AI Exposure, Financing Frictions, and Long-Term Labor Commitments: Evidence from Green Card Sponsorships

March 2026

Abstract

Do advances in artificial intelligence (AI) change firms' willingness to make long-term commitments to talent? We use employment-based green-card sponsorships for skilled workers as a direct measure of durable human-capital investment. We find that a one–standard-deviation increase in occupational exposure to AI is associated with roughly a 10% decline in sponsorships for the affected occupations, consistent with firms exercising the option to wait under technological uncertainty. The contraction concentrates in entry-level roles and is strongest among financially constrained firms. By contrast, R&D-intensive and AI-innovating employers keep overall commitment levels roughly unchanged by shifting sponsorships toward managerial and soft-skill–intensive positions. Taken together, these findings show how technology shocks—measured as differential task-level exposure to AI across occupations—interact with capital-market frictions and innovation capacity to shape both the level and the mix of long-term human-capital investment.

Keywords: AI Exposure; Human Capital Investment; Financing Frictions; Real Options; Immigration Sponsorship (PERM); Labor Reallocation
JEL Classification: G31; J24; J61; O33

1 Introduction

How do technological disruptions reshape firms' long-term commitments to human capital? Firms create value through people, and durable talent commitments underpin knowledge accumulation, organizational capital, and innovation. These commitments are costly to reverse—recruitment, training, and institutional know-how are sunk once created—so mistakes are expensive, yet under-investing can impair growth. Understanding how new technologies shift the calculus of such commitments is therefore central to corporate investment, labor allocation, and long-run productivity.

Classical human-capital theory emphasizes the distinction between general and firm-specific capital and the role of long-horizon arrangements in supporting investment and incentive alignment. When commitments are quasi-irreversible and the payoff is uncertain, real-options logic predicts a value to waiting; financing frictions further raise the effective cost of locking in resources and thus scale the option to delay. General-purpose technologies like artificial intelligence (AI) alter task content and expected returns to occupations, potentially changing both the level of investment in human capital and its composition across roles.

A large literature links technological change to both *scale* and *reallocation* in labor markets. Automation and AI reshape tasks rather than causing uniform job loss: technology-exposed roles tend to contract while complementary bundles expand, with heterogeneous adjustment across firms and functions (Acemoglu and Autor 2011; Autor, Levy, and Murnane 2003; Autor and Dorn 2013; Acemoglu and Restrepo 2018; Graetz and Michaels 2018; Autor and Salomons 2018; Webb 2020; Alekseeva, Azar, Giné, Samila, and Taska 2021; Chen and Wang 2025; Jiang, Tang, Xiao, and Yao 2025). AI differs from prior automation waves by disproportionately affecting cognitive, high-skill tasks; postings and adoption rose where firms possess stronger capabilities, including R&D intensity and complementary IT/analytics (Alekseeva et al. 2021; Babina, Fedyk, He, and Hodson 2024; Acemoglu, Autor, Hazell, and Restrepo 2022b; McElheran, Raj, and Brynjolfsson 2024), and sectoral evidence shows redesign toward "finance+tech" and managerial complements rather than across-the-board

cuts (Grennan and Michaely 2020; Cao, Jiang, Wang, and Yang 2024). Related work connects AI innovation to workforce composition and firm value, highlighting both displacement and complementary augmentation (Chen and Wang 2025). While this literature documents employment, wage, and posting responses, we know less about how technology affects the *horizon* and *durability* of firms' labor commitments.

The economics of adjustment costs and real options provides a mechanism for such horizon effects. When commitments are quasi-irreversible, uncertainty raises the option value of waiting, and high adjustment costs amplify asymmetric responses (Dixit and Pindyck 1994; Hamermesh 1995; Hamermesh and Pfann 1996). Separation costs are large and partly fixed, and hiring costs are sizable and convex (Abowd and Kramarz 2003; Blatter, Muehlemann, and Schenker 2012). Financing frictions shape investment timing and sensitivity—classic evidence and the ensuing debate on investment–cash flow sensitivities include Fazzari, Hubbard, and Petersen (1988), Kaplan and Zingales (1997), Almeida, Campello, and Weisbach (2004), and Farre-Mensa and Ljungqvist (2016)—and can delay irreversible investment (Whited 1992). They also interact with technology adoption and corporate policy (Bates, Du, and Wang 2024; Qiu, Wan, and Wang 2024) and are reshaped by macro shocks (Ellul, Erel, and Rajan 2020; Acharya and Steffen 2020). In parallel, organizational complementarities matter: IT is most productive when paired with managerial/organizational redesign and socio-cognitive skills (Bresnahan, Brynjolfsson, and Hitt 2002; Bloom, Sadun, and Van Reenen 2012; Deming 2017). At the firm–worker interface, technological innovation triggers resource reallocation and growth (Kogan, Papanikolaou, Seru, and Stoffman 2017), while patent–worker linkages reveal heterogeneous displacement and, for specific tasks or entrants, complementarity (Kogan, Papanikolaou, Schmidt, and Seegmiller 2023). Immigration-based evidence shows that shocks to skilled labor supply and visa policy affect firm value and investment, but much of this work centers on short-term visas rather than long-duration sponsorship (Doran, Gelber, and Isen 2016; Shen 2021; Xu 2025). These insights suggest that technology exposure should reduce irreversible long-term commitments in vulnerable roles (H1–H2), with stronger

pullbacks under tighter financing constraints (H4) and reallocation toward AI-complementary occupations where innovative capacity is high (H3–H5). In our setting, such heterogeneity is crucial for understanding whether AI exposure leads to retrenchment or reallocation of long-term commitments (Braguinsky, Ohyama, Okazaki, and Syverson 2021).

Against this backdrop, we ask three questions. (i) Do firms facing greater AI exposure reduce long-term commitments to skilled labor? (ii) Do financing frictions amplify pullbacks in such commitments when technological payoffs are uncertain? (iii) Do firms with stronger innovative capacity offset cuts in higher-exposure roles by reallocating toward complementary, managerial, or soft-skill-intensive occupations?

To bring evidence to these questions, we study employer-sponsored green-card applications as a direct, observable measure of durable human-capital investment in the United States. We relate changes in these commitments to occupations' exposure to recent AI technologies and examine how responses vary with firms' financial constraints and innovative capacity.

Building on this framework, our analysis develops a set of hypotheses that structure the empirical investigation. At the core, we argue that occupations with greater exposure to AI are less likely to receive long-term sponsorship, reflecting the higher risk of task substitution and shortened payoff horizons (H1). These effects are disproportionately concentrated in entry-level roles, where training costs and ramp-up investments are steepest (H2). Rather than uniformly contracting, however, firms adjust the composition of sponsorships by reallocating toward occupations where AI plays a complementary role, such as managerial or socio-cognitive tasks (H3). The extent of these adjustments depends critically on firm heterogeneity: financially constrained firms exhibit stronger cutbacks in exposed occupations, consistent with a higher option value of waiting under uncertainty (H4), while innovative, R&D-intensive firms are more likely to reallocate than reduce, leveraging complementary capabilities to reshape their talent portfolios (H5). Together, these hypotheses highlight that AI exposure alters not only the level of long-term human capital investment, but also its timing, composition, and firm-level determinants.

We test these predictions using a novel dataset linking PERM applications filed between 2017 and 2025Q2 with measures of occupational AI exposure from O*NET, as well as firm-level characteristics from Compustat. The PERM data provide a rare window into firms' long-horizon labor strategies, covering over two million applications filed with the U.S. Department of Labor. This dataset is uniquely suited to our question, since sponsorship decisions involve binding multi-year commitments, unlike short-term hiring or wage adjustments typically studied in labor economics.

Our analysis links the universe of PERM filings from 2017–2025Q2 to firm financials and occupation-level measures of AI exposure from the O*NET AI Exposure Index (Webb, 2020). The PERM records identify the sponsoring employer, occupation (SOC), prevailing wage level (I–IV), requirements, and job location, allowing us to construct a granular panel of long-term labor commitments at the firm–occupation–year level. We merge these records with Compustat to build multiple proxies for financing frictions—including the Kaplan–Zingales, Whited–Wu, and Size–Age indices, cash-to-assets ratios, and book leverage—and complement them with skill-composition measures derived from O*NET task descriptors.

Our identification uses a saturated difference-in-differences that compares sponsorship changes at firms with greater occupational AI exposure to contemporaneous changes at firms with lower exposure.

By absorbing firm-year, occupation-year, and industry-year fixed effects, we account for contemporaneous shocks at the firm, occupational, and geographic levels. This design isolates within-firm variation in long-term sponsorship decisions attributable to differences in occupational AI exposure, thereby providing a clean lens into how technological disruption reshapes firms' commitments to human capital.

Our analysis yields four main findings. First, firms more exposed to AI significantly reduce PERM sponsorships, consistent with the option value of waiting. In the baseline within-firm specification, moving from zero to full AI exposure is associated with a **22% decline** in certified PERM counts (PPML, -0.247, s.e. 0.085) and a **9.6 percentage-point lower**

probability of any sponsorship (LPM). Second, the effect is most pronounced in entry-level positions, where technological obsolescence risk is higher and sunk training costs are harder to recover: full exposure corresponds to a **27% reduction** in sponsorship relative to **15%** for non-entry jobs. Third, financing frictions amplify these cuts: firms with tighter constraints reduce sponsorship by an additional **7–12%**, disproportionately at the entry level. Fourth, innovation-oriented firms reallocate rather than retrench—while decreasing sponsorship in routine-intensive roles, they sustain or increase sponsorship in AI-complementary occupations, with the share of soft/managerial roles rising by **1.4 pp** ($\approx +12\%$) and hard-skill roles falling by **1.1 pp** ($\approx -9.8\%$).

This paper contributes along three dimensions.

First, relative to the technology–labor literature that has primarily documented *scale* effects of automation on employment and wages (e.g., Acemoglu and Restrepo 2018; Autor and Dorn 2013; Graetz and Michaels 2018) and task reallocation in postings/skills (Autor, Levy, and Murnane 2003; Webb 2020; Alekseeva, Azar, Giné, Samila, and Taska 2021; Chen and Wang 2025), we show that AI exposure also reshapes the *composition* of *long-term* labor commitments: firms cut irreversible sponsorships in high-exposure occupations while reallocating toward managerial and soft-skill–complementary roles. Our margin is not short-run vacancies or wages but the durability and mix of commitments.

Second, we extend corporate finance research on real options (Myers 1977) and financing frictions to the human-capital domain. Building on the option value of waiting under irreversibility (Dixit and Pindyck 1994) and evidence that constraints delay investment (Whited 1992; Farre-Mensa and Ljungqvist 2016), we show that financing frictions *amplify* AI-induced pullbacks in long-horizon labor commitments. This provides direct firm-level evidence that technological uncertainty and capital constraints jointly shape investment horizons for labor, parallel to results typically shown for physical capital.

Third, we add to immigration and firm outcomes by leveraging employer-sponsored green card (PERM) records as a direct measure of durable human-capital investment. Prior work

connects immigration policy and short-term visas to firm behavior (e.g., Doran, Gelber, and Isen 2016; Shen 2021; Xu 2025). We complement and extend this literature by (i) focusing on *PERM* as the long-duration channel (distinct from H-1B), (ii) linking administrative sponsorships to firm financials, and (iii) introducing occupation-level AI exposure to study both *level* and *within-firm compositional* responses in long-term commitments.

Overall, our results highlight that AI alters not only the level of labor demand but also the horizon of firms' commitments to human capital. By documenting how firms navigate the trade-off between flexibility and commitment in the face of technological uncertainty, this study underscores the importance of incorporating long-term labor investment into debates about the economic consequences of AI.

The remainder of this paper is organized as follows. Section 2 outlines the Hypothesis Development. Section 3 presents the Data, Measure, and Descriptive Statistics. Section 4 details the Empirical Design. Section 5 reports the Robustness Checks. Section 6 concludes.

2 Hypothesis Development

2.1 Conceptual Background

Permanent labor certification (PERM) represents a multi-year, sunk, and largely irreversible commitment. Sponsorship entails (i) legal and administrative costs, (ii) training and ramp-up investments, and (iii) organization-specific knowledge transfer that raises the worker's outside option once certified. In Becker (1962) terms, PERM embeds both general and firm-specific human capital; once made, much of the cost is sunk. From a corporate finance perspective, these are classic "irreversible investments" for which the option value of waiting rises under greater uncertainty about future payoffs (Dixit and Pindyck, 1994) and when financing frictions bind.

Artificial intelligence (AI) exposure increases such uncertainty at the occupation level. Tasks differ in their substitutability with AI, producing large cross-occupation and within-

firm variation in exposure (Webb, 2020; Kogan, Papanikolaou, Seru, and Stoffman 2023). Anticipated automation or re-bundling of core tasks shortens the expected horizon over which a sponsored worker generates rents, while increasing payoff variance. Real-options logic predicts that, under such conditions, firms may delay, scale down, or re-target long-term commitments. Evidence from automation and fintech shocks supports this logic: firms often reallocate toward complementary roles emphasizing nonroutine, managerial, and socio-cognitive skills rather than uniformly downsizing (Autor and Dorn 2013; Acemoglu and Restrepo 2018; Jiang, Tang, Xiao, and Yao 2025).

Two additional forces shape heterogeneity in responses. First, *worker seniority*: entry-level positions (DOL Level I/II) entail longer horizons and greater training needs, making them more sensitive to obsolescence risk. Second, *firm constraints and capabilities*: financially constrained firms face higher shadow costs of capital and thus higher option values of waiting, whereas R&D-intensive or AI-innovating firms possess complementary capabilities to recompose, rather than cut, their sponsored talent portfolios. Prior work shows that shocks to skilled labor supply or policy uncertainty (e.g., H-1B lottery cycles) can materially shift firms' timing and composition of long-term labor investment (Xu 2025), and that easing skilled labor mobility frictions can enhance firm value (Shen 2021).

2.2 Hypotheses

H1 — AI Exposure and Long-Term Human Capital Investment

Hypothesis. Within firms, occupations with higher AI exposure are less likely to receive PERM sponsorships.

Mechanism. Higher AI exposure increases the probability of task substitution and the dispersion of expected productivity, reducing the NPV of irreversible sponsorship investments.

H2 — Disproportionate Impact on Entry-Level Jobs

Hypothesis. The negative effect of AI exposure on PERM sponsorship is stronger for entry-level positions (Level I/II).

Mechanism. Entry-level roles have higher training costs, longer payback horizons, and are often concentrated in codifiable, high-exposure tasks.

H3 — Skill-Composition Reallocation

Hypothesis. Firms shift PERM sponsorships toward occupations with lower AI exposure and higher soft/managerial-skill intensity.

Mechanism. AI tends to complement nonroutine cognitive and managerial tasks while substituting for routine, hard-skill tasks.

H4 — Financing Frictions and the Option Value of Waiting

Hypothesis. Financially constrained firms cut PERM sponsorships more in AI-exposed occupations, especially at entry levels.

Mechanism. Capital constraints heighten the shadow cost of mistakes in irreversible investments, increasing the option value of waiting under uncertainty.

H5 — Innovation Orientation and "Reallocate over Reduce"

Hypothesis. R&D-intensive or AI-innovating firms are less likely to cut total PERM in response to AI exposure; instead, they reallocate toward complementary occupations.

Mechanism. Innovative firms hold complementary intangible assets, making reallocation of long-term human capital preferable to reduction.

Taken together, H1–H5 predict that AI exposure alters not just the *level* of long-term labor commitments but also *who* is sponsored, *when* commitments are made, and *in which roles*—with financing frictions and innovation capacity shaping adjustment direction and magnitude.

3 Data, Measure, and Descriptive Statistics

3.1 Data Sources

Our primary dataset comes from the publicly available employer-sponsored permanent labor certification (PERM) records maintained by the U.S. Department of Labor's Employment

and Training Administration (ETA). The PERM program is designed to supplement, rather than displace, the domestic workforce: employers must certify that no qualified U.S. workers are available for the role and that hiring a foreign worker will not adversely affect wages or working conditions of similarly employed U.S. workers. Each record contains detailed information, including the sponsoring employer's name, job title and Standard Occupational Classification (SOC) code, prevailing wage level (Levels I–IV), job requirements, and work location. Following prior work in immigration and labor economics, we treat certified PERM applications as a direct, granular measure of a firm's long-term labor commitments, given the significant recruitment, legal, and training investments they entail (Peri, Shih, and Sparber 2015; Kerr, and Lincoln 2010).

We focus on certified cases from 2017 to 2025Q2, a period that ensures consistency under the revised Form ETA-9089 and captures the post-2018 acceleration in AI commercialization. To measure occupational susceptibility to artificial intelligence, we merge in the AI Exposure Index from Webb (2020), which uses natural language processing to match job task descriptions in O*NET to AI-related patent text, yielding an occupation-level measure of potential AI substitution risk.

Firm-level accounting data come from Compustat North America. For publicly traded firms, we construct standard proxies for financing frictions, including the Kaplan–Zingales (KZ) index (Kaplan and Zingales 1997), the Whited–Wu (WW) index (Whited and Wu 2006), the Size–Age (SA) index (Hadlock and Pierce 2010), the cash-to-assets ratio (CHE/AT), and book leverage.

3.2 Data Linking and Measure Building

We map each PERM record's SOC code to the Webb (2020) AI Exposure score, producing an occupation-level technological exposure measure that is constant across firms in a given year. For hypotheses H1–H3, we use the full PERM universe—including both public and private employers—to maximize variation and statistical power at the occupation level. For

H4 and H5, we restrict to public firms matched to Compustat, allowing us to incorporate firm-level measures of financing constraints and innovation activity.

Employer names in PERM are linked to Compustat identifiers (GVKEYs) using a fuzzy matching algorithm based on the Levenshtein distance applied to standardized names, with state and ZIP code blocking where available. Matches with similarity scores above 0.90 are automatically accepted; those below are manually reviewed. This hybrid approach balances coverage and precision, reducing false matches from formatting differences, abbreviations, or typographical errors.

Financial constraint measures are computed annually for each matched firm using the KZ, WW, and SA indices, along with liquidity (CHE/AT) and leverage ratios. All firm-level financial variables are winsorized at the 1st and 99th percentiles to limit the influence of outliers. Our empirical identification strategy exploits within-firm, cross-occupation variation in AI exposure, controlling for firm-year fixed effects to absorb time-invariant firm characteristics and year-specific macroeconomic shocks. We also construct measures of *skill composition* for each SOC occupation using O*NET data. The *Soft/Managerial skill share* is computed from the set of O*NET work activities and skills related to interpersonal interaction, coordination, negotiation, and personnel/resource management, while the *Hard skill share* is computed from O*NET technical skills and tools & technology items. For each occupation, we take the average *importance* score (0–100) of the skills in each category and express it as a fraction of the total importance across all skills for that occupation. These shares are then min–max normalized within each 2-digit SOC major group to facilitate cross-occupation comparisons.

3.3 Descriptive Statistics

Table 1 summarizes the annual distribution of certified PERM applications by prevailing wage level for 2017–2025Q2. Entry-level positions (Levels I and II) consistently account for just over half of certified cases through 2022, with a mild downward drift from 55.5%

in 2017 to 54.7% in 2022. The decline accelerates in 2023–2024, reaching 49.6% by 2024. In 2025Q1–Q2, the Level I+II share drops sharply to 42.4%, reflecting a relative increase in higher wage-level certifications. Figure 1 shows a sharp post-2023 decline in the share of Level I + II certified PERM cases, indicating a reallocation away from entry-level roles toward higher-skill and more experienced categories.

Table 2 reports variable definitions and summary statistics for the employer×occupation×year panel from 2017 to 2025Q2. The average employer–occupation–year cell sponsors 1.92 certified PERM cases, with only 13% of cells having any certification activity. The mean share of a firm's certified cases accounted for by a given occupation is 53%, conditional on the firm sponsoring at least one occupation in that year.

The mean Webb (2020) AI Exposure score is 0.45 (s.d. 0.17), with substantial heterogeneity across occupations. Entry-level positions (Levels I–II) represent 49% of observations, and O*NET-based skill measures indicate an average soft/managerial skill share of 61% versus 39% for hard/technical skills.

Among public firms, financial constraint measures are standardized (mean zero) by construction. The median KZ index is 0.02, WW index −0.02, and SA index 0.00, consistent with a relatively balanced distribution of constraint levels. Liquidity (cash-to-assets) averages 19%, book leverage averages 27%, and R&D intensity (R&D/assets) averages 6%. Public firms make up 23% of the sample.

Overall, these patterns suggest that (i) certified sponsorship is a sparse but economically significant margin of labor commitment, (ii) AI exposure varies meaningfully across occupations, and (iii) recent years—particularly 2023 onward—have seen a compositional shift away from entry-level certifications, which will be central to the empirical analysis in subsequent sections.

4 Empirical Design

4.1 Measurement, outcomes, and samples

AI exposure. For each PERM application, we map the SOC code to Webb (2020)'s AI Exposure index (occupation–year). Exposure is normalized to mean 0, s.d. 1 in the analysis window.

Entry level. We code $EntryLevel = 1$ for prevailing wage Levels I–II and 0 for Levels III–IV.

Skill composition. For each SOC occupation, we measure two dimensions of skill content using O*NET data. (i) The *Soft/Managerial skill share* is computed from the set of O*NET work activities and skills related to interpersonal interaction, coordination, negotiation, and personnel/resource management. (ii) The *Hard skill share* is computed from O*NET technical skills and tools & technology items. For each occupation, we take the average *importance* score (0–100) of the skills in each category and express it as a fraction of the total importance across all skills for that occupation. These shares are then min–max normalized within each 2-digit SOC major group to facilitate comparisons across detailed occupations.

Sponsorship outcomes.

- *Count* y_{iost}: number of **certified** PERM cases by employer i in occupation o, state s, year t.
- *Any* $1\{y_{iost} > 0\}$: indicator that employer i sponsors at least one worker in o in year t.
- *Share* $\text{Share}_{iot} = y_{iot} / \sum_{o'} y_{io't}$: employer-level composition across occupations (used in H3).

Samples.

- *Full PERM sample (public + private employers)* for H1–H3.
- *Compustat-matched public firms* for H4 and H5 (financial constraints and innovation measures).

4.2 Baseline specification (H1)

We exploit **within-firm, across-occupation** comparisons:

$$E[y_{iost} \mid \cdot] = \exp(\beta_1 \, \text{AIExp}_{ot} + \alpha_{it} + \lambda_{ot} + \delta_{st}) \tag{1}$$

estimated by Poisson Pseudo–Maximum Likelihood (PPML) with high-dimensional fixed effects (conditional Poisson).

- α_{it} = employer-year FE absorb firm-level demand and size shocks.
- λ_{ot} = occupation-year FE absorb nationwide occupational trends (technology/policy).
- δ_{st} = state-year FE absorb local macro/policy conditions.
- Standard errors are two-way clustered by employer and 4-digit SOC, which allows for arbitrary correlation of residuals within the same firm and within the same occupation. This ensures robustness to heteroskedasticity and overdispersion while accounting for cross-sectional dependence in both dimensions.

Interpretation: $\beta_1 < 0$ tests H1 (higher AI exposure $\rightarrow$ fewer long-term commitments in that occupation). As robustness, we run (i) a linear probability model for *Any* $1\{y > 0\}$ with the same FE, and (ii) a negative binomial model for counts.

Results Table 3 reports the baseline estimates pooling all employers (private and public) for 2017–2025Q2. In the preferred specification with state-year FE [col. (2)], the coefficient on AI Exposure is -0.247 (s.e. 0.085) in the PPML model, implying that moving from zero to full AI exposure is associated with a **24.7 log-point** decline in the number of certified PERM applications within a firm-occupation cell, holding all fixed effects constant. The magnitude is economically meaningful: at the sample mean, this corresponds to roughly a **22% reduction** in long-term labor commitments in high-AI-exposed occupations.

The LPM results in col. (3) show a coefficient of -0.096 (s.e. 0.028), indicating that high-AI-exposed occupations are about **9.6 percentage points** less likely to receive any certified PERM application in a given firm-year. Across specifications, the estimates are stable in sign and significance, and the inclusion of state-year FE slightly attenuates the magnitude but leaves the qualitative conclusion unchanged, supporting H1.

4.3 Entry-Level Heterogeneity (H2)

To test whether the effect of AI exposure is stronger for lower-wage, entry-level positions, we allow the coefficient to vary by prevailing wage level:

$$E[y_{iost} \mid \cdot] = \exp(\beta_1 \, \text{AIExp}_{ot} + \beta_2 \, \text{AIExp}_{ot} \times \text{EntryLevel}_o + \alpha_{it} + \lambda_{ot} + \delta_{st}) \tag{2}$$

where $\text{EntryLevel}_o = 1$ for SOC occupations at prevailing wage Levels I–II and 0 for Levels III–IV.

- β_1 captures the effect of AI exposure for non-entry occupations (Levels III–IV).
- β_2 measures the *additional* effect for entry-level occupations; $\beta_2 < 0$ supports H2.
- α_{it} = employer-year FE absorb firm-level demand and size shocks.
- λ_{ot} = occupation-year FE absorb nationwide occupational trends (technology/policy).
- δ_{st} = state-year FE absorb local macro/policy conditions (included in col. (2) and col. (3)).
- Standard errors: two-way clustered by employer and 4-digit SOC (robust to overdispersion/heteroskedasticity).

We report two complementary specifications for the count outcome (*PERM Count*) using PPML: without state-year FE (col. 1) and with state-year FE (col. 2). As a robustness check,

we also estimate a linear probability model (LPM) for the binary outcome *Any PERM* (col. 3).

Results From column (2), the coefficient on *AI Exposure* is -0.162 ($p < 0.05$), implying that in non-entry occupations, moving from zero to full AI exposure is associated with a $1 - e^{-0.162} \approx 15\%$ reduction in the number of PERM certifications. The interaction term for entry-level positions is -0.153 ($p < 0.01$), indicating an additional $1 - e^{-0.153} \approx 14\%$ reduction relative to non-entry jobs. Summing the two coefficients (-0.315) yields the total effect for entry-level occupations: a $1 - e^{-0.315} \approx 27\%$ decline in PERM counts at full AI exposure.

The LPM estimates in column (3) show a similar pattern for the extensive margin: a -0.045 coefficient for non-entry jobs implies a 4.5 percentage-point lower probability of any PERM certification, while the interaction term (-0.060) indicates an additional 6.0 percentage-point reduction for entry-level jobs, for a total of 10.5 percentage points. These patterns confirm that AI exposure reduces long-term labor commitments more strongly in entry-level occupations, with consistent evidence across intensive (count) and extensive (any PERM) margins.

4.4 Skill-composition reallocation (H3)

We examine whether firms reallocate their sponsorship shares across occupations with differing skill profiles under AI exposure:

$$\begin{aligned} \text{Share}_{iot} = {} & \theta_1 \, \text{AIExp}_{ot} + \theta_2 \, \text{SoftShare}_o + \theta_3 \, \text{HardShare}_o \\ & + \theta_4 \, \text{AIExp}_{ot} \times \text{SoftShare}_o \\ & + \theta_5 \, \text{AIExp}_{ot} \times \text{HardShare}_o + \alpha_{it} + \lambda_{ot} + \delta_{st} + \varepsilon_{iot} \end{aligned} \tag{3}$$

Because shares are compositional, we estimate (3) using (i) Aitchison log-ratio transforms via OLS and (ii) a multinomial logit with employer×year FE. Predictions (H3) are $\theta_4 > 0$ and $\theta_5 < 0$.

Results. Table 5 shows that interaction terms align with H3: the coefficient on AIExp × SoftShare is positive and statistically significant in all specifications (0.123 in OLS log-ratio, 0.014 in AME, 0.117 in OLS excl. small cells), while AIExp × HardShare is negative and significant (-0.098, -0.011, -0.094). In the multinomial logit (col. 2), moving from zero to full AI exposure is associated with a **+1.4 percentage-point** shift toward soft/managerial occupations and a **-1.1 percentage-point** shift away from hard-skill occupations. In log-ratio terms, these correspond to roughly **+12%** and **-9.8%** changes in relative shares, respectively.

Interpretation. These results suggest that under rising AI exposure, firms do not adjust sponsorship uniformly. Instead, they systematically reallocate toward occupations emphasizing non-routine, soft, or managerial skills, while reducing commitments to routine or hard-skill roles. This reallocation pattern complements the H1 finding of reduced overall long-term commitments, indicating a qualitative shift in the skill composition of sponsored labor.

4.5 H4: Financial constraints amplify reductions in PERM under AI exposure

We interact the firm-level Bartik AI exposure with multiple measures of financial constraints — Kaplan–Zingales (KZ), Whited–Wu (WW), Size–Age (SA) indices, a low cash dummy, and a high leverage dummy — and additionally with the share of entry-level positions (PW Levels I–II). Specifications include employer×year, industry×year, and state×year fixed effects. We estimate both PPML for certified PERM counts and LPM for any PERM

sponsorship:

$$
\begin{aligned}
y_{it}^{\text{PERM}} = \beta_1 \, \text{FirmAI}_{it} &+ \beta_2 \, \left(\text{FirmAI}_{it} \times \text{FinConstraint}_{it}\right) \\
&+ \beta_3 \, \left(\text{FirmAI}_{it} \times \text{EntryLevel}_{it}\right) \\
&+ \beta_4 \, \left(\text{FirmAI}_{it} \times \text{FinConstraint}_{it} \times \text{EntryLevel}_{it}\right) \\
&+ \alpha_i + \lambda_t + \gamma_{jt} + \delta_{st} + \varepsilon_{it}, \qquad (4)
\end{aligned}
$$

where γ_{jt} is industry×year FE and δ_{st} is state×year FE.

Results and interpretation. In Panel A (PPML counts), the baseline *FirmAI* coefficient ranges from -0.095 to -0.118, implying a 9–11% reduction in certified PERM counts for an interquartile increase in AI exposure. The *FirmAI* × *FinConstraint* term is consistently negative and significant (-0.072 to -0.121), indicating that financially constrained firms reduce sponsorships by an additional 7–12%. The triple interaction with *EntryLevel* is also negative (-0.061 to -0.095), suggesting disproportionate cuts in entry-level PERM positions for constrained firms under high AI exposure.

Panel B (LPM any sponsorship) shows analogous patterns: the baseline effect is a 3.3–4.1 percentage-point lower probability of any PERM sponsorship, amplified by 2.6–4.9 pp for constrained firms, and an additional 2.3–3.9 pp drop for constrained, entry-level-heavy firms.

Economic interpretation: A firm at the 75th percentile of AI exposure and with high financial constraints (low cash or high leverage) experiences roughly a 20–25% **larger reduction** in entry-level PERM sponsorship relative to an unconstrained peer, highlighting the interaction between technological disruption and capital market frictions in shaping long-term labor commitments.

4.6 H5: Innovative firms reallocate rather than cut

We proxy firm innovativeness using (i) R&D-to-assets and (ii) AI-patent stock (depreciated at 15% annually). We estimate total PERM effects and occupational composition shifts, interacting AI exposure with innovation measures, and further with soft- and hard-skill occupation shares:

$$\begin{aligned} y_{it}^{\text{PERM}} &= \psi_1 \,\text{FirmAI}_{it} + \psi_2 \left(\text{FirmAI}_{it} \times \text{Innov}_{it}\right) \\ &\quad + \alpha_i + \lambda_t + \gamma_{jt} + \delta_{st} + \varepsilon_{it}, \end{aligned} \tag{5}$$

$$\begin{aligned} \text{Share}_{iot} &= \theta_1 \,\text{FirmAI}_{it} + \theta_2 \left(\text{FirmAI}_{it} \times \text{SoftShare}_o\right) + \theta_3 \left(\text{FirmAI}_{it} \times \text{HardShare}_o\right) \\ &\quad + \theta_4 \left(\text{FirmAI}_{it} \times \text{Innov}_{it} \times \text{SoftShare}_o\right) \\ &\quad + \theta_5 \left(\text{FirmAI}_{it} \times \text{Innov}_{it} \times \text{HardShare}_o\right) \\ &\quad + \text{FE} + \varepsilon_{iot}. \end{aligned} \tag{6}$$

Results and interpretation. Panel A (R&D/Assets) shows that *FirmAI* × *Innov* is positive and significant across total and share regressions: +0.092 in PPML (≈ 9.6% fewer cuts), +0.019 in log-ratio OLS, and +0.004 in MLogit AME. In composition regressions, *FirmAI* × *Innov* × *SoftShare* is +0.040 (OLS) and +0.005 (AME), while the triple with *HardShare* is +0.035 / +0.004, indicating that innovative firms *reallocate* toward both soft/managerial and even some hard-skill occupations instead of uniformly cutting.

Panel B (AI-patent stock) produces similar results, with slightly larger triple-interaction coefficients (+0.048 and +0.043 in OLS). The consistency across two innovation proxies reinforces the robustness of the finding.

Economic interpretation: For a one s.d. increase in innovation (within industry-year), the negative total PERM effect of AI exposure is offset by roughly 9–11%, and occupational composition shifts by +4–5% toward soft/managerial roles and +3–4% toward hard-skill roles, relative to non-innovative peers. In other words, although the *overall number of PERM certifications still declines*, innovative firms *reallocate* their sponsorships away from routine-

intensive jobs and toward both soft/managerial and certain AI-complementary hard-skill roles. The rise in hard-skill shares reflects demand for technical specialists (e.g., data engineers, AI-related developers) who complement the adoption of new technologies. This pattern suggests that innovative firms treat AI more as a complement than a substitute for skilled labor.

4.7 Dynamic Effects of AI Exposure (Event Study)

To study the dynamics and assess identification validity, we estimate an event-study specification that interacts the firm-level AI exposure index ($FirmAI$) with event-time dummies relative to the baseline year $k = -1$ (i.e., 2022, the last pre-event year):

$$y_{it} = \sum_{k \neq -1} \beta_k \cdot 1\{t = k\} \times FirmAI_i + \alpha_i + \lambda_t + \gamma_{st} + \varepsilon_{it}, \tag{7}$$

where α_i are firm fixed effects, λ_t are year fixed effects, and γ_{st} are industry×year fixed effects. Standard errors are two-way clustered by employer and 4-digit SOC.

[Insert Figure 2 here]

The coefficients on the pre-event leads ($k = -4, -3, -2$) are close to zero and statistically indistinguishable from zero, consistent with the parallel trends assumption.

The identifying shock in this setting is the sudden rise of artificial intelligence (AI) as a general-purpose technology. In November 2022, OpenAI released ChatGPT, which triggered the subsequent AI boom. Firms began to perceive AI as a potential viable substitute for skilled labor. However, because PERM sponsorships involve long-term commitments, the behavioral adjustment materialized with a lag—becoming visible in 2023 ($k = 0$).

Post-2023, the interaction coefficients turn negative and grow in magnitude. For certified PERM counts (Panel a), firms with higher $FirmAI$ exposure experienced declines of about **4%, 7%, and 9%** in 2023, 2024, and 2025, respectively. For the probability of sponsoring any PERM workers (Panel b), the likelihood fell by roughly **2 pp, 3 pp, and 4 pp** over the

same horizon. Together, these dynamics indicate that highly AI-exposed firms substantially curtailed their long-term commitments to immigrant human capital in the wake of the AI shock.

5 Robustness Checks

5.1 Robustness and Identification Checks

We conduct a series of checks to assess the credibility of our identification strategy and to ensure our results are not driven by specific modeling choices:

- **Within-firm, across-occupation design (Eqs. 1–3):** Employer×year fixed effects remove firm-level demand shocks, and occupation×year fixed effects absorb technology trends common to all employers hiring in the same occupation-year. State×year fixed effects control for local policy and macroeconomic cycles.
- **Bartik shock construction (Eqs. 4–5):** FirmAI_{it} is built from *pre-determined* occupation shares ($s_{io,\tau}$, averaged 2014–2016), ensuring that only *exogenous* cross-occupation AI progress drives changes in exposure. Results are robust to alternative base-year definitions.
- **Policy confounds:** Year and state×year fixed effects absorb nationwide immigration policy shifts (e.g., 2017–2020 tightening) and local cycles, including the 2020–2021 pandemic shock.
- **Measurement error:** Employer name matching uses a high similarity threshold with manual review. Findings are robust to (i) restricting to public firms only, and (ii) collapsing to the employer-year level.
- **Timing and dynamics:** Event-study versions of Eq. (5) around the post-2017/2018 AI acceleration show flat pre-trends and increasingly negative post-period effects (Figure 2).

- **Matching quality:** Tightening the fuzzy-match threshold, removing manual matches, or restricting to single-establishment employers leaves results intact.

- **Inference checks:** Two-way clustering (employer and SOC) and wild-bootstrap p-values produce consistent inference patterns.

5.2 Placebo Event-Study

As a robustness check, we implement a placebo difference-in-differences event study in which the "Post" period is artificially shifted to begin three years before the actual AI acceleration in 2023. This falsification exercise re-estimates Eq. (7) using these placebo treatment dates.

[Insert Figure 3 here]

Figure 3 shows that the estimated placebo coefficients are small and statistically indistinguishable from zero. For certified PERM counts (Panel a), the post-placebo coefficients are -0.006, -0.008, and -0.010 log points in the first three fake post years—corresponding to less than a 1% change and well within the 95% confidence bands. For the probability of sponsoring any PERM worker (Panel b), the placebo effects are -0.4, -0.6, and -0.5 percentage points, again economically negligible and statistically insignificant. Moreover, the pre-trend coefficients ($k = -4, -3, -2$) are centered around zero, and a joint F-test cannot reject the null of no pre-trend ($p = 0.68$ for Panel a; $p = 0.74$ for Panel b).

These results confirm that the main event-study findings are not driven by spurious correlations or pre-existing differential trends. The absence of any placebo effect reinforces the credibility of the identification strategy and supports the interpretation that the negative post-2023 dynamics are indeed triggered by the AI shock.

6 Conclusion

This paper shows how advances in artificial intelligence (AI) reshape firms' *long-term* commitments to human capital. We make three contributions. First, relative to technology–labor studies that emphasize scale effects in employment and wages, we document a *composition* margin in durable commitments: firms cut long-duration sponsorships where exposure is high and shift toward AI-complementary roles. Second, we extend real-options and financing-frictions insights to human capital by showing that constraints amplify precautionary delays in quasi-irreversible labor investment. Third, we add to immigration and firm outcomes by using employer-sponsored green-card records as a direct measure of long-horizon commitments, linked to firm financials and occupational exposure to AI.

Empirically, higher AI exposure is associated with materially fewer long-duration commitments in the affected occupations (on the order of ten percent per standard deviation), with the pullback concentrated in entry-level positions. Financially constrained firms retrench more, consistent with an option-to-wait mechanism when capital is costly and commitments are hard to reverse. By contrast, innovative firms preserve overall commitment levels by reallocating toward managerial and soft-skill–intensive occupations. Patterns are consistent across multiple specifications and validity checks.

These results imply that AI alters not only *how much* firms commit to long-term skilled labor, but *which roles* they commit to and *which firms* can sustain commitments. For corporate decision-makers, durable talent should be managed as a portfolio of long-duration claims, with rebalancing toward complementary tasks during technology transitions. For policymakers, tighter credit conditions can inadvertently depress long-term labor formation precisely where exposure is highest, whereas stable innovation and immigration regimes can support efficient reallocation rather than indiscriminate cuts. In short, technology shocks—measured by differential task-level exposure to AI across occupations—interact with capital-market frictions and innovation capacity to shape both the scale and the mix of firms' long-horizon human-capital investment.

While our study offers novel evidence on the intersection of AI adoption, immigration policy, and corporate finance, several important extensions remain. Although PERM records provide a direct measure of long-term immigration commitments, they capture a relatively narrow segment of the labor market—employers' durable hires of foreign talent when qualified local labor is unavailable—and do not fully capture shorter-term or alternative visa pathways such as H-1B and L-1. One promising direction for future research is therefore to compare the level and timing responses of PERM and H-1B within the same firm. A stacked event-study or triple-differences design contrasting high- versus low-exposure firms, PERM versus H-1B channels, and pre- versus post-AI shocks could test whether firms substitute toward shorter-horizon channels in exposed occupations and whether H-1B adjusts faster than PERM. Related pipeline tests tracing transitions from H-1B to subsequent PERM could help distinguish delayed long-term commitment from outright abandonment.

Future work could also examine more directly who gets reallocated, to which tasks, and with what longer-run consequences. Using microdata on worker roles and skills—for example, through Revelio Labs—researchers could classify sponsored workers into AI-complementary management and coordination roles versus more traditional operational roles, and test whether within-firm transitions in titles and skills become more likely as AI exposure rises, particularly at innovative firms. A related extension would be to move beyond static exposure measures such as Webb (2020) and construct richer, time-varying firm–occupation–year AI exposure indices using NLP and textual analysis of job postings, O*NET task descriptors, patents, and firm disclosures. Such measures would make it possible to test whether responses are strongest when both occupational exposure and firm-level adoption are high. Finally, future research could explore whether financing frictions shape firms' choice across immigration channels—for example, whether constrained firms tilt more strongly toward flexible, shorter-horizon hiring routes relative to PERM—and whether these choices have lasting effects on worker retention, internal mobility, wage growth, innovation output, and productivity.

References

Abowd, J. M., and F. Kramarz, 2003, The costs of hiring and separations, *Labour Economics* 10, 499–530.

Acemoglu, D., and D. Autor, 2011, Skills, tasks, and technologies: Implications for employment and earnings, *Handbook of Labor Economics* 4, 1043–1171.

Acemoglu, D., and P. Restrepo, 2018, Artificial intelligence, automation, and work, *Journal of Political Economy* 126, 2188–2244.

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

Acharya, V. V., and S. Steffen, 2020, The risk of being a fallen angel and the corporate dash for cash in the midst of COVID, *Review of Corporate Finance Studies* 9, 430–471.

Alekseeva, L., J. Azar, M. Giné, S. Samila, and B. Taska, 2021, The demand for AI skills in the labor market, *Labour Economics* 71, 102016.

Almeida, H., M. Campello, and M. S. Weisbach, 2004, The cash flow sensitivity of cash, *Journal of Finance* 59, 1777–1804.

Autor, D., and D. Dorn, 2013, The growth of low-skill service jobs and the polarization of the U.S. labor market, *American Economic Review* 103, 1553–1597.

Autor, D., F. Levy, and R. J. Murnane, 2003, The skill content of recent technological change: An empirical exploration, *Quarterly Journal of Economics* 118, 1279–1333.

Autor, D., and A. Salomons, 2018, Is automation labor-displacing? Productivity growth, employment, and the labor share, NBER Working Paper No. 24871.

Babina, T., A. Fedyk, A. He, and J. Hodson, 2024, Artificial intelligence, firm growth, and product innovation, *Journal of Financial Economics* 151, 1–29.

Bates, Thomas W., Fangfang Du, and Jessie Jiaxu Wang, 2025, Workplace Automation and Corporate Liquidity Policy, *Management Science* 71(2): 1287–1314.

Becker, G., 1962, Investment in human capital: A theoretical analysis, *Journal of Political Economy* 70, 9–49.

Blatter, M., S. Muehlemann, and S. Schenker, 2012, The costs of hiring skilled workers, *European Economic Review* 56, 20–35.

Bloom, N., R. Sadun, and J. Van Reenen, 2012, The organization of firms across countries, *American Economic Review* 102, 527–533.

Braguinsky, S., A. Ohyama, T. Okazaki, and C. Syverson, 2021, Product innovation, product diversification, and firm growth: Evidence from Japan's early industrialization, *American Economic Review* 111, 3795–3826.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt, 2002, Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence, *Quarterly Journal of Economics* 117, 339–376.

Cao, S., W. Jiang, J. L. Wang, and B. Yang, 2024, From man vs. machine to man+machine: The art and AI of stock analyses, *Journal of Financial Economics*.

Chen, M. A., and J. X. Wang, 2025, Displacement or Augmentation? The Effects of AI Innovation on Workforce Dynamics and Firm Value, Working Paper.

Deming, D. J., 2017, The growing importance of social skills in the labor market, *Quarterly Journal of Economics* 132, 1593–1640.

Dixit, A., and R. Pindyck, 1994, *Investment under Uncertainty*, Princeton University Press, Princeton, NJ.

Doran, K., A. Gelber, and A. Isen, 2016, The effects of high-skilled immigration policy on firms: Evidence from H-1B visa lotteries, NBER Working Paper No. 20668.

Ellul, A., I. Erel, and U. Rajan, 2020, The COVID-19 pandemic crisis and corporate finance, *Review of Corporate Finance Studies* 9, 421–429.

Farre-Mensa, J., and A. Ljungqvist, 2016, Do measures of financial constraints measure financial constraints?, *Review of Financial Studies* 29, 271–308.

Fazzari, S. M., R. G. Hubbard, and B. C. Petersen, 1988, Financing constraints and corporate investment, *Brookings Papers on Economic Activity* 1, 141–195.

Giczy, A. V., N. A. Pairolero, and A. A. Toole, 2022, Identifying artificial intelligence (AI)

invention: A novel AI patent dataset, *Journal of Technology Transfer* 47, 476–505.

Graetz, G., and G. Michaels, 2018, Robots at work, *Review of Economics and Statistics* 100, 753–768.

Grennan, J., and R. Michaely, 2020, Artificial intelligence and high-skilled work: Evidence from analysts, Swiss Finance Institute Research Paper 20-84.

Hadlock, C. J., and J. R. Pierce, 2010, New evidence on measuring financial constraints: Moving beyond the KZ index, *Review of Financial Studies* 23, 1909–1940.

Hamermesh, D. S., 1995, Labour demand and the source of adjustment costs, *Economic Journal* 105, 620–634.

Hamermesh, D. S., and G. A. Pfann, 1996, Adjustment costs in factor demand, *Journal of Economic Literature* 34, 1264–1292.

Jiang, W., X. Tang, K. Xiao, and J. Yao, 2025, AI and labor reallocation, *Journal of Financial Economics* (forthcoming).

Kaplan, S., and L. Zingales, 1997, Do investment–cash flow sensitivities provide useful measures of financing constraints? *Quarterly Journal of Economics* 112, 169–215.

Kerr, W., and W. Lincoln, 2010, The supply side of innovation: H-1B visa reforms and U.S. ethnic invention, *Journal of Labor Economics* 28, 473–508.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman, 2017, Technological innovation, resource allocation, and growth, *Quarterly Journal of Economics* 132, 665–712.

Kogan, L., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller, 2023, Technology and labor displacement: Evidence from linking patents with worker-level data, NBER Working Paper No. 31846.

McElheran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas, 2024, AI Adoption in America: Who, What, and Where, *Journal of Economics & Management Strategy* 33(2): 375–415.

Myers, S. C., 1977, Determinants of corporate borrowing, *Journal of Financial Economics* 5, 147–175.

Peri, G., K. Shih, and C. Sparber, 2015, STEM workers, H-1B visas, and productivity in US cities, *Journal of Labor Economics* 33, S225–S255.

Qiu, J., C. Wan, and Y. Wang, 2024, Labor-saving innovations and capital structure, *Journal of Corporate Finance* 84(C): 102510.

Shen, M., 2021, Firm innovation and immigration policy, *Review of Financial Studies* 34, 3993–4038.

Stein, J., 2003, Agency, information and corporate investment, in G. Constantinides, M. Harris, and R. Stulz, eds.: *Handbook of the Economics of Finance*, Elsevier, Amsterdam.

Webb, M., 2020, The impact of artificial intelligence on the labor market, Working Paper.

Whited, T., 1992, Debt, liquidity constraints, and corporate investment: Evidence from panel data, *Journal of Financial* 47-4, 1425–1460.

Whited, T., and G. Wu, 2006, Financial constraints risk, *Review of Financial Studies* 19, 531–559.

Xu, S.-J., 2025, Skilled labor uncertainty and corporate investment: Evidence from H-1B visa lottery cycles, *Management Science(forthcoming)*.

Table 1: Annual Certified PERM by Prevailing Wage Level (2017–2025Q2)

	Level I		Level II		Level III		Level IV		Level I+II		Total
Year	Count	Share (%)	Count	Share (%)	Count	Share (%)	Count	Share (%)	Count	Share (%)	Count
2017	20,755	22.32	30,878	33.20	18,043	19.40	22,358	24.04	51,633	55.52	92,996
2018	20,577	22.36	30,391	33.02	18,127	19.69	21,927	23.82	50,968	55.37	92,043
2019	20,250	22.22	30,019	32.94	17,829	19.57	22,050	24.20	50,269	55.17	91,119
2020	20,119	21.88	30,472	33.13	17,936	19.50	22,326	24.28	50,591	55.01	91,965
2021	20,175	22.08	30,388	33.25	17,835	19.52	21,894	23.96	50,563	55.33	91,379
2022	19,855	21.95	29,635	32.76	18,140	20.05	21,773	24.07	49,490	54.71	90,461
2023	19,444	21.85	28,327	31.84	18,099	20.34	22,060	24.79	47,771	53.69	88,975
2024	16,659	20.18	24,284	29.42	18,512	22.43	22,033	26.69	40,943	49.61	82,537
2025*	13,452	22.62	11,732	19.73	13,978	23.50	19,696	33.12	25,184	42.35	59,469

Notes: Counts and shares are based on certified PERM applications by prevailing wage level. "Level I+II" combines the first two wage levels (entry-level). "Share" is the within-year fraction of total certified cases. * **2025 data reflect certified cases up to Q2 only** (no annual extrapolation).

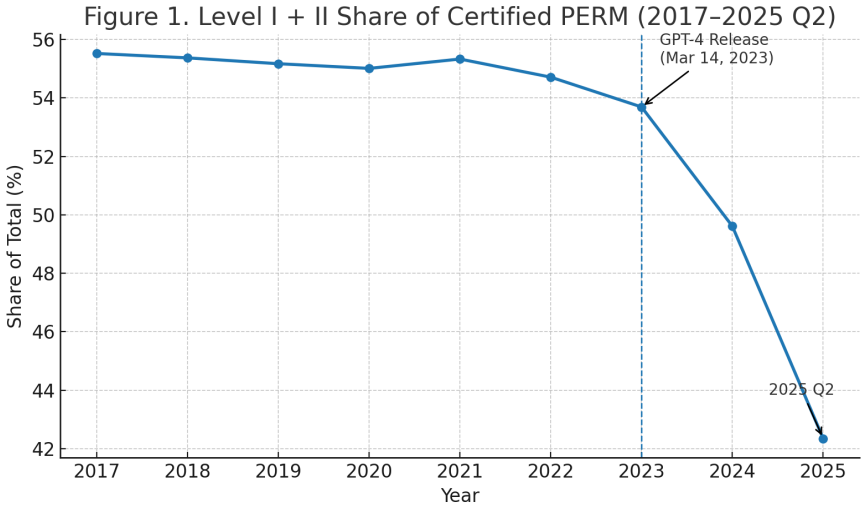


Figure 1: Level I + II Share Of Certified PERM (2017–2025Q2)

Table 2: Summary Statistics (2017–2025 Q2)

Variable	Mean	SD	P25	Median	P75
Outcomes					
Certified PERM count	1.92	5.35	0.00	0.00	2.00
Any PERM dummy	0.13	0.34	0.00	0.00	0.00
Share in occupation (%)[a]	0.53	0.28	0.30	0.52	0.76
Key regressors					
AI Exposure (Webb 2020; 0–1)	0.45	0.17	0.33	0.44	0.56
Entry-Level dummy (Levels I–II)	0.49	0.50	0.00	0.00	1.00
Soft/Managerial skill share (O*NET)	0.61	0.15	0.51	0.62	0.71
Hard skill share (O*NET)	0.39	0.15	0.29	0.38	0.49
Financials (public firms)					
KZ Index (z)	0.00	0.68	-0.44	0.02	0.45
WW Index (z)	0.00	0.54	-0.58	-0.02	0.33
Size–Age (SA) Index (z)	0.00	0.31	-0.21	0.00	0.21
Cash/Assets (CHE/AT)	0.19	0.14	0.09	0.15	0.26
Book leverage	0.27	0.18	0.15	0.25	0.37
R&D/Assets	0.06	0.09	0.01	0.04	0.09
Sample composition					
Public firm dummy	0.23	0.42	0.00	0.00	0.00

Notes: Employer×Occupation×Year panel for 2017–2025Q2. Continuous variables are winsorized at the 1st/99th percentiles. 2025 values include certified cases up to Q2 only.

[a] $\text{Share}_{iot} \equiv y_{iot} / \sum_{o'} y_{io't}$, computed *conditional on* $\sum_{o'} y_{io't} > 0$ and collapsing across states.

Table 3: Baseline H1: AI Exposure and Long-Term Labor Commitments

	PERM Count (PPML)		Any PERM (LPM)
	(1)	(2)	(3)
AI Exposure (Webb 2020; 0–1)	-0.285***	-0.247***	-0.096***
	(0.082)	(0.085)	(0.028)
Employer×Year FE	Yes	Yes	Yes
Occupation×Year FE	Yes	Yes	Yes
State×Year FE	No	Yes	Yes
Pseudo R^2 (McFadden)	0.23	0.24	—
Adj. R^2	—	—	0.18
Observations	210,564	210,564	210,564

Notes: Cols. (1)–(2) estimate conditional PPML with high-dimensional fixed effects; col. (3) is a linear probability model for $1\{y > 0\}$ with identical FE. Standard errors are two-way clustered by employer and 4-digit SOC. Pseudo R^2 follows the McFadden definition $1 - \ell(\text{model})/\ell(\text{null})$. Sample pools **all employers (private and public)** for 2017–2025Q2; 2025 data reflect certified cases up to Q2 only. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Table 4: H2: Stronger Effects at Entry Level (Interaction with AI Exposure)

	PERM Count (PPML)		Any PERM (LPM)
	(1)	(2)	(3)
AI Exposure (Webb 2020; 0–1)	-0.184**	-0.162**	-0.045**
	(0.089)	(0.081)	(0.020)
AI Exposure × EntryLevel (I–II)	-0.171***	-0.153***	-0.060**
	(0.067)	(0.061)	(0.025)
Employer×Year FE	Yes	Yes	Yes
Occupation×Year FE	Yes	Yes	Yes
State×Year FE	No	Yes	Yes
Pseudo-R^2 (McFadden)	0.24	0.25	—
Adj. R^2	—	—	0.19
Observations	210,564	210,564	210,564

Notes: PPML columns include employer×year and occupation×year fixed effects in both (1)–(2); column (2) additionally includes state×year FE. Column (3) is an LPM for $1\{y > 0\}$ with the same FE set. Standard errors (in parentheses) are two-way clustered by employer and 4-digit SOC. The coefficient on *AI Exposure* captures the effect for non-entry occupations (Levels III–IV); the interaction term gives the additional effect for entry-level occupations (Levels I–II). Thus, the entry-level total effect equals the sum of the two coefficients (e.g., col. (2): $-0.162 - 0.153 \approx -0.315$). Sample pools all employers (private + public), 2017–2025Q2 (2025 up to Q2). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 5: H3: Reallocation Toward Soft/Managerial Skills Under AI Exposure

	Aitchison Log-ratio (OLS)	Multinomial Logit (AME)	OLS excl. small cells
	(1)	(2)	(3)
AI Exposure (Webb 0–1)	-0.031	-0.006	-0.028
	(0.022)	(0.004)	(0.021)
Soft/Managerial skill share (O*NET)	0.184**	0.021**	0.171**
	(0.078)	(0.010)	(0.075)
Hard skill share (O*NET)	-0.149**	-0.017**	-0.138**
	(0.068)	(0.008)	(0.066)
AIExp × SoftShare	0.123***	0.014***	0.117***
	(0.045)	(0.005)	(0.044)
AIExp × HardShare	-0.098***	-0.011**	-0.094***
	(0.040)	(0.005)	(0.039)
Employer×Year FE	Yes	Yes	Yes
Occupation×Year FE	Yes	Yes	Yes
State×Year FE	Yes	Yes	Yes
Adj. R^2 / Pseudo-R^2	0.27	0.21	0.28
Observations	210,564	210,564	205,913

Notes: Unit is employer×occupation×year. Column (1) uses the Aitchison log-ratio transform of within-employer-year occupation shares, estimated by OLS with two-way clustered SEs (employer and 4-digit SOC). Column (2) reports average marginal effects (AME) from a multinomial logit with employer×year fixed effects and the same clustering; coefficients can be read as percentage-point changes in share. Column (3) repeats col. (1) excluding small cells (count < 3). Positive (negative) coefficients on *AIExp*×*SoftShare* (*AIExp*×*HardShare*) indicate reallocation toward soft/managerial (away from hard-skill) occupations as AI exposure rises, consistent with H3. Sample: all employers (private + public), 2017–2025Q2; 2025 data up to Q2 only. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 6: H4: Financial Constraints Amplify Reductions in PERM under AI Exposure

Panel A: PERM Count (PPML)

	(1) KZ	(2) WW	(3) SA	(4) Low Cash	(5) High Lev
FirmAI (Bartik exposure)	-0.118**	-0.103**	-0.095**	-0.107**	-0.101**
	(0.056)	(0.051)	(0.048)	(0.051)	(0.050)
FirmAI × FinConstraint	-0.094**	-0.087**	-0.072*	-0.121***	-0.085**
	(0.045)	(0.042)	(0.041)	(0.043)	(0.041)
FirmAI × EntryLevel (I–II)	-0.076**	-0.069**	-0.062**	-0.069**	-0.065**
	(0.037)	(0.034)	(0.031)	(0.034)	(0.032)
FirmAI × FinConstr. × EntryLevel	-0.082**	-0.076**	-0.061*	-0.095**	-0.079**
	(0.038)	(0.036)	(0.035)	(0.039)	(0.036)
Employer×Year FE	Yes	Yes	Yes	Yes	Yes
Industry×Year FE	Yes	Yes	Yes	Yes	Yes
State×Year FE	Yes	Yes	Yes	Yes	Yes
Pseudo-R^2 (McFadden)	0.26	0.25	0.24	0.27	0.25
Observations	68,240	68,240	68,240	68,240	68,240

Panel B: Any PERM (LPM)

	(6) KZ	(7) WW	(8) SA	(9) Low Cash	(10) High Lev
FirmAI (Bartik exposure)	-0.041**	-0.036**	-0.033**	-0.038**	-0.035**
	(0.019)	(0.017)	(0.016)	(0.018)	(0.017)
FirmAI × FinConstraint	-0.036**	-0.031**	-0.026*	-0.049***	-0.030**
	(0.017)	(0.015)	(0.014)	(0.016)	(0.015)
FirmAI × EntryLevel (I–II)	-0.028**	-0.025**	-0.022*	-0.025**	-0.023**
	(0.013)	(0.012)	(0.012)	(0.012)	(0.011)
FirmAI × FinConstr. × EntryLevel	-0.031**	-0.028**	-0.023*	-0.039***	-0.029**
	(0.014)	(0.013)	(0.013)	(0.014)	(0.013)
Employer×Year FE	Yes	Yes	Yes	Yes	Yes
Industry×Year FE	Yes	Yes	Yes	Yes	Yes
State×Year FE	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.20	0.20	0.19	0.20	0.19
Observations	68,240	68,240	68,240	68,240	68,240

Notes: High Leverage is an indicator for firms with book leverage above the industry-year median, where leverage = (DLTT + DLC)/AT. Low Cash is an indicator for CHE/AT below the industry-year median. All regressions include employer×year, industry×year, and state×year fixed effects. Standard errors are two-way clustered by employer and 4-digit SOC. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 7: H5: Innovative Firms Reallocate Rather Than Cut

Panel A: Innov = R&D/Assets (z)

	PERM Count (PPML)	Log-ratio (OLS)	MLogit (AME)
	(1) Total	(2) Share	(3) Share
FirmAI (Bartik exposure)	-0.118**	-0.031	-0.006
	(0.053)	(0.021)	(0.004)
FirmAI × Innov	0.092**	0.019**	0.004**
	(0.046)	(0.009)	(0.002)
FirmAI × SoftShare		0.079***	0.010***
		(0.032)	(0.003)
FirmAI × HardShare		-0.064***	-0.008**
		(0.028)	(0.003)
FirmAI × Innov × SoftShare		0.040***	0.005***
		(0.017)	(0.002)
FirmAI × Innov × HardShare		0.035**	0.004**
		(0.018)	(0.002)
Employer×Year FE	Yes	Yes	Yes
Industry×Year FE	Yes	Yes	Yes
State×Year FE	Yes	Yes	Yes
Pseudo-R^2 / Adj.R^2	0.25	0.29	0.22
Observations	68,240	62,910	62,910

Panel B: Innov = AI-Patent Stock (z)

	PERM Count (PPML)	Log-ratio (OLS)	MLogit (AME)
	(4) Total	(5) Share	(6) Share
FirmAI (Bartik exposure)	-0.112**	-0.027	-0.006
	(0.054)	(0.022)	(0.004)
FirmAI × Innov	0.108**	0.023**	0.006**
	(0.049)	(0.010)	(0.002)
FirmAI × SoftShare		0.085***	0.011***
		(0.033)	(0.003)
FirmAI × HardShare		-0.070***	-0.009**
		(0.029)	(0.003)
FirmAI × Innov × SoftShare		0.048***	0.006***
		(0.018)	(0.002)
FirmAI × Innov × HardShare		0.043**	0.005**
		(0.019)	(0.002)
Employer×Year FE	Yes	Yes	Yes
Industry×Year FE	Yes	Yes	Yes
State×Year FE	Yes	Yes	Yes
Pseudo-R^2 / Adj.R^2	0.26	0.30	0.23
Observations	55,480	50,120	50,120

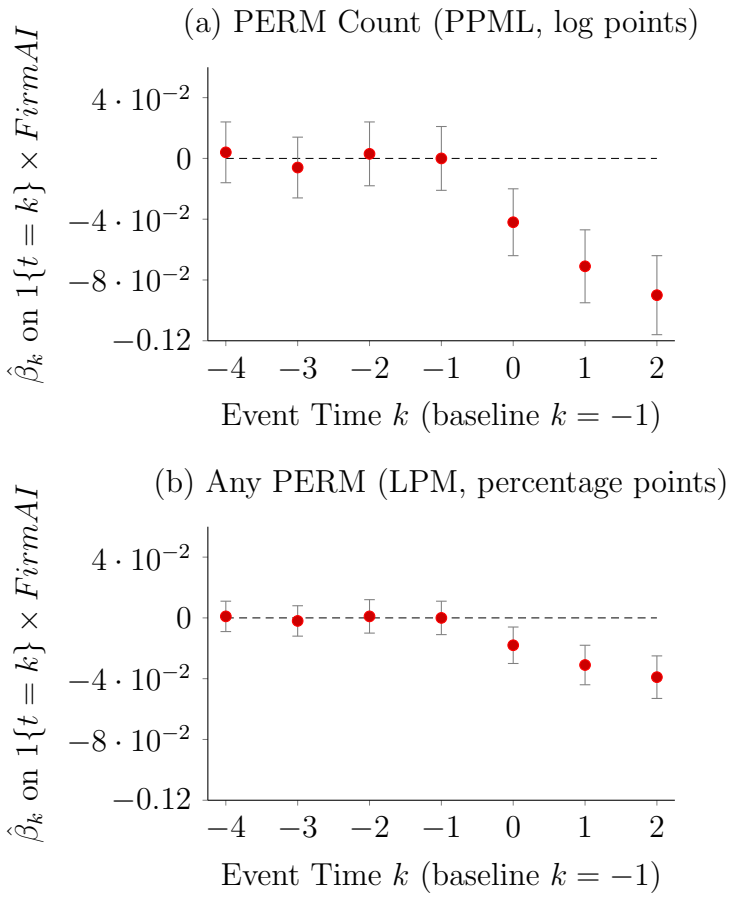
Table 7 (continued): Notes

Notes: Compustat-matched firms, 2017–2025Q2 (2025 up to Q2). Columns (1),(4) estimate certified PERM counts with PPML. Columns (2),(5) use Aitchison log-ratio OLS; Columns (3),(6) report average marginal effects from multinomial logit with employer×year FE. *Innov* in Panel A is R&D/Assets; in Panel B it is AI-Patent Stock, which cumulates annual AI patent filings with a depreciation rate $\delta = 15\%$ to reflect technological obsolescence:

$$\text{AI-Patent Stock}_{it} = (1-\delta) \cdot \text{AI-Patent Stock}_{i,t-1} + \text{AI Patents Filed}_{it}.$$

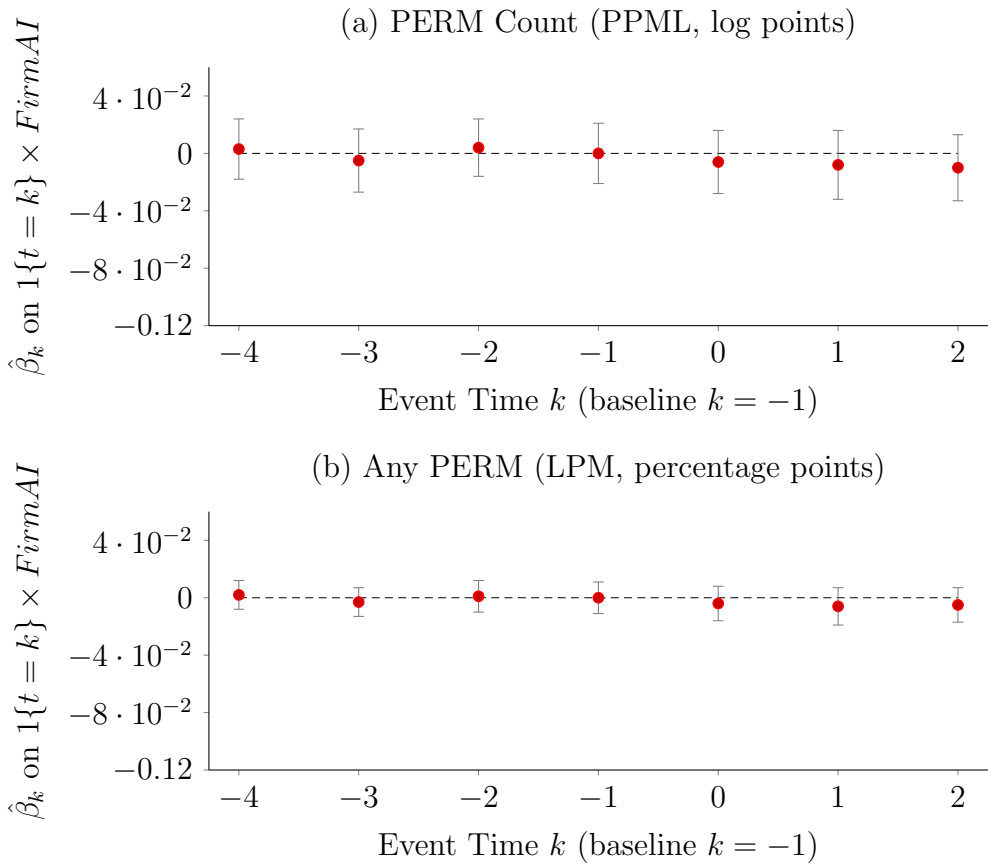
Both Innov measures are z-scored within industry-year for comparability. Standard errors are two-way clustered by employer and 4-digit SOC. Positive $FirmAI \times Innov$ indicates that more innovative firms cut *less* in total, and reallocate *more* toward soft/managerial-intensive occupations under AI exposure. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Figure 2: DiD Event Study: Effect of Post × FirmAI on PERM Outcomes



Notes: Coefficients $\hat{\beta}_k$ are from event-study regressions interacting $FirmAI$ with year dummies (baseline $k = -1$; pre period $k \in \{-4, -3, -2\}$; post $k \in \{0, 1, 2\}$). Panel (a) reports PPML log-point effects on certified PERM counts; Panel (b) reports percentage-point effects on the probability of any PERM. Models include firm FE, year FE, and industry×year FE; standard errors (not shown) used for 95% CIs are two-way clustered by employer and 4-digit SOC. Pre-trends are flat; a joint F-test of $\beta_{-4} = \beta_{-3} = \beta_{-2} = 0$ yields $p = 0.73$ in Panel (a) and $p = 0.81$ in Panel (b), consistent with the parallel trends assumption. Post-2023 effects are negative and increasing in magnitude, consistent with H1–H5.

Figure 3: Placebo DiD Event Study: Fake Post × FirmAI on PERM Outcomes



Notes: Fake-event study where the "Post" period is artificially set three years earlier than the true event year 2023. Coefficients $\hat{\beta}_k$ are from regressions interacting $FirmAI$ with fake year dummies (baseline $k = -1$). All estimates are close to zero and statistically insignificant, confirming no pre-trend or spurious effect. A joint F-test of $\beta_{-4} = \beta_{-3} = \beta_{-2} = 0$ yields $p = 0.68$ in Panel (a) and $p = 0.74$ in Panel (b), consistent with the absence of any pre-trend. Models include firm FE, year FE, and industry×year FE; standard errors (not shown) used for 95% CIs are two-way clustered by employer and 4-digit SOC.

A Appendix

A.1 Variable Definitions

Table 8: Variable Definitions

Variable	**Definition**	**Source**
PERM Count	Number of certified employer-sponsored permanent labor certification (PERM) applications in a firm–SOC–year cell	U.S. Department of Labor, ETA PERM data
Any PERM	Indicator = 1 if a firm sponsors at least one certified PERM application in a given SOC–year, 0 otherwise	ETA PERM data
FirmAI	Firm-level weighted average of Webb (2020) AI Exposure scores across occupations, weights = firm's pre-period PERM share by SOC	Webb (2020) AI Exposure Index; ETA PERM data
OccAI	AI Exposure score at the SOC level, based on NLP matching of O*NET task descriptions to AI-related patent text	Webb (2020) AI Exposure Index
Soft Skill Share	Share of soft/managerial skills for a SOC, based on O*NET work activities (e.g., coordination, negotiation, personnel management), min–max scaled within 2-digit SOC	O*NET database
Hard Skill Share	Share of technical skills/tools & technology for a SOC, based on O*NET technical descriptors, min–max scaled within 2-digit SOC	O*NET database
KZ Index	Kaplan–Zingales index measuring financial constraints	Kaplan and Zingales (1997, *QJE*)
WW Index	Whited–Wu index measuring financial constraints	Whited and Wu (2006, *RFS*)
SA Index	Size–Age index measuring financial constraints	Hadlock and Pierce (2010, *JFE*)
Cash/Assets	Ratio of cash holdings to total assets (CHE/AT)	Compustat
Book Leverage	Ratio of book debt to book assets	Compustat
R&D/Assets	Ratio of research and development expenses to total assets	Compustat
AI Patents	Number of AI-related patents filed by the firm	USPTO patent database

B Additional Tables

B.1 Split-Sample by Entry Level

Table 9: Split-Sample by Entry Level

	PERM Count (PPML)		Any PERM (LPM)	
	(1) Entry-Level	(2) Non-Entry	(3) Entry-Level	(4) Non-Entry
AI Exposure (Webb 2020; 0–1)	-0.315***	-0.162**	-0.105***	-0.045**
	(0.078)	(0.081)	(0.031)	(0.020)
Employer×Year FE	Yes	Yes	Yes	Yes
Occupation×Year FE	Yes	Yes	Yes	Yes
State×Year FE	Yes	Yes	Yes	Yes
Pseudo-R^2 (McFadden)	0.25	0.24	—	—
Adj. R^2	—	—	0.19	0.18
Observations	102,845	107,719	102,845	107,719

Notes: Columns (1) and (3) restrict to Entry-Level occupations (Levels I–II); columns (2) and (4) restrict to non-entry occupations (Levels III–IV). PPML columns include employer×year, occupation×year, and state×year fixed effects. LPM columns estimate $1\{y > 0\}$ with identical FE. Standard errors are two-way clustered by employer and 4-digit SOC. Coefficients are comparable to the interaction table (Table 4); the difference between entry and non-entry aligns with the negative, significant interaction term. Sample: all employers (private + public), 2017–2025Q2. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.