Elite Talent, Invention, and Firm Productivity in the Age of AI

Sagar Baviskar[*] Lee Branstetter[†] Cameron Drayton[‡]
Brian C. Fujiy[§] Eduard Hovy[¶] Prasanna Tambe[‖] Liujie Wu[**]

June 2026

Preliminary and Incomplete

The latest version is available here.

Abstract

We hypothesize that AI experts trained at the scientific frontier may constitute an important resource 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 individuals across time and organizational boundaries. We combine these data with Compustat records and standard panel regression methods to examine how frontier AI talent affects productivity, AI invention, and the translation of specific bodies of new scientific knowledge into firms' patents. Our results suggest a modest impact on conventionally measured firm productivity, but stronger effects on AI invention and the translation of recent science into AI invention. Examining the extent to which these relationships are causal is the subject of ongoing research, as is an exploration of these effects in Census microdata, which includes unlisted firms.

[*]Carnegie Mellon University. Corresponding author. Email: sagarbaviskar@cmu.edu
[†]Carnegie Mellon University and NBER
[‡]Carnegie Mellon University
[§]U.S. Census Bureau
[¶]University of Melbourne
[‖]Wharton School, U. Pennsylvania
[**]Carnegie Mellon University

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

In earlier eras of economic history, general purpose technologies (GPTs) such as steam power, electricity, and information technology became pervasive, improved over time, and enabled complementary innovations that reshaped industries and brought significant and persistent increases in 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 expectations are rising that it could have similar – or perhaps even stronger effects (Agrawal et al., 2021; Baily et al., 2023). Earlier GPTs expanded the ability of humans to exert physical force or store and process data. AI holds out the promise of extending human cognition, reducing uncertainty, enabling novel combinations of new ideas, and making expert-level specialized knowledge available on demand Agrawal et al. (2019a,b, 2021, 2025b). Cockburn et al. (2019) highlight AI as a technology for invention itself, altering trajectories of innovative search. 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 and complementary innovations before steep productivity gains materialize. More broadly, recent work on the economics of transformative AI emphasizes that its economic consequences will depend not only on frontier model capabilities, but also on how firms, workers, and markets reorganize around these technologies (Agrawal et al., 2025a).

The well-established consensus around the necessity of complementary investment and the implication that productivity effects may arrive with a lag has motivated 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). Another stream of research has examined early AI-related innovation by identifying AI-related patents (Cockburn et al., 2019; Webb et al., 2019; Giczy et al. 2021, 2023; Miric, Jia, and Huang, 2023; Fang, Gu, Yan, and Zhu (2026)). Baviskar et al. (2024) use NLP techniques to identify AI patents, connect these data to firm-level outcomes using Compustat and Census micro data, and 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 most find no statistically significant effects of this recruitment on firm productivity growth (Babina et al., 2024). Hampole et al. (2025) are an exception – they find that firm hiring of AI "integrators" is positively correlated with subsequent productivity growth and the generation of AI patents. New work by Babina et al. (2026) also suggests an emerging impact of AI investment on productivity in the most recent years, contingent on appropriate investment in organizational capital.

Prior research shows that recruiting star scientists can reshape organizations through knowledge transfer and peer effects (Zucker et al., 1998b; Agrawal et al., 2017). This paper extends the literature by examining how the movement of the most advanced students of elite AI scientists (whom we dub the "immortals") into industry impacts firm productivity and AI invention. 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 elite AI graduates diffuse into firms and document some stylized facts regarding their movement into the U.S. private sector workforce.

Our analysis generates three main findings. First, the hiring of elite AI graduates is associated with positive, statistically significant labor productivity gains. Second, 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. These results are robust to the inclusion of stocks of AI workers at various levels of education, skill, and responsibility. Third, we present suggestive evidence that elite AI graduates are especially important in enabling their employers to create AI inventions that build on the

specific scientific advances the elite AI graduates researched under the supervision of their mentors.

The remainder of the paper is organized as follows. We provide a brief summary of the prior literature upon which we most directly build. We then describe our data construction process. We cluster AI talent into different groups and describe their characteristics. We then present our empirical approach and results. We conclude with a discussion of limitations and plans for future research.

2 Literature Review and Conceptual Framework

2.1 The Role of Star Scientists in Transformative Innovation

A series of seminal contributions to the history and sociology of science have shown that a small group of "star scientists" account for a disproportionate share of the scientific output of entire fields (de Solla Price, 1963). These stars can accumulate prestige and resources, sometimes at the expense of equally competent peers and rivals, that amplify their scientific output, influence, and impact (Merton, 1968). They also form networks of prominent scholars that shape the research agendas of their fields, promote designated proteges, and define what constitutes legitimate inquiry within their domains (Crane, 1972). Economists Pierre Azoulay and Joshua Graff Zivin provided critical quantitative evidence for these and related ideas through careful studies of the effects of the unexpected deaths of star scientists (Azoulay et al., 2010, 2019). Fabian Waldinger and his coauthors further documented the consequences of the Nazi-era dismissal and emigration of Jewish scientists for Germany's research universities, which had previously occupied a central position in global science, and for the United States, which became the eventual destination of many scientific experts initially trained in Germany (Waldinger, 2010; Moser et al., 2014).

A related literature examines the role of star scientists in transferring frontier scientific knowledge to industrial R&D. In a series of papers that directly inspired the current study, Lynne Zucker, Michael Darby, and their coauthors argue that cutting-edge scientific knowledge is often incompletely codified and therefore not readily accessible through publications alone. Much of this knowledge remains tacit and embodied in human capital, so firms seeking to build on it often require direct collaboration with the discovering scientists or researchers trained in their laboratories (Zucker and Darby, 1996; Zucker et al., 1998a, 2002). Studying the emergence of the U.S. biotechnology industry, they show that the geographic distribution of star scientists predicted the location and formation of biotechnology firms, that localized effects reflected identifiable relationships between firms and university scientists rather

than generalized knowledge spillovers, and that collaboration with star and other leading university scientists was strongly associated with firms' patenting, product development, employment, and other measures of commercial success (Zucker et al., 1998b,a, 2002). Related work by Agrawal (2001), Agrawal and Henderson (2002), and Agrawal (2006) extends this perspective by distinguishing among the channels through which university knowledge reaches firms, demonstrating that patents capture only a limited portion of university knowledge transfer, and showing that inventor engagement can be important for commercializing inventions whose value depends on latent knowledge.

Later work suggests that star scientists may generate benefits beyond the transfer and absorption of tacit knowledge. Although their setting is university departments rather than firms, Agrawal et al. (2017) find that the principal organizational benefit of recruiting a star scientist arises through improved recruitment: the average quality of subsequent hires increases, while the aggregate productivity of incumbent scientists does not. They further show that this aggregate result masks positive effects on incumbents working in areas related to the star and negative effects on those working in unrelated areas. This research also identifies qualifications to the benefits that stars and other expert hires bring to their organizations. Kehoe and Tzabbar (2015) show that star scientists increase firm innovative productivity but can constrain the emergence of other innovative leaders, with these effects depending on the breadth of the star's expertise and the extent of collaboration. Singh and Agrawal (2011) find that hiring firms substantially increase their use of recruits' prior inventions, but much of this effect reflects recruits' continued exploitation of their own ideas, while diffusion to others remains concentrated within their immediate collaborative networks. Finally, Toole and Czarnitzki (2009) show that the value of academic scientists' human capital depends on the tasks to which it is applied: scientifically oriented human capital contributes more strongly to research and proof-of-concept activities, whereas commercially oriented human capital contributes more strongly to invention and patenting.

Prior work has also investigated the factors shaping Ph.D.-level scientists' choices between academic and private-sector careers. Stern (2004) shows that scientists are willing to accept lower wages for positions that allow them to pursue and publish an individual research agenda, consistent with a compensating differential for participation in science. Roach and Sauermann (2010) find that Ph.D. students who prefer industrial research careers tend to exhibit a weaker "taste for science," place greater weight on salary and access to resources, and express stronger interest in downstream research activities than students who prefer academic careers. Building on this evidence, Sauermann and Roach (2012) show that the attractiveness of academic research careers declines over the course of doctoral training, while nonacademic careers become relatively more attractive.

We build on this prior literature by extending its central ideas into an important new context: the emergence and diffusion of AI, a technology widely considered to be a general purpose technology. This paper focuses in particular on the role played by the doctoral advisees of the top AI scientists in raising the ability of the firms that hire them to realize productivity gains, generate AI-related inventions, and create innovations that build on recent science. We generate a new, comprehensive data set that identifies the world's top AI scientists and links them to their doctoral student and post-doctoral advisees. Using data from Revelio, we then track these students to the firms that employ them. In doing so, we extend the earlier literature on star scientists by leveraging the massive expansion of granular data on the movement of knowledge workers within and across firms enabled by providers like Revelio. Prior research generally had to infer the movement of star scientists and their mentees using only data on patents or publications. Our data allow us to track elite AI graduates' movements and their impact even in contexts where neither patents nor publications are generated. Finally, we extend the literature by moving our analysis beyond publicly traded firms. The current version of the paper uses micro data from Compustat to document the impact elite AI graduates have on the publicly traded firms that hire them. Ongoing research undertaken in collaboration with the Census Bureau has also connected these flows of experts to the Census LBD database. Unfortunately, due to Census disclosure requirements and procedures, those results cannot yet be publicly disclosed. Section 3 describes our data construction exercise. Sections 4-7 present our empirical results. In our concluding section, we place the results documented in this paper in the context of our broader research aims.

2.2 How Firms Benefit By Hiring Elite AI Graduates

First, however, we seek to clearly delineate the potential mechanisms through which the students and advisees of star scientists may benefit the firms that hire them, using a typology we call the AI Talent Diffusion Taxonomy (ATDT). Building closely on the prior research noted in the previous section, this taxonomy highlights academic lineage as a central conduit through which frontier AI knowledge is transferred, embedded, and scaled within the hiring firm. Specifically, the taxonomy highlights three interrelated but conceptually distinct mechanisms – knowledge transfer, capability augmentation, and network spillovers – that explain how the academic lineage of AI talent may shape firm-level productivity, innovation, and strategic advantage.

2.2.1 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 the advanced ideas they acquired through their training into organizational processes, patents, products, and strategies. 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.

2.2.2 Capability Augmentation

The second mechanism, capability augmentation, explains how elite AI graduates strengthen firms' internal technological and organizational capacities. A critical channel by which this happens that has been identified in the prior literature centers on the ability of the hired expert to attract additional high-quality experts in her subspecialty (Agrawal et al., 2017). The presence of the first high-level expert (or the first few experts) provides subsequent hires with a community with whom they can interact and may signal to future hires the priority placed by the hiring firm on high-level scientific expertise. Practitioners and elite scientists have used the word "legitimacy" to describe this phenomenon in conversations with us. By hiring a top expert, the firm earns "legitimacy" in the eyes of subsequent hires. The additional experts bring their own tacit knowledge and frontier skills, broadening the ability of the firm to learn from and build on recent science. 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.

The ability to recruit more expertise conferred on the firm by the hiring of a mentee of a star scientist may not be limited solely to individuals who are also graduates of elite doctoral programs. Experienced practitioners and recent college graduates may also be attracted to firms with visible investments in frontier scientific expertise. The presence of elite AI graduates might attract a higher quality of practitioner or skilled researcher to the hiring firm, and this could be an additional source of benefit. The attraction is rational if firm staff could learn frontier ideas and techniques from the elite AI graduates without

having to copmplete an AI Ph.D. under a star scientist. The presence of an elite AI graduate thus augments the capabilities of the hiring firm by attracting better workers in the lower strata of the human capital pyramid and training those workers in new ideas and techniques.

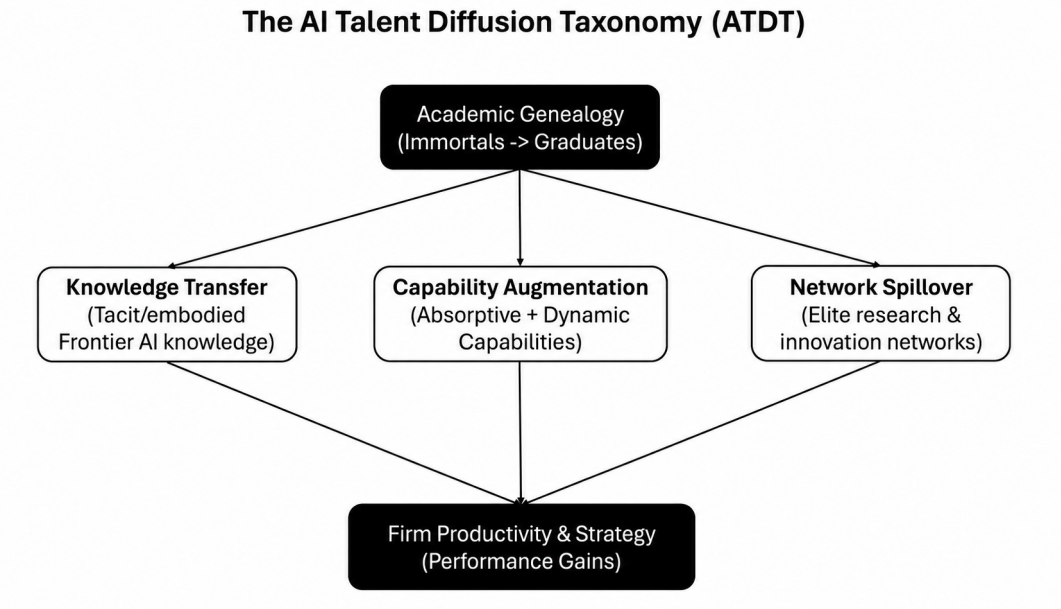


Figure 1: The AI Talent Diffusion Taxonomy (ATDT). Academic lineage channels frontier AI knowledge from elite scientists ("immortals") to their graduates, shaping firm productivity through three mechanisms: knowledge transfer, capability augmentation, and network spillovers.

2.2.3 Network Spillovers

The third mechanism, network spillovers, captures the relational 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 our 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.

3 Data

3.1 Data on Elite AI Graduates

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. The ten core AI subdomains are as follows:

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.

Illustrative examples of the individuals identified as "Immortals" include the following. Tom Mitchell of Carnegie Mellon University wrote the definitive textbook for machine learning in 1997, founded the first academic machine learning department (at CMU) in 2006, conducted pioneering research in self-supervised AI and cognitive neuroscience AI, and has supervised over 40 Ph.D. students (not all of whom we are able to track in our data), in addition to large numbers of masters students and post-docs. Geoff Hinton of the University of Toronto, often called the "Godfather of Deep Learning," is a co-recipient of the 2018 Turing Award and the 2024 Nobel Prize, popularized backpropagation algorithms for training multi-layer neural networks, and co-invented Boltzmann machines and AlexNet. His dozens of doctoral advisees include such luminaries as Ilya Sutskever (co-founder of OpenAI) and Yann LeCun, who is an immortal scholar in his own right. Andrew Ng of Stanford co-founded Google Brain, is a pioneer in large-scale unsupervised deep learning, and advised many students that have gone on to build self-driving car companies, top AI labs, and major academic departments. Fei-Fei Li is also a Stanford Professor, known as the "Godmother of AI," creator of ImageNet, a major contributer to computer vision, and the adviser of large numbers of students, including Andrej Karpathy and Olga Russakovsky. Our list also includes somewhat younger rising stars such as Graham Neubig of CMU, who has made major contributions to machine translation, code generation, and agentic AI. Given his younger age, his number of advisees is smaller, but growing and quite influential.

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, ProQuest like data sources 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 importance 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.

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, although we will not be utilizing these data in the current paper.

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	630	360	202	29	16	167	4	87	499	512
# 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	5,569	5,629	6,447	857	302	2,755	64	2,837	6,288	6,278
# of direct Immortal–corporate collaborative papers	8,553	16,026	4,582	665	987	6,859	30	4,301	12,047	7,112

Notes: AI subdomains are not mutually exclusive, and an Immortal may be assigned to multiple research domains. Consequently, counts in this table across subdomains do not sum to the total number of unique Immortals.

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 explore temporal and industry-specific heterogeneity in the returns to AI. Specifically, we examine whether hiring elite AI graduates is correlated with subsequent measurable productivity improvements, and whether these gains are broadly distributed or concentrated among a subset of firms. By pairing date on the employment of elite AI graduates with financial outcomes, the matched dataset provides the empirical foundation for testing whether AI is delivering on its promise of productivity-enhancing innovation.

3.2.1 Characteristics of Firms Hiring Elite AI graduates

We operationalize exposure to frontier AI talent as the number of graduates trained under elite AI pioneers ("immortals") employed by a firm, normalized by 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 elite AI graduate firms as those that have hired at least one graduate from an immortal lineage between 1990 and 2024.

Descriptive patterns reveal that these 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 elite AI graduates 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 with Census Bureau economists will seek to match our elite talent to a much broader set of firms, including these privately held start-ups.

In the Compustat sample, firms hiring elite AI graduates 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 Elite AI graduates Versus General AI Talent

While we hypothesize that elite AI graduates, as we have defined them, may represent an important resource for firms seeking to invent new AI technology or use frontier AI to fundamentally re-engineer their exiting products and services, we fully realize that elite AI graduates are not the only human capital resource that matters for these firms. Available evidence points to the importance of workers with varying degrees and kinds of AI-related human capital. For example, Babina et al. (2024) show that a firm's hiring of workers with a diverse range of AI-related skills leads to output gains and the generation of AI patents.

To us, this suggests the need for a "pyramid" of AI-related human capital within firms. For firms merely seeking to apply existing AI technologies to the incremental improvement of existing products, services, or procedures, a limited amount of specialized training may be sufficient. We thus imagine a "base level" of AI talent that does not even necessarily require an undergraduate degree in computer or data science. As the firm's AI-related technological ambitions grow, higher levels of skill (workers with bachelors or masters degrees or extensive

industry experience) may be required. For the firms with the most challenging goals, access to frontier scientific knowledge obtained through the hiring of a worker with elite doctoral-level training may be necessary or, at least, extremely useful. These workers represent the apex or part of the apex of the pyramid of human capital in such firms. However, even these highly talented workers cannot generate the desired outcomes alone. They need to work with the rest of the pyramid.

Since elite AI graduates, as we have defined them, constitute a relatively small part of the pyramid, and they are unlikely to exist in firms where the base of such a pyramid is entirely absent, it is important to control directly for the presence other workers if we hope to elicit even suggestive evidence that the mentees of top AI scientists are important at all, or instead represent a kind of unnecessary "trophy hire" obtained by well-funded incumbents or rich start-ups.

Before we turn to our methods for measuring the broader AI workforce within firms, we note the related work of Hampole et al. (2025), which was brought to our attention long after we began working on this project. This paper defines "AI exposure" within a firm by identifying the emergence of "AI integrators" in the firm's workforce. This is related to our notion of elite AI graduates, and we suspect there is some overlap, in that some of their AI integrators have our elite AI background. However, this will not always be the case. Our notion is defined on the basis of training and background. Their notion is defined on the basis of the roles different AI workers play within the firm. It would be of interest to study both kinds of elite AI workers, and document where these concepts overlap in the same individual, where they do not, and what difference it makes to the hiring firms. We are currently pursuing this line of research.

We construct these more broadly based measures of AI workers by combining two complementary sources of worker-level information in Revelio: standardized occupational positions and standardized skills. The position-based measure captures whether an individual held a role whose occupational content is directly or adjacently related to artificial intelligence. We use Revelio's standardized role taxonomy, specifically role k17000 v3, and manually classify AI-related roles into ordered intensity levels. Higher levels capture core AI research and engineering roles, such as machine learning scientist, AI researcher, NLP engineer, and computer vision engineer. Intermediate levels capture technical AI systems, applied machine learning, robotics, data science, and AI-adjacent analytical roles. Lower levels capture weaker but still AI-relevant roles, such as AI-facing product, support, training, or business roles.

We separately construct a skill-based measure using Revelio's standardized worker skills. We classify skill k35000 values into AI-related categories and assign each skill an intensity score. Core AI skills include machine learning, deep learning, neural networks, natural lan-

guage processing, computer vision, reinforcement learning, and related frontier methods. Broader AI skills include AI frameworks, model deployment, MLOps, ML infrastructure, cloud-based machine learning platforms, robotics and autonomy tools, and advanced statistical or data-science methods that support AI work. This skill taxonomy allows us to identify workers whose job titles may not explicitly indicate AI activity but whose reported or inferred technical skills reveal AI-related capabilities.

We then combine the position and skill measures to create alternative definitions of AI workers. In the stricter definitions, a worker must satisfy both conditions in a given firm-year: they must hold an AI-related position and have at least one AI-related skill observed by that year. For example, the most stringent definition requires a level-5 AI role and at least one level-5 AI skill. Broader definitions relax these thresholds to include level-4 and level-5 roles and skills, or level-3 through level-5 roles and skills. These intersection-based measures are designed to identify workers whose occupational assignment and technical capability both indicate AI relevance.

Finally, we construct a more lenient definition that captures any worker with an AI-related role and an AI-related skill under the broadest thresholds. For each worker-firm-year, we expand employment spells into annual observations from 1990 onward and retain observations where the worker is employed at a Compustat-linked Revelio firm and satisfies the relevant AI-role and AI-skill criteria. Skill information is applied in a time-safe manner: a skill is counted for a worker-year only if it was first reported by the end of that year, with untimestamped skills retained when Revelio does not provide a first-reported date. The resulting employee-firm-year panel provides a scalable measure of broad-based AI human capital across U.S. public firms.

3.4 AI Patent 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 clas-

sification 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 Is Science Important in AI Invention? Whose Science?

Our research presumes that science, including recent science, is an important input into industrial invention. A long literature has used patent citations to scientific papers to trace the connection between science and invention (Narin et al., 1997; Jaffe and de Rassenfosse, 2017; Roach and Cohen, 2013; Gaetani and Li Bergolis, 2015; Li et al., 2017; Fleming et al., 2019; Arora et al., 2021). A full summary of this large and still-active literature lies beyond the scope of the current paper. Prior work has revealed limitations and potential pitfalls in the use of these data, but our reading of the literature suggests that patent citations to science, when used with appropriate care, can provide useful evidence on the role of scientific knowledge in industrial R&D (Jaffe and de Rassenfosse, 2017; Roach and Cohen, 2013). Much of the research prior to 2020 relied on proprietary data sets that could not be shared publicly. Marx and Fuegi (2020) made a significant contribution to this literature by generating and sharing with the research community a public, worldwide database linking patent citations to scientific articles. Recent releases of these data include in-text as well

as front-page citations, allowing researchers to measure patent reliance on science more comprehensively.

If science – including but not limited to that undertaken by leading AI academic scientists with their graduate students – is important to AI invention, then we might expect this to be reflected in large numbers of citations made by our AI-related patents to scientific articles. We matched our AI patents to the Reliance on Science database maintained by Matt Marx at Cornell University. In some cases, we supplemented the paper references contained in these data to more accurately connect patents to the specific papers they cite. Figure 2 shows trends over time in the aggregate number of AI patents citing scientific papers. Figure 3 shows the average number of citations to science made per patent, which rises in the early years of our data sample but appears to level off in the 2000s. We interpret this as suggesting that, on average, AI patents do make significant numbers of citations to scientific papers, and the majority of AI patents make at least one.

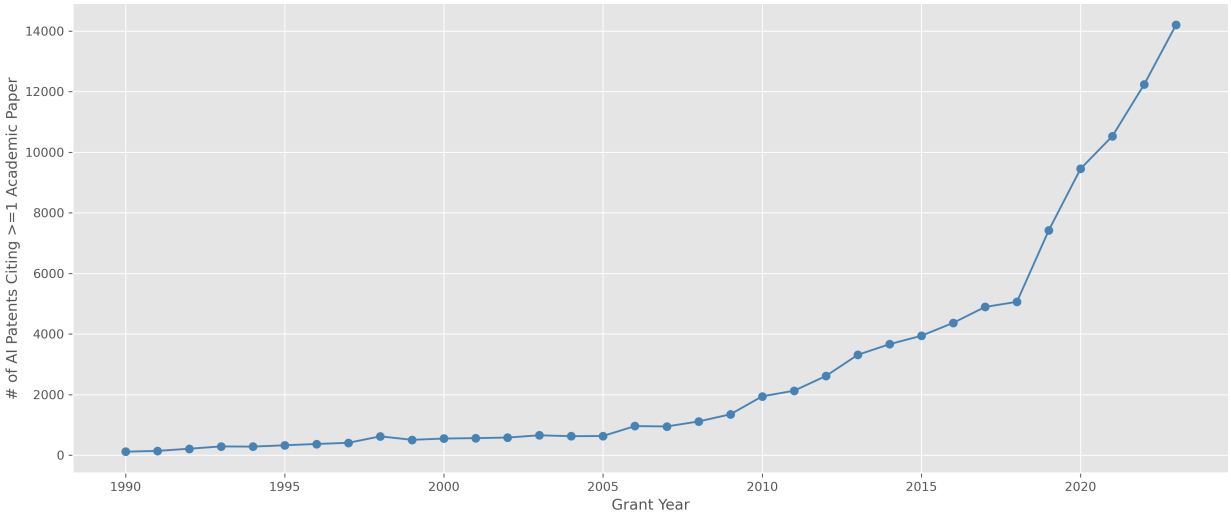


Figure 2: Aggregate Number of AI Patents Citing Scientific Papers by Year

Notes: The figure plots the annual aggregate number of AI patents that cite scientific papers. The counts are aggregated by patent grant year.

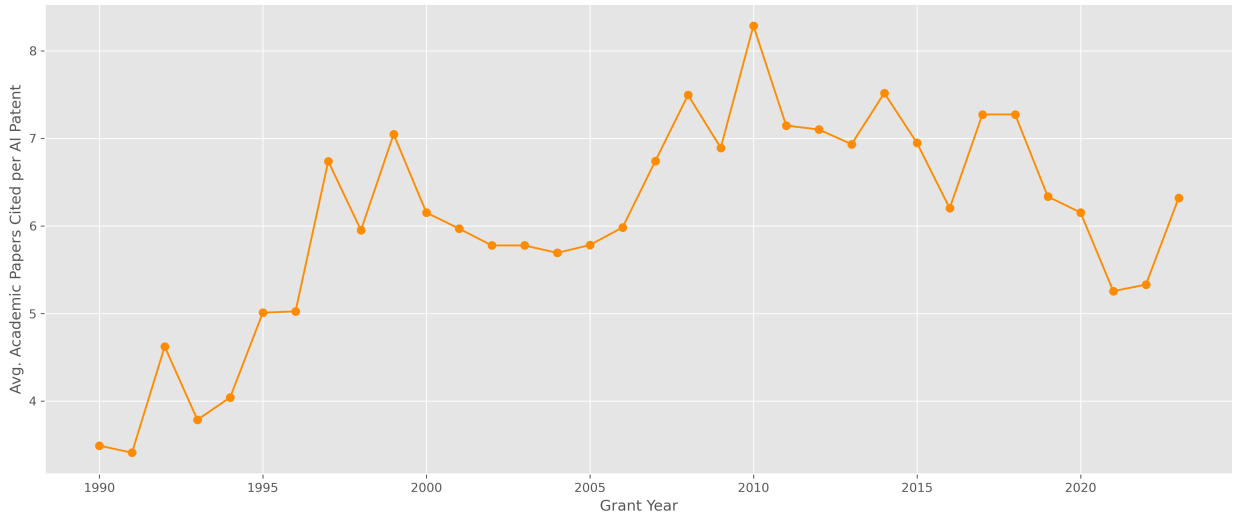


Figure 3: Average Number of Scientific Paper Citations per AI Patent by Grant Year

Notes: The figure plots the average number of citations to scientific papers made by each AI patent, aggregated by patent grant year.

Data permit us to link the papers cited by the patents to the institutional affiliations of their authors. Table 2 shows the most influential institutions in terms of the numbers of times AI patents cite their papers. The list appears reasonable in terms of the identity of institutions, but it also illustrates something important about the structure of computer science in AI-related fields: large tech companies like Microsoft and Google are extremely important generators of highly-cited science, and the patent citations to their papers are certainly not driven solely by self-citations. This may be indicative of the degree to which the boundaries between academic science and industrial R&D have blurred in the domain of AI. However, academic institutions still play a disproportionately important role in training new AI talent up to the technological frontier. Given the degree to which demand in industry for individuals with these skills appears to run well beyond supply, their movement into industry may still influence where and how AI science is translated into AI invention.

Table 2: Top 20 Institutions by Number of AI Paper Citations from AI Patents

Rank	Institution	Citations
1	Carnegie Mellon University	26,347
2	Microsoft	25,379
3	Stanford University	25,174
4	IBM	21,405
5	Massachusetts Institute of Technology	21,207
6	Google	17,699
7	University of California, Berkeley	17,594
8	University of Toronto	11,926
9	AT&T	11,236
10	University of Washington	11,168
11	SRI International	9,865
12	Cornell University	9,838
13	University of Illinois Urbana–Champaign	8,735
14	Columbia University	7,996
15	Harvard University	7,933
16	University of Maryland, College Park	7,828
17	University of Michigan–Ann Arbor	7,584
18	Georgia Institute of Technology	7,513
19	University of Southern California	7,350
20	University of Oxford	6,870

Notes: Self-citations are excluded. A self-citation is defined as a citation where the patent assignee organization and the cited paper's affiliated institution are matched to the same normalized company or institution.

5 Mobility and clustering of elite AI graduates

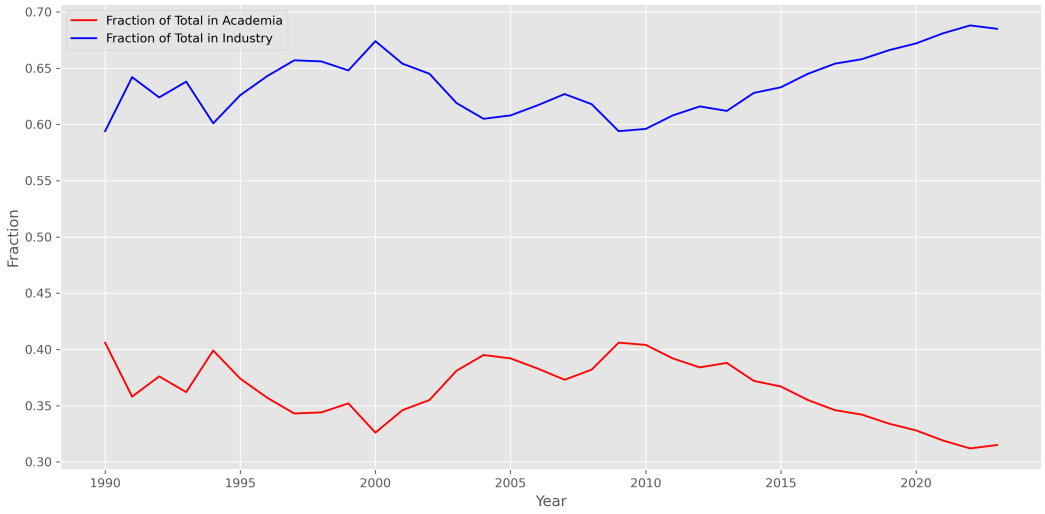
To understand the evolution of employment trajectories among AI researchers, we analyze the annual data from 1975 to 2024 on individuals who have held at least one job post-graduation. This section highlights longitudinal trends in employment distribution across sectors, with a particular focus on Academia (NAICS code 61), Industry (all the other NAICS codes), and Big Tech Firms 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 in AI 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 increasing number and size of computer science departments and the expansion of research funding during this period. This trend began to reverse in the early 2000s, as industry—particularly the technology sector—increasingly attracted 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 significantly, signaling a renewed shift away from academia.

Figure 4: Academic vs. Industry Employment Over Time

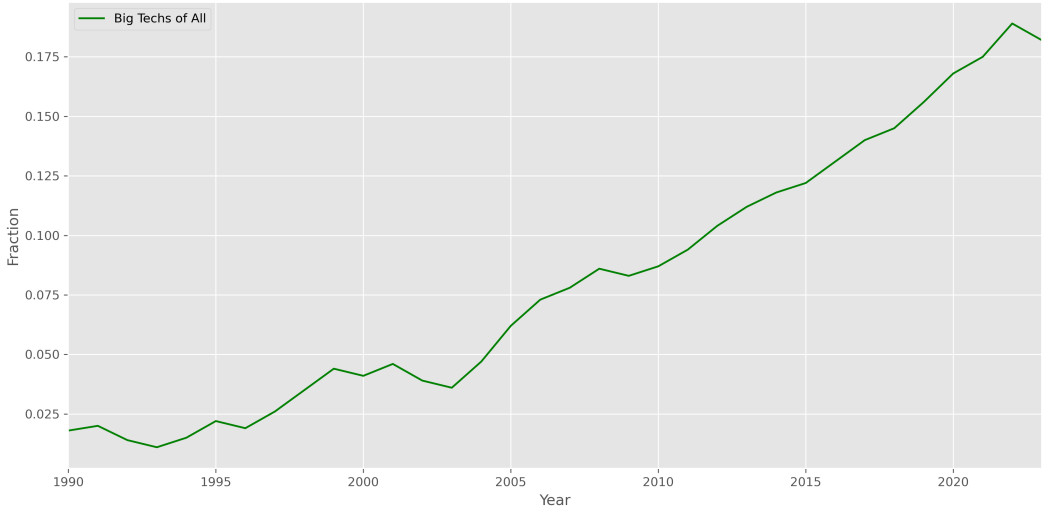


Notes: This figure chart tracks the share of AI graduates employed in academia (red) and industry (blue) from 1990 to 2024. It reveals a long-term shift: while academia's share initially grew, it peaked around 2010 and has declined since, whereas industry employment has steadily risen, reflecting increasing absorption of AI talent by the private sector.

Within industry, the so-called Big Tech Firms, have become increasingly important. In the year 2010, only about 10% of elite AI industry workers were employed in these firms, but this share rose to 20.9% by 2023. That being said, Big Tech firms still collectively account

for only about one fifth of the elite AI workers hired, including those going into academia. Put another way, 80% of elite AI graduates moves to organizations other than Big Tech, suggesting their share, while large, is neither dominant nor necessarily restrictive.

Figure 5: Rise of Big Tech Employment



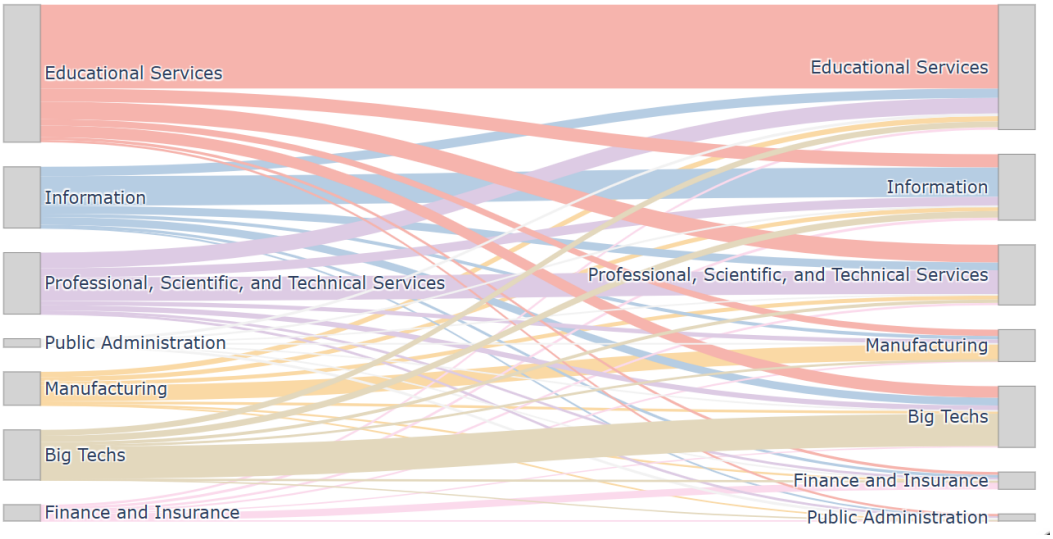
Notes: This figure charts the growing share of AI professionals working in Big Tech Firms from 1990 to 2024. It highlights a sharp and steady increase starting around 2005 and peaking above 20% in 2023. That being said, more than 80% of elite AI workers join organizations other than the Big Tech companies.

Talent mobility across industries shows us the patterns of retention and relocation in different sectors. As we can see in figure 6, 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 may indicate 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 possibly 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 may signal 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 interaction with other labor markets.

Figure 6: 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.

6 Elite AI graduates and Firm Productivity

To evaluate the impact of elite AI graduates on firm-level productivity, we take a production function approach with controls for traditional inputs, R&D, elite AI graduates (as we have defined it), and various alternative measures of the AI labor force within firms. This section outlines the empirical framework, measurement of variables, and identification strategies.

6.1 Approach

We begin by using the microdata in Compustat to estimate a Cobb-Douglas production function. We utilize the Wooldridge (2009) approach to estimation of production function parameters. This approach utilizes expenditure on materials to correct for the potential endogeneity of conventional inputs. We then take the TFP residuals derived from that regression and use them as the dependent variable of the regression reported in Table 3.

To control for firm heterogeneity and time trends, we include fixed and year effects. In addition, we include measures of the firm's stock of R&D expenditures, elite AI graduates, and general AI talent. We interpret the coefficients as indicative of the contribution of within-firm changes in these variables to firm-level TFP, but we stop short of claiming these conditional correlations have a wholly causal interpretation.

We also seek to estimate the impact of elite AI graduates on firm-level labor productivity. We take the traditional approach of regressing the log of sales per employee on capital per employee, the R&D expenditure stock per employee, firm and year effects, and our measures of elite and general AI talent. Following a line of reasoning that dates back at least to Griliches (1981), we interpret the coefficients on our measures of AI-related human capital as reflecting the excess returns associated with this input over and above the contribution of these workers to employment and the likely inclusion of at least some of their salaries in R&D. Our key variables are defined as follows:

- Y_{it}: Firm output, measured as gross deflated sales.
- 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.
- $EliteAIgraduates_{it}$: Elite talent, measured as the cumulative stock of graduates of elite AI scientists' lineages employed by firms.
- $GeneralAITalent_{it}$: General AI talent, defined in various ways.
- RnD: Firm-level R&D expenditure stock.
- μ_i, λ_t, η_s: Firm, time, and industry level (2-digit NAICS) of fixed effects.
- ϵ_{it}: Error term.

While these panel regressions cannot fully 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 and instrumental variable approaches that could further strengthen our inference regarding causal effects. At this point, we stop short of making strong causal claims.

Even if AI is a general-purpose technology, its productivity impact may 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, if any, from frontier AI talent vary by sector.

6.2 Event Study Design

Elite AI workers may contribute to firm productivity in ways that have nothing to with patented inventions. However, it would also be interesting to examine the impact of these workers within a firm on the propensity of that firm to generate AI-related patents. This would be an especially interesting result if the generation of AI patents is also related to subsequent increases in firm productivity.

To strengthen the causal interpretation of this relationship, 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:

$$Y_{it} = \sum_{k=-5, k\neq 0}^{+5} \beta_k D_{it}^k + \alpha \log(K_{it}) + \gamma \log(L_{it}) + \delta \log(XRD_{it}) + \mu_i + \lambda_t + \eta_s + \epsilon_{it}. \quad (1)$$

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.

6.3 Results

6.3.1 Productivity Effects of Elite AI Graduates

We next turn to estimation of the potential impact of elite AI graduates on firm productivity. Estimation results for TFP are provided in Table 3.

Across specifications, we see limited evidence of a statistical association between our measures of elite AI graduates and total factor productivity. Column 1 suggests a positive and statistically significant conditional correlation in the full firm sample, and column 2

suggests that it remains so, even after the inclusion of a broader measure of AI talent within the firm. However, the rest of Table 3 indicates that this is not a robust result. In the full sample, it fades depending on the exact definition of general AI talent used, and the effects are even weaker in manufacturing or the information sector, even though we might expect stronger results in the latter context.

6.3.2 Labor Productivity Effects of Elite AI Graduates

We next examine the impact of elite AI graduates on labor productivity. Results are provided in Table 4. The elasticity of labor productivity with respect to the cumulative stock of elite AI graduates does not appear to be robustly disinguishable from zero in the full sample. We find stronger results in more technology-intensive areas such as information services and manufacturing. These results suggest that the stock of elite AI graduates is correlated with labor productivity, and the results are robust to the use of lagged measures of talent stocks, and the incorporation of a wide range of different definitions of general AI talent. The estimated elasticities are consistently higher in the information sector than in manufacturing, and these differences are statistically significant. In addition, in manufacturing, the coefficients associated with elite AI graduates are generally smaller than those associated with general AI talent, whereas the reverse is true in the information sector. Finally, as we vary the definition of general AI talent across specifications using data on manufacturing firms, the coefficients are uniformly statistically significant. This is not true in the information sector. In specifications that include our measure of elite AI graduates, the significance of general AI talent depends on the definition used. Less stringent definitions are generally statistically insignificant. Given the potential for elite AI graduates to be especially impactful in the information sector, this pattern of results seems plausible.

We emphasize that we view these results as suggestive, not conclusive evidence of a causal relationship. With that caveat in mind, we consider the magnitudes of the estimated coefficients on elite AI graduates. This has an elasticity interpretation, and suggests that a 100 percent increase in the stock of elite AI graduates yields a 3-4 percent increase in labor productivity in the information sector. That effect seems small, but there are firms whose stock of elite AI graduates has increased more than 50-fold. For the firms with the greatest investment in elite AI graduate, the productivity effects could be substantial.

6.3.3 Effects of Elite AI Graduates on AI Innovation

We now examine whether hiring elite AI graduates increases firms' AI inventive activities. Here we focus on the flow of AI patents as the dependent variable. We use application

Table 3: Elite AI Graduate Stock and TFP

Panel A. Baseline Results

	Full Sample			Manufacturing			Information		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Elite AI Graduate Stock	0.0960***	0.0582**		0.0739*	0.0339		0.0929*	0.0842	
	(0.0293)	(0.0292)		(0.0386)	(0.0368)		(0.0480)	(0.0565)	
General AI Talent Stock		0.0993***			0.103***			0.0198	
		(0.0310)			(0.0359)			(0.0517)	
Lagged Elite AI Graduate Stock			0.0459			0.0315			0.0676
			(0.0285)			(0.0371)			(0.0570)
Lagged General AI Talent Stock			0.111***			0.109***			0.0419
			(0.0346)			(0.0397)			(0.0584)
R&D Stock	0.492***	0.490***	0.455***	0.662***	0.660***	0.626***	0.566***	0.565***	0.525***
	(0.0203)	(0.0203)	(0.0202)	(0.0352)	(0.0351)	(0.0355)	(0.0509)	(0.0510)	(0.0531)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	65,742	65,742	59,713	26,918	26,918	24,745	9,637	9,637	8,608
R^2	0.962	0.962	0.967	0.967	0.967	0.971	0.951	0.951	0.956

Panel B. Full Sample: Alternative Definitions of General AI Talent

	Position	Any AI	Levels 3–5	Levels 4–5	Level 5
Elite AI Graduate Stock	0.0309	0.0391	0.0421	0.0571*	0.0582**
	(0.0290)	(0.0287)	(0.0289)	(0.0296)	(0.0292)
General AI Talent Stock	0.113***	0.0935***	0.0883***	0.0878***	0.0993***
	(0.0234)	(0.0208)	(0.0207)	(0.0285)	(0.0310)
R&D Stock	0.486***	0.487***	0.488***	0.490***	0.490***
	(0.0204)	(0.0203)	(0.0203)	(0.0203)	(0.0203)
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	65,742	65,742	65,742	65,742	65,742
R^2	0.963	0.963	0.963	0.962	0.962

Panel C. Manufacturing: Alternative Definitions of General AI Talent

	Position	Any AI	Levels 3–5	Levels 4–5	Level 5
Elite AI Graduate Stock	0.0116	0.0149	0.0147	0.0297	0.0339
	(0.0368)	(0.0358)	(0.0361)	(0.0374)	(0.0368)
General AI Talent Stock	0.103***	0.0918***	0.0908***	0.0939***	0.103***
	(0.0269)	(0.0250)	(0.0250)	(0.0335)	(0.0359)
R&D Stock	0.653***	0.655***	0.655***	0.659***	0.660***
	(0.0354)	(0.0352)	(0.0352)	(0.0352)	(0.0351)
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	26,918	26,918	26,918	26,918	26,918
R^2	0.967	0.967	0.967	0.967	0.967

Panel D. Information: Alternative Definitions of NonElite AI Talent

	Position	Any AI	Levels 3–5	Levels 4–5	Level 5
Elite AI Graduate Stock	0.0624	0.0835	0.0869	0.0828	0.0842
	(0.0594)	(0.0577)	(0.0578)	(0.0586)	(0.0565)
General AI Talent Stock	0.0572	0.0164	0.0103	0.0216	0.0198
	(0.0498)	(0.0426)	(0.0422)	(0.0528)	(0.0517)
R&D Stock	0.560***	0.565***	0.565***	0.565***	0.565***
	(0.0515)	(0.0508)	(0.0508)	(0.0509)	(0.0510)
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	9,637	9,637	9,637	9,637	9,637
R^2	0.951	0.951	0.951	0.951	0.951

Notes: The dependent variable is total factor productivity (TFP), estimated using the Wooldridge (2009) proxy-variable approach. Panel A reports the baseline specifications using the Level 5 definition of general AI talent. Panels B–D examine alternative definitions of general AI talent based on AI position, any AI employee, AI Levels 3–5, AI Levels 4–5, and AI Level 5. All regressions include firm and year fixed effects. Standard errors are clustered at the firm level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4: Elite AI Graduate Stock and Labor Productivity

Panel A. Baseline Results

	Full Sample			Manufacturing			Information		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Elite AI Graduate Stock	0.0100*	0.000655		0.0352***	0.0245***		0.0396***	0.0304**	
	(0.00532)	(0.00592)		(0.00667)	(0.00755)		(0.0114)	(0.0131)	
General AI Talent Stock		0.0256***			0.0275***			0.0202	
		(0.00707)			(0.00906)			(0.0141)	
Lagged Elite AI Graduate Stock			0.00216			0.0230***			0.0257**
			(0.00538)			(0.00694)			(0.0123)
Lagged General AI Talent Stock			0.0262***			0.0254***			0.0278**
			(0.00698)			(0.00899)			(0.0142)
Capital per Employee	0.234***	0.234***	0.249***	0.205***	0.204***	0.221***	0.215***	0.215***	0.246***
	(0.00374)	(0.00374)	(0.00363)	(0.00567)	(0.00567)	(0.00555)	(0.00881)	(0.00882)	(0.00877)
R&D Stock per Employee	0.0370***	0.0369***	0.0398***	0.0320***	0.0323***	0.0389***	0.0589***	0.0585***	0.0507***
	(0.00299)	(0.00299)	(0.00280)	(0.00474)	(0.00474)	(0.00447)	(0.00768)	(0.00768)	(0.00743)
Constant	-1.003***	-1.003***	-0.908***	-1.166***	-1.166***	-1.062***	-1.106***	-1.104***	-0.979***
	(0.0118)	(0.0118)	(0.0113)	(0.0177)	(0.0177)	(0.0172)	(0.0287)	(0.0287)	(0.0281)
Observations	65,913	65,913	59,795	26,936	26,936	24,755	9,668	9,668	8,632
R^2	0.758	0.758	0.793	0.741	0.741	0.771	0.704	0.704	0.735

Panel B. Full Sample: Alternative Definitions of General AI Talent

	Sale	Position	Levels 3–5	Levels 4–5	Level 5	Any
Elite AI Graduate Stock	0.0100*	-0.00395	-0.00100	-0.00338	0.000655	-0.00118
	(0.00532)	(0.00597)	(0.00595)	(0.00597)	(0.00592)	(0.00594)
General AI Talent Stock		0.0246***	0.0182***	0.0310***	0.0256***	0.0186***
		(0.00476)	(0.00441)	(0.00630)	(0.00707)	(0.00439)
Capital per Employee	0.234***	0.234***	0.234***	0.234***	0.234***	0.234***
	(0.00374)	(0.00374)	(0.00374)	(0.00374)	(0.00374)	(0.00374)
R&D Stock per Employee	0.0370***	0.0368***	0.0369***	0.0369***	0.0369***	0.0369***
	(0.00299)	(0.00299)	(0.00299)	(0.00299)	(0.00299)	(0.00299)
Constant	-1.003***	-1.007***	-1.005***	-1.003***	-1.003***	-1.005***
	(0.0118)	(0.0118)	(0.0118)	(0.0118)	(0.0118)	(0.0118)
Observations	65,913	65,913	65,913	65,913	65,913	65,913
R^2	0.758	0.759	0.758	0.759	0.758	0.758

Panel C. Manufacturing: Alternative Definitions of General AI Talent

	Sale	Position	Levels 3–5	Levels 4–5	Level 5	Any
Elite AI Graduate STock	0.0352***	0.0181**	0.0165**	0.0188**	0.0245***	0.0161**
	(0.00667)	(0.00763)	(0.00773)	(0.00767)	(0.00755)	(0.00770)
General AI Talent Stock		0.0280***	0.0283***	0.0346***	0.0275***	0.0294***
		(0.00605)	(0.00594)	(0.00803)	(0.00906)	(0.00593)
Capital per Employee	0.205***	0.205***	0.205***	0.204***	0.204***	0.205***
	(0.00567)	(0.00567)	(0.00567)	(0.00567)	(0.00567)	(0.00567)
R&D Stock per Employee	0.0320***	0.0317***	0.0319***	0.0323***	0.0323***	0.0319***
	(0.00474)	(0.00473)	(0.00473)	(0.00474)	(0.00474)	(0.00473)
Constant	-1.166***	-1.171***	-1.168***	-1.167***	-1.166***	-1.169***
	(0.0177)	(0.0177)	(0.0177)	(0.0177)	(0.0177)	(0.0177)
Observations	26,936	26,936	26,936	26,936	26,936	26,936
R^2	0.741	0.741	0.741	0.741	0.741	0.741

Panel D. Information: Alternative Definitions of General AI Talent

	Sale	Position	Levels 3–5	Levels 4–5	Level 5	Any
Elite AI Graduate Stock	0.0396***	0.0247*	0.0316**	0.0267**	0.0304**	0.0300**
	(0.0114)	(0.0131)	(0.0130)	(0.0131)	(0.0131)	(0.0129)
General AI Talent Stock		0.0254**	0.0129	0.0264**	0.0202	0.0154
		(0.0109)	(0.00998)	(0.0133)	(0.0141)	(0.00989)
Capital per Employee	0.215***	0.216***	0.215***	0.215***	0.215***	0.215***
	(0.00881)	(0.00883)	(0.00882)	(0.00882)	(0.00882)	(0.00882)
R&D Stock per Employee	0.0589***	0.0580***	0.0584***	0.0584***	0.0585***	0.0582***
	(0.00768)	(0.00768)	(0.00769)	(0.00768)	(0.00768)	(0.00769)
Constant	-1.106***	-1.113***	-1.108***	-1.105***	-1.104***	-1.109***
	(0.0287)	(0.0288)	(0.0287)	(0.0287)	(0.0287)	(0.0287)
Observations	9,668	9,668	9,668	9,668	9,668	9,668
R^2	0.704	0.704	0.704	0.704	0.704	0.704

Notes: The dependent variable is the log of labor productivity (sales per employee). Capital and R&D stock are measured on a per-employee basis too. Panel A reports baseline estimates for the full sample, manufacturing firms, and information firms. Panels B, C and D additionally control for alternative measures of general AI talent. All specifications include firm and year fixed effects. Standard errors are clustered at the firm level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

years to date the invention of AI patents assigned to our sample firms. The current results are derived from log linear models in which we first tranfrom the dependent variable by adding 1 to all observations, then taking the logs . Alternative specifications using a 5-year moving average of AI patents and an AI patent stock are reported as robustness checks in the Appendix A.

Table 5: AI Talent and AI Patent Output

	Full Sample		Manufacturing		Information	
	(1)	(2)	(3)	(4)	(5)	(6)
Active Elite AI Graduates	0.380***		0.434***		0.356***	
	(0.0348)		(0.0478)		(0.0764)	
General AI Talent	0.105***		0.109***		0.133*	
	(0.0308)		(0.0405)		(0.0723)	
Lagged Active Elite AI Graduates		0.367***		0.427***		0.325***
		(0.0349)		(0.0488)		(0.0760)
Lagged General AI Talent		0.0956***		0.0983**		0.125
		(0.0346)		(0.0456)		(0.0775)
R&D Expenditure	0.0212***	0.0238***	0.0130*	0.0168**	0.0580**	0.0631**
	(0.00505)	(0.00535)	(0.00732)	(0.00796)	(0.0235)	(0.0256)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	61,476	55,582	26,306	24,129	9,049	8,054
R^2	0.674	0.684	0.652	0.662	0.722	0.731

Notes: The dependent variable is the annual flow of AI patents generated by each firm, where patents are dated according to their application year. Active Elite AI Graduates measures the number of incumbent elite AI graduate researchers employed by a firm in a given year. General AI Talent measures the number of other AI-related workers employed by the firm. Columns (2), (4), and (6) replace the contemporaneous AI talent variables with their one-year lags. All specifications include firm and year fixed effects. Standard errors clustered at the firm level are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

The results of this regression suggest that stocks of elite AI graduates and broader measures of the AI talent stock are both predictive of AI patenting outcomes, even in a regression with extensive fixed effects and controls for R&D investment. The elasticities are much larger than in the labor productivity regressions; magnitudes exceed those in Table 4 by a factor of ten. In addition, the coefficients associated with elite AI graduates are consistently much larger than those associated with more general measures of AI talent. This is true in all industry samples.

Consistent with a talent-driven rise in sustained AI inventive activity, the next section shows that entry into AI patenting is followed by persistent gains in firm output, reinforcing

the causal progression from frontier AI talent to innovation to productivity.

6.3.4 Event-Study Evidence: Dynamic Productivity Effects of AI Invention

Building on the preceding evidence that employing frontier AI talent may significantly expand firms' AI inventive output, this section examines the downstream performance consequences of those AI inventions. 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 invention. Treated firms are matched to structurally similar non-adopters using propensity score matching, and placebo invention years are assigned to ensure aligned event time.

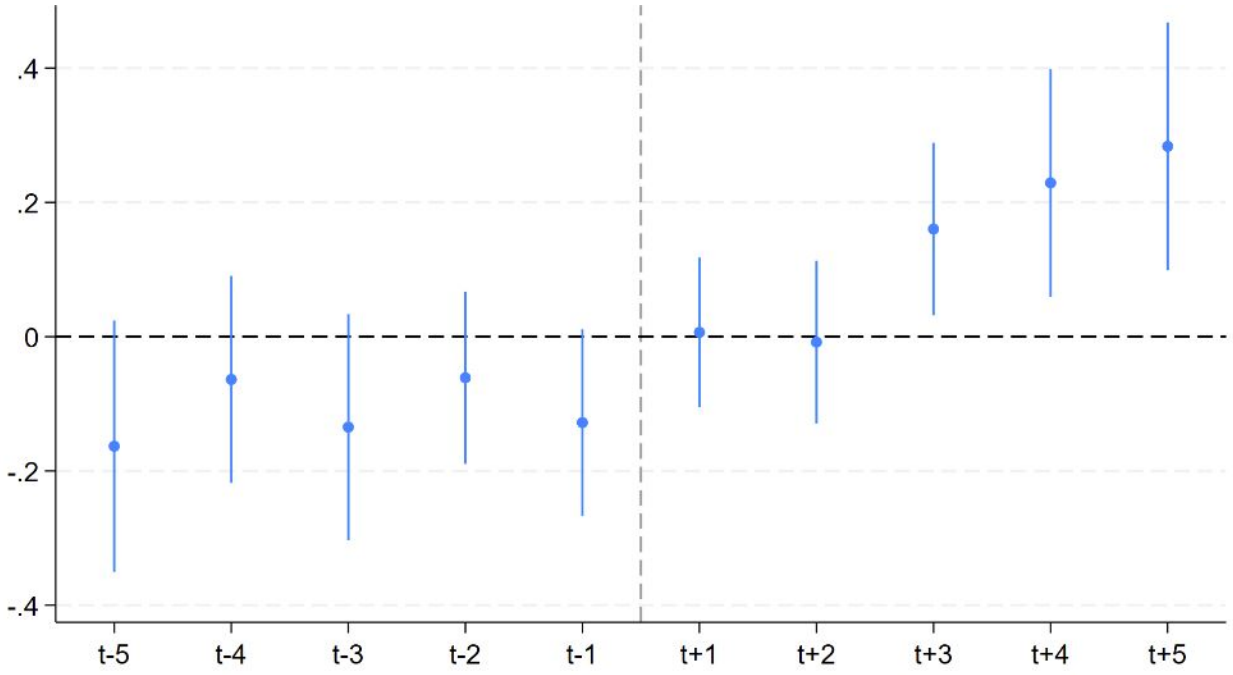


Figure 7: Dynamic effects of AI patenting on firm productivity

Notes: Coefficients and 95% confidence intervals estimated from event-time regressions with firm, and year 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 productivity relative to their matched peers.

7 Do Elite AI Graduates Help Bring Scientific Breakthroughs in AI Into Industrial Innovation?

As the previous section explicitly notes, elite AI graduates can play many roles with the firms that hire them. However, our paper has repeatedly stressed the potential role of elite AI graduates as a mechanism for transmitting frontier scientific knowledge to firms. We have already shown that AI patents frequently cite scientific papers. We used these patent citations to papers to identify a cluster of highly cited papers listed below in Table 6. This group of papers starts with Geoff Hinton's transformative "ImageNet" paper published in ACM in 2012, which he coauthored with Ilkya Sutskever. In this paper, Hinton and his coauthors demonstrated that deep convolutional neural networks could process and classify images as effectively as humans. The final paper in the cluster, the famous "Attention Is All You Need" paper published in 2017 in NEURIPS, introduces the transformer. The other 7 papers elaborate and develop critical ideas in the evolution of neural network approaches to AI, forming a kind of chain of discovery that links ImageNet to Attention is All You Need. Table 7 shows that these papers have received large numbers of citations from subsequent academic work as well as AI patents. This cluster necessarily omits many other vital academic discoveries that have undergirded the progress of the AI revolution. Nevertheless, we focus on this cluster of papers, because they appear to articulate and elaborate a series of related ideas that have proved especially important in both scientific advance and industrial R&D.

Table 6: Ten Seminal AI Papers

Year	Abbrev.	Paper	Authors
2012	AlexNet	ImageNet Classification with Deep Convolutional Neural Networks	Krizhevsky et al.
2014	Adam	Adam: A Method for Stochastic Optimization	Kingma & Ba
2014	VGG	Very Deep Convolutional Networks for Large-Scale Image Recognition	Simonyan & Zisserman
2014	GoogLeNet	Going Deeper with Convolutions	Szegedy et al.
2014	Seq2Seq	Sequence to Sequence Learning with Neural Networks	Sutskever et al.
2015	ResNet	Deep Residual Learning for Image Recognition	He et al.
2015	BatchNorm	Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift	Ioffe & Szegedy
2015	Faster R-CNN	Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks	Ren et al.
2015	U-Net	U-Net: Convolutional Networks for Biomedical Image Segmentation	Ronneberger et al.
2017	Transformer	Attention Is All You Need	Vaswani et al.

Table 7: Citation Statistics of the 10 Seminal AI Papers

Paper	Citations	Unique Immortals Citing	Share of Immortals Citing (%)
ResNet	230,673	1,300	58.2
Attention	180,324	1,464	65.6
Adam	167,172	1,602	71.7
AlexNet	129,476	1,130	50.6
VGG	112,010	965	43.2
U-Net	96,653	610	27.3
Faster R-CNN	73,117	724	32.4
GoogLeNet	47,401	662	29.6
BatchNorm	46,934	847	37.9
Seq2Seq	22,091	962	43.1
Total	1,105,851	–	–

If elite AI graduates constitute a human bridge through which these ideas flow from academia into industry, then we might expect to observe the following patterns in our data. First, we would expect that the ideas represented in the cluster of papers highlighted in Table 6 would be quickly picked up by leading academics (the Immmortal scholars) and incorporated into the research they are pursuing with their doctoral students. This is the case. We find that 83.9% percent of our leading academics cite these papers in their own

subsequently published work. Next, we would expect that companies would be much more likely to generate AI patents citing these papers after they hire advisees of leading academics who had couathored papers building on them (as evidenced by citations) during or shortly after their Ph.D. We pursue this line of reasoning through a series of exploratory regressions. We estimate a series of linear probability models where the dependent variable is the generation of at least one AI patent citing one or more of the highlighted papers. We find that this probability rises sharply with the number of elite academic children hired, but that the key coefficients level off after the fifth or the tenth hire, depending on the way we measure publication and citation. Based on this evidence, we set 5 or 10 as a rough estimate of the critical mass of elite children needed to connect to the streams of research represented by the seminal papers. We then estimate a linear probability panel in which we predict the generation of at least one AI patent citing them, controlling for firm and time effects, the log of R&D spending, and a dummy variable=1 if the focal firm has, by that year, hired at least the threshold number of elite AI graduates who cited the papers in their own graduate work.

Table 8 presents the results of that exploratory regression, using alternative definitions of graduate student exposure to the science embodied in the seminal papers. The coefficient on the "threshold" dummy variable is quite large, varying between .3 and .5. While none of this is conclusive nor can the results be given a strong causal interpretation, we nevertheless find it striking that such a strong conditional correlation exists between the presence in the firm of academic children exposed to a particular body of scientific knowledge and the generation by the hiring firm of patents which cite it. The inclusion of firm effects relaxes the concern that our results are driven solely by a concentration of hiring and patenting in a small number of productive and science-oriented firms. The inclusion of time effects relaxes the concern that the correlation is an artifact of an idea spreading simultaneously through industry and academia without the mediation of elite AI talent. In addition to providing circumstantial evidence in favor of our hypothesis, this example showcases the granular nature of our data and the degree to which we may be able to trace out the incorporation of particular AI concepts from scientific research to industrial development.

Table 8: Are Firms More Likely to Generate AI Patents Citing Seminal Papers If They Hire Elite AI Graduates Trained on Them?

Panel A. Direct AI Patent Output

	Full Sample		Manufacturing		Information	
	(1)	(2)	(3)	(4)	(5)	(6)
Coauthored AI Talent (≥ 5)	0.361***		0.428***		0.356***	
	(0.0853)		(0.152)		(0.115)	
Combined AI Talent (≥ 10)		0.400***		0.384***		0.401***
		(0.0885)		(0.107)		(0.131)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	45,793	45,793	17,985	17,985	7,347	7,347
R^2	0.334	0.334	0.311	0.308	0.424	0.427

Panel B. Combined AI Patent Output

	Full Sample		Manufacturing		Information	
	(7)	(8)	(9)	(10)	(11)	(12)
Coauthored AI Talent (≥ 5)	0.298***		0.237***		0.306***	
	(0.0654)		(0.0650)		(0.100)	
Combined AI Talent (≥ 10)		0.338***		0.183*		0.358***
		(0.0715)		(0.103)		(0.108)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	45,793	45,793	17,985	17,985	7,347	7,347
R^2	0.372	0.373	0.376	0.375	0.441	0.444

Notes: The dependent variable is an indicator for whether this company has an AI patent citing the top 10 papers. Panel A uses direct AI patent output, while Panel B uses combined AI patent output, which additionally includes second-degree citing patents. Coauthored AI Talent includes students who coauthored a Top-10-citing paper with their advisor, whereas Combined AI Talent additionally includes independent Top-10-citing students. All specifications include firm and year fixed effects and controls for the stock of elite AI graduates, the stock of general AI talent (Level 5), and log R&D stock. Robust standard errors clustered at the firm level are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

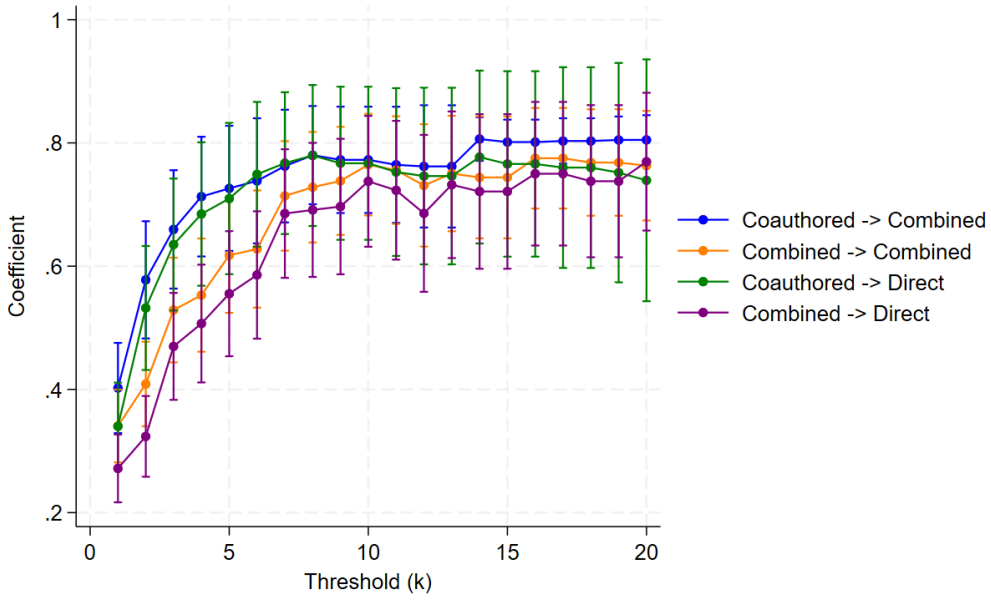


Figure 8: Threshold-response for hiring elite AI graduates.

Notes: The figure plots the estimated coefficients from linear probability model regressions in which the key explanatory variable is a threshold dummy equal to one if a firm hired at least k elite AI graduates, where $k = 1, \ldots, 20$. Panel A uses the stock of coauthored elite AI graduates, while Panel B uses the combined stock of coauthored and independently cited elite AI graduates. The dependent variables are indicators for whether the firm ever produced an AI patent citing one of the Ten Papers during 2012–2022. All regressions control for the firm's average log R&D stock over the sample period and include 2-digit NAICS industry fixed effects. Vertical bars denote 95% confidence intervals based on standard errors clustered at the firm level.

8 Conclusions

By linking firm performance data to measures of elite AI graduates, this study provides new evidence on the role these individuals may play in translating frontier AI science into AI-related innovation. We document four main results.

First, elite AI graduates appear to be an economically meaningful input, at least in some sectors. Manufacturing and information services firms that employ them exhibit higher labor productivity, and this remains true even after explicitly controlling for other measures of general AI talent.

Second, firms that employ more elite AI graduates generate significantly more AI patents. This also remains true even after explicitly controlling for other measures of general AI talent, and statistical association with AI patent output is considerably stronger in magnitude

than the association with productivity growth, suggesting that elite AI graduates play an especially important role in generating AI inventions in the firms that they hire.

Third, we show that the transition into AI patenting is itself associated with increased productivity growth, and an event study approach strengthens the causal interpretation of this statistical association.

Finally, we undertake a statistical case study that examines the role of elite AI graduates in enabling the firms that hire them to generate patents building on the specific knowledge the graduates acquired through their scientific training. This strengthens the evidence that these graduates actually function as a conduit through which frontier science is incorporated into new AI inventions.

8.1 Limitations and Directions for Future Research

That being said, our approach carries with it a number of limitations. Perhaps the most serious is the potential endogeneity of elite AI graduate employment. Our ability to obtain empirical leverage around this nettlesome issue is likely to grow as we shift our focus from Compustat data on publicly traded U.S. firms to the far more representative and expansive micro level data sets maintained by the U.S. Census Bureau, including the revenue-enhanced Longitudinal Business Database (LBD). The relatively small number of firms in the Compustat data base, their large size, and the limited engagement of many of these firms with frontier AI science almost certainly diminishes our ability to examine the impact of elite AI graduates in the contexts where their arrival is likely to make the most difference. In the LBD, we can study the impact of elite AI graduates on start-ups, privately held quantitative finance enterprises, and other firms that employ these graduates but are currently omitted from our data.

With these richer, more comprehensive data, it will become more possible to use event-study methods to analyze the impact of a transition into the initial hiring of elite AI graduates on productivity, patenting, and other outcomes of interest. It should also be more possible to leverage plausibly exogenous shocks impacting particular groups of firms, 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 us to disentangle selection from treatment effects more effectively, sharpening causal estimates of elite graduates' contributions to firm performance.

We can also sharpen our inference concerning the role of elite AI graduates in AI invention by delving more deeply into the data we already have. It is at least conceptually possible

for us identify the subset of elite AI graduates who are either directly named as inventors on AI patents or are clearly identified in the Revelio data as holding a position of research leadership in their firm. The disambiguation of inventor names is a well-known challenge in the literature, and this problem is compounded when we try to link names in Revelio to the names of the same individuals in patent documents, where they can be represented differently in ways that complicate linkage. However, better NLP techniques should allow us to make progress on this front. Related techniques could parse the text of the patents generated by the firms and assess the extent to which these patents reflect not just the expertise of the elite AI graduates but also the science that is evolving within the larger academic networks of which these graduates are a part. This could enable a quantitative assessment of the degree to which elite AI graduates are actually contributing to spillovers through networks, as we hypothesized in the description of the ATDF. More careful analysis of the textual data in Revelio describing job titles, roles, and skills could also help us flesh out a more nuanced and granular picture of how elite graduates interact with other skilled personnel within the firms that hire them in functions other than the generation of patents. In particular, it would be of interest to study the overlap between the "AI integrators" identified by Hampole et al. 2025 and our elite AI graduates – to what extent are these the same individuals – and the interactions between these groups when they are composed of different individual members.

Our current efforts are also constrained by the limited set of firm-level outcomes we currently consider. At the moment, we utilize only data on sales-based measures of productivity (TFP and labor productivity) and patents. However, some AI enterprises also publish detailed descriptions of their AI innovations. The massive compute, data, and human resources available to some of these companies have enabled them to rival and even eclipse top academic computer science departments in terms of the quantity and even quality of their academic output. The text of these papers, the innovations they describe, the citations that link them backward to the foundational work of elite AI scientists and forward to AI patents, and the authorship networks they trace out would be of great value to our efforts to better understand exactly how elite AI graduates contribute to AI innovation. Some elite academic experts frequently collaborate and coauthor with corporate lab, and these coauthored papers potentially constitute an alternative channel through which frontier science could impact and contribute to industrial innovation.

The rise of firms as generators of fundamental AI science also calls into question our focus on academia as the primary incubator of elite experts. Our current approach presumes that the way to achieve exposure to frontier knowledge is to apprentice oneself to a top AI scientist as her Ph.D. student or post-doc advisee. If firms are now taking the lead in some domains of "academic research," however, that suggests the alternative possibility that especially

talented individuals with advanced training acquired outside the circle of "immortal" experts could have a kind of post-doc experience in a firm that is working with frontier technologies, and this experience could convert these individuals to carriers of frontier knowledge in the same we presume elite AI graduates to be. This line of reasoning suggests the need to broaden our definition of elite AI graduates to include those whose formative scientific experiences occurred in firms rather than universities. Tracking where these individuals take and deploy their knowledge could substantially enrich our understanding of the role of specialized human expertise in the advance of AI invention.

Finally, we hypothesized that the initial hiring of elite AI experts might make it easier to hire subsequent members of their communities and attract a higher quality of general AI talent. In principle, these are testably hypotheses that our data allow us explore, giving us a means of assessing the degree to which elite AI experts are really delivering a capability augmentation effect and where in the firm, idea, and task space that effect is the strongest.

References

Agrawal, A. (2006). Engaging the inventor: Exploring licensing strategies for university inventions and the role of latent knowledge. *Strategic Management Journal*, 27(1):63–79.

Agrawal, A., Gans, J., and Goldfarb, A. (2019a). *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.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2019b). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2):31–50.

Agrawal, A. and Henderson, R. (2002). Putting patents in context: Exploring knowledge transfer from MIT.

Agrawal, A., McHale, J., and Oettl, A. (2017). How stars matter: Recruiting and peer effects in evolutionary biology. *Research Policy*, 46(4):853–867.

Agrawal, A. K. (2001). University-to-industry knowledge transfer: Literature review and unanswered questions. *International Journal of Management Reviews*, 3(4):285–302.

Agrawal, A. K., Brynjolfsson, E., and Korinek, A., editors (2025a). *The Economics of Transformative AI*. University of Chicago Press, Chicago.

Agrawal, A. K., Gans, J. S., and Goldfarb, A. (2025b). Genius on demand: The value of transformative artificial intelligence. In Agrawal, A. K., Brynjolfsson, E., and Korinek, A., editors, *The Economics of Transformative AI*, chapter 1. University of Chicago Press, Chicago.

Arora, A., Belenzon, S., and Sheer, L. (2021). Knowledge spillovers and corporate investment in scientific research. *American Economic Review*, 111(3):871–898.

Azoulay, P., Fons-Rosen, C., and Graff Zivin, J. S. (2019). Does science advance one funeral at a time? *American Economic Review*, 109(8):2889–2920.

Azoulay, P., Graff Zivin, J. S., and Wang, J. (2010). Superstar extinction. *The 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*, 151:103745.

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

Barrios, J. M., Fujiy, B. C., Lisowsky, P., and Minnis, M. (2025). Measurement matters: Financial reporting and productivity. NBER Working Paper 34536, National Bureau of Economic Research.

Baviskar, S. V., Alderucci, D., Branstetter, L. G., Goldschlag, N., Hovy, E., Runge, A., Tambe, P., and Zolas, N. (2024). Quantifying the impact of AI on productivity and labor demand: Evidence from U.S. census microdata. Paper presented at the NBER Summer Institute, Digital Economics and Artificial Intelligence.

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., 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., 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.

Crane, D. (1972). *Invisible Colleges: Diffusion of Knowledge in Scientific Communities*. University of Chicago Press, Chicago, IL.

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.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, 135(2):561–644.

de Solla Price, D. J. (1963). *Little Science, Big Science*. Columbia University Press, New York.

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

Fleming, L., Greene, H., Li, G.-C., Marx, M., and Yao, D. (2019). Government-funded research increasingly fuels innovation. *Science*, 364(6446):1139–1141.

Gaetani, R. and Li Bergolis, V. (2015). The economic effects of scientific shocks.

Hampole, M., Papanikolaou, D., Schmidt, L. D. W., and Seegmiller, B. (2025). Artificial intelligence and the labor market. NBER Working Paper 33509, National Bureau of Economic Research. Revised September 2025.

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

Jaffe, A. B. and de Rassenfosse, G. (2017). Patent citation data in social science research: Overview and best practices. *Journal of the Association for Information Science and Technology*, 68(6):1360–1374.

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

Kehoe, R. R. and Tzabbar, D. (2015). Lighting the way or stealing the shine? an examination of the duality in star scientists' effects on firm innovative performance. *Strategic Management Journal*, 36(5):709–727.

Li, D., Azoulay, P., and Sampat, B. N. (2017). The applied value of public investments in biomedical research. *Science*, 356(6333):78–81.

Marx, M. and Fuegi, A. (2020). Reliance on science: Worldwide front-page patent citations to scientific articles. *Strategic Management Journal*, 41(9):1572–1594.

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.

Merton, R. K. (1968). The matthew effect in science. *Science*, 159(3810):56–63.

Moser, P., Voena, A., and Waldinger, F. (2014). German jewish émigrés and US invention. *American Economic Review*, 104(10):3222–3255.

Narin, F., Hamilton, K. S., and Olivastro, D. (1997). The increasing linkage between u.s. technology and public science. *Research Policy*, 26(3):317–330.

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

Roach, M. and Cohen, W. M. (2013). Lens or prism? patent citations as a measure of knowledge flows from public research. *Management Science*, 59(2):504–525.

Roach, M. and Sauermann, H. (2010). A taste for science? PhD scientists' academic orientation and self-selection into research careers in industry. *Research Policy*, 39(3):422–434.

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

Sauermann, H. and Roach, M. (2012). Science PhD career preferences: Levels, changes, and advisor encouragement. *PLOS ONE*, 7(5):e36307.

Singh, J. and Agrawal, A. (2011). Recruiting for ideas: How firms exploit the prior inventions of new hires. *Management Science*, 57(1):129–150.

Stern, S. (2004). Do scientists pay to be scientists? *Management Science*, 50(6):835–853.

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

Toole, A. A. and Czarnitzki, D. (2009). Exploring the relationship between scientist human capital and firm performance: The case of biomedical academic entrepreneurs in the SBIR program. *Management Science*, 55(1):101–114.

Waldinger, F. (2010). Quality matters: The expulsion of professors and the consequences for PhD student outcomes in nazi germany. *Journal of Political Economy*, 118(4):787–831.

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. and Darby, M. R. (1996). Star scientists and institutional transformation: Patterns of invention and innovation in the formation of the biotechnology industry. *Proceedings of the National Academy of Sciences*, 93(23):12709–12716.

Zucker, L. G., Darby, M. R., and Armstrong, J. S. (1998a). Geographically localized knowledge: Spillovers or markets? *Economic Inquiry*, 36(1):65–86.

Zucker, L. G., Darby, M. R., and Armstrong, J. S. (2002). Commercializing knowledge: University science, knowledge capture, and firm performance in biotechnology. *Management Science*, 48(1):138–153.

Zucker, L. G., Darby, M. R., and Brewer, M. B. (1998b). 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

As a robustness check, we re-estimate the innovation regressions using the cumulative stock of AI patents. Results in Tables A1 and A2, indicating that frontier AI talent is associated with substantially greater AI patenting, whether we measure it in annual output or the stock.

Table A1: AI Talent and AI Patent Output

	Full Sample		Manufacturing		Information	
	(1)	(2)	(3)	(4)	(5)	(6)
Active Elite AI Talent	0.423***		0.495***		0.367***	
	(0.0356)		(0.0527)		(0.0705)	
Non-Elite AI Talent	0.330***		0.330***		0.370***	
	(0.0375)		(0.0540)		(0.0811)	
Lagged Active Elite AI Talent		0.421***		0.496***		0.360***
		(0.0355)		(0.0524)		(0.0694)
Lagged Non-Elite AI Talent		0.364***		0.354***		0.400***
		(0.0409)		(0.0595)		(0.0838)
R&D Expenditure	0.0115**	0.0131**	-0.00317	-0.00132	0.0398*	0.0459*
	(0.00548)	(0.00578)	(0.00866)	(0.00942)	(0.0236)	(0.0254)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	61,476	55,582	26,306	24,129	9,049	8,054
R^2	0.808	0.820	0.783	0.794	0.849	0.859

Notes: The dependent variable is the five-year moving average of AI patent output. Active Elite AI Talent measures the number of incumbent frontier AI researchers employed by a firm in a given year. Non-Elite AI Talent measures the number of other AI-related workers employed by the firm. Columns (2), (4), and (6) replace the contemporaneous AI talent variables with their one-year lags. All specifications include firm and year fixed effects. Standard errors clustered at the firm level are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

B Event-Study Evidence: Dynamic Productivity Effects of hiring elite AI talent

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

Table A2: AI Talent and AI Patent Stock

	Full Sample		Manufacturing		Information	
	(1)	(2)	(3)	(4)	(5)	(6)
Frontier AI Talent Stock	0.632***		0.767***		0.523***	
	(0.0420)		(0.0607)		(0.0882)	
Normal AI Talent Stock	0.238***		0.214***		0.305***	
	(0.0299)		(0.0422)		(0.0716)	
Lagged Frontier AI Talent Stock		0.615***		0.742***		0.510***
		(0.0415)		(0.0609)		(0.0860)
Lagged Normal AI Talent Stock		0.247***		0.226***		0.303***
		(0.0307)		(0.0435)		(0.0726)
R&D Stock	0.0456***	0.0486***	0.0120	0.0166	0.0989**	0.102**
	(0.00990)	(0.0102)	(0.0180)	(0.0194)	(0.0438)	(0.0451)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	66,172	59,996	27,001	24,793	9,693	8,647
R^2	0.852	0.862	0.847	0.855	0.878	0.886

Notes: The dependent variable is the log of AI patent stock. AI Talent Stock measures the cumulative stock of frontier AI researchers employed by a firm since 1990, while Normal AI Talent Stock measures the cumulative stock of other AI-related workers employed by the firm. Columns (2), (4), and (6) replace the contemporaneous AI talent stock variables with their one-year lags. All specifications include firm and year fixed effects. Standard errors clustered at the firm level are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

To trace this dynamic relationship, we implement an event-study design centered on the year of first hiring an AI talent, 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.

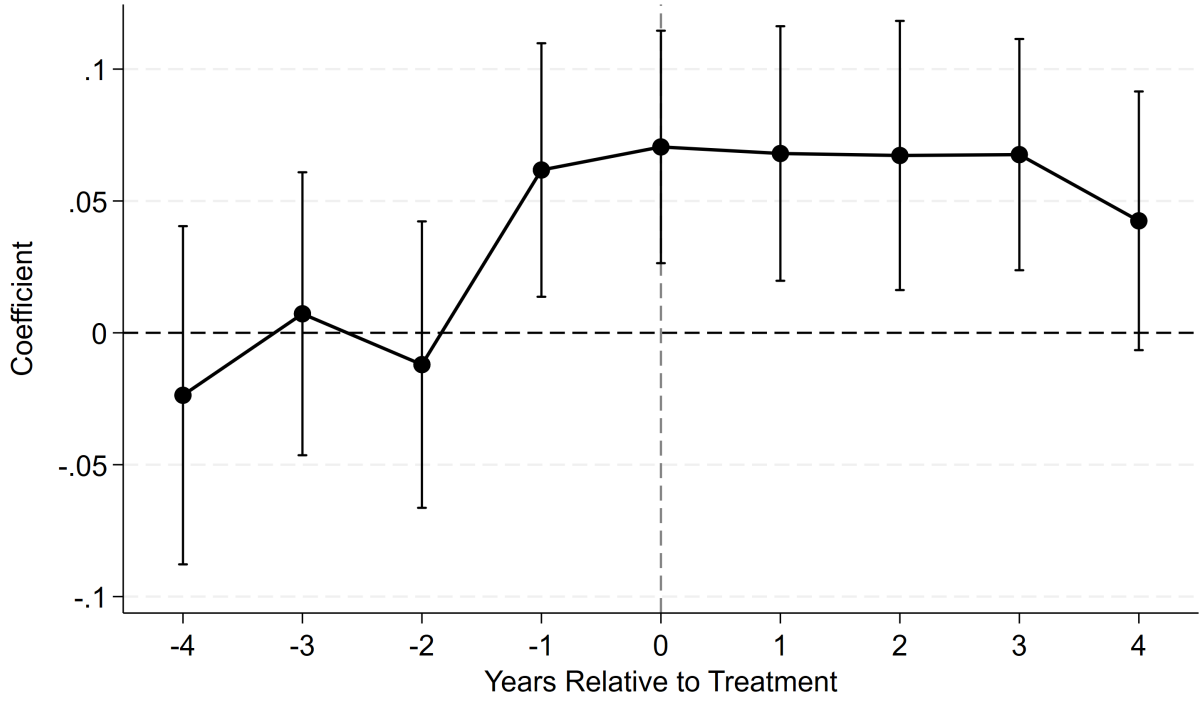


Figure A1: Effect of Hiring Elite AI Talent on Firm Value Added per Employee

Notes: The figure presents event-study estimates of the effect of hiring elite AI talent on firm value added per employee. Coefficients are plotted relative to the omitted pre-treatment period, with vertical bars indicating 95% confidence intervals.

The pre-adoption coefficients ($t = -4$ 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 hiring of an AI talent, productivity rises gradually and persistently. Five years after adoption, firms exhibit approximately 5 percent higher real value added per employee relative to their matched peers.

C AI Patents: Identification via the PaLLaFi Framework

Accurate identification of AI-related patents is fundamental to measuring technological diffusion, economic impact, and firm-level exposure to frontier innovation. Existing approaches, based on static keyword searches or predefined technology taxonomies, struggle to capture the interdisciplinary and rapidly evolving nature of AI. To overcome these limitations, we develop **PaLLaFi** (Patent Labeling via Language Models and Fine-Tuned Inference),

a multi-filter classification framework that leverages supervised machine learning models, state-of-the-art large language models (LLMs), and human-in-the-loop validation.

The framework is designed to balance three objectives that are often in tension in large-scale, longitudinal corpora:

1. Precision and generalizability through robust, fine-tuned supervised models.
2. Interpretability and domain attribution via state-of-the-art LLMs.
3. Temporal representativeness through year-wise stratified sampling to reflect evolving terminology.

Pipeline Overview

Figure A2 summarizes the five-stage pipeline:

- **Stage 1: Corpus construction and pre-filtering.** We begin with all U.S. utility patents granted from 1990–2024 (∼7.3M). Documents missing titles or abstracts are excluded; text fields serve as primary inputs.
- **Stage 2: Manual labeling and benchmarking.** Using an initial labeled set of 2,000 patents (1,000 AI; 1,000 non-AI), we benchmark:
  - *Supervised* — BERT, SciBERT, Longformer, Pat-Spectre, and PaECTER (our custom model),
  - *LLMs* — DeepSeek, LLaMA 3.1, Phi-3, Granite, Gemma-2, and Gemma-3 variants.

  PaECTER achieves the highest F1 accuracy among supervised models, while Gemma-3 (27B) demonstrates superior reasoning quality and becomes central to downstream inference.
- **Stage 3: Temporal expansion.** Year-wise stratified sampling yields a balanced dataset of ∼22,000 patents (11,000 AI; 11,000 non-AI), adjudicated through expert review and LLM-guided rationalization.
- **Stage 4: High-throughput classification.** The refined PaECTER model is applied to the full 7.3M corpus, providing fast binary classification (AI vs. non-AI).
- **Stage 5: Subdomain tagging and rationales.** Gemma-3 (27B) assigns AI subdomains (e.g., ML, NLP, Robotics, CV) and generates human-readable explanation labels. Subdomain enrichment is ongoing.

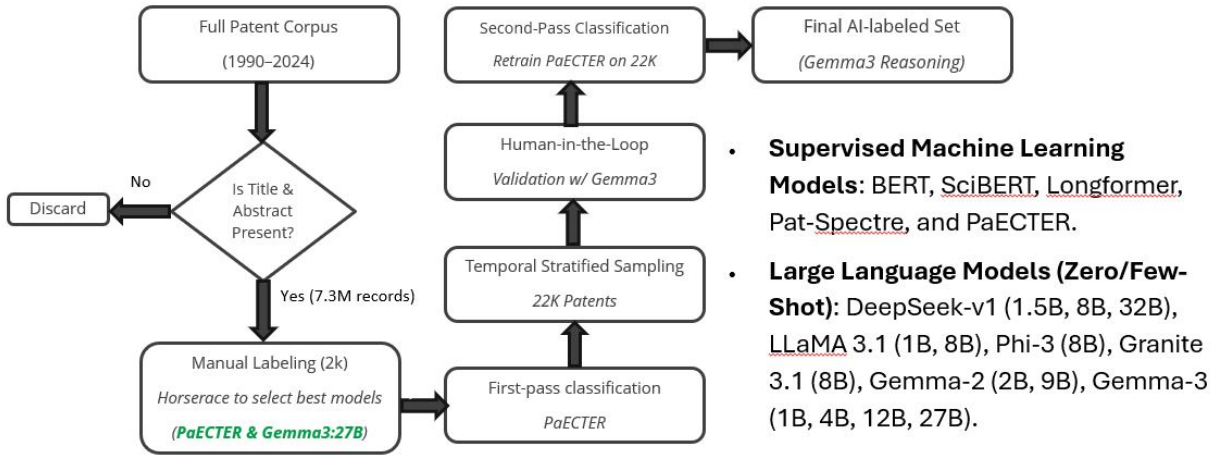


Figure A2: PALLAFI: A hybrid machine learning and LLM framework for identifying and characterizing AI-related patents across subdomains.

The resulting AI patent dataset underpins all innovation measures in the main analyses, enabling us to track the emergence of frontier technologies, the flow of elite AI talent into firms, and the productivity effects of AI adoption documented in Section 6.3.4.

D Measurement and Construction Details

D.1 Value Added

Value added Y_{it} is defined as output net of intermediate inputs, following the production function literature (e.g., De Loecker et al., 2020, Barrios et al., 2025). Because Compustat does not separately report intermediate input expenditures, we measure value added as sales minus cost of goods sold (COGS).

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.

D.5 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.

D.6 Clustering and Inference

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

E Variable Definitions

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).	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 Stock(real)	RD_{it}	Same calculating method as K_{it}	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