Elite Talent and Firm Productivity in the Age of AI

Sagar Baviskar* Liujie Wu† Cameron Drayton ‡ Lee Branstetter§
Prasanna Tambe¶ Eduard Hovy‖

October 2025

Click here for latest version.

Abstract

Artificial intelligence (AI) has advanced rapidly in recent years, and some experts see it emerging as a transformative general-purpose technology, but firms vary greatly in their ability to capture its benefits. We hypothesize that the shortage of experts trained up to the scientific frontier is an important resource constraint shaping where and how AI diffusion and innovation occur. We contribute to the study of elite human capital in this domain by constructing new data on the training of these experts and their movement from academia to industry. We use publication data to identify the most influential scientists in 10 academic subfields of AI, link these scientists to their graduate students and postdoctoral advisees who find employment in the United States, and follow the movement of these students across time and organizational boundaries. Our data suggest that elite graduates are not disproportionately absorbed by the largest technology firms ("big tech"), but are instead distributed across a wide range of sectors. We link these experts to the firms that employ them and use Compustat data and standard panel regression techniques to explore how frontier AI talent reshapes firm productivity. Firms that hire elite graduates experience sustained total factor productivity improvements of 8–13 percent, with effects that strengthen over time. Elite graduates also enhance the labor productivity of hiring firms and their presence is strongly correlated with the generation of patents describing AI-related innovations. Event study analyses point to a causal effect of AI innovation on firm productivity. The strong association between elite graduates, AI innovation, and productivity growth points to the role these individuals may play in both generating innovations and connecting firms to sources of external knowledge.

*Carnegie Mellon University. Corresponding Author. Email: sagarbaviskar@cmu.edu
†Carnegie Mellon University
‡Carnegie Mellon University
§Carnegie Mellon University and NBER
¶Wharton School of Management, University of Pennsylvania
‖Language Technologies Institute, Carnegie Mellon University, and University of Melbourne

1 Introduction: General Purpose Technologies, AI, and Elite Talent

The concept of general-purpose technologies (GPTs) has long been central to explaining technological revolutions with economy-wide impact. GPTs such as steam power, electricity, and information technology are pervasive, improve over time, and enable complementary innovations that reshape industries and raise productivity growth at the macro level (Bresnahan and Trajtenberg, 1995; Helpman, 1998). Artificial intelligence (AI) has entered a new era of rapidly advancing capabilities and the media speaks of an AI boom (Agrawal et al., 2021; Baily et al., 2023). However, productivity growth in the United States and elsewhere remains sluggish. Perhaps not surprisingly, public discourse now oscillates between optimism—heralding AI as the GPT driver of a new productivity boom—and skepticism—warning of an AI "false dawn" or a disruptive wave of job displacement and massively increased inequality (Brynjolfsson and McAfee, 2014; Susskind, 2020).

The GPT literature has long noted that the economy-wide productivity surge often arrives with a significant lag, as firms and industries undertake costly complementary investments to realize the GPT's benefits in their specific contexts (David, 1991; Jovanovic and Rousseau, 2005). The implementation of earlier general-purpose technologies (GPTs), such as electricity and information technology, began decades before economy-wide growth effects were evident in the data (Bresnahan and Trajtenberg, 1995; Helpman, 1998). Brynjolfsson et al. (2019) argue that AI's productivity impact will follow a J-curve: adoption/innovation involves adjustment costs before steep productivity gains materialize. Agrawal et al. (2019, 2021) conceptualize AI as a "prediction machine" that reduces uncertainty and enables novel recombination. Cockburn et al. (2019) highlight AI as a technology for invention itself, altering trajectories of innovative search. Collectively, this work emphasizes that AI's transformative potential depends not only on the technology but also on complementary organizational, human capital, and industry-specific investments. The work also implies that productivity effects may begin to emerge in successful early-adopter or early-innovator firms long before they diffuse to the broader economy.

The well-established consensus around the necessity of complementary investment and the implication that productivity effects may arrive with a lag has motivated different researchers to use different ways of identifying successful early adopters or early innovators, who may be already reaping AI's gains.

Government-led, large-scale surveys of AI adoption have provided valuable descriptive insights into AI diffusion, but these studies have not yet found robust productivity effects (Zolas et al., 2020; McElheran et al., 2024). A growing body of carefully implemented

randomized controlled trials (RCTs(document large productivity impacts in specific work settings, but it is hard to generalize these measured effects from the very specific contexts in which they are demonstrated (Brynjolfsson et al., 2023; Noy and Zhang, 2023). Alderucci et al. (2024) find statistically and economically significant impacts of AI innovation, as measured by patents, on firm's ex-post labor productivity growth. Results obtained with event study methods strengthen the inference that these effects are causal, but this paper does not explain why some firms are generating AI patents more quickly and effectively than others. Past efforts to study firm-level recruitment of AI workers can link these efforts to output gains and AI patents, but find no statistically significant effects of this recruitment on firm productivity growth (Babina et al., 2024).

AI is distinctive among GPTs in four respects that may heighten the heterogeneity of its impact across firms, industries, and time. First, it is unusually cross-sectoral: the same basic algorithmic approaches can be applied to domains as varied as manufacturing, logistics, finance, and healthcare, but only with significant domain-specific complements (Brynjolfsson and McElheran, 2022). Second, frontier AI knowledge is highly tacit and fast-moving, built on heuristics and architectures that challenge codification into well-established textbooks, course curricula, and widespread industrial practice. (Jordan and Mitchell, 2015). Third, AI progress relies on recombination across interdependent subfields such as machine learning, NLP, robotics, and computer vision. Fourth, diffusion is mediated by open-science platforms, benchmarks, and elite academic communities that privilege insiders (Lee and Schankerman, 2020). These features suggest that the movement of the relatively small number of experts with frontier experts with scarce cutting edge knowledge and social network ties may be critical in determining which firms realize the productivity-enhancing benefits of AI first.

Following this logic, this paper contributes to the literature by examining how the knowledge carried by the most advanced students of elite AI scientists (whom we dub the "immortals") shapes firm productivity. We assemble a novel dataset that links Compustat firm-level data with ∼2,000 elite AI scientists and ∼22,000 of their graduate students (and post-doctoral advisees) who then seek employment in the United States. This allows us to trace how frontier AI talent diffuses into firms and conditions heterogeneity in their outcomes.

Our findings suggest three broad patterns. First, embodied frontier AI knowledge brought by elite graduates generates large and persistent total factor productivity gains, underscoring the role of specialized human capital as a key enabling firms to benefit from AI. Second, frontier AI talent also boosts labor productivity. Third, and contrary to popular belief, large technology firms have not monopolized access to elite graduates so far, and sectoral patterns instead show elite talent moving into a diverse set of firms and industries. Finally, the growth of elite talent in the firm is a strong predictor of the generation of AI patents, and

the emergence of AI patents appears to be causally connected to firm productivity growth.

To interpret these findings, we develop the AI Talent Diffusion Framework (ATDF). The framework highlights three mechanisms through which frontier AI talent enhances firm productivity: knowledge transfer from elite training into firms, capability augmentation as AI expertise is applied to existing production processes, and network spillovers that expand access to complementary innovations. This perspective builds on research in the economics of innovation on human capital, technology adoption, and spillovers, clarifying why some firms translate AI into enduring performance gains while others do not.

The remainder of the paper is organized as follows. We first situate our study within the literatures on human capital, innovation, and knowledge diffusion, and introduce an AI Talent Diffusion Framework (ATDF) to interpret our findings. We then describe our data construction process, including the academic lineage and firm-level matching. We cluster the AI talents into different groups and describe their characteristics, followed by details of our empirical approach. We present results and conclude with implications for the economics of innovation and the strategy literatures.

2 Literature Review and Theoretical Framework

2.1 Human Capital and Firm Performance

Human capital is a well-established driver of firm productivity and innovation (Hitt et al., 2001; Mollick, 2012). In knowledge-intensive industries, intangible expertise can be more consequential than physical capital or even formal intellectual property. Yet most empirical work measures human capital in aggregate terms, such as education, experience, or broad skill categories, without interrogating its origins.

The resource-based view (RBV) highlights that advantage arises from resources that are valuable, rare, inimitable, and non-substitutable (Barney, 1991). From this perspective, elite-trained AI talent constitutes a uniquely potent resource: graduates trained by AI pioneers carry not only technical expertise but also cognitive models, norms, and social capital rooted in their academic lineage. Their embeddedness in frontier research communities makes them difficult to substitute or replicate. Earlier studies on university–industry spillovers show that hiring individuals from prestigious academic environments accelerates firm innovation by embedding tacit knowledge (Zucker et al., 1998; Cockburn and Henderson, 1998). We argue that academic lineage represents a distinct form of human-capital differentiation that helps explain heterogeneity in firms' ability to capture value from AI. These insights complement economics perspectives that model firm knowledge capital and managerial quality as key

drivers of productivity (Griliches, 1984; Pakes and Griliches, 1984; Bloom and Van Reenen, 2007) and also builds on findings that star scientists amplify institutional and colleague productivity (Azoulay et al., 2010).

2.2 Organizational Capabilities and Innovation

Beyond resources and individuals, strategy and organizational research emphasize firm-level processes that determine how external knowledge becomes productive. Absorptive capacity theory (Cohen and Levinthal, 1990) emphasizes the role of prior related knowledge in recognizing and exploiting external ideas. Hiring graduates from elite AI lineages increases absorptive capacity by embedding frontier expertise directly within firms. Dynamic capabilities research (Teece et al., 1997; Eisenhardt and Martin, 2000; Teece, 2007, 2018) highlights how firms adapt to shifting environments, with lineage-based talent enhancing their ability to sense and reconfigure. Finally, the complementary assets perspective (Teece, 1986) stresses that the value of an innovation depends on supporting infrastructure. In AI, such assets include data, computing resources, and organizational readiness. Our contribution is to show that lineage-based talent functions as a critical human complement that allows firms to integrate and leverage these assets effectively.

Relational perspectives emphasize that alliances, collaborations, and networks are crucial channels for knowledge transfer (Gulati, 1998; Dyer and Singh, 1998). While this research typically emphasizes inter-firm ties, internal collaboration and search-transfer frictions also shape whether knowledge moves to where it is most valuable (Hansen, 1999). We extend this view by highlighting academic lineages as relational assets. Elite AI pioneers generate cohorts of graduates who diffuse tacit knowledge across firms, embedding expertise and connecting organizations to knowledge-rich communities and entrepreneurial ecosystems (Stuart and Sorenson, 2007).

2.3 The AI Talent Diffusion Framework (ATDF)

To illustrate the mechanisms through which elite AI expertise flows from academia into firms, we use the AI Talent Diffusion Framework (ATDF). This framework builds on human capital theory (Becker, 1964), which emphasizes the role of individual skill development in shaping labor market outcomes, and the knowledge-based view of the firm (Grant, 1996), which considers knowledge as the most strategically significant resource in modern organizations. The ATDF advances these perspectives by introducing academic lineage as a central conduit through which frontier AI knowledge is transferred, embedded, and scaled within firms. Specifically, the ATDF highlights three interrelated but conceptually distinct mech-

anisms—knowledge transfer, capability augmentation, and network spillovers—that explain how the academic lineage of AI talent shapes firm-level productivity and strategic advantage.

2.3.1 Knowledge Transfer

The first mechanism, knowledge transfer, captures the direct movement of tacit and frontier knowledge from elite academic mentors to firms through their graduates. Graduates trained under pioneering AI scientists bring with them not only technical know-how but also exposure to research frontiers, methodological rigor, and conceptual frameworks shaped during their training. These graduates often operate at the cutting edge of subfields such as machine learning, natural language processing, robotics, and computer vision. When hired into firms, they act as knowledge vectors, embedding advanced ideas into organizational processes, products, and strategies.

This mechanism aligns with research on star scientists and knowledge diffusion, which shows that the mobility of elite scientists across institutional boundaries leads to measurable gains in innovation (Azoulay et al., 2010). Importantly, the knowledge carried by these experts is not generic; it is frontier-specific, situated within evolving paradigms, and often too complex to be acquired easily through market transactions or consulting relationships. Instead, it must be internalized through close interaction, a process catalyzed by the presence of elite-trained graduates within firms.

2.3.2 Capability Augmentation

The second mechanism, capability augmentation, explains how elite AI graduates strengthen firms' internal technological and organizational capacities. While knowledge transfer concerns the *content* of expertise, capability augmentation focuses on the *application* of that expertise and how it transforms organizational routines. Elite-trained graduates often bridge scientific discovery and engineering practice, designing experimentation protocols, building scalable evaluation systems, and fostering rigorous, data-driven decision-making.

Drawing on absorptive capacity theory (Cohen and Levinthal, 1990), we argue that graduates of elite academic lineages increase firms' ability to recognize, assimilate, and exploit new ideas. Over time, these contributions extend beyond individual problem-solving: they institutionalize new heuristics and methods, embedding them into firm routines. This process strengthens dynamic capabilities (Teece, 2007), allowing firms to reconfigure resources, accelerate adaptation to technological shocks, and maintain competitiveness in rapidly evolving environments. Capability augmentation is therefore inward-facing, concerned with building organizational resilience and agility from within.

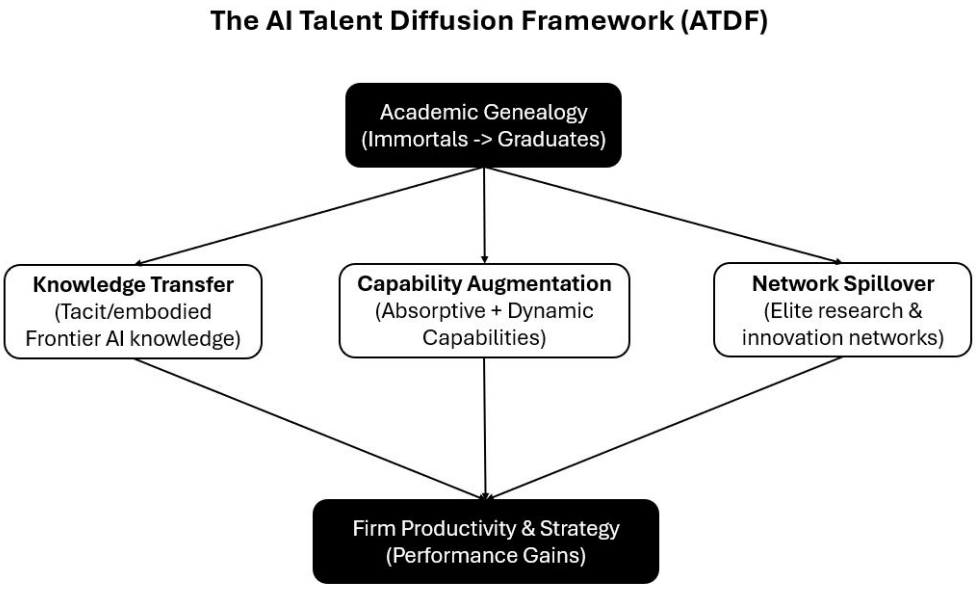


Figure 1: The AI Talent Diffusion Framework (ATDF). Academic lineage channels frontier AI knowledge from elite scientists ("immortals") to their graduates, shaping firm productivity through three mechanisms: knowledge transfer, capability augmentation, and network spillovers.

2.3.3 Network Spillovers

The third mechanism, network spillovers, captures the relational and reputational advantages firms gain by hiring graduates from elite academic lineages. Unlike traditional hires, who bring only their individual expertise, the elite scientists that are the focus of ours study connect firms to broader epistemic and innovation communities. These networks span co-authors, academic collaborators, government agencies, research consortia, and leading technology companies.

Through continued engagement in these networks, via publications, conferences, collaborations, and advisory roles, graduates create relational bridges that firms can leverage for strategic benefit. This resonates with research on inventor mobility and innovation networks, which shows that mobile talent carries access to siloed knowledge and facilitates recombination across organizational boundaries (Rosenkopf and Almeida, 2003). In the AI context, spillovers may take the form of early visibility into research breakthroughs, opportunities for cross-sector collaborations, or privileged access to public and private funding streams.

Critically, these advantages are not evenly distributed. They accrue disproportionately to firms employing graduates embedded in high-status academic lineages, and often persist even after individual employees depart. Firms known for hiring elite AI talent can develop reputations as innovation leaders, creating a self-reinforcing loop that attracts additional

top-tier candidates and strengthens their position in the AI ecosystem.

3 Data

3.1 Frontier AI Talent Data

To capture the human capital dimensions of AI innovation, we construct a novel dataset that traces the career trajectories of elite academic scientists in AI and their graduate students. This provides a unique window into how frontier knowledge is embodied in individuals and diffused into firms.

Sub-Disciplinary Structuring

The first stage of our data construction process involved segmenting the broad and rapidly evolving field of AI into ten distinct sub-disciplines. This taxonomy was informed by major conference themes, journal specializations, and topical categorizations commonly accepted in the AI research community. We extract publications in top journals and conferences associated with ten core AI subdomains:

1. Machine Learning (ML)
2. Natural Language Processing (NLP)
3. Robotics
4. Speech
5. Agents
6. Computer Vision (CV)
7. Information Retrieval (IR)
8. Human–Computer Interaction (HCI)
9. Knowledge Representation (KR)
10. General AI / Other

This framework ensured appropriate representation across AI while also enabling subfield-specific insights into mentorship and career trajectories.

Identification of Elite Scholars

For each of the ten subdomains above, we identified leading publication venues (e.g., NeurIPS, ICML, AAAI, ICLR for ML) using impact factors, expert consensus, and rankings such as CORE and Google Scholar Metrics. Using Elsevier Scopus, we compiled bibliometric data for all authors in these venues and computed their h-index, normalized by subdomain to

account for citation differences. Authors whose h-index ranked in the top ∼0.2% in each subdomain (about 2,000 in total) were designated as the *Immortals*, representing the frontier of global AI research.

Academic Lineage Mapping

Once the scientific Immortals were identified, we traced their academic descendants by combining multiple sources. Institutional websites, CVs, lab web pages, and other public records were used to validate Immortals' identities, affiliations, and lists of students. We also used ProQuest data on archived doctoral dissertations to confirm advisor–advisee relationships, recording metadata such as student names, institutions, and year of completion. This process yielded ∼30,000+ graduates in our current dataset (with 22,000 US based, and ∼15,000 identified through ProQuest alone). These individuals embody tacit knowledge, heuristics, and network ties transmitted through frontier mentorship.

Career Profiling

We then constructed longitudinal profiles of elite trained graduates' career trajectories by integrating data from the following sources:
a. LinkedIn Data: Structured search queries identified professional profiles containing education, employment, and mobility histories.
b. Revelio Labs: We matched individuals to standardized datasets with job titles, industries, and inferred skills, enabling systematic analysis of career paths.
c. Manual Validation: For ambiguous or common names, we cross-validated entries across LinkedIn and ProQuest to avoid false matches and ensure consistency.
Each individual was assigned a unique identifier linking their academic origin and employment history, providing a robust dataset to track knowledge flows.

3.1.1 Quantitative Overview of the Frontier AI Talent Dataset

Table 1 provides a quantitative overview of the scope of our dataset across ten AI subdomains.

First, the breadth of our coverage is extensive: approximately 2,000 Immortals are identified, with the largest concentrations in machine learning (320), natural language processing (208), robotics (160), information retrieval (459), and general AI—not elsewhere classified (534). These scholars are associated with nearly 460,000 papers, reflecting their centrality to the evolution of AI research.

Second, the academic lineage component documents more than 30,000 children of Immortals (22,000 in the United States). To date, subfields such as robotics (6,222), machine

learning (5,071), information retrieval (5,318), and natural language processing (4,607) account for substantial shares. In some subdomains (e.g., speech, information retrieval, and general AI) we are still completing our data collection efforts, so the results reported in this paper should be considered preliminary.

Third, the dataset also incorporates direct Immortal–corporate collaborations. Over 60,000 coauthored papers are observed, including 16,026 in machine learning, 12,047 in information retrieval, and 8,553 in natural language processing. These collaborations provide an additional dimension of knowledge transfer, alongside the movement of academic children into firms.

Taken together, the table underscores the magnitude and richness of our dataset. It demonstrates both the depth of scientific influence exerted by elite AI scholars and the breadth of diffusion channels—academic children and corporate collaborations—that form the empirical foundation of our current and future analyses.

Table 1: Counts of Immortals, academic children, and direct Immortal–corporate collaborations by AI subdomain

Data point \ Domain	**NLP**	**ML**	**Robotics**	**Agents**	**HCI**	**Speech**	**KR**	**CV**	**IR**	**Other AI**
# of Immortals	208	320	160	29	17	140	4	79	459	534
# of journals/conference venues	722	79	17	58	18	12	32	80	146	55
# of papers associated with Immortals	50,414	74,165	56,476	9,362	4,451	36,396	873	33,509	89,490	103,988
# of children of Immortals identified (so far)	4,607	5,071	6,222	836	291	2,130	57	2,610	5,318	3,771
Balance # of children yet to be tagged	0	0	0	0	0	2,232	0	0	†	†
# of direct Immortal–corporate collaborative papers	8,553	16,026	4,582	665	987	6,859	30	4,301	12,047	7,112

† Currently identifying additional children in Information Retrieval and Other AI categories.

3.2 Linking to Compustat Data

To evaluate the firm-level consequences of AI innovation, we draw on Compustat, a comprehensive financial and operational dataset maintained by S&P Global Market Intelligence. Compustat provides standardized, longitudinal data on publicly traded U.S. firms, including variables such as output, employment, R&D expenditures, capital investments, and market valuation. These measures make it an indispensable resource for linking frontier AI talent to firm productivity, strategic positioning, and financial performance.

Integrating these records allows us to analyze temporal and industry-specific heterogeneity in the returns to AI. Specifically, we examine whether frontier AI talent translates into measurable productivity improvements, and whether these gains are broadly distributed or concentrated among a subset of firms. By pairing frontier AI talent data with financial outcomes, the matched dataset provides the empirical foundation for testing whether AI is

delivering on its promise of productivity-enhancing transformation—or instead reinforcing existing disparities in technological capacity across firms.

3.2.1 Characteristics of Firms Hiring Frontier AI Talent

We operationalize exposure to frontier AI talent as the number of graduates trained under elite AI pioneers ("immortals") employed by a firm, normalized by lagged firm resources such as employment, R&D, or capital stock, depending on the outcome of interest. This measure captures the relative depth of a firm's integration of elite AI expertise rather than simply the scale of its workforce. Using the matched Compustat–academic lineage dataset, we define AI-talent firms as those that have hired at least one graduate from an immortal lineage between 1990 and 2024.

Descriptive patterns reveal that AI-talent firms are systematically different from their peers. They are generally larger in employment and assets, older, more R&D intensive, and more profitable than the average Compustat firm. They also exhibit higher labor productivity, measured as output per worker, and greater capital intensity, patterns consistent with the idea that integration of frontier AI talent may be greater among firms with the resources and organizational capacity to absorb and leverage general-purpose technologies. Of course, privately held AI start-ups are excluded from the Compustat database. Ongoing research will seek to match our elite talent to a much broader set of firms, including these privately held start-ups.

In the Compustat sample, AI-talent firms also tend to be more diversified, both geographically and across business segments, reflecting the broad applicability of AI expertise and the complementarities required for integration into production and business processes. This evidence aligns with the literature on "superstar firms," suggesting that frontier AI talent is disproportionately concentrated among firms with scale, scope, and the absorptive capacity to capitalize on GPTs.

Sectoral analysis sharpens these contrasts. In manufacturing (NAICS 33) and IT (NAICS 56), AI-talent firms are markedly larger, more capitalized, and more productive than their non-AI peers. They also display early signs of labor restructuring: lower shares of production workers and higher concentrations of knowledge-intensive roles, consistent with both automation and skill upgrading associated with AI-driven transformation.

3.3 Patents Data

We construct a large-scale longitudinal dataset of AI patents by leveraging the full corpus of patents from the United States Patent and Trademark Office (USPTO), covering the period

1990–2024. After extensive cleaning, including de-duplication, normalization of formats, and exclusion of withdrawn records, the final dataset comprises 7,256,235 unique patents. For each patent, we extract and analyze two core textual fields: the title and the abstract. These fields provide concise yet informative descriptions of technical contributions and are particularly well-suited for identifying whether a patent is AI-related or not.

To classify patents, we employ a semi-supervised pipeline. The process begins with a broad retrieval step based on AI-related CPC/USPC classifications (e.g., Class 706), generating an initial pool of patents for training and evaluation. We then introduce a novel classification framework, **PaLLaFi**, which combines supervised machine learning models, large language models, and a human-in-the-loop system. Iterative labeling focuses on high-value cases: patents clearly within or outside AI as well as ambiguous "borderline" patents where human adjudication improves model accuracy. In parallel, we are developing subdomain-level tagging (e.g., Machine Learning, Natural Language Processing, Robotics, Vision, Knowledge Representation) to capture the composition and evolution of AI innovation more precisely. This effort remains ongoing and will be incorporated in the final dataset release.

To link patents to firm-level outcomes, we match USPTO assignees to publicly traded companies in the Compustat database. We build on the patent–firm linkage methodology of Dyèvre and Seager (2023), which connects 70 years of USPTO records to Compustat firms, and extend it by implementing our own entity-resolution and name-disambiguation pipeline. This refinement addresses variation in firm names, abbreviations, and ownership structures, yielding a more accurate mapping. The resulting dataset links patents to approximately 6,000 publicly traded firms.

This integrated patent–firm dataset enables us to measure firms' exposure to AI innovation, track technological transitions over time, and evaluate how frontier AI knowledge translates into firm-level productivity outcomes. The classification framework that operationalizes the AI/non-AI distinction is discussed in detail in the appendix section.

4 Mobility and Clustering of the AI Talents

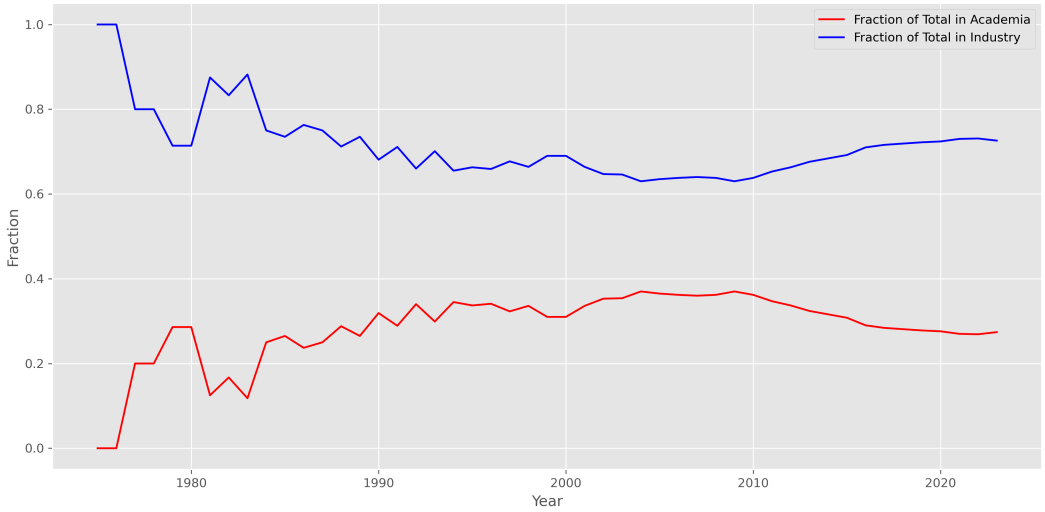
To understand the evolution of employment trajectories among AI researchers, we analyze the annual data from 1975 to 2024 on individuals who have held at least one job post-graduation. This section highlights longitudinal trends in employment distribution across sectors, with a particular focus on Academia (NAICS code 61), Industry (all the other NAICS codes), and Big Tech Firms ("777") for aggregate comparison.

The number of "children" with recorded employment has increased dramatically over the past six decades, growing from fewer than 10 individuals in the 1970s to over 5,000 by

2024. The most substantial growth occurred after 2000, corresponding with the global AI boom and expansion of both academic and industrial opportunities. This surge reflects the broader transformation of AI from a niche academic field to a central driver of technological and economic development.

In the early years of the dataset, nearly all AI professionals entered industry rather than academia. This pattern suggests that, during the formative decades of the field, the infrastructure and incentives for academic careers were relatively underdeveloped. However, beginning in the 1970s and gaining momentum through the 1980s and 1990s, academia started to absorb a growing share of the workforce. This rise likely reflects the institutionalization of computer science departments, the expansion of research funding, and the increasing prestige of academic roles during this period. This trend began to reverse in the early 2000s, as industry—particularly the technology sector—reasserted its dominance in attracting talent. The resurgence was driven by rapid commercial demand for machine learning and data science, along with the emergence of large-scale digital platforms that required pioneering AI capabilities. As a result, the fraction of workers in industry rose sharply, signaling a renewed shift away from academia.

Figure 2: Academic vs. Industry Employment Over Time

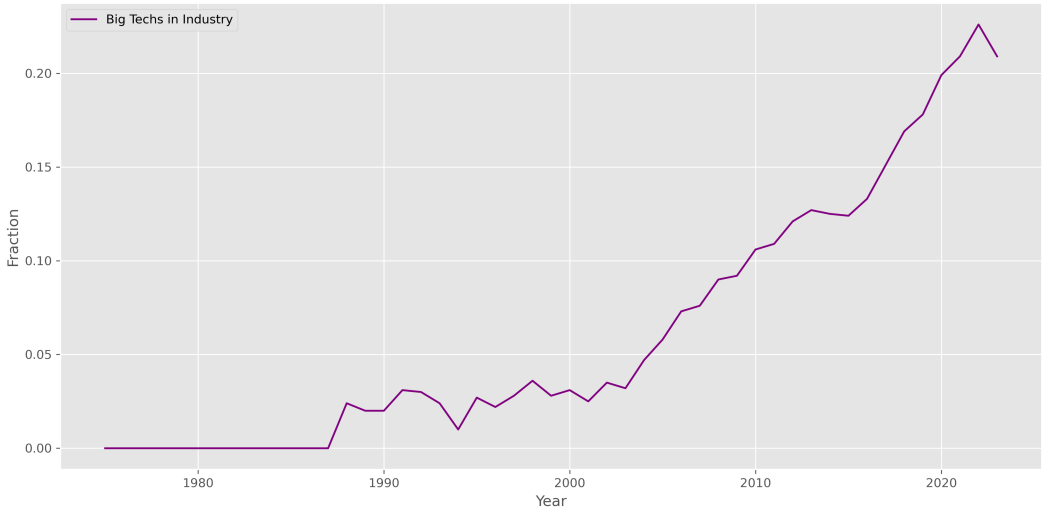


Notes: This figure chart tracks the share of AI graduates employed in academia (red) and industry (blue) from 1975 to 2024. It reveals a long-term shift: while academia's share initially grew, it peaked around 2010 and has declined since, whereas industry employment has steadily risen, reflecting increasing absorption of AI talent by the private sector.

Within industry, Big Tech Firms, have become increasingly dominant. In the year 2010, only about 10% of industry workers were employed in these firms, but this share rose dramat-

ically to 20.9% by 2023. This concentration underscores the centralizing power of big tech firms in the AI labor market. Such firms not only outcompete academia in talent acquisition but also shape the direction of research through internal priorities and resource asymmetries. The shift also raises important questions about knowledge privatization and the narrowing of research goals around commercial imperatives.

Figure 3: Rise of Big Tech Employment in Industry



Notes: This figure charts the growing share of AI professionals working in Big Tech Firms from the late 1980s to 2024. It highlights a sharp and steady increase starting around 2005, peaking above 20% in 2023, emphasizing the centralization of elite AI talent within a few dominant tech firms.

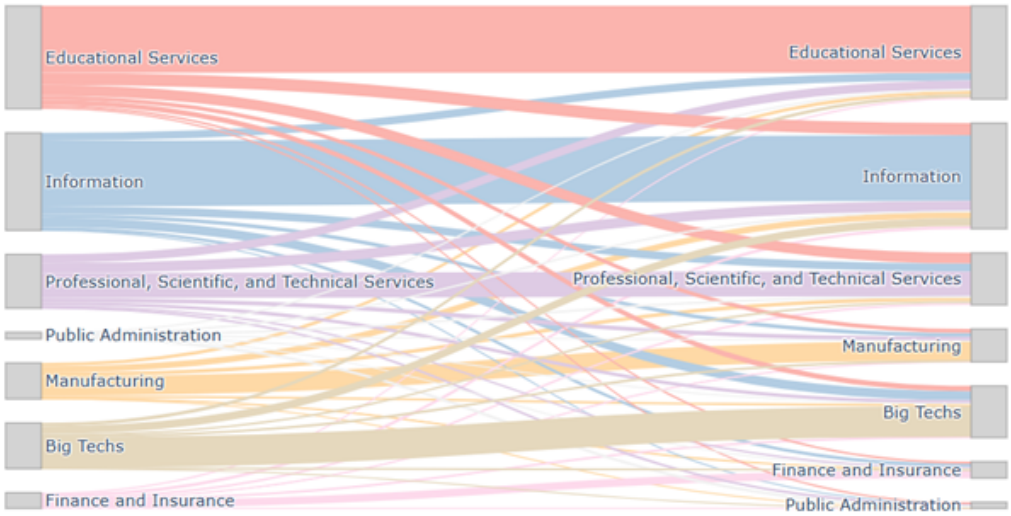
Talent mobility across industries shows us the patterns of retention and relocation in different sectors. As we can see in figure 4.

Core knowledge-intensive sectors such as Information (NAICS code 51), Education, and Big Tech firms demonstrate strong internal absorption capacities, with over 65% of transitions occurring within the same industry. This indicates a tendency toward internal labor market consolidation, where domain expertise remains locked within sectoral boundaries. At the same time, engineering and consulting professionals, especially those from manufacturing and Professional Services (NAICS code 54), are increasingly transitioning into the information sector, with nearly 20% of such movements reflecting a broader realignment toward digital-intensive roles.

Meanwhile, Professional Services appear to function as a transitional hub within the broader labor market. Not only does it maintain a significant portion of its own workforce (46%), but it also absorbs talent from a range of adjacent sectors, including Education (10%),

Manufacturing (10%), and Public Administration (14%). This signals a growing demand for hybrid skill sets, such as AI consulting, data analytics, and applied R&D, that bridge academic, industrial, and digital domains. In contrast, sectors such as Public Administration and Administrative Support (NAICS code 92) exhibit relatively low transition probabilities (mostly $< 5\%$), underscoring their limited interoperability with other labor markets.

Figure 4: Career Transitions Between Industries



Notes: This Sankey diagram visualizes inter-industry mobility, showing how AI talents transition between sectors over time. It highlights strong internal retention within Education, Information, and Big Tech, and shows significant cross-sector flows through Professional, Scientific, and Technical Services, which acts as a key intermediary hub.

4.1 Clustering Approach

Step 1: Data Preprocessing: Representing Career Trajectories

The first step involves transforming raw employment history data into a structured and machine-readable format suitable for sequential modeling. To prepare this data for sequence modeling, we construct for each individual a career path represented as a chronological sequence of position segments. Formally, we define each token representation as:

$$x_t^{(i)} = W \cdot \left[r_t^{(i)}, n_t^{(i)}, d_t^{(i)}, y_t^{(i)} \right] + b$$

where r_t is the encoded role, n_t is the encoded NAICS code, d_t is the duration in months, y_t is the job start year, and W is a learned projection matrix. Thus, the full input sequence

for individual i is:

$$X^{(i)} = \left[x_1^{(i)}, x_2^{(i)}, \ldots, x_{T_i}^{(i)}\right]$$

where T_i denotes the number of job segments for individual i, which varies across individuals. Each token $x_t^{(i)}$ corresponds to one position that individual i held in their employment trajectory.

To account for the order of jobs and allow the model to differentiate between early-career and late-career transitions, we incorporate positional information into each token too.

Step 2: Sequence Modeling with Transformer Encoder

To obtain a compact and informative representation of each individual's career trajectory, we employ a Transformer encoder architecture. Each individual's career sequence—comprising a variable number of job transitions—is treated as a sequence of tokens. The Transformer processes this sequence and outputs a fixed-dimensional embedding vector representing the entire trajectory (Vafa et al., 2022).

To enhance the contextual modeling of career trajectories, we introduce an *role and industry aware bias term* into the self-attention mechanism. This modification accounts for semantic similarities between roles and industries, enabling the model to better capture transitions occurring within related economic sectors.

In the original Transformer, the scaled dot-product attention is computed as:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V$$

In our model, we introduce two additive bias matrixes B_{role} and B_{industry} into the attention logits:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^\top}{\sqrt{d_k}} + B_{\text{role}} + B_{\text{industry}}\right)V$$

Specifically, we use a pre-trained Sentence-BERT model to embed the natural language context of each occupational role and industry. Cosine similarity between these embeddings is then used to construct the more reasonable bias matrices B_{role} and B_{industry}, enabling the model to modulate attention based on semantic closeness rather than exact identity. This allows the model to capture nuanced relationships between roles or industries that are lexically or functionally similar, even if they belong to different formal categories.

Here the bias matrix B_{role} is computed dynamically based on the similarity between occupational role embeddings. Each occupational role is mapped to a dense vector using a pre-trained Sentence-BERT model applied to the textual content R_j of that role j. The bias

matrix B_{industry} is constructed in the same manner, based on the cosine similarity between Sentence-BERT embeddings of industry descriptions R_k of industry k. These descriptions are sourced from the official NAICS definitions provided by the U.S. Census Bureau[1].

After processing the token sequence through the Transformer layers, we obtain a new sequence of contextualized embeddings:

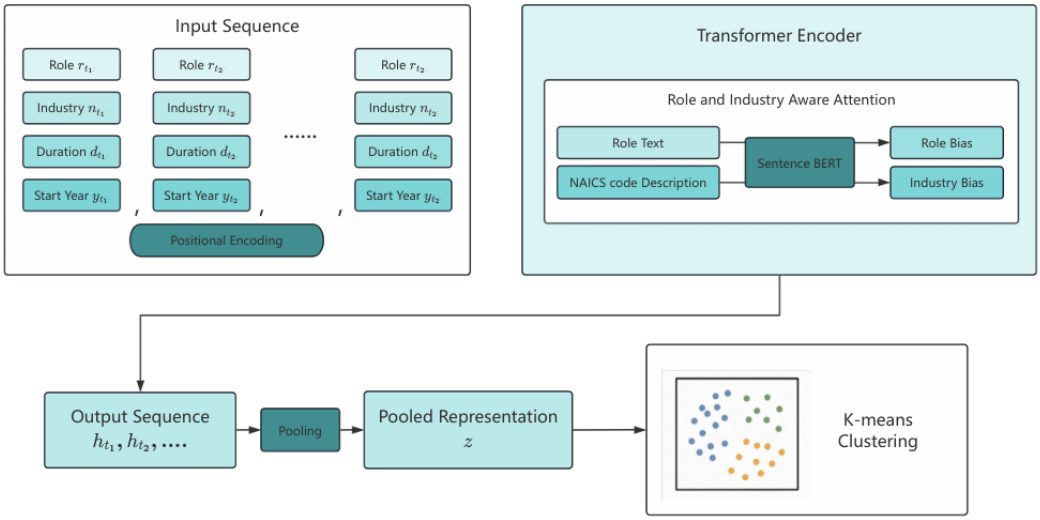
$$H^{(i)} = \left[h_1^{(i)}, h_2^{(i)}, \ldots, h_{T_i}^{(i)}\right]$$

Step 3: Clustering of Career Embeddings Using K-Means

After obtaining the high-dimensional trajectory embeddings for all individuals from the Transformer encoder, the next step is to identify groups of similar career paths using unsupervised clustering. For this purpose, we apply the K-Means clustering algorithm to map each individual's career trajectory to a cluster label, which signifies a learned career archetype.

The whole modeling process is shown in Figure 5.

Figure 5: Career Trajectory Modeling Pipeline



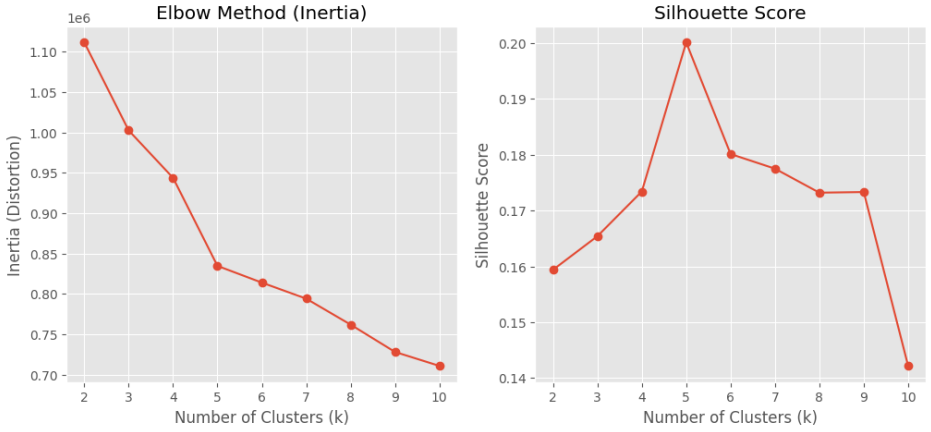
Notes: This diagram illustrates the study's methodological pipeline. Step 1 encodes each individual's career as a sequence of job attributes (industry code, role, duration, start year). Step 2 uses a Transformer encoder with role-and-industry-aware bias to generate a dense trajectory embedding. Step 3 applies K-Means clustering to group similar career paths, identifying distinct mobility archetypes in the AI labor force.

[1]Detailed description of each NAICS code is from U.S. Census Bureau - Understanding NAICS 2022, https://www.census.gov/programs-surveys/economic-census/year/2022/guidance/understanding-naics.html

4.2 Patterns of AI Talents' Movements across Industries

We apply K-Means clustering on these learned representations, selecting k = 5 based on evaluation of both the Elbow Method and Silhouette Score (Figure 6). The resulting clusters represent semantically distinct patterns of career development, characterized by differing dynamics in role transitions, industry mobility, and temporal structure.

Figure 6: Cluster Selection Evaluation



Notes: This side-by-side plot uses two standard metrics to determine the optimal number of career clusters. The Elbow Method (left) shows a sharp drop in inertia at k=5, indicating diminishing returns beyond that point. The Silhouette Score (right) peaks at k=5, suggesting the best-defined clustering also occurs at this value—justifying the use of four career trajectory clusters.

Cluster 1: Technical roles in information

Cluster 1 includes 1,491 individuals whose career trajectories are characterized by stable technical roles within the information sector. The average career start year is 2013.9 with a median of 2015. Individuals in this cluster switch jobs 3.72 times on average, suggesting moderate mobility largely confined within a single industry. 64.9% of individuals started in Information, and 78.6% were in that industry by the end of the observed period. There is limited but notable inflow from Professional, Scientific, and Technical Services (14.1%) and Education (6.4%), suggesting industrial shifts into information industry.

Cluster 2: functional transitions and cross-sector mobility, with Professional Services as a Central Hub

Cluster 2 contains 1,259 individuals and shows the highest role and industry fluidity among all clusters, with an average of 3.89 job switches per person. The average career start is

2013.1. This group is distinguished by both functional transitions and cross-sector mobility. Professionally, 36.1% start in Professional Services, followed by 15.5% in Finance and 10.2% in Education . By the latest records we have now, 41.1% remain in Professional Services, while 14.5% in Finance. For AI knowledge transfer, this group exemplifies a fluid, flexible model where expertise moves through hybrid functions and sectors.

Cluster 3: Manufacturing-Origin Professionals

Cluster 3 consists of 592 individuals, with an average start year of 2012.1. They have an average of 3.59 job transitions, reflecting a relatively stable and long-spanning career. 57.3% begin in Manufacturing, with 55.3% there at the end. Only 12.2% ended in Information, and 6.8% reach Big Tech firms in the latest records. This cluster represents experienced, technically focused professionals embedded in the industrial engineering economy, with narrow functional boundaries.

Cluster 4: Early-Career AI Engineers

Cluster 4 is composed of 863 individuals, with the shortest career timelines: an average start year of 2016.6 with the median 2018. These individuals switch jobs only 2.97 times on average, the lowest across all clusters. Strikingly, 81.7% start in Big Tech ("777") and 81.9% were there in our latest record, indicating exceptional retention and focus within big tech firms. Cluster 4 exemplifies a very young group of graduates in the knowledge ecosystem where big tech firms internalize the talent pipeline, capturing AI knowledge directly at source, upon graduation or internships.

Cluster 5: Academically Trained Researchers with Strong Retention in Education

Cluster 5 includes 1,558 individuals, making it the largest group. Their average start year is 2013.5, with 3.18 job transitions on average. It is defined by a strong academic origin and evolution of its role within the education-research ecosystem. 85.9% start in Education , and 82.8% are still there at the end. Only 2.3% reach Big Tech firms, and flows to Information (5.5%) and Professional Services (6.3%) are limited.

5 Elite AI Talent and Firm Productivity

To evaluate the impact of frontier AI talent on firm-level productivity, we estimate production functions that incorporate both traditional inputs and embodied technological knowledge.

This section outlines the empirical framework, measurement of variables, and identification strategies.

5.1 Approach

We begin with a log-linear Cobb–Douglas production function augmented with talent-based inputs:

$$\log(Y_{it}) = \alpha \log(K_{it}) + \gamma \log(L_{it}) + \delta \log(Talent_{it}) + Z'\theta + \mu_i + \lambda_t + \eta_s + \epsilon_{it} \tag{1}$$

where:

- Y_{it}: Firm output, measured as value added. Following standard practice, we compute nomimal value added as sales minus intermediate inputs which are proxied by the cost of good sold; then we deflate this nominal value added using the yearly price deflator to obtain the real value added
- K_{it}: Capital input, constructed using the Perpetual Inventory Method (PIM), with initial capital stock set equal to the firm's earliest observed capital expenditure and depreciated at 15% annually.
- L_{it}: Labor input, measured as the number of full-time employees.
- $Talent_{it}$: Embodied technological input, measured as the cumulative stock of frontier AI talent (graduates of elite AI scientists' lineages) employed by firms.
- Z: Controls, specifically firm-level R&D expenditures.
- μ_i, λ_t, η_s: Firm, time, and industry level (2-digit NAICS) of fixed effects.
- ϵ_{it}: Error term.

This structure isolates the contribution of frontier AI talent relative to traditional factors of production while leveraging high-dimensional fixed effects to mitigate confounding.

The cumulative number of graduates trained by elite AI scientists (Immortals) subsequently employed by firms. This measure captures the inflow of embodied frontier knowledge into firms. By focusing on embodied knowledge, we capture the channel of AI diffusion most closely tied to tacit, inimitable expertise.

While panel regressions cannot eliminate endogeneity concerns, our empirical design incorporates several strategies to strengthen inference. First is the use of fixed effects for firms, years and industries. Second is the use of dynamic specifications with lags in which we include one- and two-year lags of talent measures to capture delayed effects and reduce simultaneity. In ongoing research, we are exploring event study approaches that could further strengthen our inference regarding causal effects.

Given AI's status as a general-purpose technology, its productivity impact is expected

to vary by sector. We therefore estimate industry-specific models for two major contexts: (a) Advanced Manufacturing: where adoption of frontier AI talent is conditioned by physical complements, capital intensity, and longer diffusion cycles. (b) Information Services: where digital complementarities and higher absorptive capacity may facilitate more rapid assimilation. This heterogeneity analysis tests whether productivity gains from frontier AI talent are mediated by sectoral readiness and complementary assets. Our empirical strategy relies on a panel identification framework with high-dimensional fixed effects and dynamic lag structures.

5.2 Event Study Design

To strengthen causal interpretation, we implement an event study framework around firms' first entry into AI patenting. Treated firms are those that file their first AI-related patent in year $t = 0$. Each treated firm is matched to a control firm that has not patented in AI, using propensity score matching based on pre-treatment characteristics: capital stock, employment, R&D intensity, leverage, and 2-digit industry code. Control firms are assigned a placebo adoption year corresponding to their matched treated firm, ensuring comparability in event time.

We estimate the following dynamic specification over an 11-year symmetric window ($t \in [-5, +5]$), omitting the adoption year as the reference period:

$$Y_{it} = \sum_{k=-5, k\neq 0}^{+5} \beta_k D_{it}^k + \alpha \log(K_{it}) + \gamma \log(L_{it}) + \delta \log(XRD_{it}) + \mu_i + \lambda_t + \eta_s + \epsilon_{it}. \quad (2)$$

Here, D_{it}^k is an indicator for event year k relative to adoption. Coefficients for pre-treatment years ($k < 0$) test the parallel-trends assumption, while post-adoption coefficients ($k > 0$) trace the dynamic evolution of productivity following entry into AI invention. Standard errors are clustered at the firm level. The details are outlined in the appendix section.

5.3 Results

5.3.1 Productivity Effects of Frontier AI Talent

We next turn to embodied frontier knowledge, proxied by the hiring of graduates trained by elite AI scientists ("AI children"). These individuals represent a distinctive form of human capital: they carry not only advanced technical expertise but potentially also tacit knowledge, heuristics, and network ties embedded in elite academic lineages. By integrating

such talent, firms gain privileged access to frontier AI methods and capabilities that are difficult to replicate through conventional hiring channels.

Across specifications, the elasticity of TFP with respect to cumulative AI children stock ranges from 0.062 to 0.13. These estimates imply that doubling the stock of frontier AI hires increases firm value added by 6.2% to 13%, controlling for capital, labor, R&D, and fixed effects. We find stronger gains in technology-intensive industries such as manufacturing and IT. Strategically, this highlights the role of talent acquisition as a critical lever in capturing the value of general-purpose technologies.

Lagged specifications reveal that the impact of AI children endures over time. One- and two-year lags remain positive and statistically significant, suggesting that productivity gains accumulate as these individuals diffuse knowledge internally, shape organizational routines, and integrate AI into products and processes. This persistence is consistent with theories of absorptive capacity and organizational learning: the returns to frontier talent extend well beyond initial hiring, reflecting cumulative capability building.

The productivity returns to frontier AI talent are robust across industries but differ in intensity. In Information Services, the effects are particularly large, consistent with the sector's greater absorptive readiness, digital infrastructure, and shorter innovation cycles. Advanced Manufacturing also exhibits significant gains, though of slightly smaller magnitude, reflecting the integration frictions and capital intensity of physical production systems. These patterns suggest that while frontier talent is broadly valuable, sectoral absorptive capacity mediates the speed and scale of performance improvements.

Table 2: Impact of Cumulative AI Children on Total Factor Productivity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
VARIABLES	TFP	TFP	TFP	TFP	TFP	TFP	TFP Mfg	TFP Mfg	TFP Info	TFP Info	TFP Info
log_children_stock	0.279***		0.108***		0.102***		0.130***		0.0859***		
	(0.0154)		(0.00947)		(0.0108)		(0.0167)		(0.0294)		
L.log_children_stock		0.247***		0.0895***		0.0869***		0.110***		0.0721**	
		(0.0152)		(0.00981)		(0.0111)		(0.0174)		(0.0298)	
L2.log_children_stock											0.0621**
											(0.0300)
log_k_real			0.270***	0.309***	0.181***	0.209***	0.160***	0.191***	0.158***	0.174***	0.163***
			(0.00763)	(0.00885)	(0.0100)	(0.0120)	(0.0155)	(0.0185)	(0.0291)	(0.0351)	(0.0408)
log_emp			0.635***	0.615***	0.631***	0.623***	0.648***	0.634***	0.705***	0.689***	0.694***
			(0.00915)	(0.00972)	(0.0136)	(0.0143)	(0.0204)	(0.0212)	(0.0505)	(0.0550)	(0.0599)
log_xrd_real					0.132***	0.118***	0.141***	0.131***	0.0967***	0.0979***	0.0961***
					(0.00914)	(0.00954)	(0.0150)	(0.0157)	(0.0311)	(0.0332)	(0.0358)
Constant	3.981***	4.055***	2.215***	2.057***	2.239***	2.164***	2.099***	2.011***	2.336***	2.290***	2.354***
	(0.00427)	(0.00422)	(0.0272)	(0.0329)	(0.0363)	(0.0438)	(0.0534)	(0.0636)	(0.121)	(0.141)	(0.165)
Observations	65,996	61,724	63,581	59,399	41,955	39,267	19,300	18,065	5,216	4,833	4,461
R-squared	0.893	0.899	0.930	0.932	0.934	0.936	0.925	0.927	0.921	0.924	0.927

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

To complement these TFP results, we next examine whether the productivity advantages

of frontier AI talent extend to labor productivity.

5.3.2 Labor Productivity Effects of Frontier AI Talent

While Table 2 established that frontier AI talent significantly raises total-factor productivity, here we test whether these effects extend to labor productivity. This outcome captures whether hiring elite AI graduates enhances the productivity of the average worker.

Table 3: Impact of Cumulative AI Children on Labor Productivity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
VARIABLES	labor prod.	labor prod.	labor prod.	labor prod.	labor prod.	labor prod.	labor prod. mfg.	labor prod. mfg.	labor prod. info	labor prod. info	labor prod. info
log_children_stock	0.0749***		0.0883***		0.0902***		0.124***		0.0742**		
	(0.00939)		(0.00935)		(0.0107)		(0.0166)		(0.0290)		
L.log_children_stock		0.0644***		0.0762***		0.0781***		0.107***		0.0629**	
		(0.00978)		(0.00976)		(0.0111)		(0.0174)		(0.0298)	
L2.log_children_stock											0.0533*
											(0.0301)
log_k_per_emp			0.296***	0.340***	0.191***	0.225***	0.170***	0.205***	0.160***	0.180***	0.175***
			(0.00772)	(0.00867)	(0.00995)	(0.0116)	(0.0154)	(0.0179)	(0.0291)	(0.0348)	(0.0399)
log_xrd_per_emp					0.139***	0.122***	0.147***	0.136***	0.102***	0.102***	0.101***
					(0.00917)	(0.00957)	(0.0150)	(0.0157)	(0.0319)	(0.0341)	(0.0368)
Constant	3.006***	3.017***	2.035***	1.874***	2.146***	2.057***	2.014***	1.918***	2.307***	2.244***	2.279***
	(0.00316)	(0.00323)	(0.0253)	(0.0291)	(0.0335)	(0.0387)	(0.0495)	(0.0561)	(0.124)	(0.144)	(0.165)
Observations	65,942	61,675	63,581	59,399	41,955	39,267	19,300	18,065	5,216	4,833	4,461
R-squared	0.691	0.701	0.706	0.716	0.655	0.665	0.564	0.575	0.653	0.666	0.675

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Across specifications, the elasticity of labor productivity with respect to cumulative AI children stock ranges from 0.053 to 0.124. These estimates imply that doubling the stock of frontier AI hires raises the productivity of the average worker by 5.3% to 12.4%, controlling for capital, labor, R&D, and fixed effects. Again, we find stronger results in technology-intensive areas such as information services and advanced manufacturing.

Lagged specifications reveal that the labor productivity impact of AI children persists over time. One- and two-year lags remain positive and statistically significant, suggesting that efficiency gains accumulate as frontier-trained individuals diffuse tacit knowledge, shape organizational routines, and embed AI capabilities across teams. This persistence is consistent with theories of absorptive capacity and organizational learning: the productivity returns to frontier talent extend well beyond the initial hiring event, reflecting cumulative capability building that strengthens firms' long-run adaptive capacity.

From a strategic perspective, these findings highlight how elite AI hires enhance the productivity of the average worker. This aligns with resource-based and dynamic capabilities views—elite-trained individuals may embed tacit knowledge, problem-solving heuristics, and innovation routines that diffuse across teams. Sectoral splits reinforce this point: Information Services reap especially strong labor productivity gains, consistent with higher absorptive readiness, while manufacturing shows significant but more gradual effects, possibly due to integration frictions.

Taken together, these results position frontier AI talent as a VRIN-like strategic asset: rare, inimitable, and embedded in academic lineages that are not easily replicated. Unlike generic skill categories, these individuals may act as catalysts for organizational learning, cross-functional integration, and long-run capability building. Their presence may amplify the firm's ability to adapt, recombine, and scale AI knowledge, aligning with dynamic capability perspectives on sensing, seizing, and reconfiguring resources. Firms that successfully attract and embed frontier AI talent occupy a privileged position in the evolving AI landscape.

6 Discussion and Conclusions

By linking firm-level performance data to academic lineage measures of elite AI talent, this study provides new evidence on how embodied knowledge shapes productivity.

We document three main findings. First, frontier AI talent is an economically meaningful input: firms that hire elite-lineage graduates experience persistent and significant gains in total factor productivity (TFP). Second, analyses show that growth in elite graduates is associated with subsequent increases in firms' labor productivity. Third, contrary to popular belief, large technology firms have not monopolized access to elite graduates. Sectoral analyses reveal especially large effects in information services, where digital complements may accelerate adoption of AI or innovation in AI, and slower gains in manufacturing, where integration frictions may dampen the adoption and diffusion of AI-related innovations.

Together, these results advance economics of innovation by showing that the origins of talent matter: academic lineages of AI experts condition how frontier AI knowledge is absorbed and applied within firms. From a strategy perspective, our findings resonate with the resource-based view: elite AI talent functions as a VRIN-like asset—valuable, rare, inimitable, and non-substitutable—that anchors persistent heterogeneity in firm performance.

For managers, the implication is that capturing AI's value requires more than recruiting technical skills; it depends on cultivating pipelines to frontier talent and creating organizational conditions in which their capabilities can be fully utilized. For scholars, our findings suggest that lineage-based measures of human capital open new avenues for studying heterogeneity in firm performance, capability development, and adaptation to general-purpose technologies.

6.1 Limitations

While our findings are robust and broadly consistent with the view of AI as a general-purpose technology, several limitations merit caution and point to opportunities for future research.

Although the Cobb–Douglas production framework is standard, its estimation is vulnerable to endogeneity and omitted variable bias (Levinsohn and Petrin, 2003; Wooldridge, 2009). Firms anticipating future productivity gains may selectively recruit elite AI talent, generating potential reverse causality. Our design uses firm, year, and industry fixed effects, lag structures, and propensity-score matching to mitigate these concerns, but cannot fully account for unobserved, time-varying factors such as managerial quality, strategic vision, or concurrent digital investments. In strategy settings, such anticipatory capability building is itself endogenous to performance.

Because our analysis relies on Compustat, it focuses on publicly traded U.S. firms—larger, more established organizations with richer resources. This may underrepresent outcomes among private firms and startups where much frontier AI experimentation occurs. Moreover, our measures do not observe within-firm microdynamics (e.g., task reallocation, wage structures, team topology) that increasingly matter for understanding capability reconfiguration and value capture. Linking academic lineage-based talent measures to richer personnel and organizational data would illuminate these mechanisms.

Taken together, these limitations do not undermine the core contribution: documenting robust empirical relationships between frontier AI talent—embedded in elite academic lineages—and heterogeneity in firm productivity. Rather than definitive causal claims, our study provides disciplined evidence that embodied frontier knowledge is a powerful strategic lever. Future theoretical and empirical work should unpack boundary conditions (e.g., data assets, governance, complementarities), mediating processes (e.g., experimentation cultures, MLOps maturity), and dynamic interactions between elite lineage access and organizational capability development.

6.2 Future Work

Looking forward, several avenues offer promising opportunities to deepen our understanding of how frontier AI talent reshapes firm strategy and performance.

First, strengthening causal identification remains a key priority. Future work could leverage exogenous shocks, such as immigration policy changes affecting access to technical talent, shifts in university research funding, or sudden breakthroughs in AI subfields, and interact them with firm-level exposure to elite graduates. Such natural experiments would allow researchers to disentangle selection from treatment effects, sharpening causal estimates of

frontier talent's contribution to firm performance. Beyond econometric rigor, these designs would illuminate how external shocks condition firms' ability to mobilize and capture embodied frontier knowledge, a question central to strategy.

Second, expanding the scope of data would enrich theoretical and empirical insights. Incorporating U.S. Census microdata, including the Longitudinal Business Database and establishment-level records, would allow researchers to study private firms, younger ventures, and multi-unit enterprises alongside public firms. This would broaden our understanding of how organizational form, governance, and scale shape access to and utilization of elite AI talent. Such data could also enable fine-grained analysis of how frontier-trained individuals reshape labor demand within firms—across occupations, skill tiers, and wage structures—shedding light on the distributional consequences of technological change and their implications for firm boundaries, workforce strategy, and organizational design.

Third, refining production function models to account for industry-specific contexts could improve precision. Different industries exhibit distinctive capital intensities, labor compositions, and knowledge complementarities. Tailoring estimation to these realities would reveal not only whether frontier AI talent matters but also how its value unfolds differently across settings. This would allow strategy scholars to better theorize complementarities and boundary conditions of embodied knowledge.

Finally, a deeper exploration of heterogeneity is critical. The effects of frontier AI talent are unlikely to be uniform across industries, geographies, or firm sizes. Some firms may leverage elite hires to scale and entrench competitive advantages, while others may face disruption if unable to mobilize complementary assets. Future research could map these heterogeneous trajectories, identifying which other organizational capabilities, governance models, or market positions amplify or dampen the value of frontier talent. Especially important is examining how the infusion of elite AI talent restructures labor demand—whether by substituting for routine work, complementing specialized knowledge, or creating new categories of employment. These dynamics speak directly to strategic questions of capability building, resource redeployment, and long-run positioning.

References

Agrawal, A., Gans, J., and Goldfarb, A. (2019). *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Agrawal, A., Gans, J., and Goldfarb, A. (2021). *Power and Prediction: The Disruptive Economics of AI*. Harvard Business Review Press.

Azoulay, P., Graff Zivin, J. S., and Wang, J. (2010). Superstar extinction. *Quarterly Journal of Economics*, 125(2):549–589.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 154:102–132.

Baily, M., Bosworth, B., Goolsbee, A., and Syverson, C. (2023). Prospects for a productivity boom: Ai and beyond. *Brookings Papers on Economic Activity*.

Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1):99–120.

Becker, G. S. (1964). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. University of Chicago Press, Chicago.

Bloom, N. and Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *Quarterly Journal of Economics*, 122(4):1351–1408.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of Econometrics*, 65(1):83–108.

Brynjolfsson, E., Li, D., and Raymond, L. (2023). Generative ai at work. Technical Report 31161, NBER.

Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age*. W. W. Norton.

Brynjolfsson, E. and McElheran, K. (2022). The rapid adoption of data-driven decision-making. *AEJ: Applied Economics*. For cross-sectoral diffusion context.

Brynjolfsson, E., Rock, D., and Syverson, C. (2019). Artificial intelligence and the modern productivity paradox. In *The Economics of Artificial Intelligence: An Agenda*, pages 23–57. University of Chicago Press.

Cockburn, I. and Henderson, R. (1998). Absorptive capacity, coauthoring behavior, and the organization of research. *Journal of Industrial Economics*, 46(2):157–182.

Cockburn, I., Henderson, R., and Stern, S. (2019). The impact of artificial intelligence on innovation. In Agrawal, A., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*, pages 115–146. University of Chicago Press.

Cohen, W. M. and Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1):128–152.

David, P. A. (1991). The computer and the dynamo: The modern productivity paradox in a not-too-distant mirror. In *Technology and Productivity: The Challenge for Economic Policy*, pages 315–348. OECD.

Dyer, J. H. and Singh, H. (1998). The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *Academy of Management Review*, 23(4):660–679.

Dyèvre, D. and Seager, P. (2023). Linking 70 years of uspto patents to compustat. Working Paper.

Eisenhardt, K. M. and Martin, J. A. (2000). Dynamic capabilities: What are they? *Strategic Management Journal*, 21(10-11):1105–1121.

Grant, R. M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17:109–122.

Griliches, Z., editor (1984). *R&D, Patents, and Productivity*. University of Chicago Press, Chicago, IL.

Gulati, R. (1998). Alliances and networks. *Strategic Management Journal*, 19(4):293–317.

Hansen, M. T. (1999). The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly*, 44(1):82–111.

Helpman, E., editor (1998). *General Purpose Technologies and Economic Growth*. MIT Press.

Hitt, M. A., Bierman, L., Shimizu, K., and Kochhar, R. (2001). Direct and moderating effects of human capital on strategy and performance. *Academy of Management Journal*, 44(1):13–28.

Jordan, M. I. and Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260.

Jovanovic, B. and Rousseau, P. L. (2005). General purpose technologies. In *Handbook of Economic Growth*, volume 1B, pages 1181–1224. Elsevier.

Lee, R. and Schankerman, M. (2020). Knowledge diffusion and the impacts of open science. In *American Economic Review*, volume 110, pages 3091–3114.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, 70(2):317–341.

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

Mollick, E. (2012). People and process, suits and innovators. *Strategic Management Journal*, 33(9):1001–1015.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative ai. *Science*, 381(6654):187–192.

Pakes, A. and Griliches, Z. (1984). Patents and r&d at the firm level: A first look. In Griliches, Z., editor, *R&D, Patents, and Productivity*, pages 55–72. University of Chicago Press, Chicago, IL.

Rosenkopf, L. and Almeida, P. (2003). Overcoming local search through alliances and mobility. *Management Science*, 49(6):751–766.

Stuart, T. E. and Sorenson, O. (2007). Strategic networks and entrepreneurial ventures. *Strategic Entrepreneurship Journal*, 1(3-4):211–227.

Susskind, D. (2020). A world without work. *Allen Lane*.

Teece, D. J. (1986). Profiting from technological innovation. *Research Policy*, 15(6):285–305.

Teece, D. J. (2007). Explicating dynamic capabilities. *Strategic Management Journal*, 28(13):1319–1350.

Teece, D. J. (2018). Dynamic capabilities as (workable) management systems theory. *Journal of Management Studies*, 55(7):1338–1349.

Teece, D. J., Pisano, G., and Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7):509–533.

Vafa, K., Palikot, E., Du, T., Kanodia, A., Athey, S., and Blei, D. M. (2022). Career: Transfer learning for economic prediction of labor sequence data. Technical Report 4074.

Wooldridge, J. M. (2009). On estimating firm-level production functions using proxy variables. *Economics Letters*, 104(3):112–114.

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D., Buffington, C., Foster, L., and Dinlersoz, E. (2020). Advanced technologies adoption and use by u.s. firms: Evidence from the annual business survey. NBER Working Paper 28290, National Bureau of Economic Research.

Zucker, L. G., Darby, M. R., and Brewer, M. B. (1998). Intellectual human capital and the birth of u.s. biotechnology enterprises. *American Economic Review*, 88(1):290–306.

Appendix

A Effects of Frontier AI Talent on AI Innovation

This section evaluates whether firms that employ frontier AI talent—graduates of elite AI scientists ("Immortals") exhibit stronger AI inventive output. Adopting a cumulative stock measure as well as active employment of these graduates, we estimate firm-level AI patent production functions including capital, R&D, and fixed effects. Results show consistent and economically large associations: doubling the number of Immortal-trained hires is linked to 50–100 percent higher AI patenting, depending on specification. Effects are stronger in advanced manufacturing (NAICS 333–336), suggesting greater complementarity between embodied frontier talent and complex physical system innovation.

A.1 Cumulative AI Patents

Table A1: Impact of Frontier AI Talent on Cumulative AI Patenting

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES	AI Patents	AI Patents	AI Patents	AI Patents	AI Patents (Adv. Mfg.)	AI Patents (Adv. Mfg.)	AI Patents (Info)	AI Patents (Info)
log_children_stock	0.903***				0.999***		0.774***	
	(0.0140)				(0.0198)		(0.0352)	
L.log_children_stock		0.875***				0.972***		0.736***
		(0.0149)				(0.0213)		(0.0374)
log_children_personyears_cum			0.494***					
			(0.0106)					
L.log_children_personyears_cum				0.481***				
				(0.0112)				
log_k_real	0.00339	0.00803***	0.00335	0.00864***	-0.000997	0.00377	0.0725***	0.0882***
	(0.00252)	(0.00299)	(0.00273)	(0.00323)	(0.00534)	(0.00619)	(0.0123)	(0.0147)
log_rd_real	0.0306***	0.0314***	0.0486***	0.0483***	0.0316***	0.0374***	0.0418**	0.0291
	(0.00369)	(0.00413)	(0.00413)	(0.00456)	(0.00646)	(0.00733)	(0.0198)	(0.0222)
Constant	-0.0204	-0.0150	-0.0328**	-0.0291	0.0155	0.00201	-0.0789	-0.0177
	(0.0145)	(0.0166)	(0.0160)	(0.0180)	(0.0234)	(0.0272)	(0.0744)	(0.0849)
Observations	70,586	63,226	70,586	63,226	27,024	24,698	9,296	8,194
R-squared	0.841	0.852	0.813	0.829	0.840	0.848	0.863	0.875

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table A2: Impact of Frontier AI Talent on Five-Year Moving-Average AI Patents

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES	AI Patents	AI Patents	AI Patents	AI Patents	AI Patents (Adv. Mfg.)	AI Patents (Adv. Mfg.)	AI Patents (Info)	AI Patents (Info)
log_children_active	0.526***				0.619***		0.424***	
	(0.0164)				(0.0239)		(0.0357)	
L.log_children_active		0.499***				0.593***		0.383***
		(0.0173)				(0.0252)		(0.0371)
log_children_personyears			0.584***					
			(0.0189)					
L.log_children_personyears				0.555***				
				(0.0200)				
log_k_real	0.00645***	0.0107***	0.00640***	0.0105***	-0.000457	0.00410	0.0521***	0.0646***
	(0.00185)	(0.00222)	(0.00186)	(0.00224)	(0.00389)	(0.00461)	(0.00947)	(0.0114)
log_rd_real	0.0252***	0.0252***	0.0249***	0.0251***	0.0258***	0.0276***	0.0515***	0.0474***
	(0.00288)	(0.00322)	(0.00288)	(0.00322)	(0.00511)	(0.00587)	(0.0151)	(0.0170)
Constant	-0.0226*	-0.0281**	-0.0197*	-0.0250*	0.0114	0.000273	-0.102*	-0.105
	(0.0117)	(0.0134)	(0.0117)	(0.0134)	(0.0190)	(0.0224)	(0.0573)	(0.0659)
Observations	70,586	63,226	70,586	63,226	27,024	24,698	9,296	8,194
R-squared	0.799	0.810	0.799	0.809	0.781	0.792	0.837	0.847

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

These estimates indicate that frontier AI talent does not merely enhance short-run inventive activity, but materially scales firms' AI innovation efforts through absorptive capacity, capability building, and the integration of frontier scientific knowledge into applied technologies.

B Event-Study Evidence: Dynamic Productivity Effects of AI Adoption

Building on the preceding evidence that employing frontier AI talent significantly expands firms' AI inventive output (Section A), this section examines the downstream performance consequences of those AI innovations. Specifically, we analyze how the transition into AI patenting translates into changes in firm productivity over time.

To trace this dynamic relationship, we implement an event-study design centered on the year of first AI patenting, capturing the trajectory of firm output before and after AI adoption. Treated firms are matched to structurally similar non-adopters using propensity score matching, and placebo adoption years are assigned to ensure aligned event time.

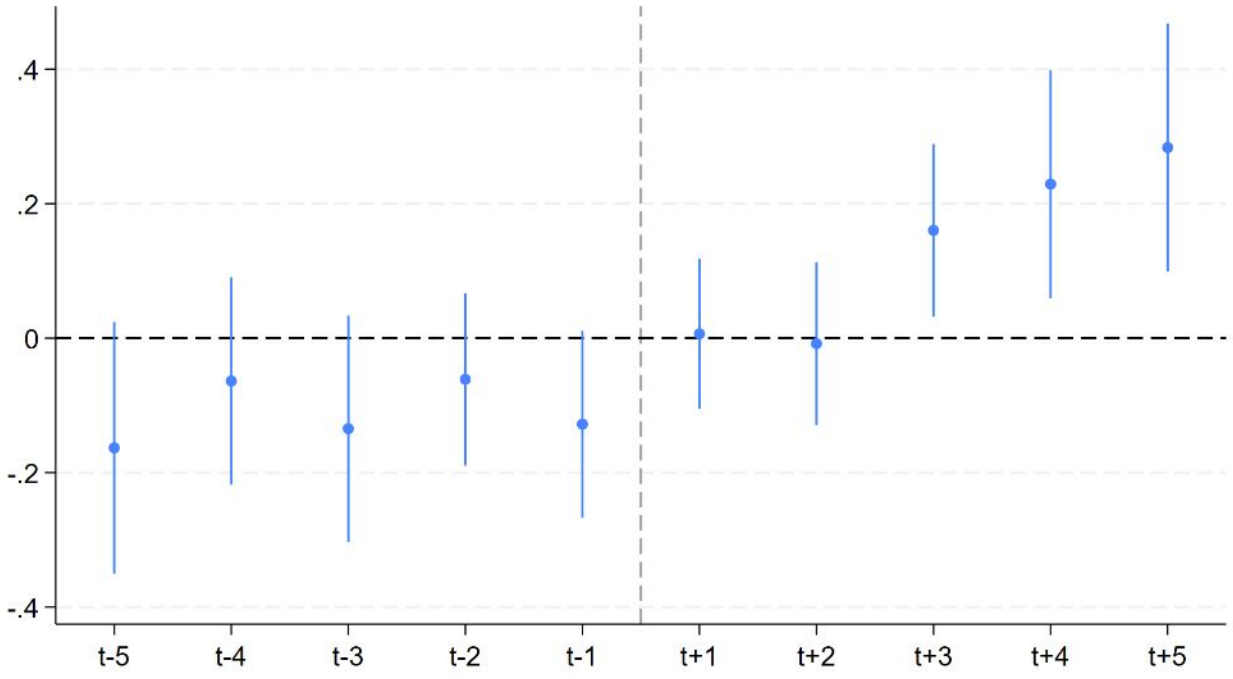


Figure A1: Dynamic effects of AI patenting on firm value added. Notes: Coefficients and 95% confidence intervals estimated from event-time regressions with firm, year, and industry fixed effects.

The pre-adoption coefficients ($t = -5$ to $t = -1$) are statistically indistinguishable from zero, confirming that treated and matched control firms follow parallel productivity trends prior to adoption. Following the first AI patent, productivity rises gradually and persistently. Five years after adoption, firms exhibit approximately 25–30 percent higher real value added relative to their matched peers.

This monotonic trajectory aligns with theories of absorptive capacity and organizational learning: the economic value of AI inventions materializes as firms accumulate complementary assets and embed new knowledge into production systems. Taken together with the AI talent results, the evidence reveals a coherent dynamic chain: frontier AI scientists spur greater AI innovation, and those innovations translate into sustained productivity growth. AI adoption thus marks a strategic inflection point that initiates long-run divergence in firm performance between adopters and non-adopters.

Table A3: Variables, Symbols, Construction, and Sources

Variable	Symbol	Definition / Construction	Source
Value added	Y_{it}	Calculated as total revenue (SALE) − cost of goods sold (COGS) − selling, general & administrative expenses (XSGA). This captures output generated net of intermediate inputs and overhead.	Compustat
Capital stock (real)	K_{it}	Constructed via Perpetual Inventory Method: $K_{it} = (1-\delta_K)K_{i,t-1} + \text{CAPX}_{it}$, with $\delta_K = 0.15$ annual depreciation. Initial K set from earliest CAPX. Deflated by investment deflator.	Compustat
Labor (employment)	L_{it}	Number of full-time employees (EMP). $\log L_{it}$ used in regressions.	Compustat
R&D (real)	XRD_{it}	Reported R&D expenditures, deflated by R&D price index; $\log(1 + XRD_{it})$ used where zeros exist.	Compustat
Frontier AI talent (cum.)	$AIChild_{it}$	Cumulative number of "AI children" (graduates of Immortals) employed by firm in year t.	Genealogy × Revelio Labs
Capital per employee	K_{it}/L_{it}	Ratio used for per-capita specifications.	Compustat
R&D per employee	XRD_{it}/L_{it}	Ratio; robustness specification.	Compustat

C Variable Definitions

D Measurement and Construction Details

D.1 Value Added

Value added Y_{it} is constructed as sale − cogs − xsga following prior strategy and productivity studies using Compustat (e.g., Hall et al., 1986; Lichtenberg, 1992). This definition ensures consistency across firms by netting out intermediate inputs and administrative overhead.

D.2 Capital Stock via PIM

Initial capital K_{i0} is set to the earliest observed CAPX; annual update $K_{it} = (1-\delta_K)K_{i,t-1} + \text{CAPX}_{it}$ with $\delta_K = 0.15$. Robustness considers $\delta_K \in \{0.10, 0.20\}$.

D.3 Deflators and Price Indices

All dollar values such as output, capital investment, and R&D etc. are deflated as described in the section.

D.4 Sample, Filters, and Winsorization

We retain public U.S. firms with non-missing Y, K, L and at least two consecutive observations.

E Additional Results and Robustness

E.1 Alternative Outcomes: TFP Proxies

We estimate Sales along with value added as outcomes; AI talent coefficients remain positive and significant, supporting productivity interpretation.

E.2 Lead Terms (Placebo)

We include L^{-1} and L^{-2} leads of AI inputs to test for reverse timing; leads are statistically null, supporting a causal ordering from AI inputs to subsequent productivity.

E.3 Clustering and Inference

Results are robust to clustering by firm and by firm×industry; wild bootstrap p-values confirm significance.