AI and Returns to Experience in Entrepreneurship

Ziqing Yan* Irisa Zhou†

May 21, 2026
Newest Version Here

Abstract

This paper examines how successive waves of AI advances have altered the value of work experience in entrepreneurship. Using employment histories from LinkedIn (2007–2019), we first exploit industry-level variation in AI exposure following the 2012 deep learning breakthrough. For this first wave, we find the aggregate effect on entry is near zero; however, for workers with prior research experience, high AI exposure raises the annual probability of founding by 58%. This response peaks for those with 5–14 years of expertise, suggesting that early AI breakthroughs complemented domain knowledge rather than ”democratizing” entry. To interpret these shifts, we develop a directed search model that quantifies how AI shocks reallocate talent by altering productivity premia across different skill sets.

1 Introduction

Artificial Intelligence (AI) is fundamentally transforming the nature of work, with profound implications for the entrepreneurial landscape. Scholars increasingly characterize AI as an "external enabler" of entrepreneurship—creating novel business opportunities even as it disrupts traditional employment structures (Fossen et al., 2024). However, AI's capacity to either replace or augment specific human skills raises a critical question: Which forms of human capital become more—or less—valuable for entrepreneurs in the age of AI?

*Department of Economics, Yale University. *ziqing.yan@yale.edu*
†Department of Economics, University of Toronto. *irisa.zhou@mail.utoronto.ca*

Classic entrepreneurship theory emphasizes the necessity of a balanced skill set. In particular, Lazear's "Jack-of-all-Trades" theory argues that individuals with a broad, multifaceted repertoire are more likely to become successful founders (Lazear, 2004). Historically, entrepreneurs had to "wear many hats"—integrating management, finance, and technical expertise—whereas specialists remained in the wage-labor market. A balanced skill set, therefore, is what we want to consider when we think of an "entrepreneur" or a "founder". We operationalize this by considering two distinct types of human capital: worker experience, which represents generalist or routine operational knowledge, and researcher experience, which represents specialized, innovation-driven expertise. The rise of AI may fundamentally alter this allocation of talent. If AI substitutes for routine cognitive tasks or lowers technical barriers to complex problem-solving, it may effectively relax the "balanced skill" constraint that has historically limited entrepreneurial entry. Therefore, our research asks: How does AI exposure change who becomes an entrepreneur?

But what is an AI? Broadly speaking, we have experienced two major AI breakthroughs recently. The first wave of deep learning around 2012, marked by the success of neural networks and the ImageNet competition, led to a surge in AI research and applications. Implementing AI tools at this stage required deep technical sophistication and specialized human capital. The second wave, starting around 2020 with the advent of large language models like GPT-3, has further accelerated AI's capabilities and adoption across industries. This era represents a transition to "General Purpose" AI. These tools are largely accessible to anyone via natural language, potentially substituting for the specialized technical barriers that once deterred prospective founders. This paper exploits both waves as distinct quasi-experiments. The second wave — the ChatGPT release in late 2022 — is our primary empirical focus, as it offers sharp variation in actual AI adoption across industries and provides much larger opportunities for a broader range of entrepreneurs. The first wave provides a comparison that illuminates the robustness and structural consistency of the pattern across distinct technological regimes.

We examine whether individual-level industry AI exposure increases the probability of transitioning into entrepreneurship. Our approach differs from the burgeoning literature on the labor market exposure of AI (Massenkoff and McCrory, 2026; Eloundou et al., 2024). While they focus on occupational displacement, we ask: If an industry is highly exposed to AI, are individuals within it more likely to found new firms? Is there a particular experience that would benefit individuals to start a new company?

The answer is non-trivial. If AI complements both operational and innovative skills, it may reinforce incumbent advantages for those who already possess a balanced skill set. Conversely, if AI substitutes for both, the "barrier to entry" for specialists might vanish, and prior experience accumulation becomes less important: anyone with the right characteristics could become an entrepreneur. If AI substitutes one and complements the other, there would be a reallocation of entrepreneurial entry toward the skill set that AI complements. We provide some answers to this question both empirically and structurally.

We use longitudinal professional profile data from Revelio Labs, covering U.S. workers observed quarterly from 2007Q1 through 2025Q1. Our main empirical strategy exploits the release of ChatGPT in 2022Q4 as a shock to the availability of general-purpose AI tools, using pre-shock cross-sectional variation in industry-level AI adoption (measured by the Anthropic Labour Market Impact score) as the treatment. This pre-determined assignment avoids reverse causality from post-shock industry sorting. In particular, we define AI exposure at the individual level using the worker's full career history before 2022Q4.

Our first finding is that the share of workers reporting a founder title rose significantly after 2022Q4, and this increase is concentrated among workers with prior research or engineering experience in high-AI-adoption industries. For researchers in high-Anthropic-exposure industries, the probability of currently being a founder grows by approximately +0.18 percentage points relative to the pre-period, with no corresponding response for non-researchers. The parallel-trends assumption is validated by flat pre-period event-study coefficients spanning eight pre-treatment quarters.

Second, this heterogeneity is not solely a new-entry phenomenon. A decomposition into new founders, serial entrepreneurs, and incumbent stayers shows that serial entrepreneurs also respond to the shock in a different timing pattern: non-researchers re-enter first, while researchers take longer. This suggests that the mechanism differs across margins: researchers leverage domain expertise to translate AI tools into new ventures, while re-entering non-researchers primarily benefit from AI lowering operational or experimentation costs.

Third, the effect is monotonically increasing in the depth of pre-shock research experience up to 14 years, consistent with a model in which workers must cross an experience threshold to exploit AI's complementarity. Non-researchers show no experience gradient.

To help us guide and interpret these findings, we develop a directed search model with multidimensional human capital and occupational choice. The model formalizes the Lazear jack-of-all-trades mechanism and yields a clear prediction: an AI-driven increase in the productivity premium of research-oriented human capital raises entrepreneurial entry, and this effect is concentrated among workers who already possess substantial research experience. Comparative statics from the calibrated model reproduce the empirically observed concentration of entrepreneurial responses among technically trained workers.

Finally, we examine whether the pattern generalizes across AI waves. Using the 2012 AlexNet shock as the first-wave experiment, we find a strikingly consistent pattern: the AI shock raised first-founding rates by approximately 58% for researchers in high-exposure industries, with no response for non-researchers. The cross-wave consistency strengthens the structural interpretation that AI's impact on entrepreneurship operates through a durable complementarity between AI and research human capital, rather than through features specific to the ChatGPT transition.

Related literature. Our work contributes to several strands of literature.

First, we build on the broad literature examining how technological change, particularly AI, is reshaping the demand for skills. A long tradition in labor economics has documented that computerization and automation tend to reduce demand for routine skills while increasing demand for non-routine cognitive and social skills (Autor et al., 2003; Autor, 2015; Autor and Thompson, 2025). Recent studies show that AI extends this trend by automating some high-skill cognitive tasks and reducing demand for certain non-AI cognitive positions (Acemoglu et al., 2022; Webb, 2019), though AI can also complement complex human work and raise productivity (Babina et al., 2024; Hampole et al., 2025). We contribute to this literature by connecting this substitution–complementarity framework to entrepreneurship, showing that AI shapes not only labor demand but also the composition of new venture formation.

Second, our work is among the first to empirically analyze the impact of AI on entrepreneurship at the micro level. Fossen et al. (2024) reviews the direct and indirect effect of AI on entrepreneurship. Gofman and Jin (2024) finds that AI professors' departures from universities reduce startup founders' AI knowledge and leads the students in affected universities to establish fewer AI startups and raise less funding. We provide new micro-level evidence that AI alters the selection of who becomes an entrepreneur, shifting the margin of entry toward experienced researchers and specialists.

Third, our research contributes to the rich literature on entrepreneurs' backgrounds and the role of human capital in new venture performance, including the seminal work of Lazear (2004). A consistent finding in prior work is that founder experience matters for entrepreneurial success. Azoulay et al. (2020) find that middle-aged entrepreneurs tend to outperform younger ones, as noted earlier. We thereby contribute to the literature by identifying a specific technological force (AI) that is altering the selection into entrepreneurship along experience lines. We show a technology-driven shift over time in who becomes an entrepreneur.

Finally, our model contributes to the literature by presenting a simple directed search framework of multidimensional skill combined with stochastic human capital accumulation à la Ljungqvist and Sargent (1998). We introduce multiple channels through which AI can impact the labor market, and, through the lens of our model, we analyze which mechanism dominates and how. We directly address the potential complementarity between AI and the researcher human capital premium, and how it affects the distribution of individuals' experience.

Section 2 introduces the theoretical framework. Section 3 describes the data and summary statistics. Section 3.2 presents the empirical strategy. Section 4 reports the main empirical results. Section 5 discusses the comparison with the GPT capability measure and presents first-wave evidence. Section 6 concludes.

2 Model

We present a directed search model to illustrate how AI affects workers' incentives to become entrepreneurs through its interaction with accumulated human capital. The model provides the theoretical framework motivating our empirical design; formal proofs and the full calibration exercise are in Appendix J.

2.1 Environment

Time is discrete with discount rate β. The economy is populated by a continuum of agents who differ in their two-dimensional experience (e_W, e_R), corresponding to **worker experience** (operational, generalist) and **researcher experience** (technical, innovative). Agents follow a perpetual-youth setup (Blanchard, 1985), exiting at rate σ, with an equal inflow of newborns drawing from a joint distribution $F(e_W, e_R)$.

There are three occupations — worker (W), researcher (R), and entrepreneur (E) — and two industries: AI-intensive (I) and non-AI-intensive (N). Firms post vacancies in submarkets indexed by (e_W, e_R, o, j), and workers direct their search across submarkets (Menzio and Shi, 2010, 2011). Workers can switch occupations at a cost $c(e_W, e_R, o, j)$.

2.2 Production and the Role of Experience Balance

Motivated by the "jack-of-all-trades" theory of entrepreneurship (Lazear, 2004), production in each occupation-industry pair depends on both dimensions of experience through a CES aggregator:

$$g_{oj}(e_W, e_R) = A_{oj}\Big(\alpha_o\, e_W^{\rho_{oj}} + (1-\alpha_o)\, e_R^{\rho_{oj}}\Big)^{1/\rho_{oj}}.$$

For workers and researchers, $\rho_{oj} \to \infty$ so that only the dominant experience dimension matters; for entrepreneurs, $\rho_{oj} \to 0$ (Cobb-Douglas), implying that a *balanced* experience pair is required. The parameter α_o governs the relative importance of worker experience: $\alpha_W > \alpha_R$, so researchers rely more on e_R while workers rely more on e_W.

2.3 Human Capital Accumulation

While employed in occupation o and industry j, agents stochastically accumulate the relevant experience dimension at rate π_j. This creates a complementarity between industry choice and the speed at which entrepreneurs build the balanced skill set required to start a firm. When unemployed, both experience dimensions depreciate at rates $\tilde{d}_o$.

2.4 Key Theoretical Implication

The model yields the following central result, which we take to the data:

Proposition 1 *An increase in the productivity premium of research-oriented human capital (A_R) in AI-intensive industries raises the equilibrium share of entrepreneurs, and this effect is concentrated among workers with higher pre-existing researcher experience.*

The intuition is straightforward. A higher research productivity premium raises the value of researcher experience. Workers with substantial accumulated e_R find that the marginal cost of transitioning into entrepreneurship — which requires a balanced (e_W, e_R) pair — falls, since they are "closer" to the balanced threshold. In contrast, workers with predominantly general work experience (e_W) gain less from the increase in A_R, so their entrepreneurial entry rate is largely unchanged. This generates the heterogeneous response to AI that we identify empirically.

The comparative statics underlying this proposition are illustrated in Appendix J. We show that raising A_R simultaneously increases the shares of both researchers and entrepreneurs, reproducing the pattern documented in the data.

3 Data and Empirical Strategy

3.1 Data

Our primary dataset comes from Revelio Labs, which is a labor market analytics company that collects and analyzes data from public LinkedIn profiles. The workers profile data contains detailed information on individuals' work history, including job titles, employers, dates of employment, location, and industries (NAICS). Revelio Labs uses machine learning and natural language processing techniques to extract and standardize information as well as predicting certain variables such as gender, race, seniority of positions, and salaries. For the current analysis, we use a 2% random sample of all U.S. LinkedIn users. As position start and ending dates are crucial in our analysis, we only keep positions with non-missing start date and impute the end date as the current date if it is missing. We also drop positions with end date earlier than start date. After cleaning, our sample contains around 0.95 million unique users and 4.3 million unique positions. Notably, user positions span 1950 to 2025Q1, yielding a total of 65 million user-quarter observations.

Identifying founders and researchers. We identify founders and researchers by searching for keywords from their job titles. The keywords for founders are:

All Founder Keywords
cofounder, co-founder, founder, owner, CEO, chief executive officer, CTO, chief technology officer

and the keywords for researchers are:

All Researcher Keywords
scientist, research, researcher, r&d, r and d, engineer, engineering, technologist, technical lead, technology lead, developer, product development, data scientist, machine learning, ai researcher, ml researcher, innovation, inventor

We exclude confounding terms such as "research assistant" or "research intern" as we are interested in post-education professional experience. To avoid misclassifying senior executives at established firms as founders, we set the founder indicator to zero if the company's founding year is five or more years before the worker's start date in that position. The researcher classification is fixed to the pre-2022Q4 period, ensuring it is not endogenous to the ChatGPT shock we study.

Measuring experience. Work experience is not directly observed in the data. We measure workers' experience by counting the number of years they worked in each position after the end date of their highest education. If the end date of the highest education is missing, we impute it as the start year of the worker's first position observed in the data. We further distinguish between *researcher experience* (years spent in researcher-classified positions) and *non-researcher experience* (remaining years). When examining the probability of becoming a founder, we focus exclusively on the worker's first founding event, restricting the sample to person-quarters in which the worker has not yet founded a company. This approach allows us to analyze how the likelihood of founding a company evolves with experience while avoiding complications from multiple founding events by the same individual.

AI exposure. We use the Anthropic Labour Market Impact score (Massenkoff and McCrory, 2026) as our primary AI exposure measure. This score captures the intensity of actual AI adoption at the industry level, reflecting how much workers in a given industry have already integrated AI tools into their work rather than a theoretical estimate of AI's potential reach. As an alternative treatment, we also use the GPT-4 capability measure from Eloundou et al. (2024), which measures the share of occupation-level tasks that LLMs could plausibly perform. Full details of both measures and their validation are in Appendix D.

Each worker's AI exposure is determined by their *pre-2022Q4 industry*, computed as the time-weighted average of the Anthropic score over all quarters in which the worker held a non-founder position with a valid industry code before 2022Q4. This pre-determined assignment avoids post-shock sorting bias: workers who moved into high-AI industries in response to ChatGPT would contaminate a time-varying treatment assignment. We exclude Agriculture (NAICS 11), Construction (NAICS 23), and Government (NAICS 92) throughout, as these sectors have markedly different labor market characteristics. Workers with missing industry codes are retained under a placeholder code.

Summary statistics. Figure 1 documents the unconditional time series of founder prevalence from 2007 to 2025. It exhibits significant seasonality, but the trend was easy to see. The founder share was broadly flat at roughly 7.3% through 2014–2019, began rising gradually from 2020 onward, and accelerated sharply after the public release of ChatGPT in late 2022. By early 2025, the founder share had reached approximately 7.8%. This visual pattern provides the motivation for focusing on the second wave of AI, and the formal event study in Section 4 controls for quarter fixed effects to isolate this acceleration from secular trends.

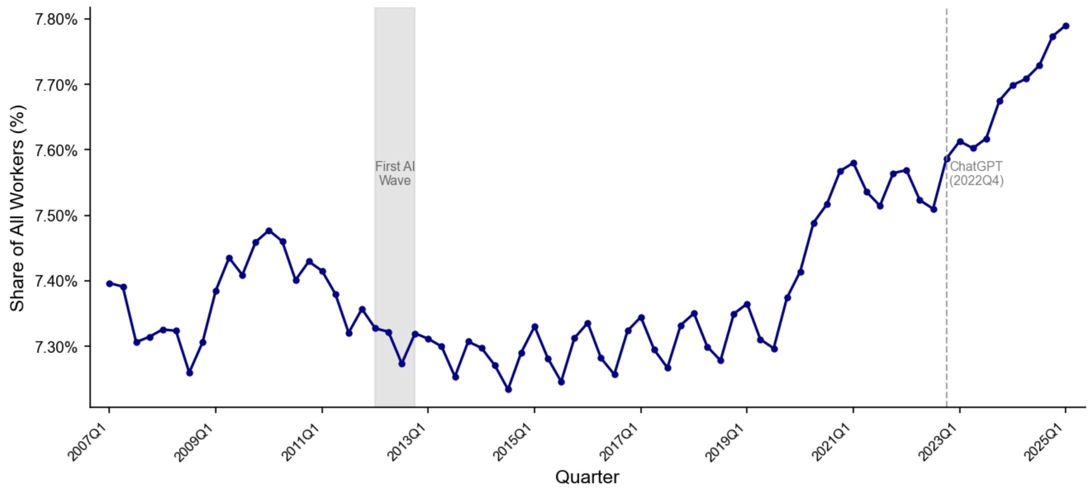

Figure 1: Share of workers currently reporting a founder or entrepreneur title, by quarter (2007Q1–2025Q1). Grey band: first AI wave (2012Q1–2012Q4). Dashed vertical line: ChatGPT launch (2022Q4).

Figure 2 plots the share of workers holding a researcher or scientist title over the same period. The series follows a U-shaped trajectory, declining to a trough precisely at the first AI wave in 2012–2013 before rising monotonically through 2023. Crucially, the researcher share shows no discontinuity at the 2022Q4 ChatGPT cutoff, alleviating concerns that the pre-treatment researcher classification is itself endogenous to the second-wave shock we study.

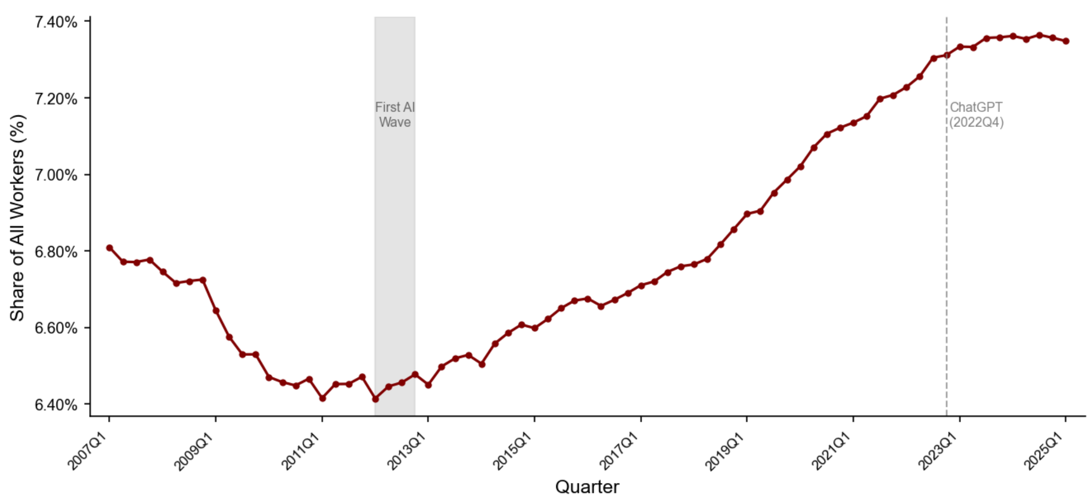

Figure 2: Share of workers currently reporting a researcher or scientist title, by quarter (2007Q1–2025Q1). Grey band: first AI wave (2012Q1–2012Q4). Dashed vertical line: ChatGPT launch (2022Q4).

3.2 AI Shock

We exploit the public release of ChatGPT in November 2022 as an exogenous shock to the availability of general-purpose AI tools. Unlike the first wave of AI adoption (anchored to the 2012 AlexNet breakthrough), which required deep technical expertise to deploy, the ChatGPT release democratized access to large-language-model capabilities through a natural-language interface, with broad and immediate implications for knowledge workers across industries. We therefore use 2022Q4 as our treatment period, defining $\text{Post}_t = \mathbf{1}[t \geq 2022\text{Q4}]$.

3.3 Treatment Construction

Each worker's AI exposure is fixed at their *pre-2022Q4 industry*, defined as the time-weighted average of the Anthropic Labour Market Impact score over all quarters in which the worker held a non-founder, non-missing industry affiliation before 2022Q4. We classify workers into a binary treatment indicator $\text{High}_i = \mathbf{1}[\text{Anthropic score}_i > \text{median}]$, where the median is taken across workers in the estimation sample. An important distinction of our definition AI exposure is that it is at the individual level. This is important distinction because workers' histories of employment shape their future employment choices. As a robustness check, we also use the continuous Anthropic score; results are reported in Appendix H.

Fixing treatment to the pre-2022Q4 industry is critical for identification. Workers who moved into high-AI industries *in response* to the ChatGPT shock would contaminate a time-varying industry assignment with the very outcome we are studying. The predetermined industry assignment breaks this reverse causality: High_i is a function only of pre-shock career histories. Full details of the treatment construction are in Appendix D. We exclude Agriculture (NAICS 11), Construction (NAICS 23), and Government (NAICS 92) throughout, as these sectors have markedly different labor market characteristics and are not the focus of our analysis. Workers with missing industry codes are retained under a placeholder code so no observations are dropped.

Two outcomes are studied throughout: (i) the probability of *currently* holding a founder title (stock measure: captures both new entrants and incumbents), and (ii) the probability of a worker's *first* founding event conditional on not having founded before (flow measure: captures new entry). All outcomes are scaled to percentage points.

3.4 Difference-in-Differences

We estimate a difference-in-differences (DiD) specification comparing workers whose pre-2022Q4 industry is in the high-AI-exposure group to those in the low-exposure group, around the ChatGPT shock. The baseline estimating equation is:

$$y_{it} = \beta\,(\text{Post}_t \times \text{High}_i) + \alpha_i + \alpha_j + \alpha_t + \varepsilon_{it}, \tag{1}$$

where i indexes workers and t indexes calendar quarters (2020Q4–2025Q1). High_i is the binary treatment fixed at the worker's pre-2022Q4 industry. α_i denotes worker fixed effects, α_j 2-digit industry fixed effects, and α_t quarter fixed effects (which absorb Post_t itself). Standard errors are clustered by worker throughout.

The coefficient β captures the differential change in founding propensity post-2022Q4 for workers in high-AI-exposure industries relative to those in low-exposure industries. We further split the sample by *prior researcher status* — workers who held a researcher, scientist, or engineer title before 2022Q4 versus those who did not — to examine whether the AI effect on entrepreneurship is concentrated among workers with technical expertise. This heterogeneity is the focus of the analysis and is consistent with the model's Proposition 1.

3.5 Dynamic DiD and Parallel Trends

To test the parallel trends assumption and trace the timing of the treatment effect, we estimate an event study specification:

$$y_{it} = \sum_{k \neq -1} \beta_k \left(\mathbf{1}[t - t^* = k] \times \text{High}_i \right) + \alpha_i + \alpha_j + \alpha_t + \varepsilon_{it},$$

where $t^* = $ 2022Q4, k runs from -8 to $+9$ quarters relative to the shock, and $k = -1$ (2022Q3) is the excluded base period. The pre-period coefficients $\hat{\beta}_k$ for $k < 0$ test whether founding trajectories were already diverging between high- and low-exposure industries before 2022Q4; flat pre-period estimates support the parallel trends assumption. We estimate this equation separately for researchers and non-researchers and plot both series together to visualize differential dynamics.

A pooled triple-difference (DDD) specification directly estimates the excess researcher response:

$$y_{it} = \sum_{k\neq -1}\Big[\gamma_k\big(\mathbf{1}[t-t^*=k]\times \text{High}_i \times \text{Researcher}_i\big) + \beta_k\big(\mathbf{1}[t-t^*=k]\times \text{High}_i\big) + \delta_k\big(\mathbf{1}[t-t^*=k]\times \text{Researcher}_i\big)\Big] + \alpha_i + \alpha_j + \alpha_t + \varepsilon_{it},$$

where γ_k is the triple-interaction coefficient capturing the *excess* response of researchers in high-AI industries, period by period.

3.6 Robustness

We assess robustness along three dimensions. First, we estimate a simple before/after specification relating each outcome to Post_t with worker and 2-digit industry fixed effects (quarter fixed effects excluded since Post_t would be collinear), to establish the aggregate pre-versus-post shift prior to the DiD analysis. Second, we replace the binary treatment with the continuous Anthropic score to check sensitivity to the treatment threshold; results are reported in Appendix H. Third, to further validate the parallel trends assumption, we examine the re-entrant subsample — prior founders who were not active founders as of 2022Q3 — whose behavioral response provides an additional falsification test with different expected timing.

4 Empirical Results

4.1 Motivating Evidence

Table 1 establishes that the share of workers reporting a founder title rose significantly after the ChatGPT release, and that this increase reflected genuine within-worker transitions rather than a compositional shift.

Table 1: Post-ChatGPT Change in Founder Outcomes

	(1) All	(2) Researchers	(3) Non-Researchers
Panel A: Currently a Founder			
Post 2022Q4	0.170***	0.243***	0.158***
	(0.008)	(0.024)	(0.008)
Control mean	7.548	3.855	8.109
Industry FE	✓	✓	✓
Quarter FE			
Worker FE	✓	✓	✓
N	13,168,882	1,711,816	11,457,066
Panel B: First Founding Event			
Post 2022Q4	0.066***	0.082***	0.063***
	(0.001)	(0.004)	(0.001)
Control mean	0.048	0.059	0.046
Industry FE	✓	✓	✓
Quarter FE			
Worker FE	✓	✓	✓
N	11,868,171	1,587,210	10,280,961

Notes: OLS linear probability model; outcomes scaled to percentage points. Column (1): all workers. Column (2): researchers (held a researcher/engineer title before 2022Q4). Column (3): non-researchers. All columns include worker and 2-digit industry fixed effects; quarter fixed effects are excluded as `post` is collinear with a full set of quarter dummies. *Panel A*: outcome is the probability of currently holding a founder title (full sample). *Panel B*: outcome is the probability of a first founding event (at-risk sample, i.e. workers who had not previously founded). Standard errors clustered by worker. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The aggregate post-ChatGPT increase in founder prevalence is approximately +0.17 percentage points (Panel A, column 1), on a pre-period mean of roughly 7.55%, a relative increase of about 2.3%. The comparable estimate from the first-founding flow (Panel B, column 1) is also positive and statistically significant, ruling out the possibility that the stock effect is driven entirely by incumbents retaining their titles. Columns (2) and (3) reveal the first sign of the heterogeneity that is the paper's central focus: researchers (column 2) exhibit larger post-period increases than non-researchers (column 3). This motivates our difference-in-differences design, which exploits cross-sectional variation in industry AI exposure intensity.

Table 2: AI Exposure and Founder Outcomes: DiD Estimates

	(1) All	(2) All	(3) Researchers	(4) Non-Researchers
Panel A: Pr(Founder) (pp)				
Post × High (Anthropic)	0.044**	0.035*	0.119**	0.016
	(0.020)	(0.018)	(0.050)	(0.019)
Control mean (Low)	5.12	5.12	3.47	5.38
Industry FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Worker FE		✓	✓	✓
N	9,675,530	9,673,076	1,536,767	8,136,309
Panel B: Pr(First Founder) (pp)				
Post × High (Anthropic)	−0.003	0.007**	0.026***	0.003
	(0.003)	(0.003)	(0.008)	(0.003)
Control mean (Low)	0.076	0.049	0.054	0.048
Industry FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Worker FE		✓	✓	✓
N	8,918,697	8,915,201	1,424,628	7,490,573

Notes: OLS. Treatment: above-median Anthropic industry AI exposure score, fixed to pre-2022Q4 quarters. Column (1): full sample, industry and quarter FEs only. Columns (2)–(4): worker, 2-digit industry, and quarter fixed effects. Column (2): all workers. Column (3): researchers (held a researcher/engineer title before 2022Q4). Column (4): non-researchers. *Panel A*: outcome is the probability of currently holding a founder title (×100). *Panel B*: outcome is the probability of a first founding event (×100), restricted to the at-risk sample (workers who had not yet founded). Standard errors clustered by worker. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4.2 Main Results

Table 2 presents the difference-in-differences estimates from equation (1), exploiting cross-sectional variation in pre-ChatGPT industry AI exposure as measured by the Anthropic Labour Market Impact score.

The aggregate DiD coefficient (column 2) is positive and precisely estimated for the founder-stock outcome (Panel A), but near zero for the first-founding flow (Panel B). The key finding lies in the researcher sub-sample: workers who held a research or engineering title before 2022Q4 and worked in high-Anthropic-exposure industries became significantly more likely to be active founders after ChatGPT (column 3, Panel A), and the same group shows a significant increase in first founding events (column 3, Panel B). By contrast, non-researchers (column 4) show no significant differential response in either panel. The contrast between columns (3) and (4) is the paper's central finding: the ChatGPT shock raised entrepreneurial activity exclusively for workers with prior technical expertise. This pattern is consistent with Proposition 1: AI productivity gains in AI-intensive industries are converted into new venture formation specifically by workers who brought technical human capital to the threshold at which entrepreneurship becomes optimal. Robustness to a continuous treatment specification is reported in Appendix H.

4.3 Robustness: Pre-Trends and Industry Dynamics

Event studies. Figure 3 plots the dynamic DiD coefficients from separate regressions for researchers and non-researchers. The pre-period estimates (quarters $k = -8$ through $k = -1$) are flat and close to zero for both groups, supporting the parallel-trends assumption. After 2022Q4, the researcher series diverges steadily upward, reaching approximately +0.18 percentage points by $k = +9$ (2025Q1, significant from $k = +3$ onward), while the non-researcher series remains flat throughout.

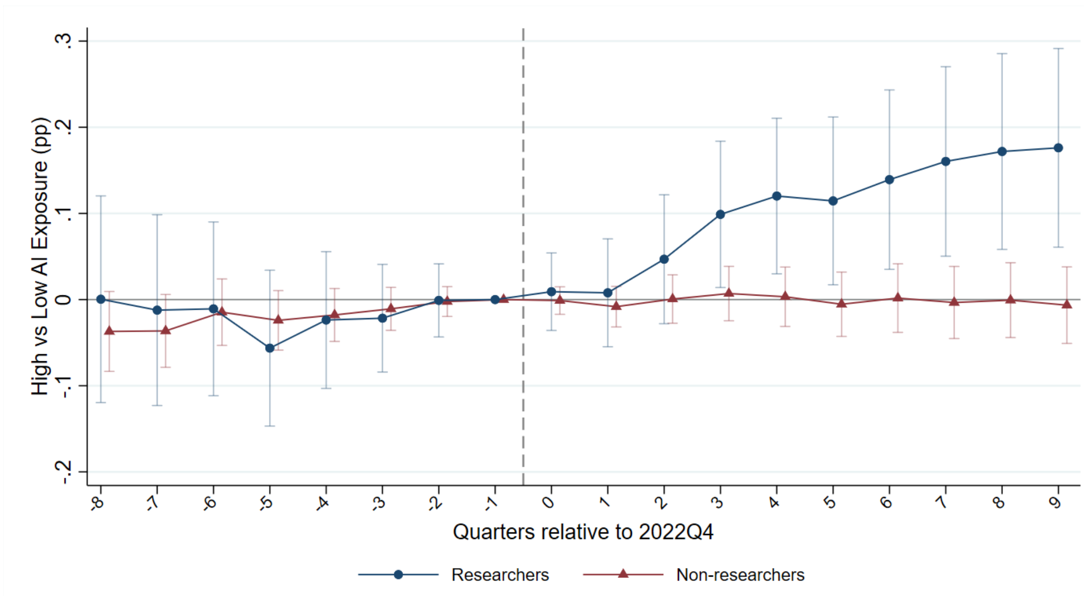

Figure 3: Dynamic DiD: Probability of being a founder around the ChatGPT release. Treatment: above-median Anthropic industry AI exposure; base period: 2022Q3. Fixed effects: worker, 2-digit industry, quarter. Standard errors clustered by worker. Navy circles: researchers; maroon triangles: non-researchers.

Figure 4 presents the first-founding event study on the at-risk sample. Pre-trends are flat for both groups. The researcher series shows a modest and gradually accumulating positive post-period path that is individually insignificant in each quarter — consistent with the low per-quarter first-founding rate — but directionally aligned with the significant pooled DiD estimate in Table 2. Non-researchers exhibit no response.

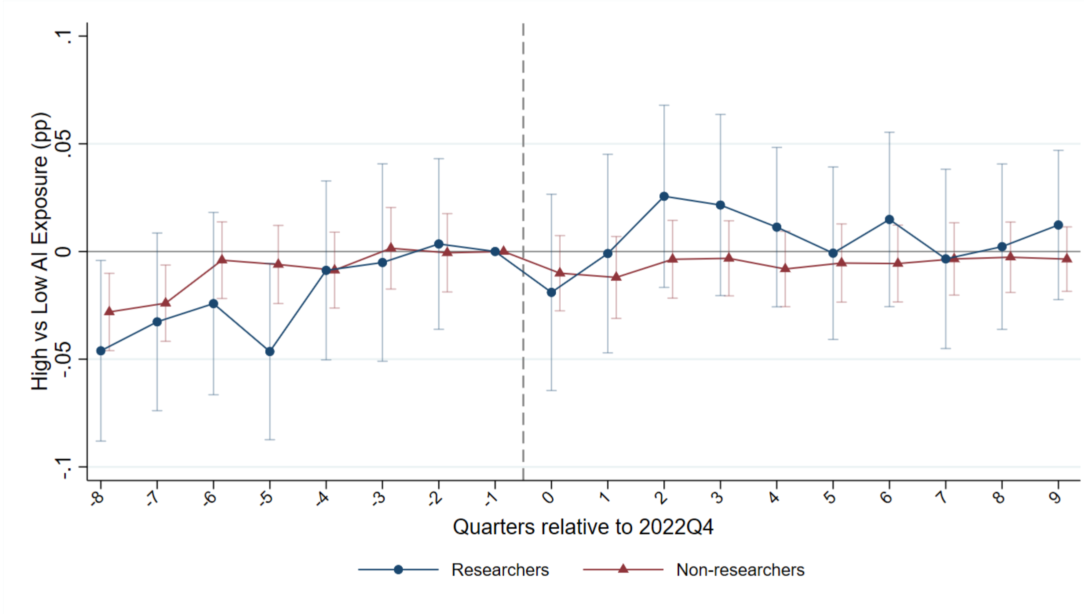

Figure 4: Dynamic DiD: Probability of first founding around the ChatGPT release (at-risk sample only). Treatment, fixed effects, and SE as in Figure 3.

Re-entrants. To decompose the founder-stock effect, we isolate workers who had previously founded but were not active founders as of 2022Q3 ("re-entrants" — dormant serial entrepreneurs). Figure 5 presents the DiD event study for this group. Both researchers and non-researchers show flat pre-trends, but the post-period dynamics differ strikingly from the main sample. Non-researchers respond quickly, with positive effects visible from $k = +2$ (2023Q2) onward, while researchers exhibit a longer lag, with the effect emerging around $k = +4$ (2023Q4). This reversal in timing is consistent with the model: re-entering non-researchers already have business-formation experience and primarily needed an enabling technology, while researchers spend more time evaluating technical opportunities before committing.

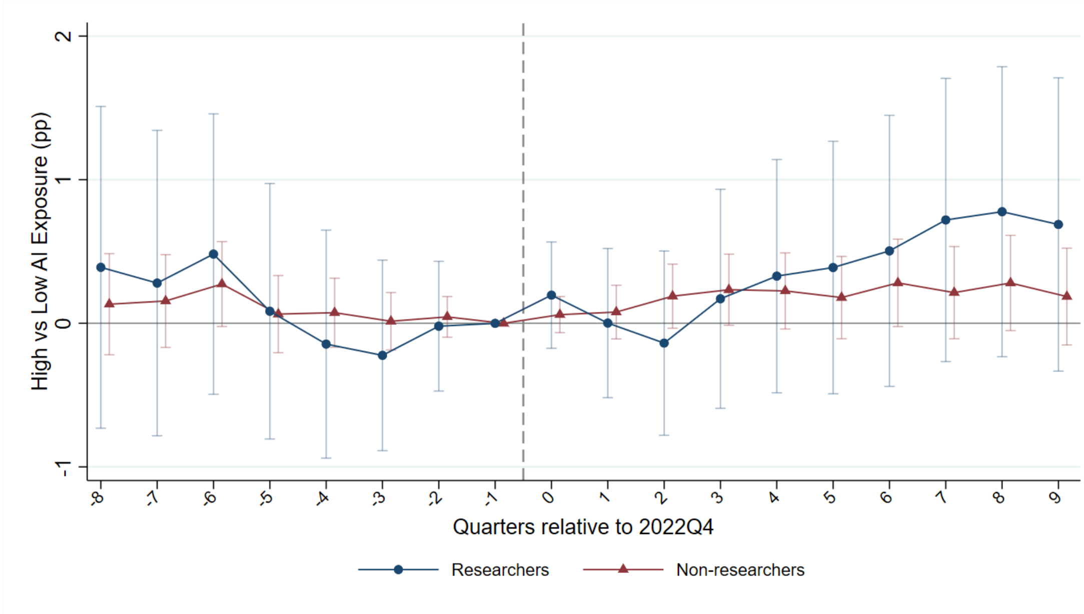

Figure 5: Dynamic DiD: Probability of re-entering founding activity, restricted to prior founders (at-risk = 0) who were not active founders in 2022Q3. Treatment, fixed effects, and SE as in Figure 3.

Figure 6 presents the triple-difference (DDD) event study for the re-entrant sample. The plotted coefficients $\hat{\gamma}_k$ measure the *excess* re-entry response of researchers in high-AI industries relative to non-researchers. Pre-period $\hat{\gamma}_k$ close to zero confirm that the two groups were on parallel trends before 2022Q4. The positive post-period triple interaction confirms that, even within the re-entrant sample, researchers in AI-intensive industries disproportionately returned to active founding. DDD event studies for the full and at-risk samples are in Appendix I.

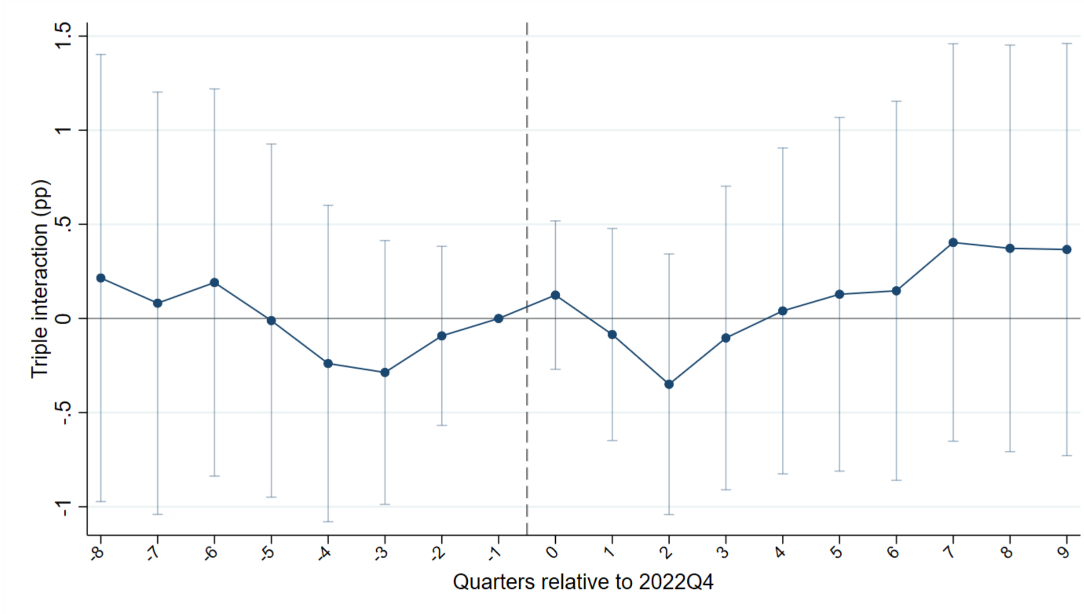

Figure 6: Triple-difference (DDD) event study for re-entrant founding. Sample: prior founders with Anthropic coverage. The plotted coefficient $\hat{\gamma}_k$ measures the excess response of researchers in high-AI industries relative to non-researchers. Base period: 2022Q3. Fixed effects: worker, quarter, 2-digit industry. SE clustered by worker.

4.4 Experience Heterogeneity

Proposition 1 predicts that the AI effect on entrepreneurship should be increasing in the stock of pre-existing researcher experience, since workers closer to the balanced-skills threshold face a lower marginal cost of entry. Figures 7 and 8 test this prediction non-parametrically by estimating the DiD coefficient separately within four work-experience bins (0–4, 5–9, 10–14, and 15+ years of cumulative pre-2022Q4 experience).

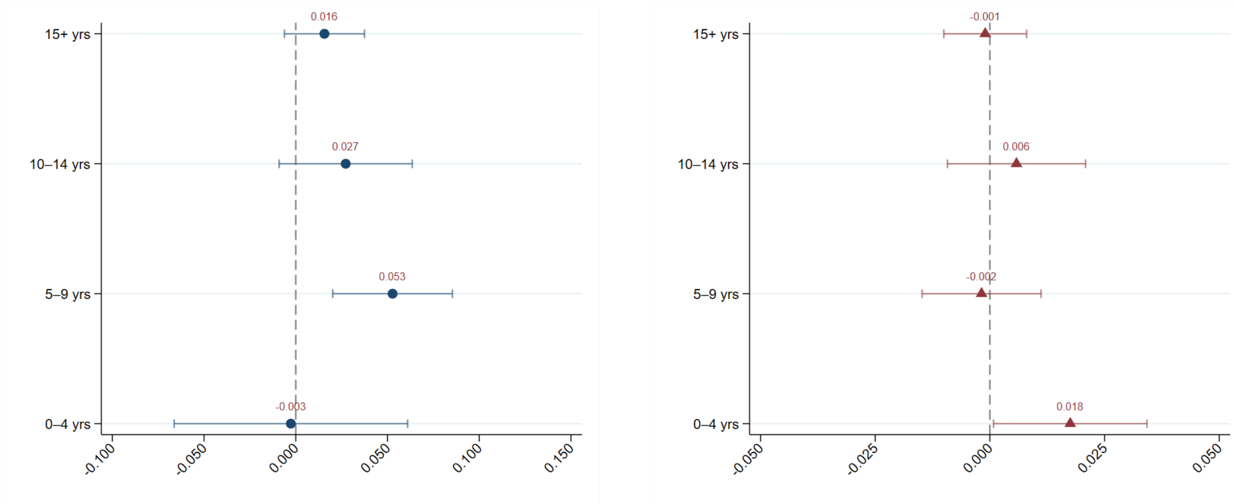

(a) Researchers (b) Non-Researchers

Figure 7: DiD coefficient on Post × High-AI-Exposure for the probability of first founding, by experience bin and researcher status. Left panel: researchers; right panel: non-researchers. Fixed effects: worker, quarter, 2-digit industry. SE clustered by worker. **The two panels use different x-axis scales: the researcher panel spans −0.10 to +0.15 pp and the non-researcher panel spans −0.05 to +0.05 pp.** Markers show 95% confidence intervals.

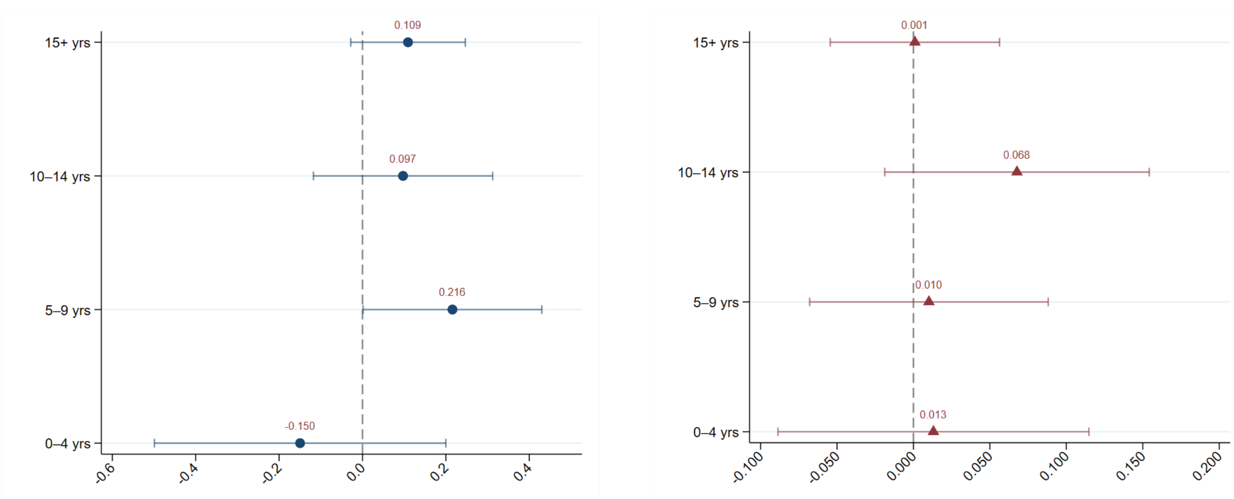

(a) Researchers (b) Non-Researchers

Figure 8: DiD coefficient on Post × High-AI-Exposure for the probability of currently being a founder (founder stock), by experience bin and researcher status. Left panel: researchers; right panel: non-researchers. Fixed effects: worker, quarter, 2-digit industry. SE clustered by worker. **Panels use different x-axis scales: the researcher panel spans −0.60 to +0.50 pp and the non-researcher panel spans −0.10 to +0.20 pp, reflecting the considerably wider dispersion of researcher coefficients in the founder-stock outcome.** Markers show 95% confidence intervals.

For researchers, the DiD effect is near zero for workers with fewer than five years of experience, grows through the 5–9 and 10–14 bins, and then declines modestly for the most experienced group (15+ years). This hump-shaped pattern is consistent with the model: workers with very little experience have not yet accumulated enough researcher human capital to clear the entrepreneurial threshold, while the most experienced researchers were likely already sorting into entrepreneurship before the shock. For non-researchers, the estimated effects are indistinguishable from zero across all experience bins for both outcomes. The experience gradient is exclusively a researcher phenomenon, confirming that it is the *interaction* of researcher human capital with AI exposure — not general work experience — that drives entrepreneurial entry.

5 Discussion

5.1 Anthropic Adoption Measure vs. GPT Capability Measure

A key identification decision is the choice of AI exposure measure. The Anthropic Labour Market Impact score captures *actual AI adoption* within industries, while the GPT-4 capability index (Eloundou et al., 2024) measures the *potential capability* of LLMs to automate or augment occupation-level tasks. Comparing the two measures provides a window into the mechanism: it is adoption — translating capability into deployment — that drives the entrepreneurial margin.

Figures 9, 10, and 11 present the dynamic DiD event studies using the GPT capability measure as the treatment, mirroring the Anthropic analysis in Section 3.5.

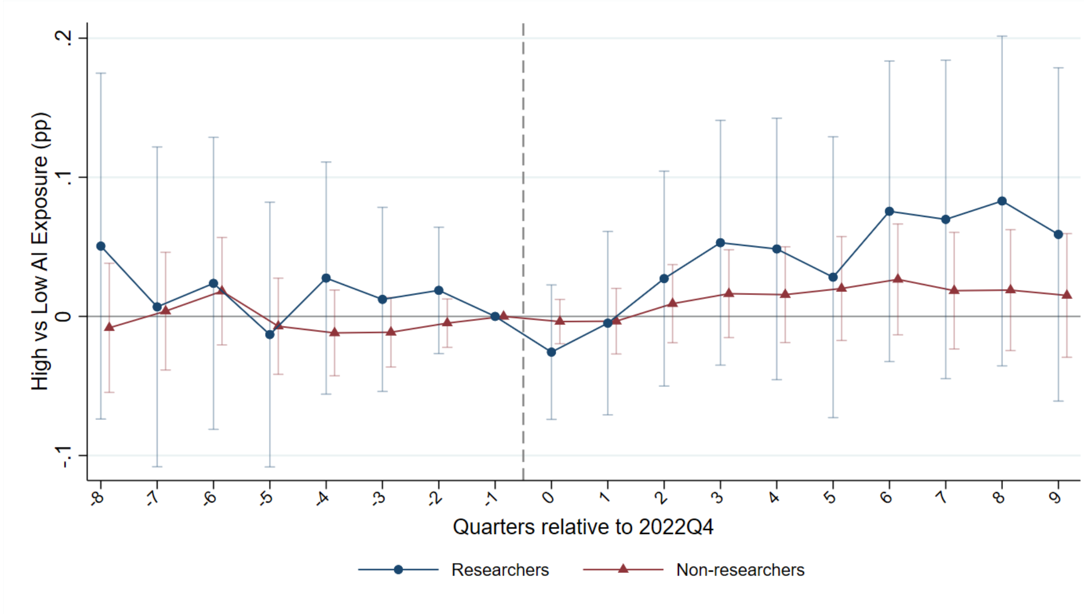

Figure 9: Dynamic DiD: Probability of being a founder around the ChatGPT release. Treatment: above-median GPT industry AI exposure score (Eloundou et al., 2024); base period: 2022Q3. Fixed effects: worker, 2-digit industry, quarter. SE clustered by worker. Navy circles: researchers; maroon triangles: non-researchers.

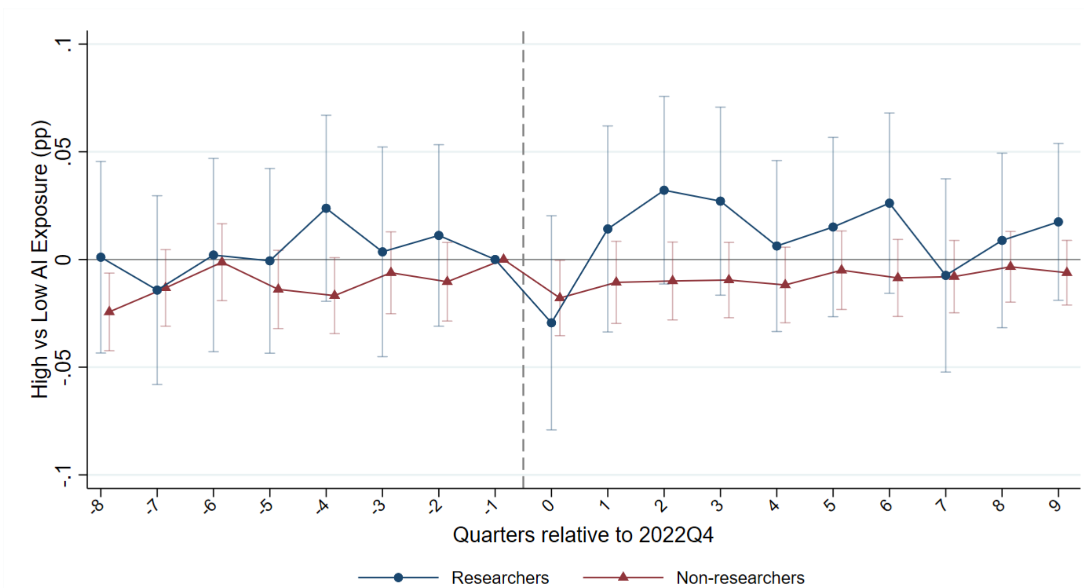

Figure 10: Dynamic DiD: Probability of first founding around the ChatGPT release (at-risk sample). Treatment: above-median GPT industry AI exposure. Fixed effects and SE as in Figure 9.

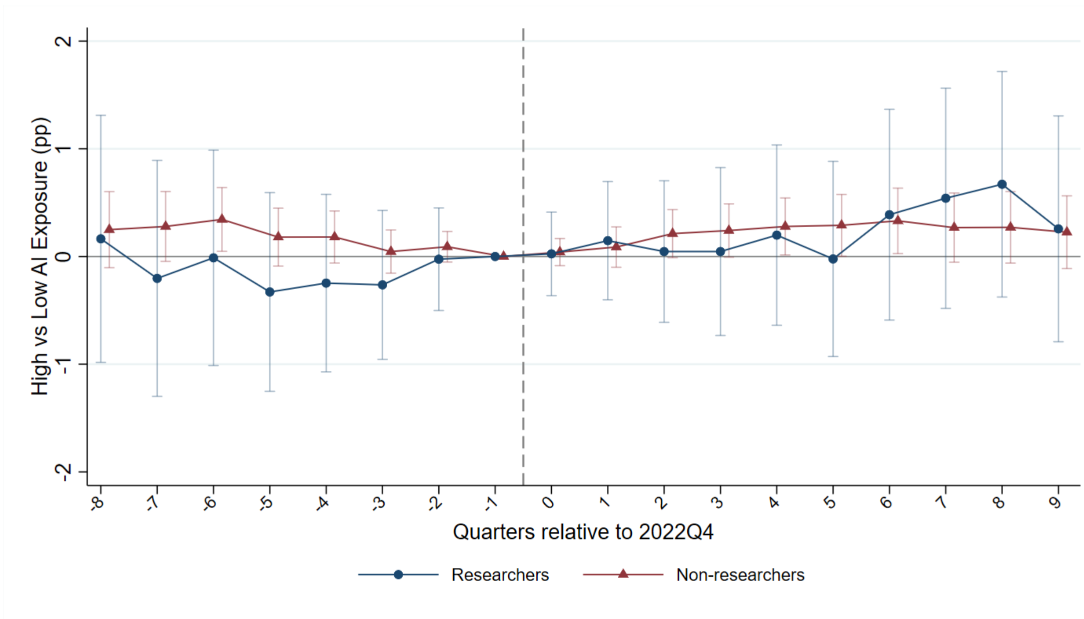

Figure 11: Dynamic DiD: Probability of re-entering founding activity among prior founders (at-risk = 0), using GPT industry AI exposure as treatment. Fixed effects and SE as in Figure 9.

The GPT measure produces a directionally consistent but significantly attenuated pattern. For the founder-stock outcome (Figure 9), researcher coefficients are positive after 2022Q4 (roughly 0.05–0.08 pp) but individually insignificant at conventional levels. Non-researchers are flat throughout. For first-founding and re-entry, neither group shows a discernible trend.

The contrast is informative. The GPT capability measure captures which industries could, in principle, be affected by LLM-era AI; the Anthropic measure captures which industries *actually deployed* AI tools at scale. The fact that the Anthropic measure yields far stronger and more robust estimates suggests that it is the realized adoption of AI — not merely its theoretical reach — that unlocks the entrepreneurial margin for researchers. Workers respond to tools they actually use, not to the abstract potential encoded in capability indices. This finding aligns with the broader literature distinguishing potential AI exposure from actual automation (Autor et al., 2003; Acemoglu et al., 2022) and reinforces our interpretation that the mechanism operates through lowering the practical barriers to venture creation.

5.2 First Wave of AI: Replication and Comparison

The previous sections document the entrepreneurial response to the second wave of AI. To assess whether this pattern is a new phenomenon or echoes an earlier structural break, we briefly present results from the first wave, defined by the AlexNet deep learning breakthrough in 2012. The full analysis and robustness checks are in Appendix G; here we highlight the key parallels and contrasts.

Figure 12 shows the raw unconditional trends in first-founding probability by AI exposure group around 2012. High-AI-exposure industries diverged upward relative to low-exposure industries after 2012, while mid-exposure industries tracked closely with the low group, providing the visual motivation for the DiD design.

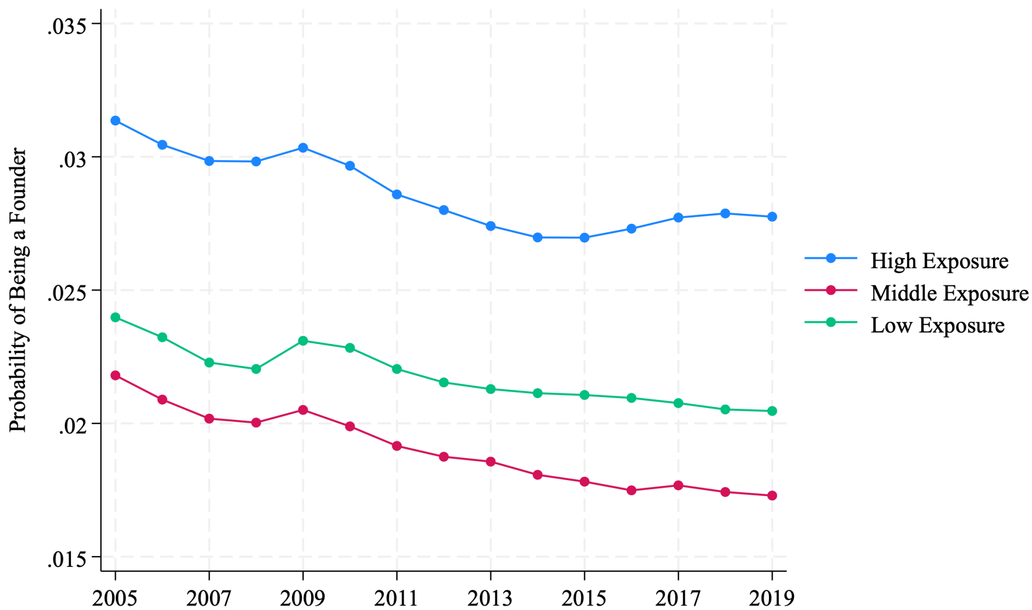

Figure 12: Annual probability of first founding by AI exposure group, 2007–2019. Groups defined by AI Industry Impact Exposure score (Felten et al., 2021). Dashed vertical line: 2012 (AlexNet shock).

Table 3 presents the main DiD estimates for the first wave: the aggregate effect is near zero, but researchers in high-AI-exposure industries show a 58% increase in the annual first-founding rate (column 3), while non-researchers show no significant response (column 4). The event study in Figure 13 confirms flat pre-trends and a gradual post-2012 divergence concentrated among researchers.

Table 3: First Wave of AI and Entrepreneurship: Main DiD Estimates

	Pr(First Founding)			
	(1)	(2)	(3)	(4)
	No Worker FE	Worker FE	Researchers	Non-researchers
Post × High Ind	0.0000	0.0002	0.0015**	-0.0000
	(0.0003)	(0.0004)	(0.0006)	(0.0004)
Control mean	0.0089	0.0026	0.0026	0.0026
Industry FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Worker FE		✓	✓	✓
N	5203609	5156612	735991	4420621

Notes: LinkedIn worker profiles 2007–2019, at-risk sample. Outcome: probability of first founding in year t (×100). Treatment: above-median AI Industry Impact Exposure score (Felten et al., 2021), fixed to workers' modal pre-2012 industry. All columns include year and 2-digit industry fixed effects; columns (2)–(4) add worker fixed effects. Column (3): researchers; column (4): non-researchers. SE clustered by modal pre-2012 4-digit industry (≈300 clusters). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

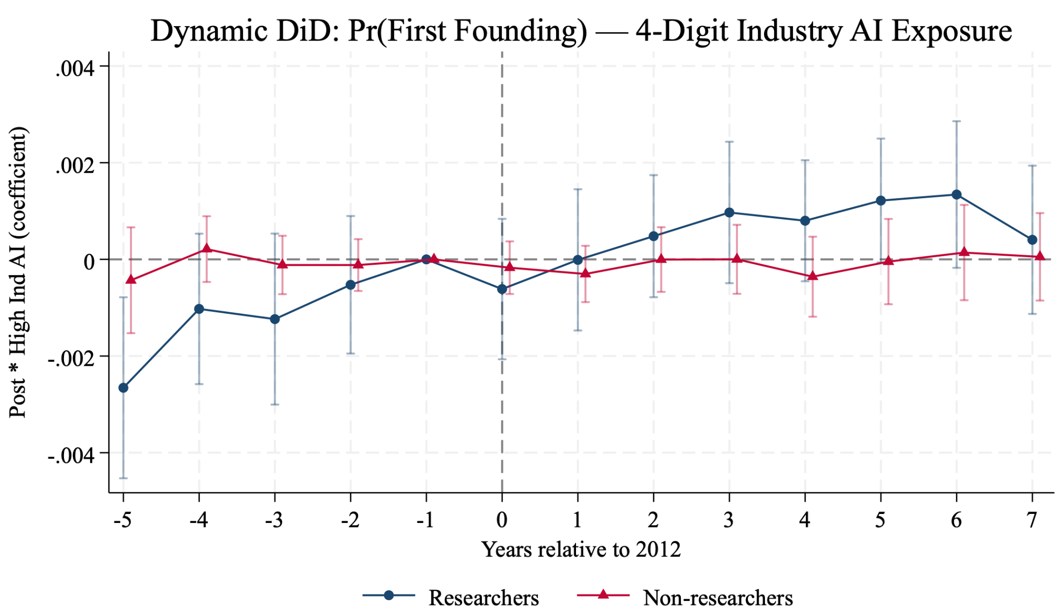

Figure 13: Dynamic DiD: Probability of first founding, researchers vs. non-researchers, first wave (AlexNet, 2012). Baseline year: 2011. Navy circles: researchers; maroon triangles: non-researchers. 95% CIs shown. SE clustered by modal pre-2012 4-digit industry.

The qualitative pattern across both waves is strikingly consistent: in each case, the AI shock raises entrepreneurial entry specifically among workers with technical research experience, while leaving non-researchers unaffected. The magnitude is comparable — a 58% increase in the baseline founding rate for researchers in the first wave — and the event-study dynamics show the same gradual post-shock divergence between researchers and non-researchers. This cross-wave consistency strengthens the interpretation that the mechanism is structural rather than wave-specific: AI-driven increases in the productivity premium of research human capital persistently lower the skill-gap barrier to entrepreneurship for technically trained workers.

The two waves differ in one important respect: the first-wave effect is concentrated in the *first founding* margin (the at-risk sample), reflecting an increase in new venture creation; the second-wave effect is also visible in the *stock* of active founders, suggesting that the ChatGPT shock additionally sustained founders who might otherwise have exited. This difference is consistent with the generative-AI wave reducing ongoing operational costs for active founders — a channel beyond the initial entry decision that was less available in the more specialized first wave.

6 Conclusion

This paper studies the impact of Artificial Intelligence on the returns to experience in entrepreneurship. We distinguish between two main types of experience: general work experience, which is repetitive and prone to substitution by AI, and research experience, which is otherwise hard to substitute. We confirm that AI increases the rate of entrepreneurship and the share of researchers via a difference-in-differences analysis using LinkedIn data. We also find that the average work experience before becoming a founder has increased. Importantly, research experience is more valuable, and we document a shift towards more experienced researchers starting their own businesses. Through the lens of our quantitative model with multidimensional skill and human capital accumulation, we find that AI primarily increases the productivity premium for research experience, biasing the labor market towards individuals with such backgrounds. This has important implications for understanding how AI is reshaping the entrepreneurial landscape and the value of different types of experience in this context.

Our work is among the first to empirically investigate the intersection of AI and entrepreneurship through multidimensional skill frameworks. This is an active area of research, and we aim to carefully estimate the model parameters and provide quantitative results on how AI impacts entrepreneurship and the returns to experience.

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.

Autor, D. and N. Thompson (2025, jun). Expertise. Working Paper 33941, National Bureau of Economic Research.

Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of economic perspectives 29*(3), 3–30.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics 118*(4), 1279–1333.

Azoulay, P., B. F. Jones, J. D. Kim, and J. Miranda (2020). Age and high-growth entrepreneurship. *American Economic Review: Insights 2*(1), 65–82.

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

Blanchard, O. J. (1985). Debt, deficits, and finite horizons. *Journal of political economy 93*(2), 223–247.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). GPTs are GPTs: An early look at the labor market impact potential of large language models. *Science 384*(6702), 1306–1308.

Felten, E., M. Raj, and R. Seamans (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal 42*(12), 2195–2217.

Fossen, F. M., T. McLemore, A. Sorgner, et al. (2024). Artificial intelligence and entrepreneurship. *Foundations and Trends® in Entrepreneurship 20*(8), 781–904.

Gofman, M. and Z. Jin (2024). Artificial intelligence, education, and entrepreneurship. *The Journal of Finance 79*(1), 631–667.

Hampole, M., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2025). Artificial intelligence and the labor market. Technical report, National Bureau of Economic Research.

Lazear, E. P. (2004). Balanced skills and entrepreneurship. *American Economic Review 94*(2), 208–211.

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

Massenkoff, M. and P. McCrory (2026, mar). Labor market impacts of AI: A new measure and early evidence. Report, Anthropic.

Menzio, G. and S. Shi (2010). Block recursive equilibria for stochastic models of search on the job. *Journal of Economic Theory 145*(4), 1453–1494.

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

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.

Appendices

A Data Description

A.1 Sample Construction

Our primary dataset comes from Revelio Labs, which processes public professional profile data. The worker profile data contain detailed information on individuals' work histories, including job titles, employers, dates of employment, location, and industry codes (NAICS). Revelio Labs applies machine learning and natural language processing to extract and standardize information, and also imputes variables such as gender, race, seniority, and salaries.

For the analysis, we use a sample of U.S. LinkedIn users. Position start and end dates are essential for constructing quarterly spell observations; we keep only positions with a non-missing start date and impute a missing end date as the current date. Positions with an end date earlier than the start date are dropped. After cleaning, the sample for the first-wave analysis contains approximately 0.95 million unique users and 3.8 million unique positions; the second-wave (Revelio GenAI) sample extends the time horizon through 2025Q1.

A.2 Identifying Founders and Researchers

Founders. We identify founders and entrepreneurs by searching for the keyword list reported in Section 3.

To avoid misclassifying senior executives at established firms as founders, we set the founder indicator to zero if the company's founding year is five or more years before the worker's start date in that position.

Researchers. Workers are classified as having prior research experience if any of the researcher keywords reported in Section 3 appear in their job titles. We exclude "research assistant" and "research intern" to focus on post-education professional experience. For the second-wave analysis, the researcher classification is fixed to the pre-2022Q4 period (`wr_pre_s2`), ensuring it is not endogenous to the ChatGPT shock.

A.3 Measuring Work Experience

Work experience is not directly observed. We proxy it by counting the years worked in each position after the worker's highest-education end date. When the education end date is missing, we impute it as the start year of the worker's first observed position. We further distinguish *researcher experience* (years in researcher-classified positions) from *non-researcher experience* (remaining years).

A.4 AI Exposure Measures

First wave (Felten-Raj-Seamans). We use the AI Industry Impact Exposure (AIIE) scores constructed by Felten et al. (2021). These are occupation-level AI exposure (AIOE) scores, weighted by industry employment shares to yield a 4-digit NAICS industry exposure score. For the first-wave analysis, each worker's AI exposure is determined by their modal pre-2012 4-digit NAICS industry — the industry in which they spent the most time during 2007–2011 — to avoid post-shock sorting bias. AIIE ranks and the stability of the classification under alternative employment weights are documented in Appendix C.

Second wave (Anthropic and GPT). For the second-wave analysis, we use (i) the Anthropic Labour Market Impact score, which captures actual AI adoption intensity at the industry level, and (ii) the GPT-4 industry exposure score from Eloundou et al. (2024), which measures the theoretical capability of LLMs to automate or augment occupation-level tasks. Each worker's exposure score is computed as the time-weighted average of the relevant industry score over all pre-2022Q4 quarters in which the worker held a non-founder position with a valid industry code.

B First Wave: Empirical Strategy Details

B.1 AI Shock

We exploit the diffusion of AI capabilities as an exogenous source of variation across industries and over time. The key event is the victory of AlexNet at the ImageNet Large Scale Visual Recognition Challenge in late 2012. This deep convolutional neural network achieved an error rate of 15.3%, more than 10 percentage points lower than the runner-up, demonstrating the practical power of deep learning for image recognition and catalyzing broad AI adoption across industries. We define $\text{Post}_t = \mathbf{1}[t \geq 2012]$.

B.2 Treatment Construction

Each worker's AI exposure is assigned based on their modal pre-2012 4-digit NAICS industry (the industry in which they accumulated the most position-years during 2007–2011). The AIIE score for that industry is fixed as the worker's treatment and classified into a binary indicator $\text{High}_i = \mathbf{1}[\text{AIIE}_i > \text{median}]$. Workers with no pre-2012 record are assigned their modal industry across all years and flagged with a separate indicator. Agriculture (NAICS 11) and Construction (NAICS 23) are excluded throughout.

B.3 Estimating Equations

The baseline DiD specification is:

$$\Pr(\text{First Founding})_{it} = \alpha + \beta\,(\text{Post}_t \times \text{High}_i) + \lambda_i + \gamma_t + \delta_{n(i,t)} + \varepsilon_{it}, \quad (2)$$

where λ_i are worker fixed effects, γ_t year fixed effects, and $\delta_{n(i,t)}$ current-year 2-digit NAICS fixed effects. Standard errors are clustered by modal pre-2012 4-digit NAICS ($\approx$300 clusters).

The event study specification replaces Post_t with year-relative dummies:

$$\Pr(\text{First Founding})_{it} = \alpha + \sum_{\substack{k=-5 \\ k\neq -1}}^{7} \beta_k\,\mathbf{1}[t-2012=k] \times \text{High}_i + \lambda_i + \gamma_t + \delta_{n(i,t)} + \varepsilon_{it}, \quad (3)$$

with baseline year 2011 ($k = -1$).

For robustness, we also estimate a version adding industry-specific linear time trends $\theta_{n(i)} \cdot t$:

$$\Pr(\text{First Founding})_{it} = \beta\,(\text{Post}_t \times \text{High}_i) + \lambda_i + \gamma_t + \delta_{n(i,t)} + \theta_{n(i)} \cdot t + \varepsilon_{it}. \quad (4)$$

C AIIE Weight Stability: 2019 vs. 2010 Employment Weights

C.1 Motivation

The AI Industry Impact Exposure (AIIE) scores from Felten, Raj, and Seamans (2021) are constructed as employment-weighted averages of occupation-level AI exposure (AIOE) within each 4-digit NAICS industry, using 2019 BLS Occupational Employment Statistics (OES) employment shares. Because our sample ends in 2019 and the treatment event is in 2012, a concern arises: the 2019 employment composition may itself reflect the AI-driven occupational recomposition we are studying. Industries that responded most strongly to AI may have absorbed more AI-exposed workers by 2019, making their measured AIIE higher precisely because they were treated. If so, AIIE would not be a fully predetermined measure of exposure.

C.2 Method

We address this concern by reconstructing AIIE using 2010 BLS OES employment data (the May 2010 National Industry-Specific file), which predates the 2012 AI shock. For each 4-digit NAICS industry n, the re-weighted measure is:

$$\widehat{\text{AIIE}}_{2010}(n) = \sum_{o} \frac{\text{emp}_{2010}(o, n)}{\sum_{o'} \text{emp}_{2010}(o', n)} \cdot \text{AIOE}(o) \tag{5}$$

using the same AIOE scores from Appendix A of Felten et al. (2021). The 4-digit values are then aggregated to 2-digit NAICS using 2010 employment-weighted means and standardised (mean 0, SD 1) to match the original AIIE scale. The 2010 OES uses 2010 SOC codes; 736 of the 774 AIOE occupations (95.1%) match directly, with the 38 unmatched codes being occupations introduced in the 2018 SOC revision.

C.3 Results

Figure A1 plots the standardised AIIE under 2019 vs. 2010 employment weights across 2-digit NAICS industries. The two measures are highly correlated: the Spearman rank correlation is $\rho = 0.967$ ($p < 0.001$). More importantly, *no industry changes its High-/Low/Middle classification* when switching from 2019 to 2010 weights. The industries that rank as high-exposure under the original measure (Information, Finance, Professional & Technical Services, Management of Firms) remain high-exposure under the pre-shock weighting, and vice versa for low-exposure industries.

The near-perfect rank agreement between pre- and post-shock employment weights indicates that the potential endogeneity of the 2019 weights is empirically negligible. The structural features captured by AIOE — the intrinsic AI-complementarity of occupational task requirements — dominate the industry-level exposure ranking regardless of whether 2010 or 2019 employment shares are used. We therefore proceed with the original Felten et al. AIIE without modification.

D 4-Digit Industry Coverage and Treatment Validation

This appendix documents the construction and validity of the modal pre-2012 4-digit NAICS industry treatment variable. We present six descriptive figures that address coverage, dominance of the modal assignment, and industry sorting patterns around the 2012 AI shock.

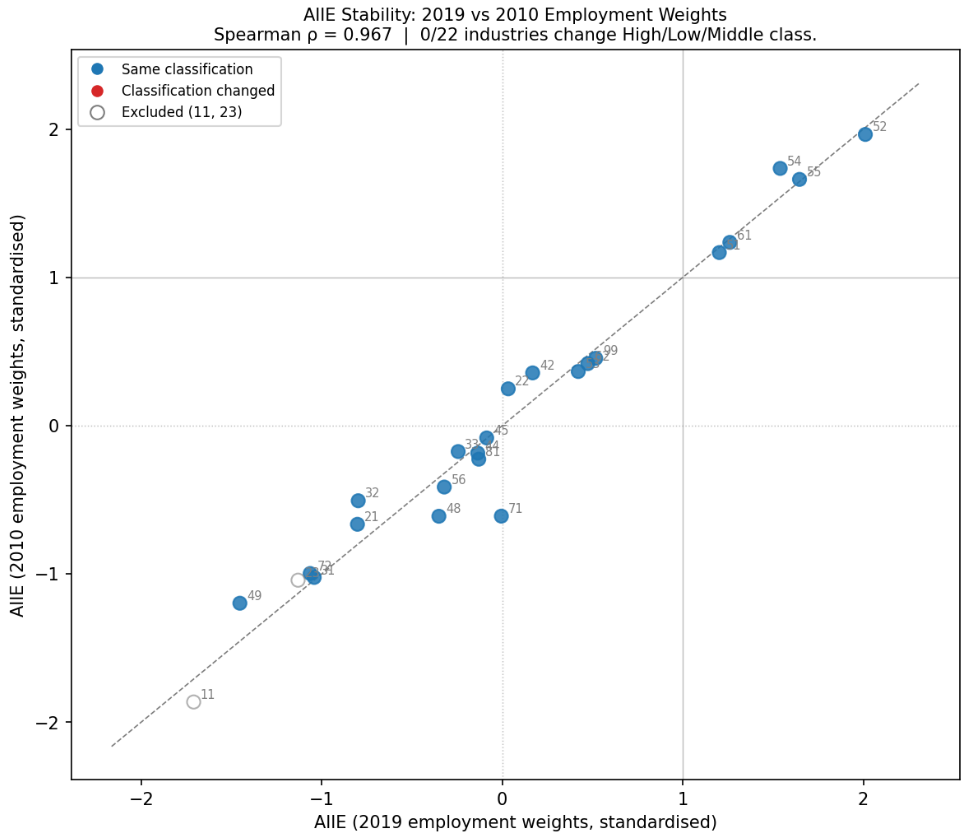

Figure A1: AIIE Stability: 2019 vs. 2010 Employment Weights

Notes: Each point is a 2-digit NAICS industry. The x-axis is the AIIE from Felten et al. (2021) standardised using the mean of 4-digit values. The y-axis is the AIIE reconstructed using May 2010 BLS OES employment weights and the same AIOE scores. Vertical and horizontal dashed lines mark the High (≥ 1) and Low (< 0) classification thresholds. No industry changes classification. Spearman $\rho = 0.967$.

AIIE distribution across sectors. Figure A2 shows the distribution of AIIE scores by 2-digit NAICS sector. Information (NAICS 51), Finance and Insurance (52), and Professional, Scientific, and Technical Services (54) cluster in the high-exposure range (AIIE ≥ 1). Retail Trade (44–45), Accommodation and Food Services (72), and Health Care (62) populate the low-exposure range. This cross-sector variation is the primary source of identification in our DiD design.

Modal industry dominance. A valid modal assignment requires that a worker's modal industry is a reliable signal of their primary industry of attachment, not an artifact of near-ties across many industries. Figure A3 plots the distribution of modal dominance — the share of a worker's pre-2012 position-years spent in their modal industry. The mean dominance is 0.61 and the median is 0.62, indicating that the typical worker spends the majority of their pre-2012 career in a single industry. Less than 10% of workers have a

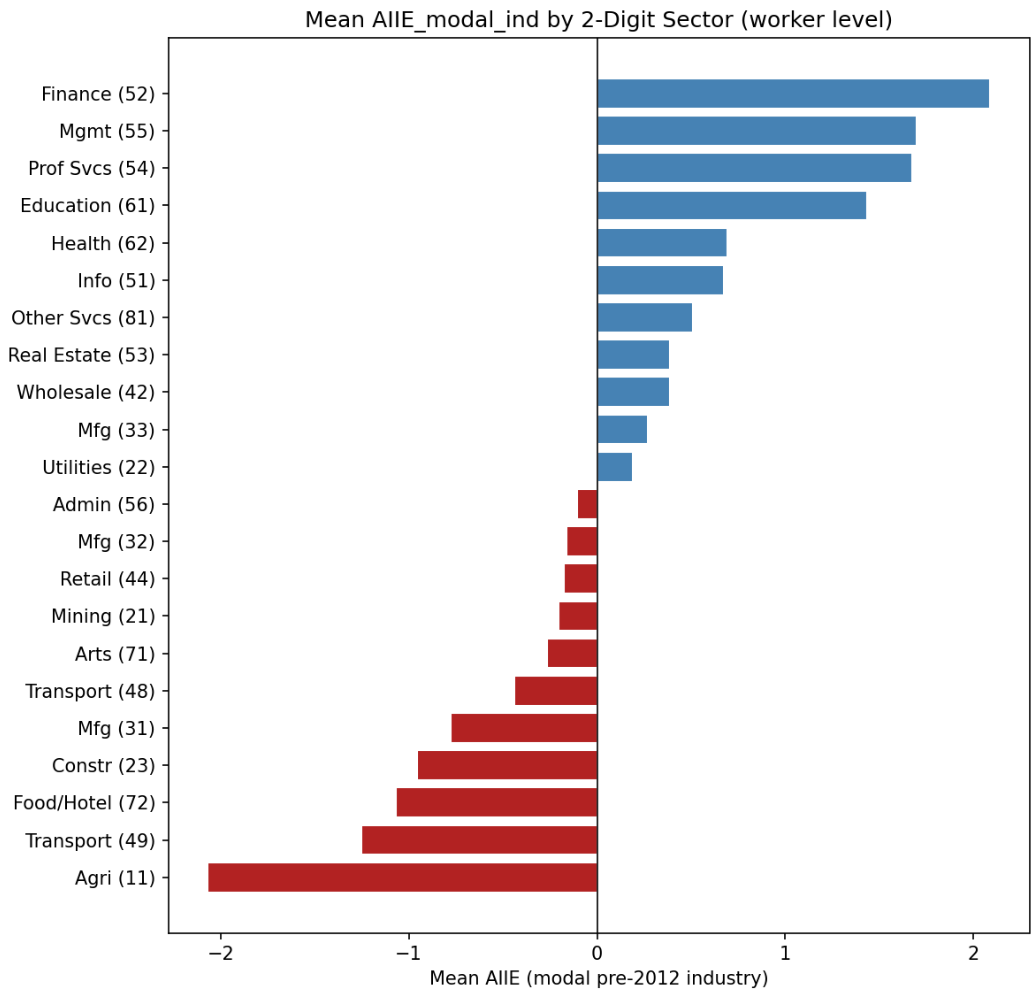

Figure A2: AIIE Score Distribution by 2-Digit NAICS Sector
Notes: Each point is a 4-digit NAICS industry. Box plots show the interquartile range of AIIE within each 2-digit sector. Dashed lines mark the High (≥ 1) and Low (< 0) treatment thresholds.

modal share below 0.40, suggesting near-ties are uncommon.

Year-to-year industry switching. If workers systematically move into high-AI industries *after* 2012 in response to the AI wave, a time-varying industry assignment would be contaminated by the outcome of interest. Figure A4 plots the annual rate of year-to-year 4-digit NAICS industry switches. Switching rates are stable before and after 2012 and show no discontinuity at the treatment year, supporting the assumption that our fixed modal assignment is not confounded by post-AI sorting.

Pre/post industry composition. Figure A5 directly tests for post-2012 compositional shifts by comparing the industry distribution of workers in the pre-period (2007–2011)

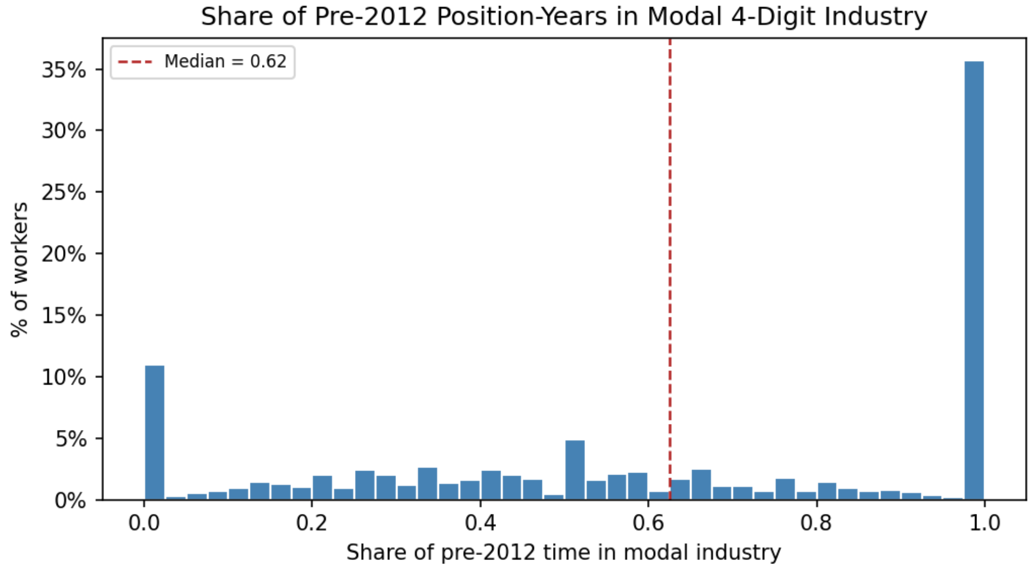

Figure A3: Distribution of Modal Industry Dominance (Pre-2012)
Notes: Modal dominance = share of pre-2012 position-years in worker's modal 4-digit NAICS industry. Sample restricted to workers with at least one pre-2012 observation.

and post-period (2012–2019). The distributions are nearly identical, confirming that the aggregate industry mix of our sample does not change after the AI shock.

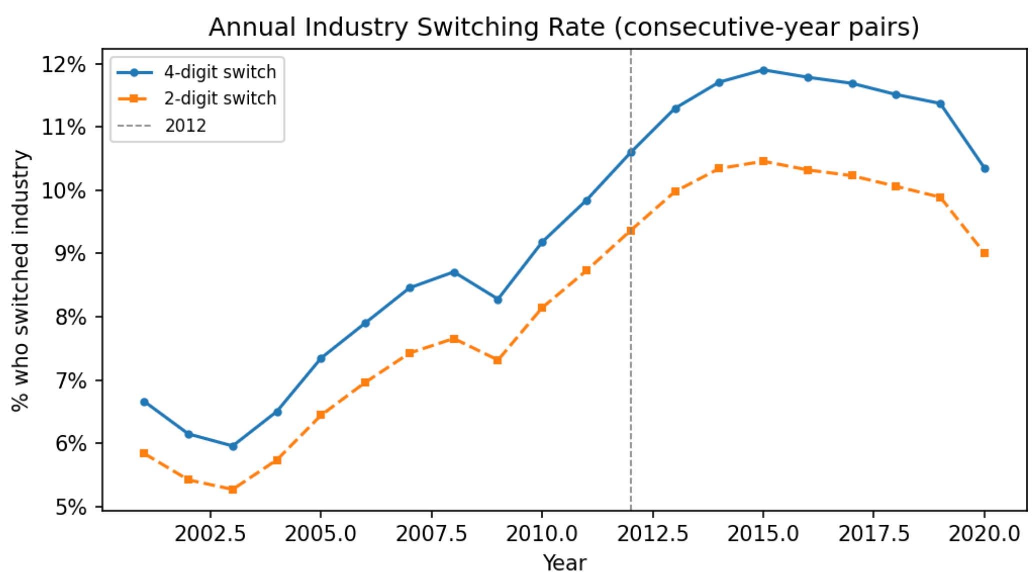

Figure A4: Year-to-Year Industry Switching Rate, 2007–2019
Notes: Share of person-year observations in which the worker's current 4-digit NAICS industry differs from the prior year. Dashed vertical line marks 2012.

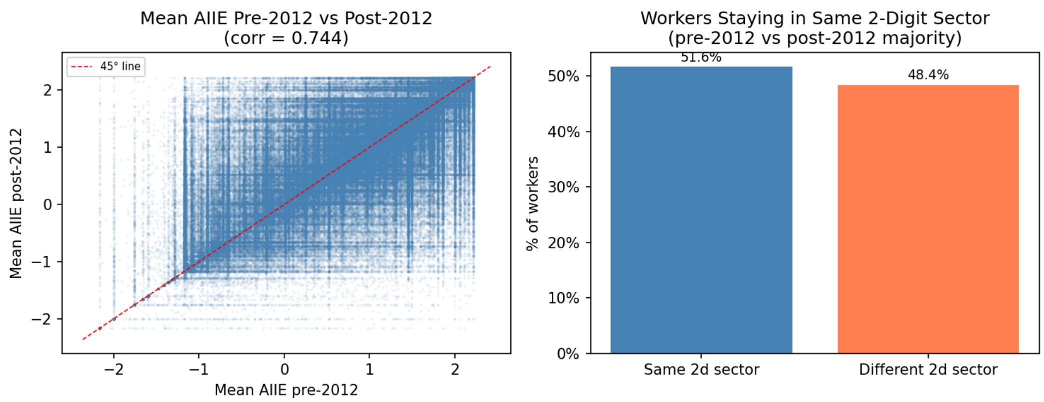

Figure A5: Industry Composition Pre- vs. Post-2012

Notes: Share of person-year observations in each 2-digit NAICS sector, averaged over the pre-period (2007–2011) and post-period (2012–2019).

AIIE coverage by year. Figure A6 plots the AIIE match rate by year. Coverage is stable across the sample period at approximately 57–58% of all position-years (77–78% of position-years with any LinkedIn industry code), with no systematic trend before or after 2012. Unmatched observations arise from LinkedIn NAICS codes that do not correspond to any FRS code, even at the 3-digit level; these are excluded from the analysis sample.

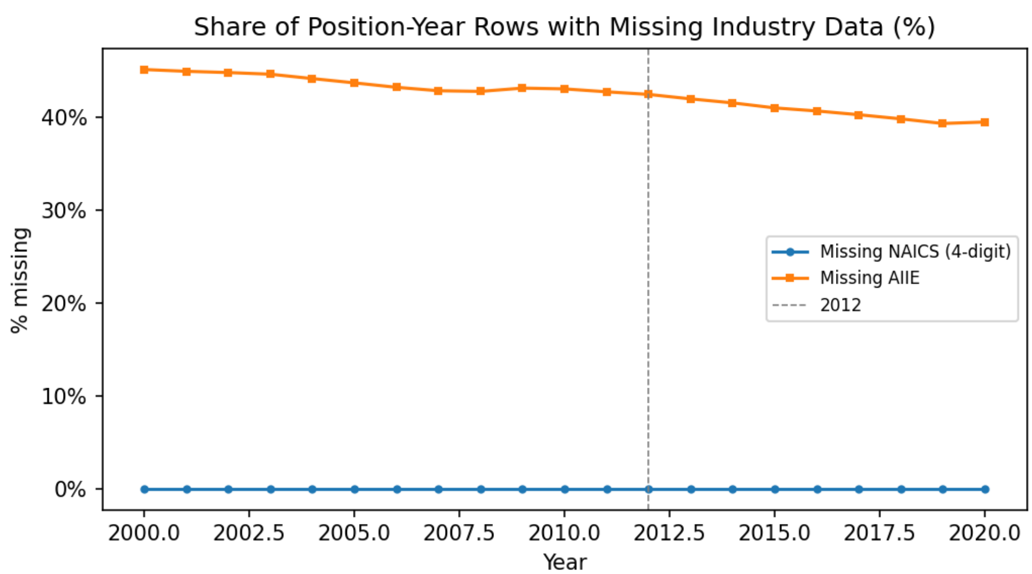

Figure A6: AIIE Coverage Rate by Year

Notes: Share of person-year observations with a valid AIIE match, by calendar year.

Workers observed in multiple industries per year. Figure A7 plots the fraction of workers observed in more than one 4-digit NAICS industry in a given year (due to overlapping position spells). This fraction is below 15% in most years and stable over time, confirming that the modal assignment is well-defined for the vast majority of workers and is not sensitive to concurrent multi-industry spells.

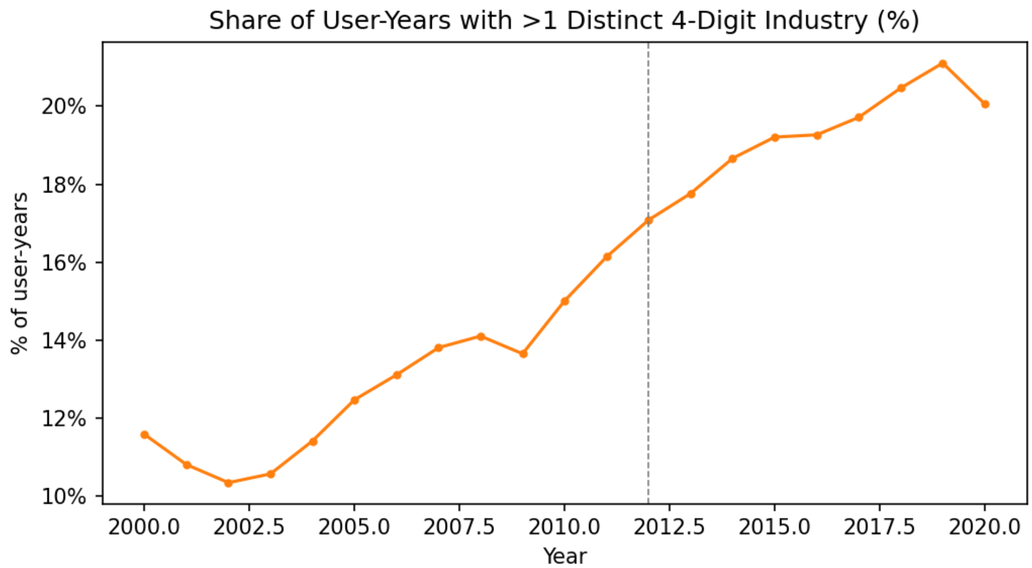

Figure A7: Share of Workers with Multiple Industries in a Given Year
Notes: Share of workers for whom at least two distinct 4-digit NAICS codes are observed within a calendar year (due to overlapping employment spells). Dashed vertical line marks 2012.

E Alternative Treatment Specifications

This appendix examines the robustness of the main findings to alternative ways of classifying AI exposure. The baseline specification uses a binary indicator for above-median AIIE (High_i). We consider two alternatives: a three-category High/Middle/Low classification and a continuous AIIE measure.

E.1 High/Middle/Low Classification

Table A1 replaces the binary High indicator with a three-way split: High (AIIE ≥ 1), Middle ($0 \leq$ AIIE < 1), and Low (AIIE < 0), with Low as the baseline. This finer classification tests whether the researcher effect is driven by the top of the AIIE distribution or is present more broadly across above-median industries.

The Middle category shows no significant differential effect relative to the Low baseline in any column or sample. The researcher-specific effect is concentrated in the High group: $\hat{\beta}_{\text{High}} = 0.0014^{*}$ for researchers (column 2), consistent with Table 3. This con-

Table A1: AI Exposure and First Founding: High/Middle/Low Treatment

	Pr(First Founding)		
	(1)	(2)	(3)
	All workers	Researchers	Non-researchers
Post × High Ind	0.0001	0.0014*	-0.0002
	(0.0004)	(0.0007)	(0.0004)
Post × Middle Ind	-0.0003	-0.0003	-0.0003
	(0.0004)	(0.0007)	(0.0004)
Control mean (Low)	0.0027	0.0028	0.0027
Industry FE	✓	✓	✓
Year FE	✓	✓	✓
Worker FE	✓	✓	✓
N	5156612	735991	4420621

Baseline group: Low AI exposure (AIIE < 0).

Notes: As Table 3 columns (2)–(4). Treatment groups: High (AIIE ≥ 1), Middle (0 ≤ AIIE < 1), Low (AIIE < 0, baseline). All columns include worker, year, and current-year 2-digit NAICS industry fixed effects. Standard errors clustered by modal pre-2012 4-digit NAICS code. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

firms that the founding response is concentrated among workers in the most AI-intensive industries, not merely above-average ones.

E.2 Continuous AIIE

Table A2 replaces the binary treatment with the continuous AIIE score, interacted with Post_t. This specification exploits the full cross-industry variation in AI exposure rather than discretising at the median.

The pattern is consistent with Table 3: the aggregate and non-researcher estimates are near zero and insignificant, while the researcher estimate is positive and marginally significant ($\hat{\beta} = 0.0005^{*}$ per unit AIIE). The lower precision relative to the binary specification reflects reduced statistical power from the continuous measure, which assigns small treatment weights to industries near the median. The directional consistency across both specifications reinforces the robustness of the researcher-specific finding.

F Occupation-Level AI Exposure: Robustness

The main analysis uses industry-level AIIE as treatment, comparing workers whose pre-2012 modal 4-digit NAICS industry is above vs. below the median AI exposure score. A potential concern is that AIIE captures sector-level confounders beyond AI — e.g., that high-AIIE industries were already more dynamic or entrepreneurship-prone for reasons

Table A2: AI Exposure and First Founding: Continuous AIIE Treatment

	Pr(First Founding)		
	(1)	(2)	(3)
	All workers	Researchers	Non-researchers
Post × AIIE	-0.0000	0.0005*	-0.0001
	(0.0002)	(0.0003)	(0.0002)
Control mean	0.0027	0.0029	0.0027
Industry FE	✓	✓	✓
Year FE	✓	✓	✓
Worker FE	✓	✓	✓
N	5156612	735991	4420621

Notes: As Table 3 columns (2)–(4), replacing the binary High_i indicator with the continuous AIIE_i score (modal pre-2012 4-digit NAICS). The coefficient is the change in first-founding probability per one-unit increase in AIIE after 2012. All columns include worker, year, and current-year 2-digit NAICS industry fixed effects. Standard errors clustered by modal pre-2012 4-digit NAICS code. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

unrelated to AI technology. In this appendix we use the occupation-level AI exposure score (AIOE) from Felten et al. (2021) as an alternative treatment, which identifies off within-industry variation in workers' task-level AI exposure.

F.1 Treatment Construction

Each worker is assigned their *modal pre-2012 6-digit SOC occupation*: the occupation in which they accumulated the most position-year observations in years 2007–2011. The corresponding AIOE score is fixed as the worker's occupation-level treatment variable, AIOE_i. A binary indicator $\text{High Occ}_i = \mathbf{1}[\text{AIOE}_i > \text{median}]$ is defined analogously to the industry-level treatment. Workers with no pre-2012 occupation record are assigned their modal occupation across all years and flagged with a separate indicator.

F.2 Results

Table A3 presents the baseline DiD estimates using occupation-level treatment, including worker, year, and current 2-digit NAICS industry fixed effects. Because the treatment now varies within industries (across occupations), the 2-digit industry fixed effects control for any sector-level confounders while leaving the occupation-level treatment variation intact.

The qualitative pattern mirrors the industry-level results: the researcher-specific effect is positive and marginally significant ($\hat{\beta} = 0.0034^{*}$, column 3), while the non-researcher estimate is near zero and insignificant (column 4). The aggregate binary estimate (column

Table A3: Occupation-Level AI Exposure and Founding: Robustness

	Pr(Founder)			
	(1)	(2)	(3)	(4)
	Binary	Continuous	Researchers	Non-researchers
Post × High Occ	0.0011		0.0034*	0.0007
	(0.0010)		(0.0019)	(0.0010)
Post × AIOE		0.0007		
		(0.0005)		
Control mean	0.0391	0.0351	0.0181	0.0421
Industry FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Worker FE	✓	✓	✓	✓
N	3393107	3393107	449555	2943552

Notes: As Table 3 but treatment is based on the worker's modal pre-2012 6-digit SOC occupation. Column (1): binary above-median AIOE indicator. Column (2): continuous AIOE. Columns (3)–(4): binary treatment for researchers and non-researchers separately. All columns include worker, year, and current-year 2-digit NAICS industry fixed effects. Standard errors clustered by modal pre-2012 6-digit SOC occupation. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

1) is 0.0011 and insignificant. This within-industry identification supports the interpretation that the researcher effect reflects AI complementarity at the task level, not sector-level selection into high-AI industries.

G First Wave of AI: Full Empirical Results

This appendix reports the complete empirical analysis for the first wave of AI (AlexNet shock, 2012). The main DiD results and the headline event study are presented in Section 5.2 of the main text; here we report the robustness and heterogeneity analyses.

G.1 Robustness: Industry-Specific Linear Time Trends

A potential concern is that high-AI industries were already on a steeper founding trajectory before 2012, and the DiD estimate reflects this pre-existing divergence. Table A4 adds industry-specific linear time trends to the baseline specification. The aggregate and non-researcher estimates remain near zero. Crucially, the researcher coefficient falls only modestly — from 0.0015 to 0.0011 — and remains statistically significant. The researcher result is not driven by differential pre-existing trajectories.

Table A4: First Wave DiD with Industry-Specific Linear Time Trends

	Pr(First Founding)		
	(1)	(2)	(3)
	All workers	Researchers	Non-researchers
Post × High Ind	0.0001	0.0011**	-0.0001
	(0.0003)	(0.0005)	(0.0003)
Control mean	0.0026	0.0026	0.0026
Industry FE	✓	✓	✓
Year FE	✓	✓	✓
Worker FE	✓	✓	✓
Industry × Trend	✓	✓	✓
N	5156612	735991	4420621

Notes: As Table 3 columns (2)–(4), augmented with modal pre-2012 2-digit NAICS industry × linear time trend. SE clustered by modal pre-2012 4-digit industry. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G.2 Experience Heterogeneity

Table A5 estimates triple-interaction specifications examining whether the AI effect grows with the depth of pre-2012 research experience. The triple interaction with researcher experience is positive and highly significant (column 2): each additional year of pre-2012 research experience amplifies the post-2012 founding response to AI exposure. In contrast, the interactions with total experience (column 1) and non-researcher experience (column 3) are near zero and insignificant. It is specifically the depth of *research* experience — not general tenure — that drives the heterogeneous response.

Table A5: First Wave DiD: Interaction with Pre-AI Experience

	Pr(First Founding)		
	(1)	(2)	(3)
	Total Exp (2011)	Research Exp (2011)	Non-Research Exp (2011)
Post × High Ind	0.0001	0.0001	0.0002
	(0.0005)	(0.0004)	(0.0005)
Post × High Ind × Exp (2011)	0.0000		
	(0.0000)		
Post × High Ind × Research Exp (2011)		0.0004***	
		(0.0002)	
Post × High Ind × Non-Research Exp (2011)			-0.0000
			(0.0000)
Control mean	0.0026	0.0026	0.0026
Industry FE	✓	✓	✓
Year FE	✓	✓	✓
Worker FE	✓	✓	✓
N	5156612	5156612	5156612

Experience fixed as of 2011 (pre-AI baseline). Post x Exp (2011) included but not shown; Exp (2011) and High x Exp (2011) absorbed by worker FE.

Notes: LinkedIn 2007–2019, at-risk sample. Outcome: probability of first founding. Experience fixed as of 2011; main effects and interactions with High_i are absorbed by worker fixed effects. Only the Post × High_i × Exp(2011) triple interaction is identified. All columns include worker, year, and 2-digit industry fixed effects. SE clustered by modal pre-2012 4-digit industry. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G.3 DiD by Experience Bin

Figure A8 estimates the baseline DiD coefficient separately for four experience bins (0–4, 5–9, 10–14, 15+ years), for researchers and non-researchers separately. The AI effect for researchers follows a hump-shaped pattern across experience bins: near zero for workers with fewer than 5 years, growing through the 5–14 bin, and declining modestly for the most experienced group. The non-researcher effect is indistinguishable from zero across all bins.

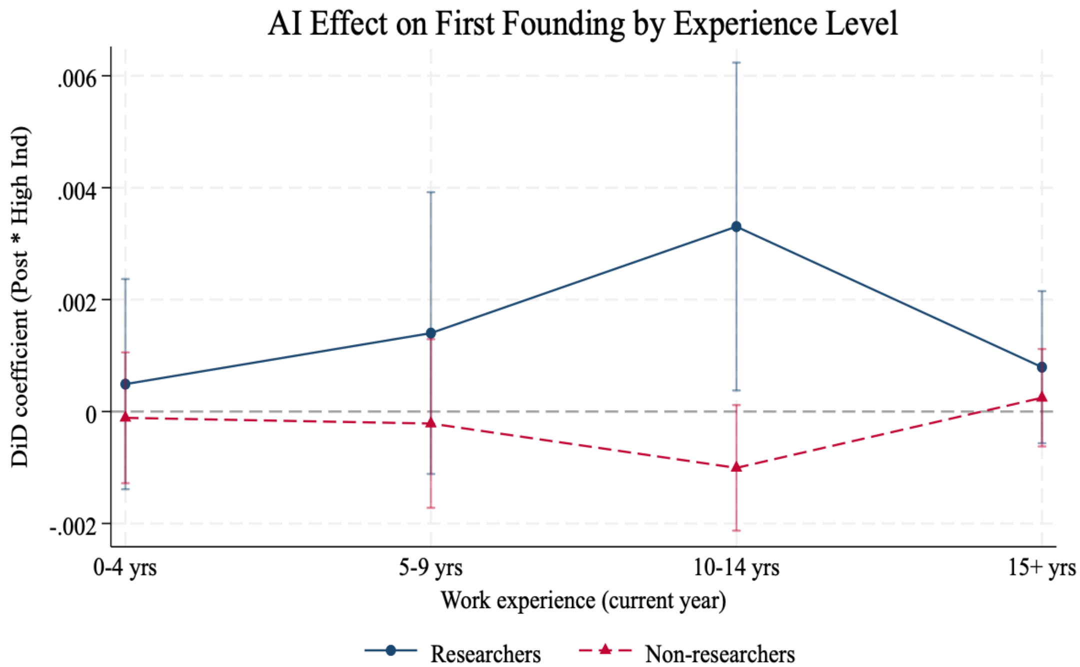

Figure A8: First wave DiD coefficient by experience bin: researchers vs. non-researchers. Each point is a separate DiD estimate for the given experience bin. 95% CIs shown. SE clustered by modal pre-2012 4-digit industry.

H Continuous Anthropic Treatment Specification

This appendix reports the robustness of the main DiD results to using the continuous Anthropic Labour Market Impact score as the treatment variable, rather than the median-split binary indicator. In the continuous specification, the raw pre-period average Anthropic score replaces the High_i indicator in equation (1). Column (1) additionally controls for the level of exposure as a regressor (since worker fixed effects are absent). Columns (2)–(4) absorb the level through worker fixed effects.

Table A6: Continuous Anthropic Treatment: Probability of Being a Founder

	Pr(Founder) (pp) — Anthropic Industry Exposure (Continuous)			
	(1)	(2)	(3)	(4)
	All	All	Researchers	Non-Researchers
Post × Anthropic Ind Exposure	0.022**	0.010	0.044*	0.002
	(0.010)	(0.009)	(0.026)	(0.010)
Control mean	5.16	5.16	3.74	5.43
Industry FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Worker FE		✓	✓	✓
N	9675530	9673076	1536767	8136309

Notes: Treatment is the continuous pre-2022Q4 Anthropic Labour Market Impact score, interacted with Post. Column (1) includes the exposure level as an additional regressor; columns (2)–(4) absorb it via worker fixed effects. All other details as in Table 2. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A7: Continuous Anthropic Treatment: Probability of First Founding

	Pr(First Founder) (pp) — Anthropic Industry Exposure (Continuous			
	(1)	(2)	(3)	(4)
	All	All	Researchers	Non-Researchers
Post × Anthropic Ind Exposure	-0.002	0.003**	0.010**	0.001
	(0.002)	(0.002)	(0.004)	(0.002)
Control mean	0.078	0.054	0.062	0.053
Industry FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Worker FE		✓	✓	✓
N	8918697	8915201	1424628	7490573

Notes: Same as Table A6 but restricted to the at-risk sample and the outcome is the probability of first founding. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The continuous treatment specification confirms that the Anthropic binary findings reflect a genuine dose-response gradient rather than a threshold artefact: the interaction coefficient in the researcher column (column 3) is significant at the 5–10% level for both outcomes, while the aggregate (column 2) and non-researcher (column 4) estimates remain near zero.

I Triple-Difference Event Studies: Full and At-Risk Samples

This appendix presents the triple-difference (DDD) event studies for the full sample (founder stock) and at-risk sample (first founding). The DDD for the re-entrant sample appears in the main text (Figure 6). The DDD specification is described in Section 3.2; the plotted coefficient $\hat{\gamma}_k$ measures the excess response of researchers in high-AI industries relative to non-researchers in each period.

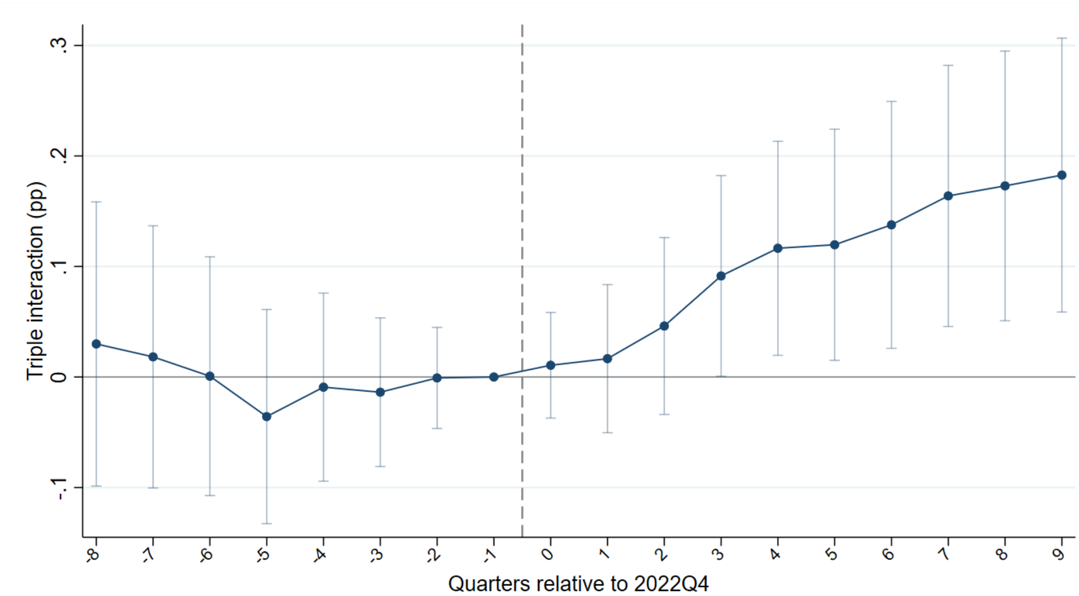

Figure A9: DDD event study: differential probability of being a founder for researchers in high-Anthropic-AI-exposure industries relative to non-researchers (full sample). Base period: 2022Q3. Fixed effects: worker, 2-digit industry, quarter. SE clustered by worker. A growing positive post-period $\hat{\gamma}_k$ indicates that the aggregate founder-stock effect is disproportionately concentrated among researchers in high-AI industries.

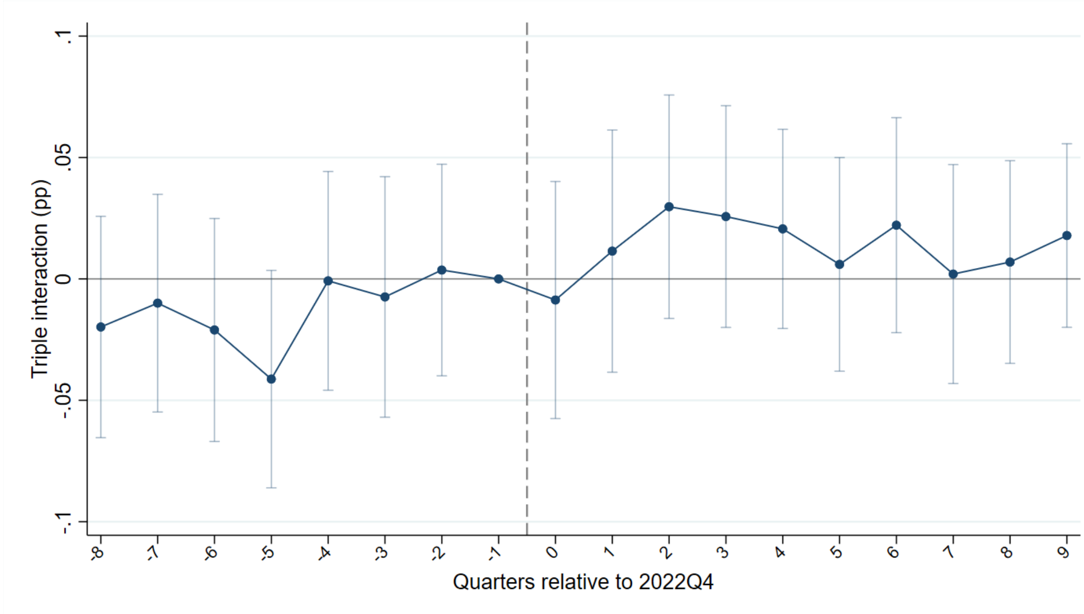

Figure A10: DDD event study: differential probability of first founding for researchers in high-Anthropic-AI-exposure industries relative to non-researchers (at-risk sample). Base period: 2022Q3. Fixed effects: worker, 2-digit industry, quarter. SE clustered by worker. The individually imprecise quarter-by-quarter estimates are consistent with the significant pooled DiD result in Table 2.

Pre-period triple-interaction coefficients $\hat{\gamma}_k$ are close to zero for both figures, confirming parallel trends in the relative researcher/non-researcher response before the ChatGPT shock. After 2022Q4, the founder-stock triple interaction (Figure A9) shows a gradually growing positive trend, confirming that the overall increase in founder prevalence documented in Table 2 is driven disproportionately by researchers in high-AI industries. The at-risk DDD (Figure A10) is noisier due to the low base rate of first-founding events, but the post-period pattern is directionally consistent.

J Model Calibration and Comparative Statics

This appendix presents the full model solution, calibration exercise, and comparative statics underlying the illustrative results summarized in Section 2.

J.1 Value Functions

Let $\omega \equiv (e_W, e_R, o, j)$ denote the state. For vacant jobs, firms post a vacancy with cost k_{oj} associated with a promised lifetime utility x:

$$V_t(\omega) = -k_{oj} + (1-\sigma)\, q(\theta_t(\omega))\, (J_t(\omega) - x).$$

Free entry implies $q(\theta(\omega)) = k_{oj}/(J_t(\omega) - x)$.

For unemployed individuals conditional on not switching:

$$U_t(\omega) = b_{oj} + \max_x \Big[(1-\sigma)\beta\, \mathbb{E}\{p(\theta_t(\omega_{t+1}))(x - U_{t+1}(\omega_{t+1})) + U_{t+1}(\omega_{t+1})\}\Big].$$

Including endogenous switching, $\mathcal{U}_t(\omega) = \max_{o',j'}\{U_t(\omega) - c(\omega)\}$.

For employed matches:

$$J_t^{\text{Act}}(\omega) = y_t(\omega) + (1-\sigma)\beta\, \mathbb{E}\Big[\delta U_{t+1}(\omega') + (1-\delta)(J_{t+1}(\omega') + \lambda_e R(\omega'))\Big],$$

where $R(\omega) = \max_{o,j,x}\{p(\theta(\omega))(x - J_t(\omega)) - c(\omega)\}$ is the on-the-job search value. The final value function accounts for endogenous separation: $J_t(\omega) = \max\{J_t^{\text{Act}}(\omega), \mathcal{U}_t(\omega)\}$.

Block Recursive Equilibrium. We focus on Block Recursive Equilibrium (BRE) (Menzio and Shi, 2010, 2011), in which market tightness $\theta(\omega)$ is completely determined by matching efficiency and vacancy costs, and individual optimal policy functions depend only on (e_W, e_R), not on the aggregate distribution.

J.2 Solution Strategy

The equilibrium is solved by value function iteration:

1. Precompute flow values for each state; build the human capital transition matrix.

2. Initialize U_0 and J_0 using separation and unemployment values.

3. Given U_0 and J_0, compute net switching values and optimal policies.

4. Repeat until convergence.

Figure A11 illustrates the optimal occupation choices across the (e_W, e_R) space. Workers sort into entrepreneurship when both experience dimensions are sufficiently high, consistent with the Lazear jack-of-all-trades framework.

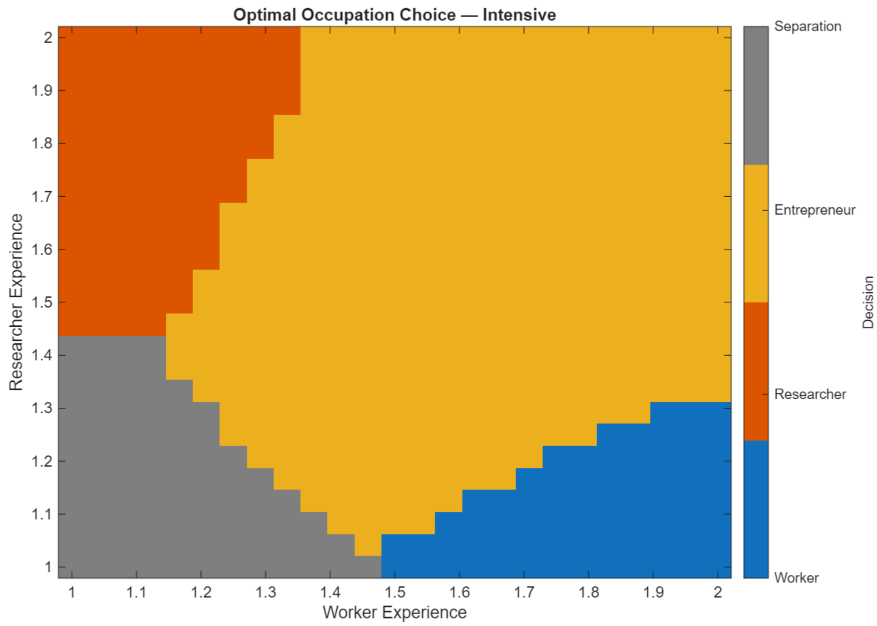

(a) AI-intensive industries

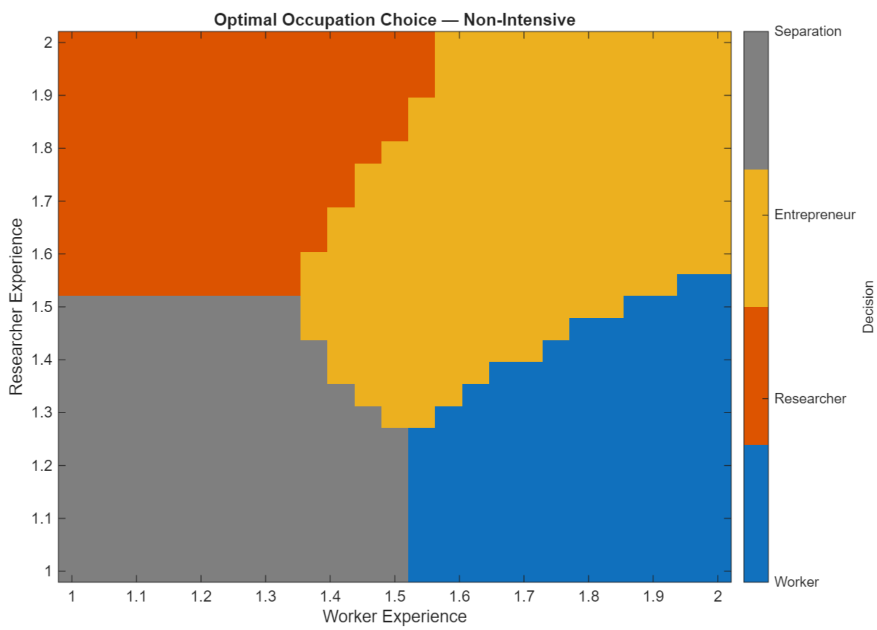

(b) Non-AI-intensive industries

Figure A11: Optimal occupation choices across the experience space.

After solving the policy functions, the steady-state distribution is obtained by building and inverting the transition matrix. Figure A12 shows the resulting experience distributions for newborns, employed workers, and all individuals.

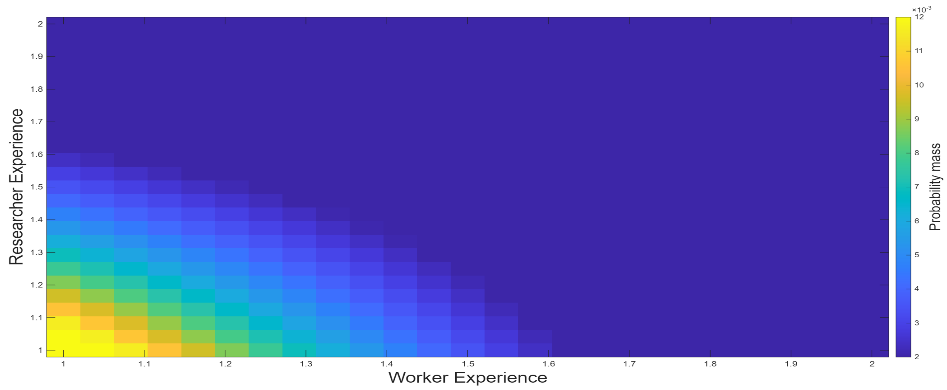

(a) Newborn experience distribution

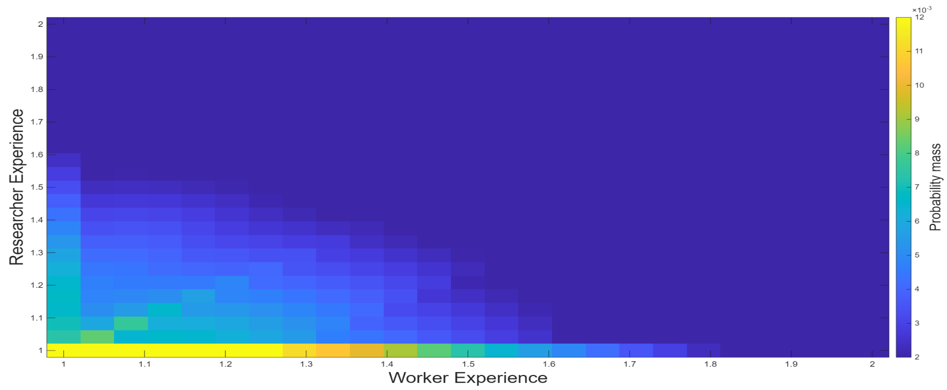

(b) Employed individuals

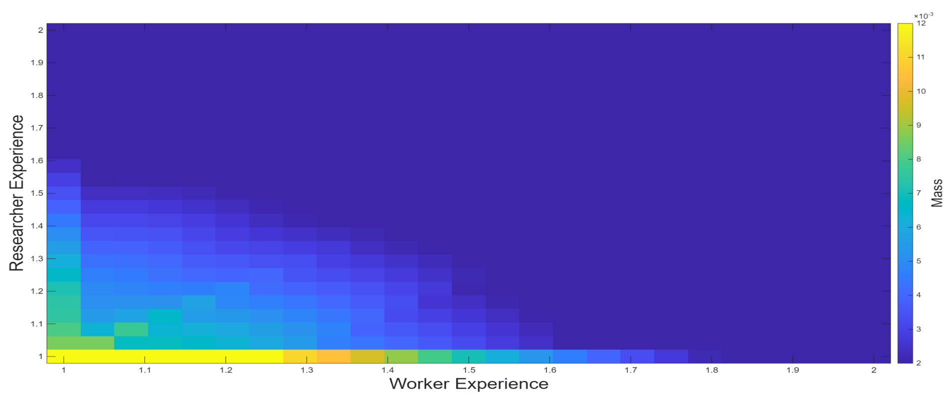

(c) All individuals

Figure A12: Baseline calibration: experience distributions.

J.3 Comparative Statics

Occupational productivity premium. Figure A13 shows how occupational shares respond to changes in key parameters. Panel (a) varies the worker productivity premium A_W; panel (b) varies the researcher premium A_R; panel (c) varies the human capital accumulation probability π; panel (d) varies matching efficiency γ.

The key comparative static for the paper's mechanism appears in panel (b): an increase in A_R simultaneously raises the shares of both researchers and entrepreneurs, with the effect on entrepreneurship operating through workers who already have high e_R. This reproduces the pattern in the data and validates Proposition 1.

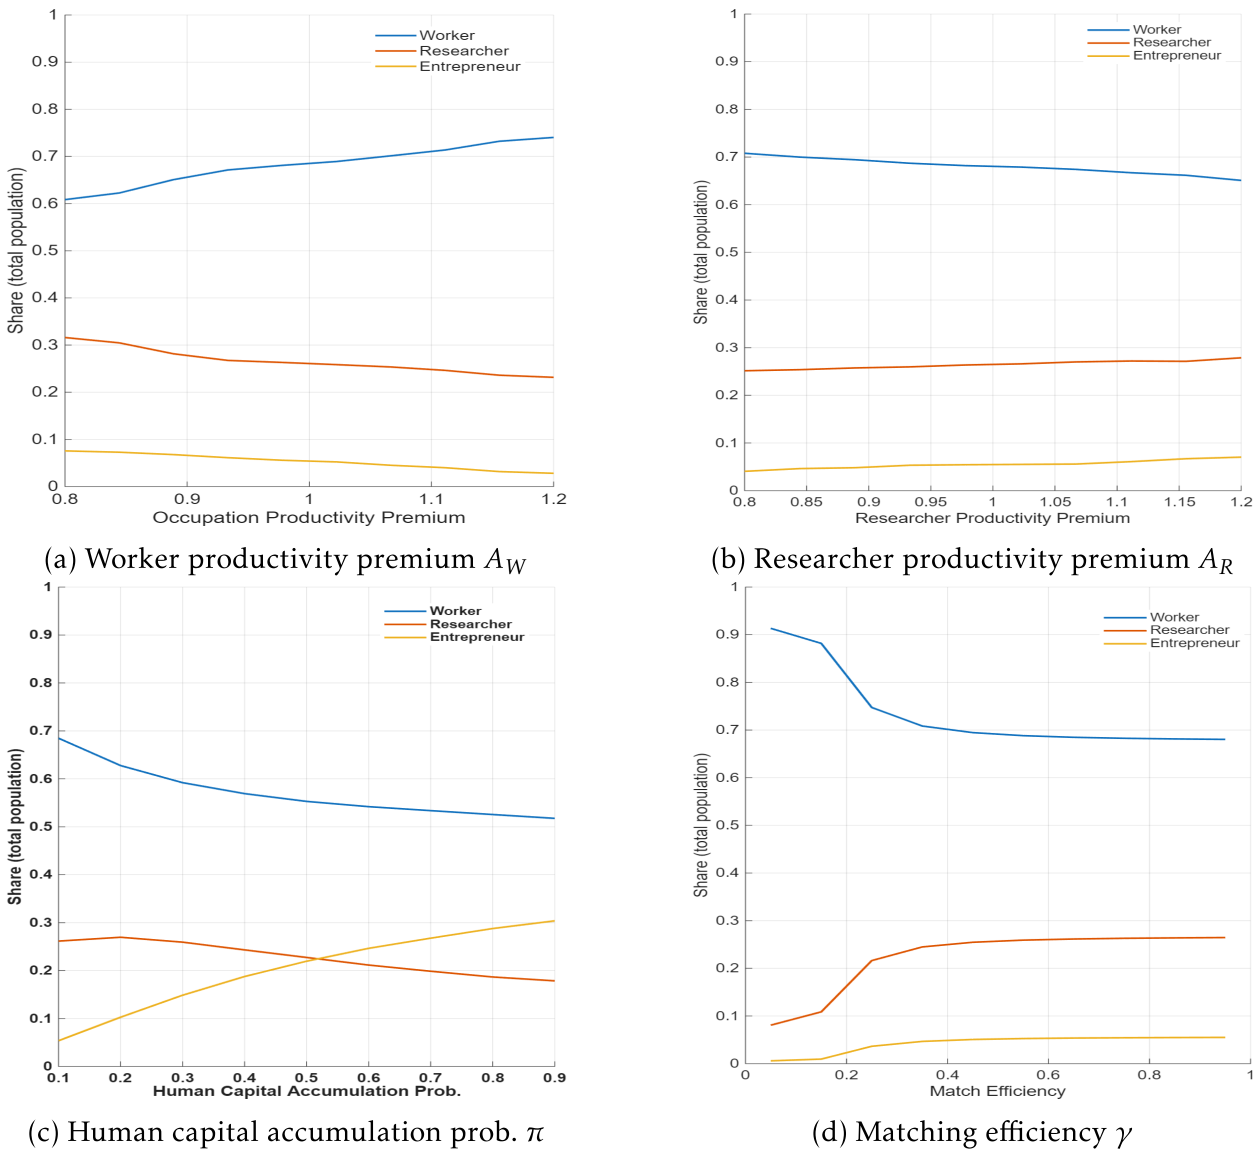

(a) Worker productivity premium A_W

(b) Researcher productivity premium A_R

(c) Human capital accumulation prob. π

(d) Matching efficiency γ

Figure A13: Comparative statics: occupational shares vs. key parameters.

Effect on experience distribution. Figure A14 shows that a higher A_R shifts the steady-state distribution of experience toward the upper-left region of the (e_W, e_R) space, reflecting a reallocation of workers toward research-intensive occupations and entrepreneurship.

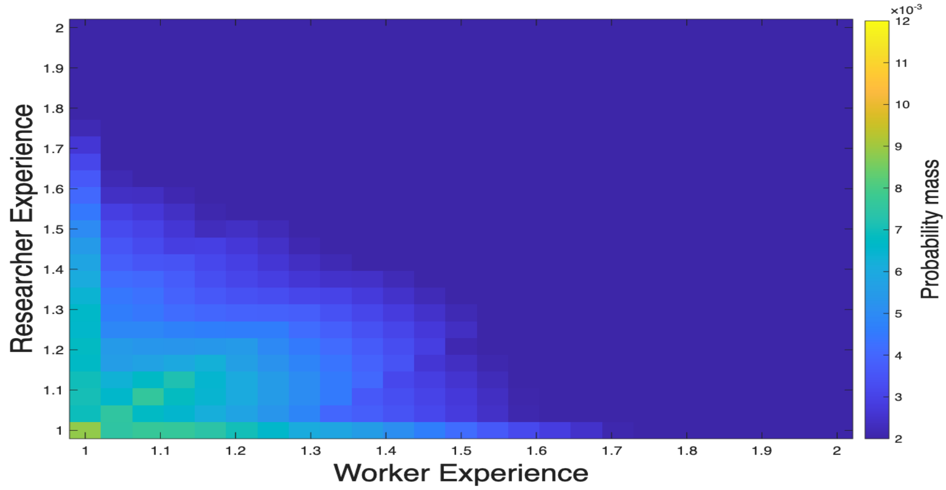

(a) Employed individuals (higher A_R)

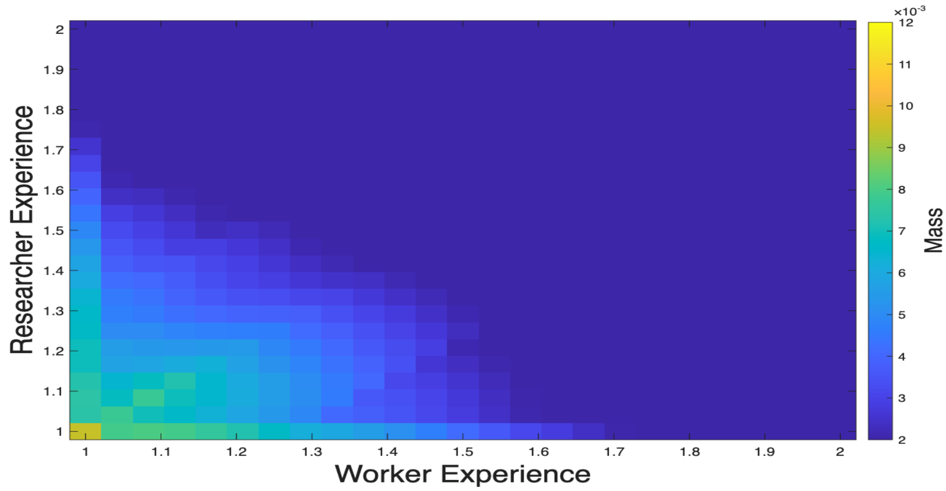

(b) All individuals (higher A_R)

Figure A14: Effect of higher researcher productivity premium on experience distributions.

J.4 Model Timing

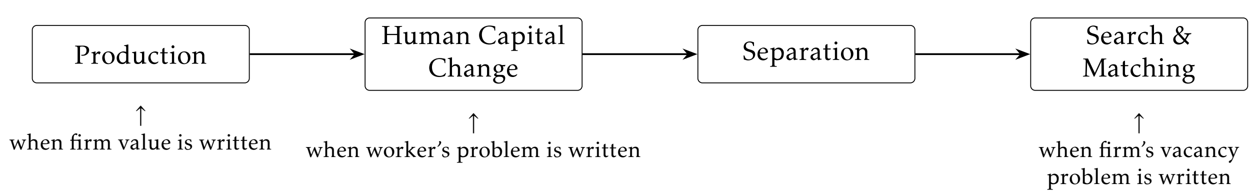

Figure A15: Timing of the Model