automation rate of 50.5 percent from the ABS.

The scale parameter $\sigma_{h,e}$ links exposure to total human capital and is identified by the correlation between wages and exposure (0.919) from CPS and Hampole et al. (2025). The standard deviation of exposure σ_h at birth is set to match the cross-sectional standard deviation of 0.051 from the same source. The probability of human capital growth p_h is the only source of on-the-job wage variation in the model and is identified by the probability of wage gains for exposed workers (20.9 percent).

Natural experiment identification. We interpret the unexpected permanent reduction in the automation cost by δ_κ as the model analogue of the AI adoption event described in Section 2.2 and discussed theoretically in Section 3.4.2. This experiment is simulated outside the steady-state sample, so the pre-shock moments remain representative of the pre-2021 economy.

Two moments discipline this experiment. The post-shock layoff rate is computed as the product of the adopter share (13.52 percent from the 2023 ABS, Table B1) and the reported employment decrease rate (3.04 percent), yielding 0.4 percent. Larger reductions in κ produce

more full automation and hence more layoffs, identifying δ_κ. The post-shock wage differential between exposed and non-exposed workers (2.0 log points from Table 2) contains information on the production complementarity ρ. A high degree of substitutability (high ρ) shifts firms' task allocation toward the most productive set of tasks. Because exposed workers also have higher total human capital, the wage gains of this group increase with ρ, identifying the parameter from the regression coefficient. One caveat is that the model regression cannot include firm fixed effects, as there is no concept of a multi-worker firm in our framework.

Table 5 reports the target moments, model fit, and the parameter each moment primarily identifies. Table 6 collects all calibrated and estimated parameter values.

Table 5: Target Moments and Model Fit

Target Moment	Source	Data	Model	Parameter	Value
Steady-state moments					
Unemployment rate	BLS	0.062	0.039	ψ: Setup cost	1.181
Partial automation rate	ABS	0.217	0.252	κ: Cost of automation	0.186
Replacement rate	JLS (2005)	0.880	0.963	z_0: Home production	1.687
Retraining rate	JLS (2005)	0.166	0.138	χ: Cost of retraining	0.148
Avg. time to automate	Revelio & BTOS	0.199	0.143	p_d: Data learning probability	0.452
Immediate automation rate	ABS	0.505	0.536	c_0: Price of data	0.170
SD of exposure	Hampole et al. (2025)	0.051	0.032	σ_h: SD of HC at birth	0.718
Corr. of wages & exposure	CPS & Hampole et al. (2025)	0.919	0.849	$\sigma_{h,e}$: HC–exposure scale	0.967
Prob. of wage gain (exposed)	CPS & Hampole et al. (2025)	0.209	0.131	p_h: HC growth prob.	0.000
Natural experiment moments					
Layoffs post-shock	ABS	0.004	0.010	δ_κ: Cost drop	0.098
Wage differential post-shock	ABS & Tax data	0.020	0.049	ρ: Complementarity	0.673

Notes. Data column reports empirical targets. Model column reports simulated moments at estimated parameter values. Arrow indicates the parameter primarily identified by each moment. JLS = Jacobson et al. (2005). Average time to automate conditions on non-immediate adopters. Layoff rate = adopter share $\times$ employment decrease rate.

Table 6: Model Parameters

Parameter	Description	Value
Estimated		
ψ	Setup cost	1.181
κ	Per-period cost of automation	0.186
z_0	Home production level	1.687
χ	Cost of retraining	0.148
p_d	Data learning probability	0.452
c_0	Price of data	0.170
σ_h	SD of total human capital at birth	0.718
$\sigma_{h,e}$	Scale of total HC and exposure at birth	0.967
p_h	Human capital growth probability	0.000
δ_κ	Drop in cost of automation	0.098
ρ	Production complementarity	0.673
Calibrated		
β	Discount factor	0.990
σ	CRRA coefficient	2
γ	Matching function parameter	1.6
λ	Death rate (quarterly)	0.005
δ	Exogenous separation rate	0.034

Notes. Estimated parameters (top panel) are obtained by minimizing the distance between simulated and empirical moments in Table 5. Calibrated parameters (bottom panel) are set to standard values or direct data counterparts.

4.4 Numeric Characterization

Figure 19 plots four key policy functions of the estimated model. These objects summarize the decisions that firms and workers make conditional on the state variables – human capital, data, and wages – and they make visible the mechanisms that drive the counterfactual results in the next section.

Panel A plots the job-finding rate as a function of automatable human capital (h_a) and non-automatable human capital (h_n), holding the piece-rate fixed at $\phi = 0.6$. The job-finding rate is highest for workers with high levels of both types of human capital, reflecting the fact that these workers generate the most productive matches and therefore attract the most vacancy posting. Workers with low h_n and high h_a – the most exposed to full automation – face lower job-finding rates, because firms anticipate that these matches will be short-lived: the worker's

non-automatable contribution is small, so the match is likely to end in full automation or separation. This pattern is a direct consequence of the free-entry condition: firms post fewer vacancies in submarkets where the expected match surplus, net of automation risk, is low.

Panel B plots the firm's initial data purchase policy as a function of the piece-rate (ϕ), holding automatable human capital fixed at its upper bound ($h_a = \bar{h}$). At low piece-rates, the firm retains a large share of output and is willing to purchase substantial data upfront to accelerate automation. As the piece-rate rises, the firm's share of output falls and the return to upfront data investment declines, so firms purchase less data at entry and instead rely on on-the-job learning. At sufficiently high piece-rates, the firm purchases no data at all – the match is not productive enough to justify the fixed cost, and the firm either learns slowly through experience or does not automate. This policy function illustrates the experience-good mechanism: when data is costly, firms substitute toward on-the-job learning, which requires keeping the worker employed.

Panel C plots the automation policy as a function of human capital (h_n on the horizontal axis, h_a on the vertical axis), holding the piece-rate at $\phi = 0.6$ and the data stock at $d = 0.8$. The color indicates the automation state: values of 2 correspond to partial automation, values of 3 correspond to full automation, and values of 1 (or below 2) correspond to no automation. The figure reveals a clear partition of the human capital space. Workers with high h_n and moderate h_a are partially automated – the firm uses AI to handle the automatable tasks while retaining the worker for the non-automatable tasks, reflecting the complementarity between the two types of human capital ($\rho = 0.67$). Workers with very high h_a and low h_n are fully automated – their non-automatable contribution is too small to justify continued employment given the available data stock. Workers with low h_a are not automated at all, regardless of h_n, because the gains from automation are too small to cover the per-period cost κ. This panel makes visible the key prediction of the model: partial automation is the dominant outcome for workers with balanced skill profiles, and full automation is confined to a narrow region of the state space.

Panel D plots the automation policy as a function of the data stock (d) and the piece-rate (ϕ), holding automatable human capital fixed at $h_a = \bar{h} = 0.8$. At low data levels, no automation occurs regardless of the piece-rate – the firm has not yet accumulated enough information on the worker's tasks to make automation worthwhile. As data accumulates, partial automation becomes optimal for matches with moderate piece-rates: the firm has learned enough to deploy AI on automatable tasks while retaining the worker. At high data levels and low piece-rates, full automation emerges – the firm has sufficient data and retains enough of the output to justify replacing the worker entirely. The transition from no automation to partial automation to full automation as d increases illustrates the experience-good dynamics at the core of our model:

automation is a process, not a one-time decision, and its pace depends on the rate at which the firm accumulates data on the worker's tasks.

Together, these four panels illustrate why the job creation margin dominates in our counterfactuals. Partial automation occupies a large and stable region of the policy space, sustained by the complementarity between non-automatable human capital and AI-generated data. Full automation is confined to a narrow corner – workers with very high exposure, sufficient data, and low piece-rates. As the cost of automation falls, the partial automation region expands, but the full automation region does not, because firms continue to value the non-automatable contribution of workers. This is the mechanism through which labor persists despite AI.

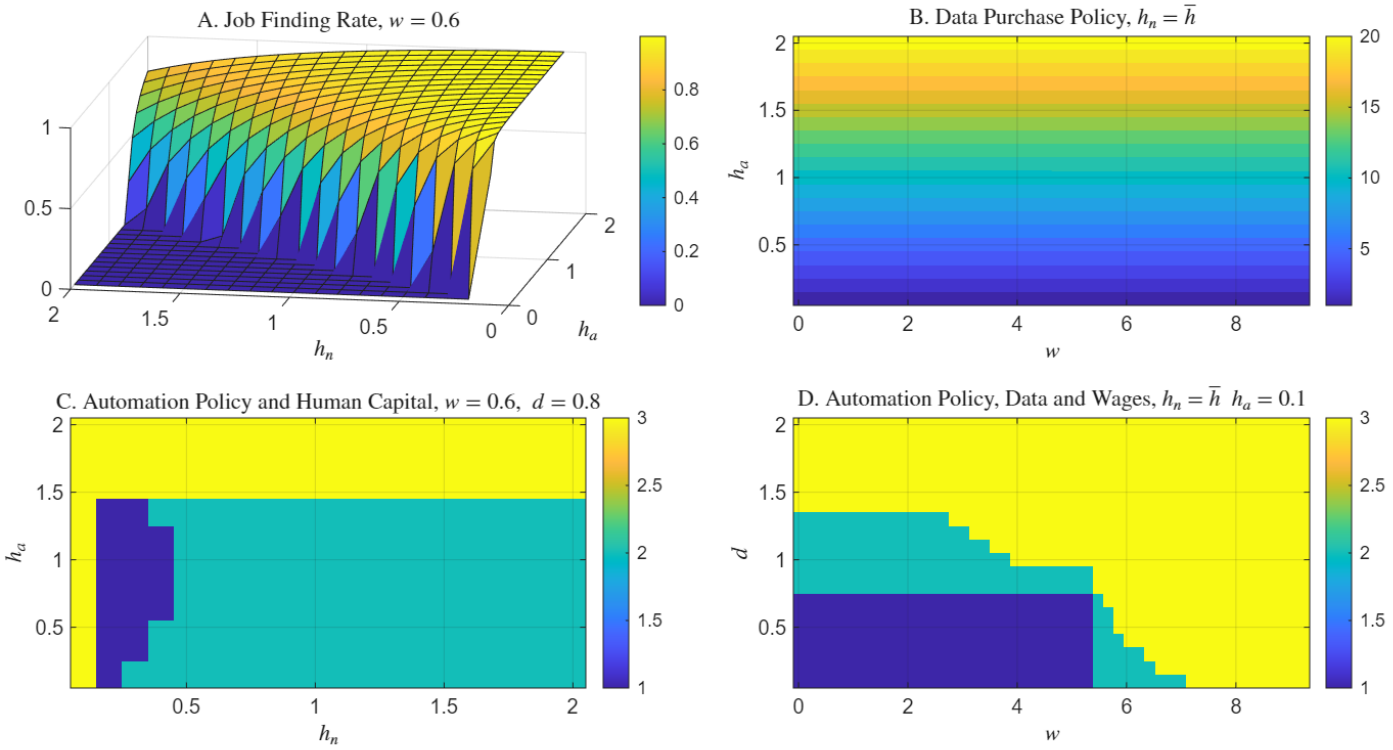


Figure 19: **Policy functions of the estimated model.**

Notes. Panel A: job-finding rate as a function of automatable (h_a) and non-automatable (h_n) human capital, at piece-rate $\phi = 0.6$. Panel B: firm's initial data purchase as a function of the piece-rate ϕ, at $h_a = \bar{h}$. Panel C: automation policy as a function of h_n (horizontal) and h_a (vertical), at $\phi = 0.6$ and $d = 0.8$; color scale: 1 = no automation, 2 = partial automation, 3 = full automation. Panel D: automation policy as a function of data stock d (horizontal) and piece-rate ϕ (vertical), at $h_a = \bar{h} = 0.8$; color scale as in Panel C. All panels evaluated at estimated parameter values (Table 6).

5 Counterfactuals

We use the estimated model to ask how the labor market responds to large reductions in the cost of AI. These exercises are motivated by the rapid pace of AI development since 2021 and by the widespread concern that continued cost declines will displace large numbers of workers. We examine three channels through which AI becomes cheaper or more effective: reductions in the per-period cost of automation (κ), reductions in the price of external data (c_0), and decreases in the speed of data accumulation (p_d). For each channel, we compare the new steady state at 95 percent and 50 percent of the baseline parameter value to the estimated baseline. Table 7 reports the results.

Table 7: Steady-State Moments Under Alternative AI Cost Scenarios

	Baseline	κ		c_0		p_d	
		95%	50%	95%	50%	95%	50%
Unemployment rate (%)	3.9	4.0	3.9	4.0	4.3	3.9	4.0
Partial automation rate (%)	25.0	24.9	96.1	25.0	26.1	24.9	26.0
Full automation rate (%)	0.2	0.2	0.2	0.2	0.5	0.2	0.2
Time to partial automation (yrs)	0.140	0.139	0.000	0.140	0.000	0.160	0.000
Retraining rate (%)	13.8	13.8	13.8	13.8	21.9	13.8	13.8
Average wages	1.756	1.761	1.918	1.755	1.743	1.756	1.758
Δ wages (%)	—	+0.3	+9.2	−0.1	−0.7	0.0	+0.1

Notes. Each column reports the steady-state moments under the indicated parameter variation. Columns labeled 9 and 50% reduce κ or c_0 (or increase p_d) to that fraction of its baseline value. Rates are shares of employed workers. Δ wages reports the percent change in average wages relative to baseline.

Reducing the cost of automation. The most striking result is what happens when the per-period cost of automation falls by half. Partial automation rises from 25 to 96 percent of employed workers – nearly universal adoption – yet the unemployment rate is unchanged at 4 percent and the full automation rate remains at zero. The time to partial automation drops to zero, reflecting the fact that at low enough costs, firms no longer want to accumulate data before adopting (note that they continue to accumulate data during partial automation, by assumption). The retraining rate is unchanged. Average wages rise by 9.1 percent, from 1.76 to 1.92. This result is a direct consequence of our model's key mechanism: because AI is more productive when trained on in-house data, firms prefer to keep workers and partially automate

them rather than fully replace them. The complementarity between non-automatable human capital and data ($\rho = 0.67$, indicating complements) means that workers remain valuable even as AI handles an increasing share of their automatable tasks. Partial automation raises the productivity of the match, and because workers receive a piece-rate of output, these productivity gains flow directly to wages. In our model, the nature of work changes – skill-weights shift toward non-automatable tasks – but jobs do not disappear, and workers are paid more.

Reducing the price of data. Halving the price of external data (c_0) produces a qualitatively different response. Because cheaper data allows firms to bypass the in-house learning process, full automation becomes more attractive: the full automation rate rises from zero to 1 percent, while the unemployment rate is unchanged at 4 percent. The retraining rate increases substantially, from 14 to 22 percent, as displaced workers invest in non-automatable skills. Average wages fall modestly by 1.1 percent, from 1.76 to 1.74. This channel captures the concern that off-the-shelf AI models, trained on freely available data, could substitute for workers without requiring firm-specific experience. Even so, the decline in wages is small, because the partial automation margin continues to create new and more productive task configurations for workers with sufficient non-automatable human capital.

Reducing the speed of learning. Lowering the learning probability p_d – so that firms accumulate data on worker tasks more slowly – modestly increases the partial automation rate and eliminates delays in adoption. The unemployment rate, full automation rate, retraining rate, and average wages are all unchanged. Firms no longer wait for data accumulation before automating.

Benefits to labor. Across all three channels, the central finding is that the job creation margin dominates. Even under extreme reductions in the cost of automation, workers are not largely displaced. Instead, AI reshapes the content of jobs: workers spend less time on automatable tasks and more time on tasks that AI cannot replicate. When automation becomes cheaper, this reallocation raises the productivity of partially automated matches, and average wages rise by 9 percent. The workers who benefit most are those with high total human capital – who also tend to be the most exposed to AI – because partial automation raises their marginal product. Workers at the bottom of the exposure distribution are largely unaffected, as their jobs are neither attractive enough to automate nor productive enough to justify the costs of AI adoption. The workers who face the greatest risk are those with high automatable but low non-automatable human capital, particularly when external data is cheap. For these workers,

retraining – investing in non-automatable skills during unemployment – is the primary margin of adjustment, and the retraining rate rises from 14 to 22 percent under the cheap-data scenario.

The only channel that produces wage losses is cheaper external data, and even there the decline is just over 1 percent. By contrast, cheaper automation per se produces large wage gains with no increase in displacement. These results hinge on our model's treatment of AI as an experience good. If AI could be adopted without in-house data – if all the relevant training data were freely available and external to the firm – then the job destruction margin would be substantially larger. It is precisely because firms need workers to generate the data that makes AI productive that the complementarity between labor and AI persists even as adoption costs fall. In this sense, labor is necessary for AI, and this necessity protects workers from the displacement that purely inspection-good models of automation would predict.

6 Conclusion

This paper develops the hypothesis that AI models must be trained on data generated by workers, and that this complementarity between labor and data fundamentally alters AI's effects on the labor market. We build this argument in three steps.

First, we use the Annual Business Survey linked to worker tax records to document the importance of in-house R&D, worker talent, and data quality for AI adoption, and we provide the first estimates, to our knowledge, of the worker-level wage effects of firm-reported AI adoption in the United States. We find that exposed workers gain on average relative to non-exposed workers within the same firm, but that these gains mask significant heterogeneity: at firms that shed workers in response to AI, high-exposure workers experience meaningful wage losses. These patterns hold both for AI exposure and for an alternative treatment based on occupation-level exposure to in-house R&D, and they are robust to demanding specifications that include firm-by-year, occupation-by-year, and income-bin-by-year fixed effects.

Second, we develop a directed search model in which AI investment is an experience good: firms must accumulate internal data on worker tasks before automation becomes productive. We allow firms to adjust the task content of jobs as they partially or fully automate, endogenously giving rise to new work, and we allow workers to retrain by investing in non-automatable skills. The model is tractable and block recursive, permitting us to compute transition paths in response to aggregate shocks.

Third, we use the model to evaluate how the labor market responds to large reductions in the cost of AI. Our central finding is that the job creation margin wins, even under extreme cost declines. Halving the per-period cost of automation drives partial automation to near-universal

levels, raises average wages by 9 percent, and leaves unemployment unchanged, because the complementarity between non-automatable human capital and AI-generated data makes firms prefer to keep workers and reallocate their tasks rather than replace them outright. The channel that poses the greatest risk to labor is not cheaper automation per se, but cheaper external data – the possibility that off-the-shelf AI models could substitute for firm-specific learning. Even in this scenario, average wages decline by only 1 percent, and the retraining margin absorbs much of the adjustment.

References

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics 40*(S1), S293–S340.

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 (2024). Automation and rent dissipation: Implications for wages, inequality, and productivity. Technical report, National Bureau of Economic Research.

Althoff, L. and H. Reichardt (2026). Task-specific technical change and comparative advantage.

Autor, D., C. Chin, A. Salomons, and B. Seegmiller (2024). New frontiers: The origins and content of new work, 1940–2018. *The Quarterly Journal of Economics*, qjae008.

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), 1553–1597.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics 151*, 103745.

Baslandze, S., Z. Edwards, J. R. Graham, T. McClure, B. Meyer, M. Sparks, S. Waddell, and D. Weitz (2026). Artificial intelligence, productivity, and the workforce: Evidence from corporate executives. *Productivity, and the Workforce: Evidence from Corporate Executives (March 13, 2026)*.

Beraja, M., D. Y. Yang, and N. Yuchtman (2023). Data-intensive innovation and the state: Evidence from ai firms in china. *The Review of Economic Studies 90*(4), 1701–1723.

Beraja, M. and N. Zorzi (2025). Inefficient automation. *Review of Economic Studies 92*(1), 69–96.

Bick, A., A. Blandin, and D. J. Deming (2026). The rapid adoption of generative ai. *Management Science*.

Bonney, K., C. Breaux, C. Buffington, E. Dinlersoz, L. S. Foster, N. Goldschlag, J. C. Haltiwanger, Z. Kroff, and K. Savage (2024). Tracking firm use of ai in real time: A snapshot from the business trends and outlook survey. Technical report, National Bureau of Economic Research.

Braxton, J. C. and B. Taska (2025). Technological change and insuring job loss. *Review of Economic Dynamics*, 101308.

Brynjolfsson, E., B. Chandar, and R. Chen (2025). Canaries in the coal mine?: Six facts about the recent employment effects of artificial intelligence.

Brynjolfsson, E., D. Li, and L. Raymond (2025). Generative ai at work. *The Quarterly Journal of Economics 140*(2), 889–942.

Dechezleprêtre, A., D. Hemous, M. Olsen, and C. Zanella (2019). Automating labor: evidence from firm-level patent data. *Available at SSRN 3508783*.

Dell'Acqua, F., E. McFowland III, E. R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard business school technology & operations mgt. Unit working paper* (24-013).

Den Haan, W. J., G. Ramey, and J. Watson (2000). Job destruction and propagation of shocks. *American economic review 90*(3), 482–498.

Dillon, E. W., S. Jaffe, N. Immorlica, and C. T. Stanton (2025). Shifting work patterns with generative ai. Technical report, National Bureau of Economic Research.

Farboodi, M. and L. Veldkamp (2021). A model of the data economy. Technical report, National Bureau of Economic Research Cambridge, MA, USA.

Freund, L. B. and L. F. Mann (2025). Job transformation, specialization, and the labor market effects of ai.

Gathmann, C., F. Grimm, and E. Winkler (2024). Ai, task changes in jobs, and worker reallocation. Technical report, CESifo Working Paper.

Hampole, M., D. Papanikolaou, L. Schmidt, and B. Seegmiller (2025). Artificial intelligence and the labor market. Technical report, Mimeo.

Hémous, D. and M. Olsen (2022). The rise of the machines: Automation, horizontal innovation, and income inequality. *American Economic Journal: Macroeconomics 14*(1), 179–223.

Hosseini Maasoum, S. M. and G. Lichtinger (2025). Generative ai as seniority-biased technological change: Evidence from us resume and job posting data. *Available at SSRN 5425555*.

Humlum, A. and E. Vestergaard (2025). Large language models, small labor market effects. Technical report, National Bureau of Economic Research.

Hyman, B. G., B. Lahey, K. Ni, and L. Pilossoph (2025). How retrainable are ai-exposed workers? Technical report, National Bureau of Economic Research.

Jacobson, L., R. LaLonde, and D. G Sullivan (2005). Estimating the returns to community college schooling for displaced workers. *Journal of Econometrics 125*(1), 271–304.

Jaimovich, N., I. Saporta-Eksten, H. Siu, and Y. Yedid-Levi (2021). The macroeconomics of automation: Data, theory, and policy analysis. *Journal of Monetary Economics 122*, 1–16.

Jones, C. I. and C. Tonetti (2020). Nonrivalry and the economics of data. *American Economic Review 110*(9), 2819–2858.

Kogan, L., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2023). Technology and labor displacement: Evidence from linking patents with worker-level data. Technical report, National Bureau of Economic Research.

Korinek, A. and D. Suh (2024). Scenarios for the transition to agi. Technical report, National Bureau of Economic Research.

Lazear, E. P. (2009). Firm-specific human capital: A skill-weights approach. *Journal of political economy 117*(5), 914–940.

Lise, J. and F. Postel-Vinay (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review 110*(8), 2328–2376.

Ljungqvist, L. and T. Sargent (1998). The european unemployment dilemma. *Journal of Political Economy 106*(3), 514–550.

McElheran, K., J. F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. Foster, and N. Zolas (2024). Ai adoption in america: Who, what, and where. *Journal of Economics & Management Strategy 33*(2), 375–415.

Menzio, G. and S. Shi (2011). Efficient search on the job and the business cycle. *Journal of Political Economy 119*(3), 468–510.

Moll, B., L. Rachel, and P. Restrepo (2022). Uneven growth: automation's impact on income and wealth inequality. *Econometrica 90*(6), 2645–2683.

Schaal, E. (2017). Uncertainty and unemployment. *Econometrica 85*(6), 1675–1721.

Veldkamp, L. and C. Chung (2024). Data and the aggregate economy. *Journal of Economic Literature 62*(2), 458–484.

Wang, P. and T.-N. Wong (2026). Artificial intelligence and technological unemployment. *Journal of Monetary Economics 158*, 103905.

7 Proofs

Some notation:

$$A(\mathbf{h}) = \max_{s\in[0,1]} \left((s \cdot h_n)^\rho + ((1-s) \cdot h_a)^\rho\right)^{1/\rho}$$

$$B(h_n, d) = \left(h_n^\rho + d^\rho\right)^{1/\rho}$$

Lemma 1.

- Let $\rho < 1$ and compute the first and second derivatives of $g(s;\mathbf{h}) = \left((s \cdot h_n)^\rho + ((1-s) \cdot h_a)^\rho\right)^{1/\rho}$, by s:

$$g'(s;\mathbf{h}) = \left((s \cdot h_n)^\rho + ((1-s) \cdot h_a)^\rho\right)^{\frac{1-\rho}{\rho}} \cdot \left[s^{\rho-1}h_n^\rho - (1-s)^{\rho-1}h_a^\rho\right]$$

$$g''(s;\mathbf{h}) = \frac{1-\rho}{\rho}\left((s \cdot h_n)^\rho + ((1-s) \cdot h_a)^\rho\right)^{\frac{1-2\rho}{\rho}} \cdot \left[s^{\rho-1}h_n^\rho - (1-s)^{\rho-1}h_a^\rho\right] - (1-\rho) \cdot \left[s^{\rho-2}h_n^\rho + (1-s)^{\rho-2}h_a^\rho\right]$$

  as g is differentiable everywhere in $(0,1)$ and has a unique s^* that solves $g'(s^*,\mathbf{h}) = 0$, given by

$$s^* = \frac{h_n^{\frac{\rho}{1-\rho}}}{h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}}}$$

  attaining a value of

$$g(s^*,\mathbf{h}) = \left(h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}}\right)^{\frac{1-\rho}{\rho}}$$

  as the first bracket of g'' is equal to zero at s^* (it is the FOC), we can conclude that $g''(s^*,\mathbf{h}) < 0$. Therefore, s^* is a local maximum, and the only one in $(0,1)$. Finally, wlog let $h_n \geq h_a > 0$, notice that

$$g(s^*,\mathbf{h}) = \left(h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}}\right)^{\frac{1-\rho}{\rho}} > \left(h_n^{\frac{\rho}{1-\rho}} + 0\right)^{\frac{1-\rho}{\rho}} = h_n = \max\{h_n, h_a\}$$

  so, $g(s^*,\mathbf{h}) > \max\{g(1,\mathbf{h}), g(0,\mathbf{h})\}$, and s^* is a global maximum.

- Let $\rho > 1$, $s \in (0,1)$ and wlog $h_n \geq h_a > 0$. Then,

$$\left((s \cdot h_n)^\rho + ((1-s) \cdot h_a)^\rho\right)^{1/\rho} \leq h_n\left(s^\rho + (1-s)^\rho\right)^{1/\rho} < h_n$$

  as $s^\rho + (1-s)^\rho < 1$. Therefore, $s^* \in 0, 1$ and $g(s^*;\mathbf{h}) = \max\{h_n, h_a\}$.

- The same logic follows for $\rho = 1$ when $h_a \neq h_n$. However, for $h_a = h_n$, any $s \in [0,1]$ is optimal. This doesn't break that for $\rho \geq 1$, $g(s^*; \mathbf{h}) = \max\{h_n, h_a\}$. ■

Solving the Deterministic Model. Consider the problem of the worker searching for a tenured job and let $f(\mathbf{h}) = (h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}}$ and $f^P(\mathbf{h}) = (h_n^\rho + h_a^\rho)^{1/\rho}$

$$V^T = \max_{0 \leq \phi \leq 1} \phi \cdot \begin{cases} f(\mathbf{h}) + \dfrac{\beta}{1-\beta} f(\mathbf{h}), & \text{if } \phi^{NP}(\mathbf{h}) < \phi \leq \widehat{\phi}^N(\mathbf{h}) \\ f(\mathbf{h}) + \dfrac{\beta}{1-\beta} f^P(\mathbf{h})(1-\kappa), & \text{if } \phi \leq \widehat{\phi}^P(\mathbf{h}) \end{cases}$$

where ϕ^{NP} = makes na and pa indifferent:

$$\phi^{NP}(\mathbf{h}) = \begin{cases} \max\left\{1 - \dfrac{\widetilde{c} h_a}{f^P(\mathbf{h})(1-\kappa) - f(\mathbf{h})}, 0\right\}, & \text{if } f(\mathbf{h}) < f^P(\mathbf{h})(1-\kappa) \\ 0, & \text{if otherwise} \end{cases}$$

notice that if $f(\mathbf{h}) \geq f^P(\mathbf{h})(1-\kappa)$, there is no ϕ that makes partially automating optimal for the firm (it would require $\phi > 1$ and negative profits), thus the whole domain $[0,1]$ is given to non-automating firms. $\widehat{\phi}^P(\mathbf{h})$ is the highest piece-rate a worker $\mathbf{h}$ can search for that guarantees tenure and partial automation once matched:

$$\widehat{\phi}^P(\mathbf{h}) = \min\left\{\phi^{NP}(\mathbf{h}), \phi^P(\mathbf{h}), \phi^{PF}(\mathbf{h})\right\}$$

where ϕ^P exhausts the pa firm's ex-ante profits and ϕ^{PF} makes it indifferent between pa and fa in the following period, i.e.,

$$\phi^P(\mathbf{h}) = \max\left\{1 - \frac{\psi(\mathbf{h})(1-\beta) + \widetilde{c} h_a \beta}{f(\mathbf{h})(1-\beta) + f^P(\mathbf{h})(1-\kappa)\beta}, 0\right\},\ \phi^{PF}(\mathbf{h}) = \max\left\{1 - \frac{h_a}{f^P(\mathbf{h})}, 0\right\}$$

similarly, $\widehat{\phi}^N(\mathbf{h})$ is the upper bound of the na firm,

$$\widehat{\phi}^N(\mathbf{h}) = \min\left\{\phi^N(\mathbf{h}), \phi^{NF}(\mathbf{h})\right\}$$

where

$$\phi^N(\mathbf{h}) = \max\left\{1 - \frac{\psi(\mathbf{h})(1-\beta)}{f(\mathbf{h})}, 0\right\},\ \phi^{NF}(\mathbf{h}) = \max\left\{1 - \frac{h_a(1-\widetilde{c}-\kappa)}{f(\mathbf{h})}, 0\right\}$$

Claim. if $\widehat{\phi}^N(\mathbf{h}) \leq \phi^{NP}(\mathbf{h})$, then $\widehat{\phi}^N \cdot \left[f(\mathbf{h}) + \frac{\beta}{1-\beta} f(\mathbf{h})\right] \leq \widehat{\phi}^P \cdot \left[f(\mathbf{h}) + \frac{\beta}{1-\beta} f^P(\mathbf{h})(1-\kappa)\right]$

Proof:

As $\widehat{\phi}^N(\mathbf{h}) \leq \phi^{NP}(\mathbf{h})$, then $f(1-\widehat{\phi}^N) < f^P(1-\kappa)(1-\widehat{\phi}^N)$, and therefore $f < f^P(1-\kappa)$. So, it's clear that the case of $\widehat{\phi}^P = \phi^{NP}$ satisfies the statement. Consider now $\widehat{\phi}^P = \phi^P$. Then, if $\widehat{\phi}^N = \phi^N$, then $(1-\phi^N)f^P(1-\kappa) - \widetilde{c}h_a > f(1-\phi^N) = (1-\phi^P)f^P(1-\kappa) - \widetilde{c}h_a$, which implies that $\phi^P > \phi^N$. If $\widehat{\phi}^N = \phi^{NF}$, then $\psi < f(1-\phi^{NF}) = f(1-\beta) + \beta h_a(1-\widetilde{c}-\kappa) < f(1-\phi^P)(1-\beta) + \beta[f^P(1-\phi^P)(1-\kappa) - \widetilde{c}h_a] = \psi$, which is a contradiction. Consider, alternatively, that $\widehat{\phi}^P = \phi^{PF}$. Then, if $\widehat{\phi}^N = \phi^N$, we have $f(1-\phi^{PF})(1-\beta) + \beta[f^P(1-\phi^{PF})(1-\kappa) - \widetilde{c}h_a] < \psi = f(1-\phi^N)(1-\beta) < f(1-\phi^N)(1-\beta) + \beta[f^P(1-\phi^N)(1-\kappa) - \widetilde{c}h_a]$, which implies that $\phi^N < \phi^{PF}$. Finally, the case of $\widehat{\phi}^N = \phi^{NF}$ is immediate. ■

Therefore, we can rewrite V^T as

$$V^T = \max\left\{\widehat{\phi}^N\left[(h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}} + \frac{\beta}{1-\beta}(h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}}\right], \widehat{\phi}^P\left[(h_n^{\frac{\rho}{1-\rho}} + h_a^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}} + \frac{\beta}{1-\beta}(h_n^{\rho} + h_a^{\rho})^{\frac{1}{\rho}}(1-\kappa)\right]\right\}$$

and, the problem of the unemployed searching for all jobs becomes

$$U(\boldsymbol{h}) = z(\boldsymbol{h}) + \beta \max_{0\leq\phi\leq 1} \begin{cases} V^T, & \text{if } \phi \leq \max\{\widehat{\phi}^N, \widehat{\phi}^P\} \\ \phi \cdot f(\mathbf{h}) + \beta U(\boldsymbol{h}), & \text{if } \max\{\phi^{NF}, \phi^{PF}\} < \phi \leq \phi^F \\ \frac{z(\boldsymbol{h})}{1-\beta}, & \text{otherwise.} \end{cases}$$

Lets first tackle the nontenured option:

$$U(\boldsymbol{h}) = \phi^F(\mathbf{h})f(\mathbf{h}) + \beta z(\boldsymbol{h}) + \beta^2 U(\boldsymbol{h}) \Rightarrow U(\boldsymbol{h}) = \frac{\phi^F(\mathbf{h})f(\mathbf{h}) + \beta z(\boldsymbol{h})}{1-\beta^2}$$

Claim. if $\phi^F(\mathbf{h}) \leq \max\{\widehat{\phi}^N, \widehat{\phi}^P\}$, then $\frac{\phi^F(\mathbf{h})f(\mathbf{h}) + \beta z(\boldsymbol{h})}{1+\beta} \leq V^T(\mathbf{h})$

Proof:

if $\widehat{\phi}^F(\mathbf{h}) \leq \max\{\widehat{\phi}^N, \widehat{\phi}^P\}$, then $\frac{\widehat{\phi}^F f(\mathbf{h}) + \beta z(\boldsymbol{h})}{1+\beta} \leq \frac{V^T + \beta z(\boldsymbol{h})}{1+\beta} \leq \max\{V^T, z(\boldsymbol{h})\}$. ■

Finally, we can write the full problem as

$$U(\boldsymbol{h}) = z(\boldsymbol{h}) + \frac{\beta}{1-\beta}\max\left\{V^T(\mathbf{h}), \frac{\phi^F(\mathbf{h})f(\mathbf{h}) + \beta z(\mathbf{h})}{1+\beta}, z(\mathbf{h})\right\}$$

Using the definitions for the present value of a tenured worker's wage (automated or not)

$$w^N(\mathbf{h}) = \widehat{\phi}^N(\mathbf{h}) \cdot \frac{f(\mathbf{h})}{1-\beta} = \frac{1}{1-\beta} \min\left\{f(\mathbf{h}) - \psi(\mathbf{h})(1-\beta), f(\mathbf{h}) - h_a(1-\tilde{c}-\kappa)\right\}$$

$$w^P(\mathbf{h}) = \widehat{\phi}^P(\mathbf{h}) \cdot \left[f(\mathbf{h}) + \frac{\beta}{1-\beta} f^P(\mathbf{h})(1-\kappa)\right]$$

$$= \frac{1}{1-\beta} \min\Big\{(1-\beta)(f(\mathbf{h}) - \psi(\mathbf{h})) + \beta(f^P(\mathbf{h})(1-\kappa) - h_a\widetilde{c}),$$

$$(f(\mathbf{h}) + \beta f^P(\mathbf{h})(1-\kappa))\left(1 - \frac{\widetilde{c}h_a}{f^P(\mathbf{h})(1-\kappa) - f(\mathbf{h})}\right), \left\{f^P(\mathbf{h}) - h_a\right\}\left[\frac{f(\mathbf{h})}{f^P(\mathbf{h})}(1-\beta) + (1-\kappa)\beta\right]\Big\}$$

and finally, the present value of the non-tenured (pace-setter) is given by

$$w^F(\mathbf{h}) = \phi^F(\mathbf{h}) \cdot f(\mathbf{h}) = f(\mathbf{h}) - \widetilde{\psi}(\mathbf{h})$$

and, therefore we can write the problem of the unemployed searching as

$$U(\boldsymbol{h}) = z(\boldsymbol{h}) + \beta \max\left\{w^N(\mathbf{h}), w^P(\mathbf{h}), \frac{w^F(\mathbf{h}) + \beta z(\mathbf{h})}{1-\beta^2}, \frac{z(\mathbf{h})}{1-\beta}\right\}$$

Proposition 2. Let $\kappa > \kappa'$, consider an unemployed worker with $\mathbf{h}$ and define $\Delta w = w(\mathbf{h}, \kappa') - w(\mathbf{h}, \kappa)$. In this proof we index the equilibrium objects by the aggregate shock too, κ. Let $\widetilde{\kappa}^N$ and $\widetilde{\kappa}^P$ be s.t. the marginal and non-marginal present value of wages are equal:

$$f(\mathbf{h}) - h_a(1-\tilde{c}-\widetilde{\kappa}^N) = f(\mathbf{h}) - (1-\beta)\psi(\mathbf{h})$$

as the derivative of the left hand side by κ is positive, non-automation is preferred for $\kappa \geq \widetilde{\kappa}^N$.

$$\left\{f^P(\mathbf{h}) - h_a\right\}\left[\frac{f(\mathbf{h})}{f^P(\mathbf{h})}(1-\beta) + (1-\widetilde{\kappa}^P)\beta\right] = (1-\beta)(f(\mathbf{h}) - \psi(\mathbf{h})) + \beta(f^P(\mathbf{h})(1-\widetilde{\kappa}^P) - \widetilde{c}h_a)$$

Similarly, as the derivative of both sides by κ are negative, with the right hand side being of greater magnitude, marginal automation is preferred for $\kappa < \widetilde{\kappa}^P$.

1. A non-automated worker that stays non-automated has no wage change as $(1-\beta)\Delta w = f(\mathbf{h}) - (1-\beta)\psi(\mathbf{h}) - (f(\mathbf{h}) - (1-\beta)\psi(\mathbf{h})) = 0$.

2. A non-automated worker who becomes marginally non-automated has wage losses. Notice that κ' must be s.t. $(1-\beta)w^N(\mathbf{h}, \kappa) = f(\mathbf{h}) - h_a(1-\tilde{c}-\kappa') < f(\mathbf{h}) - h_a(1-\tilde{c}-\widetilde{\kappa}^N) = f(\mathbf{h}) - (1-\beta)\psi(\mathbf{h}) = (1-\beta)w^N(\mathbf{h}, \kappa)$. So, $\Delta w < 0$.

3. A marginally non-automated that stays marginally non-automated has wage losses: $(1-\beta)\Delta w = h_a(\kappa' - \kappa) < 0$

4. A pace-setter that remains a pace-setter has wage gains: $(1-\beta)\Delta w = \dfrac{\beta}{1-\beta} h_a(\kappa - \kappa') > 0$.

5. A marginally automated that stays marginally automated has wage gains $(1-\beta)\Delta w = \beta(f^P(\mathbf{h}) - h_a)(\kappa - \kappa') > 0$

6. A partially automated that stays partially automated has wage gains: $(1-\beta)\Delta w = \beta f^P(\mathbf{h})(\kappa - \kappa') > 0$.

7. A partially automated that becomes marginally automated has wage gains, as $(1-\beta)w^P(\mathbf{h},\kappa) = (1-\beta)(f(\mathbf{h}) - \psi(\mathbf{h})) + \beta(f^P(\mathbf{h})(1-\kappa) - \widetilde{c}h_a) < (1-\beta)(f(\mathbf{h}) - \psi(\mathbf{h})) + \beta(f^P(\mathbf{h})(1 - \widetilde{\kappa}^P) - \widetilde{c}h_a) = \left\{f^P(\mathbf{h}) - h_a\right\}\left[\frac{f(\mathbf{h})}{f^P(\mathbf{h})}(1-\beta) + (1-\widetilde{\kappa}^P)\beta\right] \leq (1-\beta)w^P(\mathbf{h},\kappa')$.

■

Proposition 3. for a non-automated worker to prefer to become a pace-setter with a drop in $\kappa' < \kappa$, the following needs to hold

$$f(\mathbf{h}) - h_a \max\{\psi(1-\beta), 1 - \widetilde{c} - \kappa'\} < \frac{f(\mathbf{h}) - \psi h_a + \frac{\beta}{1-\beta} h_a(1 - \widetilde{c} - \kappa') + \beta z(\mathbf{h})}{1+\beta}$$

which can be rewritten in terms of exposure, e,

$$f(e) < z + e\frac{1}{1-\beta}(1 - \widetilde{c} - \kappa' + \frac{1-\beta}{\beta}[m(1+\beta) - \psi])$$

where $m = \max\{\psi(1-\beta), 1-\widetilde{c}-\kappa'\}$, and $f(e) = (e^{\frac{\rho}{1-\rho}} + (1-e)^{\frac{\rho}{1-\rho}})^{\frac{1-\rho}{\rho}}$. Notice that $f(0) > z$, and under our assumptions $f(1) < z + \frac{1}{1-\beta}(1 - \widetilde{c} - \kappa' + \frac{1-\beta}{\beta}[m(1+\beta) - \psi])$. So, by the IVT, there must exist some $\bar{e}$ s.t.

$$f(\bar{e}) = z + \bar{e}\frac{1}{1-\beta}(1 - \widetilde{c} - \kappa' + \frac{1-\beta}{\beta}[m(1+\beta) - \psi])$$

moreover, as $f'(e) \leq 1$, under the assumption of $1/2 \leq \rho$,

$$f'(e) = (e^{\frac{\rho}{1-\rho}} + (1-e)^{\frac{\rho}{1-\rho}})^{\frac{1-2\rho}{\rho}}\left(e^{\frac{2\rho-1}{1-\rho}} - (1-e)^{\frac{2\rho-1}{\rho}}\right) \leq (e^{\frac{\rho}{1-\rho}} + (1-e)^{\frac{\rho}{1-\rho}})^{\frac{1-2\rho}{\rho}} e^{\frac{2\rho-1}{1-\rho}}$$

$$= \left(1 + \left(\frac{1}{e} - 1\right)^{\frac{\rho}{1-\rho}}\right)^{\frac{1-2\rho}{1-\rho}} \leq 1$$

as this last function is increasing in e due to $1 - 2\rho < 0$, we have that

$$f'(e) \leq \frac{1}{1-\beta}(1 - \widetilde{c} - \kappa' + \frac{1-\beta}{\beta}[m(1+\beta) - \psi])$$

and, therefore, pace-setting is preferred to a non-automated job iff $\bar{e} < e$. Finally, notice that a partially automated (marginal or not) worker would never change his search policy to pace-setting as

$$w^P(\mathbf{h})(1-\beta) - \frac{f(\mathbf{h}) - \psi h_a + \frac{\beta}{1-\beta} h_a(1 - \widetilde{c} - \kappa') + \beta z(\mathbf{h})}{1+\beta}$$

$$= \beta \min\left\{(f_P(\mathbf{h}) - \frac{h_a}{1+\beta})(1-\kappa) + K_1, (f_P(\mathbf{h}) - h_a - \frac{h_a}{1+\beta})(1-\kappa) + K_2\right\}$$

where K_1, K_2 are comprised of terms independent of κ. As we can see, a partially automated worker at κ would have this difference being positive, which implies that the same would be true under $\kappa' < \kappa$. Needless to say, it would still be optimal for a firm to automate him as the condition $\hat{\phi}^P(\mathbf{h}, \kappa) < \phi^{NP}(\mathbf{h})$ can be written as

$$\frac{\widetilde{c} h_a}{f_P(\mathbf{h})(1-\kappa) - f(\mathbf{h})} \leq \max\left\{\frac{\psi(\mathbf{h})(1-\beta)}{f(\mathbf{h})}, \frac{h_a}{f^P(\mathbf{h})}\right\}$$

which will continue to hold under $\kappa' < \kappa$. ■

Proposition 4. Let $\kappa > \kappa'$, and consider an employed worker with $\mathbf{h}$.

i) If he is non-automated or marginally automated, with a piece rate of $\phi(\mathbf{h})$,

$$\phi(\mathbf{h}) = 1 - \frac{h_a}{f(\mathbf{h})} \max\{\psi(1-\beta), 1 - \widetilde{c} - \kappa\}$$

a necessary condition for him to be partially automated

$$(1 - \phi(\mathbf{h}))f(\mathbf{h}) < (1 - \phi(\mathbf{h}))f^P(\mathbf{h})(1-\kappa') - \widetilde{c} h_a \Leftrightarrow \kappa' < \kappa^P(\mathbf{h}) := 1 - \frac{f(\mathbf{h})}{f^P(\mathbf{h})} - \frac{\widetilde{c} h_a}{1 - \phi(\mathbf{h})} \frac{1}{f^P(\mathbf{h})}$$

$$= 1 - \frac{f(\mathbf{h})}{f^P(\mathbf{h})}\left[1 + \frac{\widetilde{c}}{\max\{\psi(1-\beta), 1 - \widetilde{c} - \kappa\}}\right]$$

and a necessary condition for him to be fully automated is

$$(1-\phi(\mathbf{h}))f(\mathbf{h}) < h_a(1-\widetilde{c}-\kappa') \Leftrightarrow \kappa' < \kappa^F(\mathbf{h}) := 1-\widetilde{c}-(1-\phi(\mathbf{h}))\frac{f(\mathbf{h})}{h_a}$$

$$= 1-\widetilde{c}-\max\{\psi(1-\beta), 1-\widetilde{c}-\kappa\}$$

therefore, $\kappa^P(\mathbf{h}) \geq \kappa^F(\mathbf{h})$ iff

$$f(\mathbf{h}) \leq \max\{\psi(1-\beta), 1-\widetilde{c}-\kappa\} f^P(\mathbf{h})$$

which is the same condition for pa to be preferred to fa

$$0 \leq \left(f^P(\mathbf{h})(1-\phi(\mathbf{h})) - h_a\right)(1-\kappa) = \left(\frac{f^P(\mathbf{h})}{f(\mathbf{h})}\max\{\psi(1-\beta), 1-\widetilde{c}-\kappa\} - 1\right)(1-\kappa)h_a$$

so, if $\kappa' \leq \max\left\{\kappa^P(\mathbf{h}), \kappa^F(\mathbf{h})\right\}$, then this condition determines if the worker will be fully or partially automated. Notice that it can be written in terms of exposure $e = h_a/(h_n + h_a)$:

$$\left(e^{\frac{\rho}{1-\rho}} + (1-e)^{\frac{\rho}{1-\rho}}\right)^{\frac{1-\rho}{\rho}} \leq \max\{\psi(1-\beta), 1-\widetilde{c}-\kappa\}\left(e^\rho + (1-e)^\rho\right)^{\frac{1}{\rho}}$$

therefore, workers with $e > e^H$ and $e < e^L$ will be fully automated. Finally, stayers, i.e., $\max\{\kappa^P(\mathbf{h}), \kappa^F(\mathbf{h})\} < \kappa'$, won't face wage changes. Meanwhile, if $\kappa' \leq \kappa^P(\mathbf{h})$ and $e^L \leq e \leq e^H$, then the worker will be partially automated and have wage gains, as

$$(1-\phi(\mathbf{h}))f(\mathbf{h}) < (1-\phi(\mathbf{h}))f^P(\mathbf{h})(1-\kappa') - \widetilde{c}h_a \Rightarrow f(\mathbf{h}) < f^P(\mathbf{h})(1-\kappa')$$

ii) If he is partially automated or marginally automated, he can't be fired as the condition $(1-\phi)f^P(\mathbf{h}) \geq h_a$ is independent of κ (same reason for why a mg automated worker will remain so), and he will have wage gains as $\phi f^P(\mathbf{h})(1-\kappa') > \phi f^P(\mathbf{h})(1-\kappa)$. ■

8 ABS Appendix Tables

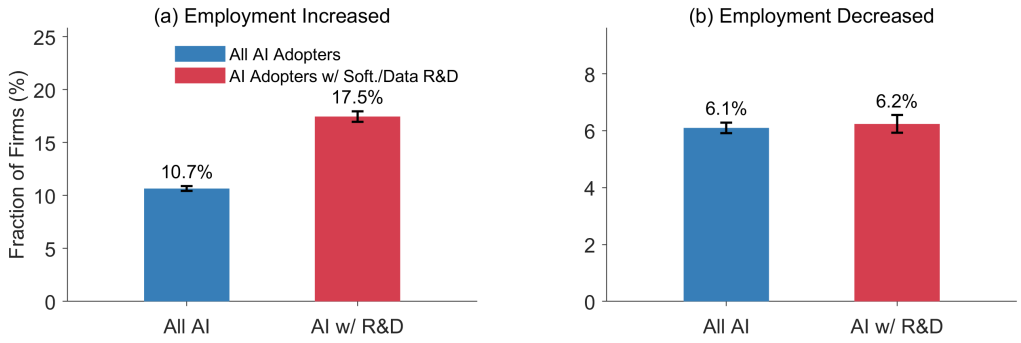


Figure 20: **Employment effects of AI adoption by software/data R&D status, 2023 ABS (reference years 2020–2022).**

Notes. Blue bar: all AI adopters ($N \approx 16{,}000$); red bar: AI adopters with software/data R&D ($N \approx 6{,}000$). Panel (a): fraction reporting employment increased; panel (b): fraction reporting employment decreased. Whiskers show ± 1 SE. *Employment question:* "What were the effects of adopting or using AI on the following? The number of workers employed by this business." *R&D question:* "Did this business have any of the following innovation activities: in-house development and purchase of computer software; collection and analysis of data in proprietary databases; activities to upgrade IT systems."

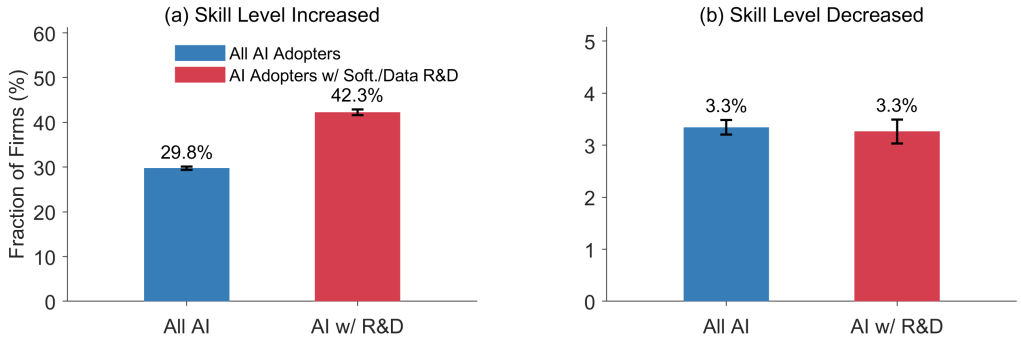


Figure 21: **Workforce skill-level effects of AI adoption by software/data R&D status, 2023 ABS (reference years 2020–2022).**

Notes. Sample and R&D definition as in Figure 20. Panel (a): fraction reporting skill level increased; panel (b): fraction reporting skill level decreased. Whiskers show ± 1 SE. *Survey question:* "What were the effects of adopting or using AI on the following? The skill level of workers employed by this business."

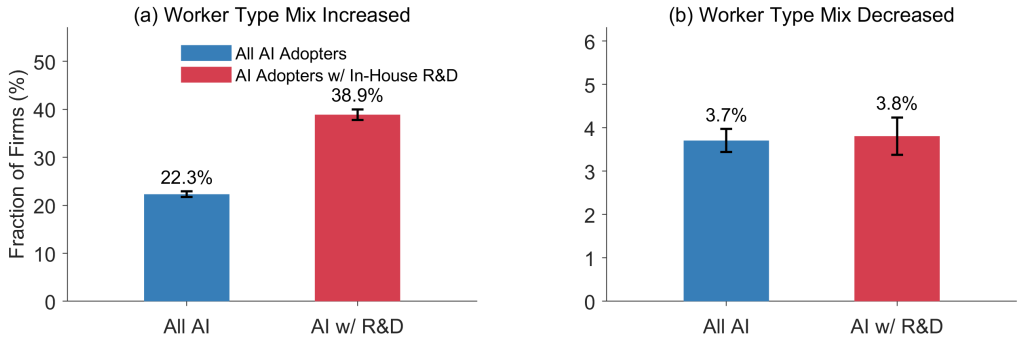


Figure 22: **STEM workforce composition effects of AI adoption by in-house R&D status, 2022 ABS (reference year 2021).**

Notes. Sample and R&D definition as in Figure 5. Panel (a): fraction reporting STEM composition increased; panel (b): fraction reporting STEM composition decreased. Whiskers show ±1 SE. *Survey question:* "To what extent did this business's use of AI technology impact your workforce? The scientific, technological, engineering, and mathematical (STEM) skills of workers employed at this business."

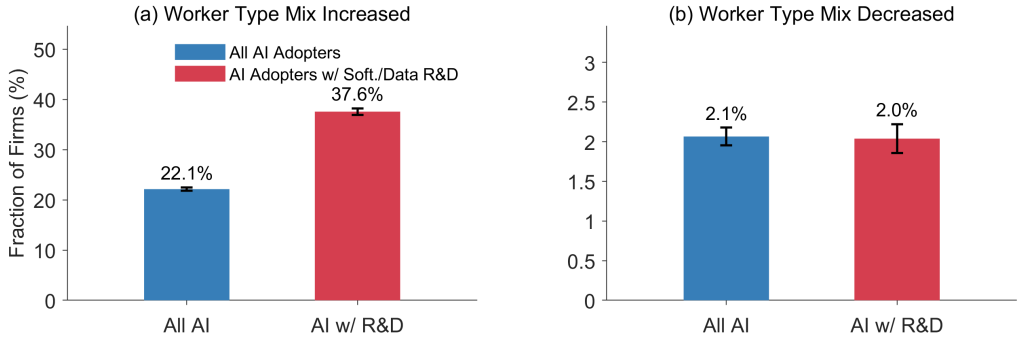


Figure 23: **STEM workforce composition effects of AI adoption by software/data R&D status, 2023 ABS (reference years 2020–2022).**

Notes. Sample and R&D definition as in Figure 20. Panel (a): fraction reporting STEM composition increased; panel (b): fraction reporting STEM composition decreased. Whiskers show ±1 SE. *Survey question:* "What were the effects of adopting or using AI on the following? The scientific, technological, engineering, and mathematical (STEM) skills of workers employed by this business."

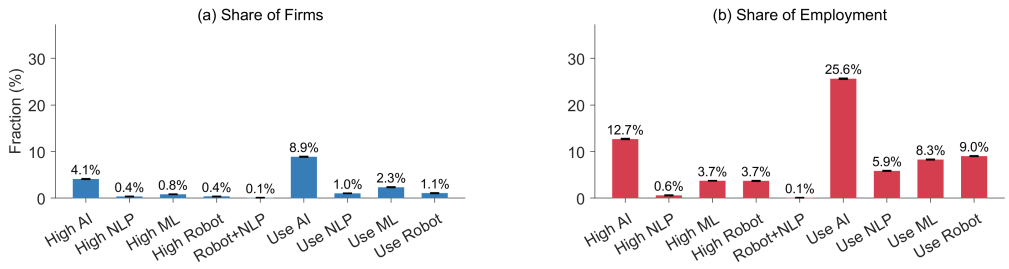


Figure 24: **ML and LLM Adoption Rates 2019 ABS (reference years 2018-2016).**

Table B1: AI Adoption and Workforce Effects, 2023 Annual Business Survey.

Outcome	Firm %	Emp. %	N
Any AI use	2.89	13.52	460,000
Labor-saving intent	36.77	62.56	16,000
Employment increased	10.66	18.76	16,000
Employment decreased	6.10	3.04	16,000
No employment change	83.24	78.20	16,000

Notes: Sample is 2023 ABS respondents (reference years 2020–2022). The top row is unconditional ($N \approx 460{,}000$); remaining rows are conditional on AI adoption ($N \approx 16{,}000$). Census AI definition: "Artificial intelligence is a branch of computer science and engineering devoted to making machines intelligent. Intelligence is that quality that enables an entity to perceive, analyze, determine response and act appropriately in its environment." Labor-saving intent: firms reporting "To automate tasks performed by labor" as their reason for AI adoption. Employment effects: "During the three years 2020 to 2022, what were the effects of adopting or using Artificial Intelligence on the following? The number of workers employed by this business."

Table B2: AI Intensity and In-House R&D, 2022 Annual Business Survey.

Outcome	Firm %	Emp. %	N
Any AI use	4.60	21.74	179,000
In-house AI R&D	0.99	11.10	179,000
In-house AI R&D | high-intensity AI use	49.76	77.08	2,000
In-house AI R&D | any AI use	18.33	49.53	13,000

Notes: Sample is 2022 ABS respondents (reference year 2021). Top panel rows are unconditional ($N \approx 179{,}000$); bottom panel rows are conditional on AI intensity subsamples. AI definition: "Artificial intelligence (e.g., machine learning, planning, reasoning, and decision making)." High-intensity AI users responded "A lot" to the AI usage question. R&D question: "During 2021, did this business perform or fund R&D on the following technologies?"

Table B3: Barriers to AI Adoption, 2019 Annual Business Survey.

Barrier	Firm %	Emp. %	N
Too expensive	23.94	30.50	2,000
Lack of data access	9.38	21.44	2,000
Poor data quality	7.95	35.69	2,000
Lack of talent	10.88	35.67	2,000
No reported barriers	26.89	15.66	2,000
Other barriers	20.96	n.a.	2,000

Notes: Sample is 2019 ABS respondents that tested AI between 2016 and 2018 ($N \approx 2{,}000$). Survey question: "During the three years 2016 to 2018, indicate which factors adversely affected the adoption or utilization of the following technologies to produce goods or services. Select all that apply." Remaining responses are collapsed into "Other barriers."

Table B4: Employment Effects of AI Adoption by In-House R&D Status, 2022 Annual Business Survey.

Outcome	All AI Adopters (%)	AI Adopters w/ R&D (%)
Employment increased	15.86	24.72
Employment decreased	7.94	8.43
No change	76.20	66.86
N	≈5,000	≈2,000

Notes: Sample is 2022 ABS AI-adopting firms (reference year 2021). Survey-weighted means. R&D group: firms that also performed in-house R&D on AI-related technologies. Employment question: "To what extent did this business's use of AI technology impact your workforce? The number of workers employed at this business." R&D question: "During 2021, did the business perform or fund R&D on the following technologies?"

Table B5: Employment Effects of AI Adoption by Software/Data R&D Status, 2023 Annual Business Survey.

Outcome	All AI Adopters (%)	AI Adopters w/ R&D (%)
Employment increased	10.66	17.46
Employment decreased	6.10	6.24
No change	83.24	76.30
N	≈16,000	≈6,000

Notes: Sample is 2023 ABS AI-adopting firms (reference years 2020–2022). Survey-weighted means. R&D group: firms with software/data R&D activity. Employment question: "What were the effects of adopting or using AI on the following? The number of workers employed by this business." R&D question: "Did this business have any of the following innovation activities: in-house development and purchase of computer software; collection and analysis of data in proprietary databases; activities to upgrade IT systems."

9 LEHD Appendix Tables

Table B6: Skill-Level and STEM Workforce Effects of AI Adoption by In-House R&D Status, 2022 Annual Business Survey.

Outcome	All AI Adopters (%)	AI Adopters w/ R&D (%)
Workforce skill level		
Increased	30.15	46.53
Decreased	4.87	5.10
STEM workforce composition		
Increased	22.31	38.90
Decreased	3.71	3.81
N	≈5,000	≈2,000

Notes: Sample is 2022 ABS AI-adopting firms (reference year 2021). Survey-weighted means. Sample sizes and R&D definition as in Table B4. Skill question: "To what extent did this business's use of AI technology impact your workforce? The skill level of workers employed at this business." STEM question: "The scientific, technological, engineering, and mathematical (STEM) skills of workers employed at this business."

Table B7: Skill-Level and STEM Workforce Effects of AI Adoption by Software/Data R&D Status, 2023 Annual Business Survey.

Outcome	All AI Adopters (%)	AI Adopters w/ R&D (%)
Workforce skill level		
Increased	29.79	42.27
Decreased	3.34	3.27
STEM workforce composition		
Increased	22.13	37.56
Decreased	2.07	2.04
N	≈16,000	≈6,000

Notes: Sample is 2023 ABS AI-adopting firms (reference years 2020–2022). Survey-weighted means. Sample sizes and R&D definition as in Table B5. Skill question: "What were the effects of adopting or using AI on the following? The skill level of workers employed by this business." STEM question: "The scientific, technological, engineering, and mathematical (STEM) skills of workers employed by this business."

Table C1: Outcome: Employment Change Indicator

	(1)	(2)	(3)	(4)	(5)
Outcome	Emp. Change	Emp. Change	Emp. Change	Emp. Change	Emp. Change
$\tau_{-2} \times T_i$	-0.0107	-0.00958	-0.0116*	-0.0107*	-0.00734
	(0.00682)	(0.00663)	(0.00636)	(0.00623)	(0.00538)
$\tau_{-1} \times T_i$	0.00202	0.00565	0.000608	0.00387	0.00173
	(0.00667)	(0.00631)	(0.00681)	(0.00648)	(0.00572)
$\tau_0 \times T_i$	0.00799	0.00896	0.00872	0.00973	0.00835
	(0.00761)	(0.00753)	(0.00760)	(0.00756)	(0.00681)
$\tau_{+1} \times T_i$	0.0248***	0.0110	0.0261***	0.0131	0.0120
	(0.00867)	(0.00841)	(0.00912)	(0.00895)	(0.00835)
$\tau_{+2} \times T_i$	0.0196**	0.00792	0.0233***	0.00919	0.00828
	(0.00771)	(0.00775)	(0.00883)	(0.00879)	(0.00828)
R^2	0.463	0.470	0.493	0.501	0.497
N	455K	455K	455K	455K	455K
Worker FE	Y	Y	Y	Y	Y
Firm × Year FE	Y	Y	Y	Y	Y
Inc. Bin × Year FE	N	Y	N	Y	Y
Occ. × Year FE	N	N	N	N	N
Lagged outcome	N	N	Y	Y	Y
ABS weights	Y	Y	Y	Y	N

Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. $T_i = \mathbf{1}(e_i > e^{90})$ is the high-AI-exposure indicator (above 90th percentile). $\tau = 0$ = 2021 (AI adoption year).

Table C2: Outcome: Low Earnings Indicator

	(1)	(2)	(3)	(4)	(5)
Outcome	Low Earnings	Low Earnings	Low Earnings	Low Earnings	Low Earnings
$\tau_{-2} \times T_i$	0.00105*	0.000858	0.00104*	0.000834	0.000940*
	(0.000563)	(0.000522)	(0.000575)	(0.000537)	(0.000537)
$\tau_{-1} \times T_i$	0.00320***	0.00167***	0.00318***	0.00164***	0.00173***
	(0.000857)	(0.000560)	(0.000875)	(0.000581)	(0.000575)
$\tau_0 \times T_i$	-0.00105	0.00172	-0.00111	0.00167	0.000916
	(0.00160)	(0.00153)	(0.00167)	(0.00155)	(0.00122)
$\tau_{+1} \times T_i$	5.38e-05	0.00276	4.19e-05	0.00273	0.00283
	(0.00246)	(0.00260)	(0.00246)	(0.00260)	(0.00289)
$\tau_{+2} \times T_i$	-0.00200	0.000795	-0.00203	0.000755	0.000616
	(0.00200)	(0.00180)	(0.00203)	(0.00182)	(0.00179)
R^2	0.418	0.427	0.418	0.427	0.427
N	455K	455K	455K	455K	455K
Worker FE	Y	Y	Y	Y	Y
Firm × Year FE	Y	Y	Y	Y	Y
Inc. Bin × Year FE	N	Y	N	Y	Y
Occ. × Year FE	N	N	N	N	N
Lagged outcome	N	N	Y	Y	Y
ABS weights	Y	Y	Y	Y	N

Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. $T_i = \mathbf{1}(e_i > e^{90})$ is the high-AI-exposure indicator (above 90th percentile). $\tau = 0$ = 2021 (AI adoption year).

Table C3: Outcome: $\Delta \ln w_2$ — RD Exposure Baseline

	(1)	(2)	(3)	(4)	(5)
Outcome	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$
$\tau_{-2} \times T_i^{RD}$	-0.000280	-0.000780	0.00171	0.00127	0.00132
	(0.00781)	(0.00779)	(0.00695)	(0.00692)	(0.00573)
$\tau_{-1} \times T_i^{RD}$	0.00928	0.00905	0.00797	0.00755	0.00484
	(0.00701)	(0.00699)	(0.00669)	(0.00662)	(0.00601)
$\tau_0 \times T_i^{RD}$	-0.000203	-0.00114	0.000981	-0.000234	-0.00300
	(0.00949)	(0.00942)	(0.00893)	(0.00882)	(0.00804)
$\tau_{+1} \times T_i^{RD}$	0.0153	0.0145	0.0161*	0.0154	0.0200**
	(0.00954)	(0.00940)	(0.00958)	(0.00948)	(0.00913)
$\tau_{+2} \times T_i^{RD}$	0.0215**	0.0199**	0.0222**	0.0206**	0.0188**
	(0.00859)	(0.00870)	(0.00943)	(0.00948)	(0.00853)
$\tau_{-2} \times T_i^{AI}$	0.00552	0.00342	0.00526	0.00292	-0.000607
	(0.00850)	(0.00888)	(0.00743)	(0.00766)	(0.00691)
$\tau_{-1} \times T_i^{AI}$	0.0134	-0.00214	0.0165*	2.15e-05	-0.00234
	(0.00949)	(0.00841)	(0.00989)	(0.00870)	(0.00799)
$\tau_0 \times T_i^{AI}$	0.0104	0.0208*	0.0157	0.0205*	0.0165*
	(0.0109)	(0.0114)	(0.0101)	(0.0106)	(0.00928)
$\tau_{+1} \times T_i^{AI}$	0.00330	0.00477	0.000563	0.00748	-0.000750
	(0.0134)	(0.0143)	(0.0138)	(0.0147)	(0.0151)
$\tau_{+2} \times T_i^{AI}$	0.0300***	0.0192*	0.0292***	0.0217**	0.0151
	(0.00974)	(0.0101)	(0.0102)	(0.0104)	(0.00927)
R^2	0.276	0.286	0.336	0.345	0.340
N	455K	455K	455K	455K	455K
Worker FE	Y	Y	Y	Y	Y
Firm×Year FE	Y	Y	Y	Y	Y
Inc. Bin×Year FE	N	Y	N	Y	Y
Occ.×Year FE	N	N	N	N	N
Lagged outcome	N	N	Y	Y	Y
ABS weights	Y	Y	Y	Y	N

Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. $T_i^{RD} = \mathbf{1}(e_i^{RD} > e_{RD}^{90})$ is the high-RD-exposure indicator (above 90th percentile). T_i^{AI} controls for AI exposure. $\tau = 0$ = 2021 (AI adoption year).

Table C4: Outcome: $\Delta \ln w_2$ — RD Exposure Triple-Difference

	(1)	(2)	(3)	(4)	(5)
Outcome	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$	$\Delta \ln w_2$
$\tau_{-2} \times T_i^{RD}$	0.00204	-0.000909			
	(0.00819)	(0.00887)			
$\tau_{-1} \times T_i^{RD}$	0.00777	0.00723			
	(0.00806)	(0.00831)			
$\tau_0 \times T_i^{RD}$	0.00735	0.00383			
	(0.0102)	(0.0104)			
$\tau_{+1} \times T_i^{RD}$	0.0173*	0.0140			
	(0.0103)	(0.0103)			
$\tau_{+2} \times T_i^{RD}$	0.0242**	0.0210*			
	(0.0117)	(0.0108)			
$\tau_{-2} \times T_i^{RD} \times J_j^+$	-0.00378	-0.00504	-0.00507	-0.00564	-0.00750
	(0.0144)	(0.0176)	(0.0159)	(0.0156)	(0.0154)
$\tau_{-1} \times T_i^{RD} \times J_j^+$	0.00292	0.0122	0.00166	0.00250	0.00271
	(0.0176)	(0.0190)	(0.0177)	(0.0164)	(0.0162)
$\tau_0 \times T_i^{RD} \times J_j^+$	-0.0266	-0.0217	-0.0287	-0.0313	-0.0259
	(0.0199)	(0.0220)	(0.0206)	(0.0209)	(0.0198)
$\tau_{+1} \times T_i^{RD} \times J_j^+$	0.00575	0.00899	0.00769	0.00555	0.00123
	(0.0231)	(0.0231)	(0.0208)	(0.0209)	(0.0201)
$\tau_{+2} \times T_i^{RD} \times J_j^+$	0.000197	0.00174	0.00325	0.00367	0.00719
	(0.0206)	(0.0196)	(0.0192)	(0.0189)	(0.0182)
$\tau_{-2} \times T_i^{RD} \times J_j^-$	0.00928	0.0201	0.00409	0.00544	0.0185
	(0.0220)	(0.0266)	(0.0212)	(0.0212)	(0.0167)
$\tau_{-1} \times T_i^{RD} \times J_j^-$	-0.000772	-0.00284	-0.00295	-0.00559	-0.0120
	(0.0266)	(0.0256)	(0.0255)	(0.0259)	(0.0265)
$\tau_0 \times T_i^{RD} \times J_j^-$	-0.0302	-0.0223	-0.0385	-0.0368	-0.0301
	(0.0329)	(0.0264)	(0.0271)	(0.0258)	(0.0262)
$\tau_{+1} \times T_i^{RD} \times J_j^-$	-0.0388	-0.0205	-0.0558**	-0.0528*	-0.0505*
	(0.0429)	(0.0345)	(0.0277)	(0.0270)	(0.0291)
$\tau_{+2} \times T_i^{RD} \times J_j^-$	-0.0375*	-0.0269	-0.0508***	-0.0536***	-0.0443**
	(0.0216)	(0.0192)	(0.0184)	(0.0188)	(0.0200)
$\tau_{-2} \times T_i^{AI}$	0.00514	0.00324			
	(0.00744)	(0.00886)			
$\tau_{-1} \times T_i^{AI}$	0.0165*	-0.00190			
	(0.00990)	(0.00844)			
$\tau_0 \times T_i^{AI}$	0.0155	0.0207*			
	(0.0101)	(0.0114)			
$\tau_{+1} \times T_i^{AI}$	0.000777	0.00496			
	(0.0138)	(0.0143)			
$\tau_{+2} \times T_i^{AI}$	0.0293***	0.0193*			
	(0.0103)	(0.0101)			
R^2	0.336	0.286	0.342	0.352	0.347
N	455K	455K	455K	455K	455K
Worker FE	Y	Y	Y	Y	Y
Firm×Year FE	Y	Y	Y	Y	Y
Inc. Bin×Year FE	N	Y	N	Y	Y
Occ.×Year FE	N	N	Y	Y	Y
Lagged outcome	Y	N	Y	Y	Y
ABS weights	Y	Y	Y	Y	N

Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. $T_i^{RD} = \mathbf{1}(e_i^{RD} > e_{RD}^{90})$ is the high-RD-exposure indicator (above 90th percentile). T_i^{AI} controls for AI exposure. $\tau = 0 = 2021$ (AI adoption year).