Buy or Build? How AI Shapes Corporate Acquisitions and Innovation*

Sebastiao Oliveira† Vahid Shadram‡

June 27, 2026

Abstract

We construct a novel classification of firms' exposure to AI to show that following a shock to the returns of AI capabilities, acquirers with greater AI exposure significantly increase acquisitions of small, young, and AI-*producing* targets, as opposed to AI-*adopting* targets, while their non-AI acquisitions decline. The response is strongest among firms least able to build AI in-house. These acquisitions also raise innovation, with gains extending well beyond AI into acquirers' core domains. Our results indicate a *reallocation* rather than an expansion in deal-making, and show how firms use acquisitions to adjust to AI-driven technological change.

Keywords: Acquisitions, Artificial Intelligence, Generative AI, innovation, M&A.
JEL Classification: D22, G31, G34.

*We thank Heitor Almeida, Dan Bernhardt, Janice Eberly and Spyridon Lagaras for comments. All errors remain our own. A previous version of this paper was circulated under the title "How AI Shapes Firm Investment: Evidence from Mergers and Acquisitions".

†Department of Economics, University of Illinois at Urbana-Champaign. E-mail: so28@illinois.edu.
‡Department of Economics, University of Illinois at Urbana-Champaign. E-mail: shadram2@illinois.edu.

1. Introduction

U.S. and global mergers and acquisitions (M&A) activity accelerated in 2025. According to PitchBook, deal counts increased by roughly 12% year-over-year and total deal value rose by 37%. This resurgence has been led by strategic acquisitions, with artificial intelligence (AI) being repeatedly cited as the motivation behind the increased deal-making. At the core of this acquisition wave is the realization that AI, particularly recent advances in generative models, has hugely expanded the range of tasks that can be automated or augmented by software, with far-reaching implications for firm productivity, organization, and competitive advantage.

This paper asks whether firms buy or build AI capability when its returns rise, and whether buying it fosters innovation. When the economy-wide returns to AI increase, firms must decide whether to develop specialized knowledge and human capital in-house or source these capabilities externally through market transactions. Frontier AI capabilities—state-of-the-art models, the scarce talent that builds them, and the proprietary data and computational resources they require—cannot be assembled in-house within a short window. A sudden increase in the returns to AI should therefore tilt firms toward *buying* rather than *making*, especially among firms whose internal build option is weakest. These acquisitions should also lead to broad innovation gains that extend well beyond AI itself, as firms may deploy the acquired AI as an input to advance their own core products and processes. Yet no existing study has examined how sudden increases in the returns to AI reshape the composition and dynamics of M&A activity, or its consequences for innovation.

We provide novel evidence that following a major advance in generative AI, firms with pre-existing exposure to AI significantly reallocate their acquisition activity toward AI-related targets. This reallocation reflects a shift in the composition of M&A activity rather than an expansion in overall deal-making, and is concentrated in acquisitions of AI-*producing* targets rather than the broad population that uses AI. The effects are concentrated among acquisitions involving young and small targets and among acquirers least able to develop AI internally, suggesting that firms use M&A to rapidly access frontier AI technologies.

We show that this reallocation leads to more innovation. Firms that acquire AI-*producing* targets significantly increase their subsequent patenting, both in AI and—more importantly—in non-AI technologies, whereas firms that acquire AI-*adopting* targets show no such gains. The increase is, in fact, larger for non-AI than for AI patents, and concentrates in technologically AI-adjacent domains rather than in unrelated fields. Consistent with AI behaving as a general-purpose

technology, the acquired capabilities diffuse across the firm's broader innovation portfolio. These patterns indicate that firms acquire AI-*producing* targets to internalize AI capabilities that enhance innovation across a wide set of technologies, not just AI technologies.

One challenge in documenting these effects is measuring the nature of firms' exposure to AI. Standard industry classifications do not cleanly capture AI exposure, making it difficult to distinguish firms that produce AI technologies from those that merely adopt them. To overcome this challenge, we develop a new firm-level AI classification using PitchBook's verticals and emerging spaces, which describe the sectors and technologies in which firms operate. This classification allows us not just to distinguish between AI and non-AI firms but also to identify firms as AI producers or AI adopters, which is crucial for identifying and understanding acquisition strategies.

A second challenge is identifying a clean, plausibly exogenous AI shock. Isolating an economy-wide AI shock is difficult because new technology adoption typically evolves gradually and is often correlated with firms' underlying acquisition strategies and contemporaneous market conditions. To address this challenge, we exploit the release of ChatGPT in November 2022 as a sudden and unanticipated AI shock (Eisfeldt et al., 2025). The event was unanticipated, as neither firms nor users had advance information about the timing or capabilities of the release, limiting the scope for strategic anticipation. Moreover, ChatGPT represented a discrete jump in AI capabilities, particularly in natural language understanding and open-ended reasoning, sharply increasing the marginal value of AI-related assets across a broad set of industries. Finally, the technology diffused unusually rapidly, reaching mass adoption within months, supporting the interpretation of the release as a common, economy-wide productivity shock rather than a gradual technological trend.

We begin by documenting five facts. First, M&A activity by acquirers that are AI firms increases by approximately 40% between 2020Q1 and 2025Q3. Second, acquirers that are AI firms increase AI-related deals (deals in which the target is an AI-exposed firm) in lieu of other deal types: among acquirers that are AI firms, AI deals increase by about 65%, while non-AI deals by the same acquirers fall by roughly 17%. Third, the increase in M&A activity is concentrated in acquisitions of AI-*producing* targets, whereas acquisitions of AI-*adopting* targets remain comparatively flat, suggesting that firms exposed to AI use M&As primarily to access external AI-*producing* innovation. Fourth, the acquired producers are frontier AI firms (concentrated in generative AI and machine learning) whose technologies overlap with the domains in which buyers themselves patent. Fifth, AI adopters' own post-shock patenting concentrates in the complementary digital infrastructure needed to deploy AI, not in core AI itself.

Having established these novel key facts, we next gauge the impact of AI on M&A activity. To do this, we employ a matched difference-in-differences approach that aligns acquirers that are AI firms with acquirers that are non-AI firms based on observable pre-shock characteristics and restricts comparisons to firms within the same industry. The core idea underlying our empirical design is that the release of ChatGPT affected firms asymmetrically depending on their pre-existing exposure to AI. In particular, acquirers with greater AI exposure experienced a larger shift in acquisition behavior following the AI shock. Following the release of ChatGPT, acquirers that are AI firms are 23% more likely to acquire AI-exposed targets. Moreover, these acquirers do not significantly increase total deal value across either AI or non-AI targets, indicating a strategic reallocation toward AI-related acquisitions rather than an increase in overall acquisition spending.

To further characterize the treatment effects, we estimate the dynamic responses of M&A activity to the AI shock. We find that the effects become statistically significant approximately one year after the shock and peak after three years, with no evidence of pre-trends. This delayed response is consistent with the time required for firms to assess the strategic implications of AI, identify suitable targets, and complete the M&A process, which typically involves lengthy search, due diligence, and integration phases. These findings support the validity of our identification strategy and highlight the long-term impact of AI on firms' acquisition strategies.

We next study the economic mechanism behind the reallocation we document by analyzing which firms AI-exposed acquirers buy. The make-versus-buy mechanism predicts that when in-house development becomes slow or costly relative to acquisition, firms substitute away from *making* and toward *buying* the capability (Phillips and Zhdanov, 2013; Bena and Li, 2014; Frésard et al., 2020). To test this, we classify targets according to the role of AI in their core activities. Firms for which AI primarily serves as an input to enhance products are classified as AI adopters, whereas firms whose primary output involves the development of AI technologies are classified as AI producers. Make-versus-buy implies that AI-exposed acquirers buy the capability they cannot quickly build—the technology held by AI-*producing* firms—rather than the complementary assets of firms that merely *adopt* AI. This is precisely what we find. Following the AI shock, we observe a significant 56.43% increase in the probability of acquiring AI-*producing* targets, but no change in acquisitions involving AI adopters.

The framework also implies that acquirers buy capability rather than scale: the targets should be the small, young innovators whose value is the technology and human capital they embody. Consistent with this, we show that acquisitions are concentrated among young and small targets.

For instance, the probability of acquiring small targets increases by 25.9% following the release of ChatGPT. Similar findings obtain for young targets. This result provides direct evidence that firms respond by reorganizing how they acquire innovation, not just how much they invest or acquire.

We further show that the reallocation response is strongest among the firms least able to build AI in-house: AI adopters turn to AI-producing targets, and in a within-adopter triple-difference the buy response is significantly weaker for adopters with more pre-shock patenting. For instance, moving from the 25th to the 75th percentile of pre-shock patenting among AI-adopter buyers reduces the post-shock acquisition response by 40%. This evidence indicates that internal capability substitutes for acquisition, and does so only for AI, where capabilities cannot be assembled in-house within a short window. We are the first to show that the make-versus-buy margin shapes firms' response to an AI shock.

Turning to innovation outcomes, we combine our M&A panel with patent records from the USPTO PatentsView database and classify patents as AI patents and non-AI patents. We show that acquirers of AI-producing targets significantly increase their post-acquisition patenting, while acquirers of AI-adopting targets do not. These gains extend well beyond AI itself: they are nearly twice as large for non-AI as for AI patents and concentrate in the acquirer's core, AI-adjacent domains. These findings suggest that the innovation associated with these deals is generated by the acquirer, which deploys the acquired AI as an input to advance its own core technologies, rather than reflecting the patenting of the acquired targets themselves.

We conclude by showing that the acquired AI producers are *complementary capability suppliers* rather than *substitutes*: the post-shock increase is concentrated in complementary, non-competitor targets, indicating that AI-driven M&A operates through synergy rather than consolidation. These findings are consistent with Bowen et al. (2023)—the acquired producers are precisely the complementary segment their framework associates with selling out, while the disruptive frontier remains independent.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 describes the data. Sections 4 and 5 present the identification strategy and main results, respectively. Section 6 provides additional results and robustness tests. Section 7 concludes.

2. Related Literature

Our paper is related to the literature on the make-versus-buy margin and the role of M&A in reallocating innovation across firms. Phillips and Zhdanov (2013) show that large firms often acquire small innovators rather than win an in-house R&D race, so an active acquisition market shapes who innovates and how. Bena and Li (2014) document that acquirers and targets sort on complementary innovation characteristics—low-R&D, patent-rich firms acquiring high-R&D innovators—and that technological overlap drives both deal incidence and post-merger innovation. Frésard et al. (2020) show that innovation incentives influence integration and acquisition decisions through vertical linkages. More broadly, Acemoglu et al. (2018) show that aggregate growth depends on reallocating resources from low-type to high-type firms, highlighting that the selection of which firms innovate is central to productivity growth. We contribute to this literature in three ways. First, rather than studying cross-sectional patterns, we exploit a sudden, economy-wide shock to the returns of a specific capability—generative AI—that sharply raises the cost of *making* relative to *buying*, providing a clean test of the make-versus-buy margin. Second, our novel producer–adopter taxonomy identifies the capability being sourced and lets us show that the response is strongest among the firms least able to build it in-house. Third, we show that these acquisitions raise subsequent innovation, with gains differentially accruing in acquirers' core technological domains rather than in AI, itself.

Our paper also contributes to the literature studying how economic and technological shocks reshape M&A and reallocate corporate assets toward more productive uses. Mitchell and Mulherin (1996) show that takeover and restructuring activity clusters by industry and is driven by industry-level shocks. Maksimovic and Phillips (2001) document an active market for corporate assets in which mergers and asset sales move plants to more efficient users, while Maksimovic and Phillips (2002) show that firms allocate resources across industries in line with productivity and industry fundamentals. Jovanovic and Rousseau (2002) formalize this logic in a Q-theory of mergers, in which high-productivity firms acquire and redeploy the assets of lower-productivity ones, and Schoar (2002) examines the productivity consequences of such reallocation. More recent work studies technology-driven deals directly: Ma et al. (2025) show that firms with greater technological capabilities tend to acquire less tech-savvy firms and subsequently increase technology investment at the target. Fang et al. (2025) show that, between 2009 and 2019, AI-exposed firms that are interested in data acquisitions are more likely to acquire data-rich companies.

Our work offers several advances and insights to this literature. First, our analysis exploits a new firm- and deal-level classification of AI exposure based on PitchBook verticals and emerging spaces, which allows us to distinguish between AI-*producing* and AI-*adopting* firms on both the acquirer and target sides of M&A transactions. Second, rather than relying on endogenous cross-sectional heterogeneity in AI capabilities, we provide the first evidence on how firms adjust their acquisition behavior in response to an exogenous and economy-wide increase in the returns to AI induced by the release of ChatGPT. Third, while prior work emphasizes data complementarities and target-side investment, we document a shift in the *composition* of acquisitions—toward small, young, AI-producing targets rather than AI adopters—and show that it raises the acquirer's own innovation.

Another strand of related literature examines the nature of acquired startups and the boundary between remaining independent and selling out. Bowen et al. (2023) show that startups in rapidly evolving (frontier) technology areas tend to stay independent and go public, whereas those in stable areas complement existing technologies and sell out. Our evidence shows that the acquired AI producers are complementary, non-substitute capability suppliers—precisely the segment Bowen et al. (2023) associate with selling out—so AI-driven M&A operates through synergy rather than consolidation, in contrast to a killer-acquisition motive.

Finally, we add to a growing literature studying how investments in AI relate to firm and industry outcomes (D'Acunto et al., 2019; Chen et al., 2019; Fuster et al., 2022; Acemoglu et al., 2022; Yang, 2022; Babina et al., 2023; Babina et al., 2024; Gofman and Jin, 2024; Marchesi and Tang, 2025; Jansen et al., 2025; Eisfeldt et al., 2025; Gupta et al., 2025). We use ChatGPT's release as a sudden and unanticipated increase in the returns to AI-related assets to study a distinct corporate-finance margin: the composition of acquisitions and the reallocation of deal activity toward AI-related targets. We provide novel evidence that firms reallocate external investment toward AI-related assets through M&A.

3. Data

3.1 PitchBook M&A Data

We construct an acquirer–quarter–level panel of U.S. mergers and acquisitions using data from PitchBook, which tracks M&A deals involving startups.[1] We focus on completed M&A deals between the first quarter of 2020 (2020Q1) and the third quarter of 2025 (2025Q3), as this period allows for a clean comparison of M&A activity before and after the widespread diffusion of artificial intelligence technologies. We exclude the fourth quarter of 2025 because our dataset only includes deals up to November 6, 2025.

Our main sample includes transactions classified by PitchBook as *Merger/Acquisition*, in which a company acquires at least a controlling stake in another company's capital stock. The final sample comprises 3,992 unique deals and 2,610 unique buyers. We exclude deals classified as *Merger of Equals*, defined as combinations of two or more companies into a single entity in which control is shared or indeterminate, and *Corporate Divestiture*, defined as transactions in which any investor type acquires a business division or subsidiary previously part of an operating company.

We use M&A data from PitchBook rather than the more traditional SDC/Refinitiv M&A data for two main reasons. First, PitchBook's granular, multi-tag technology taxonomy—its Verticals and Emerging Spaces—lets us construct a novel classification that not only distinguishes AI from non-AI firms but also separates AI *producers* from AI *adopters*. This producer–adopter distinction is crucial for our identification: it is what lets us characterize firms' acquisition strategies, whether they buy AI capability they cannot build or merely acquire users of it, rather than simply documenting that AI firms transact with one another. Coarser industry classifications such as SIC or NAICS, available in SDC, cannot support this distinction. Second, PitchBook also provides the private-firm characteristics and competitor-relationship data underlying our matching, heterogeneity, and complementarity analyses. Importantly, while PitchBook's strength is its coverage of private targets, our findings are not driven by them: the results are robust to restricting the sample to public acquirers, the segment covered by SDC, indicating that our conclusions do not hinge on PitchBook-specific private-deal coverage.

[1] We use the terms "buyer" and "acquirer" interchangeably to refer to firms that acquire another firm.

3.2 Classification of AI Firms

A central challenge in studying the relationship between AI and corporate outcomes is to determine which firms are exposed to artificial intelligence. One way to do this is to use traditional industry classifications such as SIC or NAICS codes, which assign each firm to a single sector based on its primary line of business. For our purposes, such classifications are too coarse to distinguish companies that develop AI technologies from the ones that adopt it, and they do not help capture the multi-dimensional nature of modern technology firms that simultaneously operate across several domains. This leads us to exploit PitchBook's granular activity taxonomy, which categorizes each firm's activities under two complementary schemes: *Verticals*, which provide descriptions of specialized technological sectors, and *Emerging Spaces*, which capture narrower, rapidly evolving technology niches. An advantage of using these classifications is that a single firm can be assigned multiple verticals and emerging spaces, which reflects the breadth of its technological activities. We exploit this granularity to decompose each firm's technological profile and separately identify the role of AI in its activities—whether as a primary output, an enabling input, or neither.

Classification of industry verticals and emerging spaces. We classify AI exposure of firms in three steps at the vertical and emerging-space level before we aggregate them to the firm level.

Collection of category descriptions. We first collect the official descriptions provided by PitchBook for each of its industry verticals and emerging spaces.[2] These descriptions are publicly available on PitchBook's website and summarize the core technologies, products, and business activities associated with each category. Table A.1 in Appendix A presents two illustrative examples. The descriptions significantly vary in the degree at which artificial intelligence features in the category's core activities. For example, for the Artificial Intelligence and Machine Learning vertical, the description focuses on the development of intelligent systems, learning algorithms, and AI applications. However, Construction Technology—AI-related tools such as construction robots and automated survey drones—appear alongside non-AI solutions such as project management software and equipment marketplaces, which shows that AI serves as one of several enabling technologies rather than as the primary output.

LLM-assisted classification. In the next step, we use OpenAI model "gpt-5.2-2025-12-11,"

[2] See https://pitchbook.com/what-are-industry-verticals and https://pitchbook.com/blog/what-are-emerging-spaces for details.

accessed December 2025 to assist in the initial classification of each vertical and emerging space into one of three groups: 1) AI-producing, 2) AI-adopting, or 3) non-AI. If a firm is classified as AI-producing, AI-adopting, or both, it is also labeled as AI. We provide the model with the full PitchBook description for each category (one by one to improve accuracy) and instruct it to determine whether within the category, artificial intelligence constitutes the primary technology being developed, an enabling technology employed, or is largely absent. The model is asked to return a single classification label for each description. For example, the agent labels as AI-producing those categories centered on the development of machine learning, natural language processing, generative AI, or related AI technologies. Categories that primarily operate in another industry—such as construction, agriculture, healthcare, or financial services—but explicitly employ AI for automation, optimization, forecasting, recommendation, etc. are classified as AI-adopting. Lastly, categories where AI does not feature in the PitchBook description such as Cannabis, Oil&Gas, or Indoor Farming are classified as non-AI.

Manual Review. After we get all the LLM-generated classifications, we manually reviewed them for consistency and plausibility. Table A.2 reports all verticals and emerging spaces.

Aggregation to firm-level classification. For each firm in the dataset, we decompose its reported verticals and emerging spaces into individual categories and aggregate the category-level classifications to assign a firm-level AI status. A firm is classified as an *AI producer* (*AI adopter*) if at least one of its assigned categories is classified as AI-producing (AI-adopting). Therefore, a firm can be classified as both AI adopter and AI producer depending on the categories they have in both verticals and emerging spaces. Finally, firms with no AI-related classifications are labeled as *non-AI firms*. We symmetrically apply these classifications to both buyer and target firms.[3] Table A.3 provides illustrative examples mapping well-known firms to their classifications. For example, firms such as OpenAI and Anthropic are classified as AI firms under our scheme.

External validation of AI classification. To validate our PitchBook-based AI classification against an independently constructed measure, we compare it to the firm-level generative AI exposure scores developed by Eisfeldt et al. (2025). Their measure quantifies each company's

[3]Several emerging-space categories—such as Generative AI, LLM Agents, and AI Neoclouds—may have been introduced after the ChatGPT release. To ensure that treatment status is not defined by post-shock taxonomic additions, we re-estimate all specifications excluding emerging spaces entirely and defining AI exposure using only PitchBook's industry verticals. The results are unaffected.

exposure to generative AI according to the content of each firm's workforce's task. They use detailed occupation-level data from LinkedIn employment structures and decompose firm-level exposure into a *core* component, which captures the share of workers whose primary job functions overlap with generative AI capabilities, and a *supplemental* component, which has to do with the share of workers whose peripheral tasks can be enhanced by generative AI. Of note, their measure is constructed entirely from labor market data and is therefore independent of PitchBook's technology-based classification system.

We match this measure to our acquirer panel and get 232 public acquirers for which both our AI classification and the Eisfeldt et al. (2025) firm-level exposure scores are available. Firms classified as AI acquirers under our scheme exhibit significantly higher core GenAI exposure ($r = 0.23$, $p < 0.001$), which shows that our vertical-based classification identifies firms whose workers perform tasks that overlap more closely with the primary capabilities of generative AI.

At the same time, AI acquirers display significantly *lower* supplemental GenAI exposure ($r = -0.31$, $p < 0.001$). This pattern is economically intuitive: firms that already have substantial in-house AI capability have less scope for generative AI to complement peripheral worker tasks, precisely because many of those tasks are already automated or integrated into the firm's existing AI infrastructure. The total exposure measure is also positively correlated with AI-acquirer status ($r = 0.16$, $p = 0.014$), although the magnitude is attenuated because the core and supplemental components move in opposite directions.

A comparison of means across the two measures reveals similar patterns. AI acquirers exhibit 5% higher total GenAI exposure than non-AI acquirers (0.404 versus 0.385; $t = -2.44$, $p = 0.016$), with the difference driven primarily by the core component, which is 8.3% higher (0.357 versus 0.330; $t = -3.58$, $p < 0.001$). In contrast, AI acquirers' supplemental exposure is 14.6% lower (0.047 versus 0.055; $t = 4.93$, $p < 0.001$).

Overall, these results suggest that our PitchBook classification and the Eisfeldt et al. (2025) exposure measure, while constructed from entirely different data sources and methodologies, identify overlapping but conceptually distinct dimensions of firms' relationship to generative AI. Our classification captures firms that actively operate in AI-related technology markets, as identified by PitchBook's industry taxonomy. The Eisfeldt et al. (2025) measure captures the extent to which a firm's workforce performs tasks that are exposed to generative AI capabilities.

3.3 Sample Construction

After we implement this classification, our sample features 1,617 acquirers that are either AI producers or AI adopters and 993 acquirers that are neither AI producers nor AI adopters. At the deal level, we observe transaction values when available. Deals with missing transaction values are excluded from our analysis on deal value. For each acquirer-quarter, we construct the total number of completed deals, as well as counts of acquisitions of AI-*producing* targets, AI-*adopting* targets, and non-AI targets. Of 3,992 deals, 2,357 are deals in which the target is an AI firm, and the rest involve non-AI targets. If a firm did not make any deal in a given quarter, we assign a zero to that quarter.[4] We also construct total deal values and deal values by target type. Deal values are winsorized at the 1st and 99th percentiles. Table 1 provides firm characteristics and deal descriptive statistics. Acquirers that are AI firms tend to be younger and smaller than acquirers that are non-AI firms. They are also more likely to engage in M&A activity after the release of ChatGPT.

Table 2 reports deal counts by acquirer type and target type before and after the ChatGPT release, and reveal a reallocation pattern in the composition of AI acquirers' deals. Post-shock acquisitions of AI *producing* targets by AI acquirers rise from 296 to 437, an increase of 47.6%, while acquisitions of AI *adopting* targets by the same firms decline from 594 to 559, a 5.9% decrease. Acquisitions of non-AI targets also fall, from 392 to 315, a 19.6% decline. These patterns indicate that AI-exposed acquirers did not just increase M&A activity across the board; rather, they shifted the composition of their deals toward firms that develop frontier AI capabilities and away from firms that merely use AI as an input. Among non-AI acquirers, the corresponding changes are modest: AI-producing target acquisitions are roughly stable, from 103 to 108 (4.9%), while both AI-adopting and non-AI target acquisitions decline. The contrast between the two acquirer groups, reallocation toward AI producers among AI acquirers, and broad contraction among non-AI acquirers, provides the descriptive foundation for our main analysis.

3.4 Patent Data

Data source. To study firms' innovation outcomes, we combine our M&A panel with patent records from the USPTO PatentsView database. We obtain disambiguated patent assignee records, Cooperative Patent Classification (CPC) codes, and both application filing dates and grant dates for all patents granted through 2025. Following Lerner and Seru (2022), we date each patent by its

[4] Restricting to firms that existed before 2022Q4 yields qualitatively similar estimates.

application date rather than its grant date, as the filing date closely approximates the timing of the underlying invention activity and avoids the variable lags introduced by the patent examination process. This is an important convention for our setting, as we study how M&A activity responds to a discrete technological shock: dating patents by grant date would introduce mechanical delays that blur the timing of firms' innovative responses to the release of ChatGPT in 2022.

Matching procedure and panel construction. We link patent records to our PitchBook acquirer panel using company names. We standardize names in both datasets by converting to lowercase, removing punctuation, stripping common legal suffixes (e.g., Inc., Corp., LLC, Ltd., L.P.), and collapsing whitespace. Then we perform exact matching between cleaned assignee names in PatentsView and cleaned buyer names in PitchBook. After we match the datasets, we get 711 matched buyer firms, representing approximately 27% of all acquirers in the M&A sample. Among these firms, 388 are classified as AI firms and 323 as non-AI.

We next aggregate patents to the buyer-quarter level by counting the number of distinct patent applications filed by each matched buyer in each calendar quarter. We merge these counts into the buyer-quarter panel and assign zero to buyer-quarter observations in which the matched firm filed no patent applications. The resulting panel contains 711 unique firms with 16,353 buyer-quarter observations.

Classification of AI patents. We use two approaches to classify patents as AI and non-AI. Our baseline classification uses Cooperative Patent Classification (CPC) codes assigned by the USPTO. We define a patent as an AI patent if any of its assigned CPC codes fall within the following subclasses: G06N (computing arrangements based on specific computational models, encompassing machine learning, neural networks, and knowledge-based systems), G06V (image or video recognition and understanding, encompassing computer vision), G10L (speech analysis, synthesis and recognition), and G06F18 (pattern recognition). We follow a conservative baseline that focuses on subclasses most directly linked to AI methods. In robustness analyses, we expand the definition to include G06F40 (natural language processing and computational linguistics), which has become increasingly central to AI following the rise of large language models. We assign the rest of patents as non-AI.

In another robustness analysis, we use the USPTO Artificial Intelligence Patent Dataset (AIPD 2023) developed by Pairolero et al. (2025). The AIPD uses machine-learning classifiers

trained on BERT for patents to identify AI-related inventions directly from patent text—including titles, abstracts, and claims—rather than relying only on examiner-assigned CPC codes. We merge the AIPD to our matched patent sample using patent identifiers and retain only patents covered by the AIPD release. Because the current AIPD release contains AI classifications only for patent documents published through 2023, coverage of our patent sample falls sharply after 2023. We therefore restrict all AIPD-based analyses to observations through 2023Q4 (one year after ChatGPT release) to ensure that every patent in the sample receives an explicit USPTO AI classification. One advantage of this classification is that it provides multiple prediction thresholds, which lets us further assess whether our results are sensitive to the stringency of the AI patent definition. We use the `Predict93`, which labels a patent as AI-related when the model-assigned probability of AI content exceeds 93%.

3.5 Empirical Regularities

Using this dataset and our novel AI classification, we establish five empirical facts about the relationship between AI, M&As, and innovation. To the best of our knowledge, our paper is the first to document these facts.

Fact 1: Increase in M&A Deals by acquirers that are AI firms. We first examine M&A activity over time and across acquirer types. Panel A of Figure 1 shows a significant increase in M&A deals by acquirers that are AI firms. For example, in 2020Q1, such firms acquired around 100 firms; in 2025Q3, this number was approximately 140, representing a 40% increase. At the same time, there is a sustained decline in M&A deals by acquirers that are non-AI firms, from around 80 in 2020Q1 to 35 deals in 2025Q3.

Fact 2: Increase in M&A Deals is Concentrated in Deals Involving Targets who are AI firms. We next decompose deals into those in which the target is an AI firm and those in which the target is a non-AI firm. Panels B and C of Figure 1 show that acquirers that are AI firms primarily engage in deals in which the target is an AI firm rather than other deal types. Specifically, we document an increase of 65% in deals in which the target is an AI firm by acquirers that are AI firms, together with a decline of 17% in acquisitions of non-AI firms by the same acquirers.

Fact 3: Increase in M&A Deals is Concentrated in AI-Producing Targets. The granularity of our data lets us distinguish between firms that are AI producers and AI adopters. Figure 2 plots M&A deals involving acquirers that are AI-*adopting* firms and acquirers that are AI-*producing* firms. It shows that the increase in M&A activity is concentrated in acquisitions of AI-*producing* targets. Panel A shows that among acquirers that are AI-*producing* firms, acquisitions of AI-*producing* targets increase from 10 deals in 2020Q1 to 44 deals in 2025Q3, representing an increase of over 300%. Panel B shows a similar pattern for acquirers that are AI-*adopting* firms, with acquisitions of AI-*producing* targets almost doubling from 16 to 29 deals. In contrast, acquisitions of AI-*adopting* targets remain roughly flat, suggesting that firms exposed to AI engage in M&A primarily to access external AI-*producing* innovation.

Fact 4: Post-shock acquisitions target frontier AI capabilities that complement buyers' existing innovation. Tables A.4–A.6 show that targets concentrate in frontier AI: the leading verticals are Artificial Intelligence & Machine Learning (32.1%), SaaS (20%), and Big Data (10.6%), and the leading emerging spaces are Generative AI (43.4%) and LLM Agents (9%). This concentration intensified after the shock—SaaS (+8.6 p.p.) and AI/ML (+4.5 p.p.) rise the most, while TMT (–11 p.p.) and Big Data (–2 p.p.) fall—indicating a shift away from hardware- and perception-oriented AI toward software- and language-based capabilities. These targets' technologies also overlap with buyers' own innovation: Table A.7 shows that, after linking the patent-matched targets to buyers' post-shock patenting, both concentrate in Computing and Software (63.2% of target vs. 47.3% of buyer patents) and Telecommunications (20.7% of target vs. 28.6% of buyer patents). These findings suggest that firms acquired frontier AI capabilities that complement, rather than diversify away from, their existing portfolios.

Fact 5: AI adopters' post-shock patenting concentrates in complementary digital infrastructure rather than core AI. We next examine the technological composition of patents filed by AI-adopting firms after the ChatGPT release. Importantly, we distinguish patents classified as *non-AI* under our CPC-based definition, capturing innovation activity outside of core AI technologies. If AI adopters were simply developing their own AI capabilities internally, we would expect their post-shock patenting to concentrate in AI-adjacent CPC classes (e.g., machine learning). Instead, the distribution of post-shock non-AI patents by AI adopters reveals a significant concentration in digital and information-processing technologies that are complementary to

AI deployment. Table A.8 shows that 44.6% of these patents fall within Computer Technology, 29.8% belong to Digital Communication, and 9.2% are classified as IT Methods for Management.

We also examine the firms driving these patents and their representative patent titles, which reinforces this interpretation. The largest contributors include enterprise software, cloud computing, cybersecurity, and data infrastructure firms such as Salesforce, Snowflake, Pure Storage, VMware, Block, Wiz, Intuit, Fortinet, Rubrik, Twilio, and Zscaler. Patent titles reveal innovation in database transaction management, across-control systems for data lakes, cloud-platform service routing, automated log-data redaction, secure cross-deployment messaging, payment authorization systems, and database replication, which are plausibly technologies that form the software and data infrastructure required to apply AI at scale.

4. Empirical Design

Having established these core features, we now gauge the impact of AI on M&As. The central challenge in estimating the effect of AI on M&A activities is that firms with higher exposure to AI technologies may systematically differ from firms with little or no exposure, reflecting pre-existing differences in innovative capabilities and strategic orientation that are correlated with acquisition behavior. To address these systematic differences, we combine a technological, economy-wide shock with a matching strategy that aligns buyers that are AI firms with buyers that are non-AI according to their observable pre-shock attributes. The idea underlying our empirical design is that this shock affected firms asymmetrically, depending on the pre-existing exposure of these firms to AI.

4.1 Release of ChatGPT as an Exogeneous AI shock

Following Eisfeldt et al. (2025) and others, we use the ChatGPT 3.5 release date as a plausibly exogenous AI shock. There are three features of this event that are essential for our identification strategy. First, there was no anticipation of the ChatGPT 3.5 release. Prior to November 2022, neither firms nor users had any public information regarding the exact release date of this platform or how capable and productive this version would be.[5] Consequently, buyers could not time M&A deals or strategically decide in advance which

[5]See OpenAI's public announcement introducing ChatGPT, which describes the release as intended to receive user feedback: https://openai.com/index/chatgpt/.

activities to undertake involving firms operating in AI.

Second, the release of ChatGPT 3.5 represented a discontinuous jump in AI capabilities and in the productivity gains it could generate. Unlike earlier generations of machine learning and deep learning tools, generative AI models, especially ChatGPT 3.5, employ natural language modeling and open-ended reasoning, which have broadly expanded the applicability and potential of tasks across different industries. This discrete technological advance sharply increased the marginal value of AI-related capabilities for a wide range of firms, rather than incrementally improving a narrow set of applications.

Third, while past technologies such as the World Wide Web or social media platforms like Instagram or Facebook diffused over several years, ChatGPT reached one million users within days and over one hundred million users within two months of its release.[6] This unusually rapid adoption implies that firms were exposed to the productivity implications of generative AI almost immediately, strengthening the interpretation of the release of ChatGPT 3.5 in 2022Q4 as a common and economy-wide productivity shock.

Overall, these three features support our interpretation of the release of ChatGPT as a sudden and unanticipated increase in the economic returns to AI-related assets and capabilities. Importantly, this shock affects firms asymmetrically according to their pre-existing exposure to AI technologies, motivating our matched difference-in-differences design that compares AI-exposed acquirers with otherwise similar non-AI acquirers around the 2022Q4 cutoff.

4.2 Empirical Specification

We begin with a standard two-way fixed effects DiD specification:

$$Y_{it} = \alpha_i + \delta_t + \beta\big(\text{Treated}_i \times \text{Post}_t\big) + \varepsilon_{it}, \tag{1}$$

where α_i are acquirer fixed effects and δ_t are quarter fixed effects. The outcome variables, Y_{it}, are the probability of engaging in an M&A deal (dummy variable that equals one if the number of deals is greater than one and zero otherwise), the log of the number of deals, and the log of the deal value. Treated_i is an indicator equal to one if the buyer i is an AI firm (either AI producer,

[6] See Reuters (2023), which documents the rapid adoption of ChatGPT and reports that it reached 100 million monthly active users within two months of launch: https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/.

AI adopter or both) and zero otherwise, and Post_t is a dummy variable equal to one in case $t \geq$ 2022Q4 and it is zero otherwise. Therefore, β measures the average post-shock change in Y_{it} for acquirers that are AI firms relative to acquirers that are non-AI. Standard errors are clustered at the acquirer level to account for potential serial correlation in outcomes within firms over time.

To strengthen identification, we construct a matched control group that is similar to the treated firms according to their M&A activities and firm characteristics prior to the shock. The goal of the matched DiD is to approximate the counterfactuals of acquirers that are AI firms in the absence of the AI shock by comparing them to acquirers that are non-AI firms with similar observable attributes prior to the shock. Our matching approach addresses concerns that the divergence that happened after the shock may be due to the differential selection into AI exposure, deal-making intensity or size rather than the causal effect of AI.

To do this matching, we first define a vector of pre-treatment covariates, X_i, that includes: pre-shock acquisition intensity (total number of deals), pre-shock deal scale (log total deal value), firm size (log employment), and pre-shock AI exposure in acquisitions (share of AI-target deals). We then apply nearest-neighbor matching with replacement. This allows us to match each treated buyer to one or more non-AI acquirers with similar X_i. Moreover, to ensure that comparisons are drawn within similar environments, we impose exact matching within the broad industry sectors that PitchBook provides. Industry sectors include Information Technology, Business Products and Services (B2B), Healthcare, Consumer Products and Services (B2C), Financial Services, Materials and Resources, and Energy. This further reduces any likelihood of having treated-control differences as a result of sectoral differences rather than the shock. Table 3 provides a comparison of the treated and control groups before and after matching, showing that the matching procedure substantially improves balance and eliminates statistically significant differences between the two groups.

We investigate the validity of the parallel trends assumption, which is key in the DiD method, estimating an event-study specified as follows:

$$Y_{it} = \alpha_i + \delta_t + \sum_{k \neq -1} \beta_k \left(\text{Treated}_i \times \mathbf{1}\{t - t_0 = k\} \right) + \varepsilon_{it}, \tag{2}$$

where $t_0 =$ 2022Q4 and k indexes event time in quarters. α_i denotes acquirer fixed effects, and δ_t denotes quarter fixed effects. The outcome variables are the same as those studied in Equation (1). Treated_i is an indicator equal to one if buyer i is an AI firm (either an AI producer or an AI adopter) and zero otherwise.

Panel A of Figure 3 plots the results for the probability of engaging in an AI deal. We find no evidence of pre-trends. We obtain similar results when estimating the dynamics for the probability of engaging in acquisitions involving AI-*producing* targets (Panel B of Figure 3), young targets (Panel C of Figure 3), small targets (Panel D of Figure 3), and young and small targets (Panel E of Figure 3). In summary, we find no evidence of pre-trends, providing strong support for our identification strategy.[7]

5. Main Results

5.1 The Impact of AI on M&A Activity

We first document the impact of the AI shock on M&A activity. Panel A of Table 4 reports the results.[8] After the release of ChatGPT, acquirers that are AI firms are 1.2 p.p. more likely to acquire targets that are also AI firms. This effect is economically significant: relative to the baseline probability of 5.2%, it represents a 23% increase in the likelihood of engaging in AI deals. Panel A of Table 4 also documents a 0.8% increase in the number of deals involving AI-firm targets. At the same time, column (3) of Panel A shows no significant effect on deal value, indicating that while acquisitions of AI-firm targets increase, total acquisition expenditures with those deals remain unchanged.

Panel A of Figure 3 plots the dynamic effects for the probability of engaging in an AI deal. We find that the effects peak three years after the shock. This delayed response is consistent with the time required for firms to assess the strategic implications of AI, identify suitable targets, and complete the M&A process, which typically involves long search, due diligence, and integration phases. It also highlights the long-term consequences of AI to firms' acquisition strategies.

Panel B of Table 4 shows that acquirers that are AI firms do not increase non-AI acquisitions following the AI shock. The change in the number of deals involving targets that are non-AI firms is economically and statistically insignificant. Importantly, we again find no significant effect on deal value, indicating that the increase in acquisitions of AI firms does not reflect an overall *expansion* in M&A activity, but rather a *reallocation* towards acquisitions of AI firms. Panel C of Table 4 further confirms this pattern by documenting a significant increase in the share of deals involving AI targets among all acquisitions undertaken by acquirers that are AI firms. Overall,

[7] In the Internet Appendix, we provide additional pre-trends analysis in Section B.

[8] Our baseline sample excludes corporate divestitures, corresponding to 46 firm-quarter observations out of 48,438. Including them leaves all coefficient estimates virtually unchanged.

these results are distinct from general effects of AI on investment and point to a new compositional shift in M&A activity toward AI-related targets, indicative of a strategic reallocation.

5.2 How AI Reallocates Acquisitions Across Target Types

In this section, we ask which firms AI-exposed acquirers buy, and use the answer to identify the economic mechanism behind the reallocation documented. We organize the analysis around the make-versus-buy margin, and derive three predictions from it: that acquisitions reallocate toward AI-*producing* targets (Prediction 1), toward young and small targets (Prediction 2), and that AI-*adopting* acquirers, those least able to build generative AI in-house, turn to AI-producing targets (Prediction 3). We test each prediction in turn.

A firm that wants to add a capability can either build it internally or acquire it from another firm. This make-versus-buy margin is central to a large literature on the boundaries of the firm and the role of M&A in reallocating innovation: cash-rich firms with weak internal innovation pipelines optimally outsource R&D to smaller innovators and then acquire the winners rather than win an in-house R&D race (Phillips and Zhdanov, 2013); acquirers and targets sort on *complementary innovation characteristics* rather than on scale per se, with low-R&D, patent-rich firms buying high-R&D innovators (Bena and Li, 2014); and whether a firm integrates a supplier of innovation through acquisition depends on the stage and appropriability of that innovation (Frésard et al., 2020). The common prediction is that, when in-house development becomes slow or costly relative to acquisition, firms should substitute away from *making* and toward *buying* the capability.

The release of ChatGPT offers a clean test of this margin. Generative AI capabilities—frontier models, the scarce talent that builds them, and the proprietary data and computational resources they require—cannot be assembled in-house within a short window. The sudden, unanticipated rise in the returns to these capabilities should therefore make *buying* sharply more attractive than *making* for firms whose internal-development option is slowest and most expensive. We preserve the core identification strategy and treatment status from before, adjusting only the matching procedure to incorporate the pre-shock covariates relevant to each dimension.

Acquisitions reallocate toward AI-*producing* targets. We first study acquisitions involving AI-*producing* and AI-*adopting* targets. This allows us to provide novel evidence on whether post-shock acquisitions are reallocated toward AI-*producing* or AI-*adopting* targets. Make-versus-buy

implies that AI-exposed acquirers buy the capability they cannot quickly build—the technology held by AI-*producing* firms—rather than the complementary assets of firms that merely *adopt* AI.

Consistent with this prediction, Panel A of Table 5 shows an increase in both the probability and the number of deals involving *AI-producing* targets after the release of ChatGPT. The probability that an AI-firm acquirer engages in a deal with an *AI-producing* target rises by 56.43% relative to the baseline.[9] By contrast, Panel B shows no significant effect on deals involving *AI-adopting* targets. This contrast is exactly what the make-versus-buy mechanism predicts and is hard to reconcile with a generic increase in dealmaking: the reallocation is concentrated in targets that *supply* the scarce, hard-to-build capability, not in those that simply use it. It also echoes Frésard et al. (2020), in which the firms acquired through integration are those holding the relevant, appropriable innovation.

Acquisitions reallocate toward young, small targets. The framework further implies that acquirers buy *capability rather than scale*: the targets should be the small, young innovators whose value is the technology and human capital they embody, precisely the firms Phillips and Zhdanov (2013) show large firms acquire instead of out-innovating, and the high-R&D innovators that Bena and Li (2014) identify as targets. To test this idea, we use target age to classify firms in our sample as "young" or "old", and target employment to classify firms as "small" or "large". Firms with an age (number of employees) below the median are classified as "young" ("small"), while firms with an age (number of employees) above the median are classified as "old" ("large").[10]

Panel C of Table 5 documents a strong increase in acquisitions of young targets: AI-firm acquirers have a 25.9% higher probability of engaging in a deal with a young target after the shock. Panels D and E show similar results for small targets and for targets that are both small and young. Consistent with acquirers buying capability rather than capacity, we again find no effect on deal value. This null is itself informative: a neoclassical scale or demand shock would predict *larger* transactions as assets move to higher-productivity owners, whereas the make-versus-buy mechanism predicts small deals in which acquirers pay to internalize specialized capabilities embodied in small teams. Our evidence favors the latter. Further, these findings provide the first direct evidence that AI leads firms to reorganize how they acquire

[9]The results are not driven by a handful of mega-acquirers: excluding the 10 or 20 most active buyers leaves the probability and count estimates virtually unchanged.

[10]Qualitatively identical results obtain if, instead of the median split above, we define young/old and small/large targets using the top and bottom terciles of target age and employment (see Section 6 for details).

innovation, rather than merely how much they invest or acquire.

AI-*adopting* acquirers drive into AI-producing targets. AI-*adopting* firms use AI but do not build it in-house, so they have fewer internal-development options: if the returns to AI capability jump, the only fast way to add that capability is to buy it. The mechanism we propose predicts that AI-adopting acquirers, in particular, should increase their acquisitions of AI-producing targets after the shock. We test this by interacting the treatment with a pre-shock indicator for whether the acquirer is itself an AI adopter or an AI producer.[11]

Table 6 shows that AI-adopting acquirers significantly increase acquisitions of AI-producing targets following the shock, consistent with firms that cannot build generative AI internally turning to the acquisition market to obtain it. For example, the probability of engaging in a deal involving an AI-producing target rises by 0.7 percentage points, a 61% increase relative to the pre-shock baseline of 1.15%. AI-producing acquirers also increase such acquisitions by similar magnitudes, which is expected as even firms that produce AI buy the complementary capabilities they do not already possess, and acquisition is faster than scaling internal development within the short window after the shock. In both cases, acquisitions of AI-*adopting* targets are unaffected, confirming that the reallocation is specific to targets that supply the scarce capability. Overall, these findings support the make-versus-buy mechanism that the firms least able to build AI in-house are pulled into the market for AI-producing targets.

The buy response is concentrated among adopters least able to build in-house. Of course, AI-producing acquirers also increase their acquisitions, so the adopter/producer split alone does not isolate the make-versus-buy channel. We therefore sharpen the test *within* adopters, asking whether the buy response is weaker among those better able to build AI in-house. We proxy build capability with pre-shock patenting and estimate a triple difference, $\text{Post}_t \times \text{AI Adopter Buyer}_i \times \log(1 + \text{Pre-Shock Patents}_i)$, on the sample of patent-filing AI-adopter (non-producing) and non-AI buyers. We restrict the test to adopters because an adopter's innovative capability captures its ability to develop AI internally. The make-versus-buy

[11] The make-versus-buy literature typically proxies a firm's in-house innovative capability with R&D intensity. We do not observe R&D for the firms in our sample as our main dataset comes from PitchBook. We instead measure the make margin directly through our producing-versus-adopting taxonomy, which we view as better suited to the question at hand: R&D intensity captures generic innovative effort, whereas the producer/adopter distinction isolates the specific capability the make-versus-buy margin turns on here—whether a firm can develop generative AI internally.

reasoning predicts that more-capable adopters substitute away from buying.[12]

Table 7 shows that the post-ChatGPT increase in AI acquisitions is significantly smaller for adopters with more pre-shock patents: the triple-interaction coefficient is -0.017 ($p < 0.05$) for the probability of acquiring an AI target and -0.010 ($p < 0.05$) for the number of such deals. The magnitude is economically large: moving from the 25th to the 75th percentile of pre-shock patenting among AI-adopter buyers reduces the post-shock acquisition response by 3.1 percentage points, which is about 40% of the pre-shock baseline acquisition probability of 7.8%. Crucially, the same effect is absent for acquisitions of *non-AI* targets (-0.005 and -0.003, both insignificant). This placebo is important because it shows that if pre-shock patents simply proxied size or a general propensity to acquire, capable firms would pull back on non-AI deals as well. They do not—the dampening is specific to the AI capability that ChatGPT made scarce, exactly as make-versus-buy predicts. Panel B, which splits AI targets into producing and adopting firms, points in the same direction but is noisier once the sample is divided, and we read it as suggestive rather than conclusive.

Overall, this result confirms that the reallocation toward AI targets is driven by the firms with the weakest internal build option, and capability substitutes for buying precisely where buying is the only fast route to AI. This is, to our knowledge, the first direct evidence to date that AI reorganizes *how* firms acquire innovation along the make-versus-buy margin.

5.3 Innovation Outcomes

Having shown that AI-exposed firms reallocate their acquisitions toward AI-producing targets, we now ask whether internalizing AI capabilities translates into more innovation. By raising the returns to AI capability, ChatGPT induces AI-exposed firms to acquire AI-producing targets; the acquired capabilities are then integrated into the buyer's operations; also, because AI is a general-purpose technology, the resulting innovation gains should appear not only in AI-related patenting but across the firm's broader innovation portfolio. In this subsection, we test the idea that firms acquire AI-producing targets to internalize AI capabilities that enhance innovation across a wide set of technologies, *not just* AI technologies.

Acquiring AI producers raises innovation; acquiring AI adopters does not. Table 8 reports triple-difference estimates of post-acquisition patenting, interacting the

[12]The results are qualitatively similar when including AI buyers in general (adopting and producing).

post-shock period, an indicator for AI buyers, and an indicator for buyers that acquire at least one AI target of the relevant type.

Panel A shows that AI buyers acquiring AI-*producing* targets significantly increase their patenting after the shock, and the effects are economically large. The triple-interaction coefficient of 0.538 on non-AI patents implies a 71% increase relative to the pre-shock mean. The effect on AI patents is roughly half as large (a 34% increase), indicating that the innovation response is concentrated in the acquirer's *non-AI* domains rather than in AI itself, consistent with the acquirer deploying the acquired AI as a general-purpose input to advance its core technologies. Panel B shows that acquisitions of AI-*adopting* targets produce no significant change (0.188 and 0.144, both insignificant). This contrast mirrors the rest of our results: it is the buying of AI *technology*, not the buying of firms that merely use AI, that raises subsequent innovation—exactly what the make-versus-buy mechanism implies if acquisitions are a route to capabilities a buyer cannot build in-house.

The gains extend well beyond AI. We next document that the innovation response is not confined to AI patents. In Panel A of Table 8, we show that the increase is roughly twice as *large* for non-AI patents as for AI patents, indicating that acquiring AI capability raises innovation most in the firm's core, non-AI domains. The domain-level estimates in Table 9 sharpen this evidence. We find that among non-AI domains, the gains are concentrated in technologically AI-adjacent fields—Computing and Software increases by 40% and Physics increases by 24%—and are absent in unrelated fields such as Chemistry and Human Necessities. Among AI subtypes, the effect loads on Machine Learning and Neural Networks (which increases by 27%), the core general-purpose layer of the technology—rather than on narrower applications such as computer vision or speech. Acquisitions of AI-adopting targets (Table 10) show the same qualitative pattern but weaker and noisier, consistent with the aggregate nulls in Panel B. Overall, these results are consistent with AI behaving as a general-purpose technology: the acquired capability diffuses into the domains where it complements the firm's existing innovation, raising output broadly rather than only in AI itself.

5.4 Complementary Capabilities

Bowen et al. (2023) show that startups innovating in rapidly evolving (frontier) technology areas tend to substitute for existing technologies and exit by going public, remaining independent, whereas startups in stable areas tend to complement existing technologies and exit by selling

out. In this section, we argue that the acquired AI producers are not disruptive substitutes, but *complementary capability suppliers*. Consistent with Bowen et al. (2023), we find that a positive shock to the returns to AI capability raises incumbents' acquisitions precisely in the complementary segment their theory associates with selling out.

Acquisitions are of complementary capabilities, not substitutes. To separate acquisitions of direct competitors from acquisitions of complementary capabilities, we classify each AI-producing target acquisition into one of three mutually exclusive categories using PitchBook's firm-relationship and vertical-classification data. We label a target a *substitute* if the buyer and target appear in each other's competitor set in PitchBook's *CompanyCompetitorRelation* database; coverage is extensive, with 91.7% of AI-producing acquirers and 76.5% of AI-producing targets appearing in the competitor network at least once. Among non-substitute pairs, we measure relatedness as the Jaccard similarity of the buyer's and target's vertical sets (e.g., SaaS, FinTech, Big Data, Digital Health)—shared verticals divided by total unique verticals across the pair—and classify a non-substitute acquisition as *complementary* if its overlap is above the sample median and *unrelated* if below.

Panel A of Table 11 shows that the post-shock increase is concentrated entirely in *complementary* acquisitions (0.003, $p < 0.05$), with no significant effect for substitutes or unrelated targets. This finding indicates that AI-exposed buyers respond to the AI shock by acquiring complementary AI capabilities, not by absorbing direct competitors. This matters because it shows AI reallocation operates through synergy rather than consolidation: incumbents acquire complementary AI capabilities to expand their innovation frontier, not to remove disruptive entrants.

Acquired targets map to the domains where buyers later innovate. If these are capability-augmenting acquisitions, the targets' technologies should line up with the areas in which buyers subsequently patent. We build a bridge between the two. For each AI-producing target acquired by an AI buyer, we collect its PitchBook verticals, emerging spaces, keywords, and industry classifications, and map these descriptors into the broad patent domains used in our innovation analysis (Computing/Software, Human Necessities, Chemistry, Physics, and others)—for example, SaaS, Generative AI, Cybersecurity, and CloudTech map to Computing/Software, HealthTech to Human Necessities, and AI-powered Drug Discovery to Chemistry. We then form buyer-quarter counts of acquisitions in each target domain.

Panel B of Table 11 shows that the post-shock increase in acquisitions of AI-producing targets is concentrated in the Computing/Software domain (0.004, $p < 0.01$), with a smaller effect in Human Necessities—precisely where these same buyers go on to increase their patenting (Section 5.3). The correspondence between the domains of acquired targets and the domains of subsequent innovation supports a capability-augmentation interpretation: buyers acquire complementary technologies and build on them, rather than acquiring to neutralize competitors.

6. Additional Results and Robustness

This section discusses additional results and rules out alternative channels that might explain our main findings.

Placebo Test. Several macroeconomic events coincided with the period surrounding the ChatGPT release, including the Federal Reserve's tightening cycle beginning in March 2022, compressed tech valuations through 2022–2023, and the enactment of the CHIPS Act and Inflation Reduction Act in summer 2022. To assess whether our results reflect pre-existing divergence driven by these or other contemporaneous forces, we conduct a placebo test in which we reassign the shock date to 2021Q4—one year before the actual release of ChatGPT and re-estimate the baseline difference-in-differences specification using this placebo cutoff. Figure B.1 shows no significant differential change in acquisition activity between AI and non-AI firms at the placebo date, indicating that our results are driven by AI and not other confounding factors.

Target Age, Size, and AI-Producing Status. AI-producing targets are disproportionately young and small. To address whether the age and size results just repackage the producer-target result rather than identify a distinct margin, we estimate a specification that interacts the treatment with targets that are simultaneously AI-producing *and* young, small, or both. Table B.1 reports the results. The post-shock increase is concentrated precisely in young, small, AI-producing targets: the probability and number of such acquisitions rise significantly across all three panels (e.g., 0.005 and 0.004, $p < 0.05$, for young AI-producing targets), while the deal value does not increase. The persistent no effect on deal-value shows the age/size dimension is not merely inherited from producer status but carries independent content, namely that buyers acquire capability rather than scale.

In our baseline specification, we use median splits to define young/old and small/large firms, which raises the concern that our findings reflect the specific threshold rather than a robust pattern across the distribution. To address this, we re-estimate the results using alternative cutoffs: young targets are defined as firms in the lowest tercile of target age, and small targets as firms in the lowest tercile of target employment. Table B.2 shows that our results are qualitatively similar, confirming that our findings are not driven by the choice of median splits.

Alternative AI Patent Definitions. Our baseline patent analysis classifies AI patents using CPC subclasses directly tied to core AI methods (G06N, G06V, G10L, G06F18). We assess the sensitivity of all patent-based results—including the DiD specifications and the triple-DiD with acquisition exposure—to two alternative AI patent definitions. First, we expand the baseline CPC definition to include G06F40 (natural language processing and computational linguistics), which has become increasingly central to AI following the rise of large language models. This expanded definition reclassifies a subset of previously non-AI patents as AI-related, altering both the aggregate AI/non-AI patent counts and the composition of non-AI patent domains. Table B.3 shows that our results are very similar to our baseline findings.

Second, we employ the USPTO Artificial Intelligence Patent Dataset (AIPD 2023) developed by Pairolero et al. (2025), which uses machine-learning classifiers trained on BERT for Patents to identify AI-related inventions from patent text rather than examiner-assigned CPC codes. Because the current AIPD release covers documents published through 2023, we restrict all AIPD-based analyses to observations through 2023Q4. We then re-estimate the patent specifications using the `Predict93`, the USPTO AIPD classification that identifies AI patents with a predicted probability of at least 93%. Table B.4 confirms that results are similar to the baseline under this specification.

Matching Without Pre-Shock AI Acquisition Share. Our baseline matching procedure includes the pre-shock share of deals targeting AI firms as a covariate, which could raise concerns about conditioning on an outcome-related variable. Specifically, this covariate selects control firms that were already acquiring AI targets before the shock, potentially attenuating the estimated treatment effect. To address this, we re-estimate the matching without the pre-shock AI acquisition share and reconstruct the matched sample using only firm size, firm age, pre-shock deal activity, and sector. Table B.5 shows that the results are strongly robust to this alternative specification.

Venture Capital Supply Channel Our baseline results are consistent with a demand-side interpretation in which acquirers actively seek AI capabilities after the ChatGPT shock. However, an alternative explanation is that the post-shock expansion of venture capital funding in AI increased the supply of acquisition targets, mechanically raising the volume of AI-related deals regardless of acquirer demand. To address this, we control for the share of each buyer's quarterly acquisitions involving VC-backed targets. If the results were driven by an influx of VC-funded startups becoming available for acquisition, this control should attenuate the estimated treatment effect. Table B.6 documents that our results remain quantitatively and statistically similar to the baseline, indicating that the post-shock increase in AI-target acquisitions is not explained by differential exposure to the venture-capital supply channel.

Section 174 R&D Amortization. A potential concern is that our results reflect the implementation of Section 174 of the Tax Cuts and Jobs Act rather than the emergence of generative AI. Beginning in tax year 2022, Section 174 required firms to capitalize and amortize domestic research and development expenditures over five years, replacing the prior regime under which these costs could generally be expensed immediately. This change increased the effective tax burden and reduced after-tax cash flows for R&D-intensive firms, particularly in software and technology sectors. Because AI firms tend to be R&D-intensive, industries with greater exposure to Section 174 may have experienced changes in acquisition activity—either as acquirers facing tighter cash flows or as targets becoming more vulnerable—for reasons unrelated to the release of ChatGPT.

To address this concern, we construct an industry-level measure of Section 174 exposure using quarterly Compustat data for publicly traded acquirers. For each firm-quarter prior to the implementation of Section 174 (before 2022Q1), we compute R&D intensity as the ratio of R&D expenditures to total assets. We then aggregate these firm-level measures to the PitchBook industry classification level. Because Compustat covers only public firms, the resulting exposure measures are estimated from the public-acquirer subsample and subsequently assigned to all firms—public and private—within the corresponding PitchBook industry.

We construct two alternative exposure measures using PitchBook industry classifications. Specifically, we compute the median pre-2022 R&D intensity among Compustat firms within each PitchBook industry sector and, alternatively, within each more granular industry group. Both measures are standardized to have mean zero and unit variance. We merge these industry-level exposure measures into our matched sample and augment the baseline difference-in-differences

model with a triple interaction between treatment status, a post-Section 174 indicator (equal to one beginning in 2022Q1), and industry-level Section 174 exposure.

Were our baseline results driven by the R&D amortization shock rather than generative AI, then including the Section 174 interaction should substantially attenuate or eliminate the estimated ChatGPT effect. Table B.7 shows that the ChatGPT coefficient remains stable across all specifications, and the Section 174 interaction is economically small, indicating that differential exposure to the tax treatment of R&D expenditures does not account for the observed increase in AI-related acquisitions.

7. Conclusion

This paper shows how generative AI reshapes the ways in which firms acquire innovation. Exploiting the release of ChatGPT in 2022Q4 as a sudden, unanticipated shock to the returns on AI, together with a matched difference-in-differences design, we find that AI-exposed acquirers reallocate their M&A toward AI-*producing*, young, and small targets, while their non-AI acquisitions do not rise, indicating reallocation rather than an expansion of deal-making. Consistent with a make-versus-buy mechanism, the response is strongest among acquirers least able to build AI in-house. These acquisitions also raise innovation: firms that buy AI producers patent more afterward, with gains extending primarily beyond AI into their core technological domains, consistent with AI serving as a general-purpose technology. Moreover, the acquired producers are complementary capability suppliers rather than substitutes, indicating that AI-driven M&A operates through synergy rather than consolidation.

Overall, our analysis shows how firms use acquisitions to adjust to AI-driven technological change, and its consequences for innovation. The innovation associated with these deals is generated by the acquirer, which deploys the acquired AI as an input to advance its own core technologies, rather than reflecting the patenting of the acquired targets themselves. When the returns to a capability rise faster than firms can build it, the market for corporate control reallocates that capability to the firms able to deploy it most productively—and to generate new innovation from it. These findings suggest that M&A is a central mechanism through which firms, and the broader economy, absorb and build on transformative, general-purpose technologies like AI.

References

Acemoglu, Daron, Akcigit, Ufuk, Alp, Harun, Bloom, Nicholas, and Kerr, William (2018). "Innovation, reallocation, and growth". In: *American Economic Review* 108.11, pp. 3450–3491.

Acemoglu, Daron, Autor, David, Hazell, Jonathon, and Restrepo, Pascual (2022). "Artificial intelligence and jobs: Evidence from online vacancies". In: *Journal of Labor Economics* 40.S1, S293–S340.

Babina, Tania, Fedyk, Anastassia, He, Alex, and Hodson, James (2023). *Firm investments in artificial intelligence technologies and changes in workforce composition*. Vol. 31325. National Bureau of Economic Research.

— (2024). "Artificial intelligence, firm growth, and product innovation". In: *Journal of Financial Economics* 151, p. 103745.

Bena, Jan and Li, Kai (2014). "Corporate innovations and mergers and acquisitions". In: *The Journal of Finance* 69.5, pp. 1923–1960.

Bowen, Donald E, Frésard, Laurent, and Hoberg, Gerard (2023). "Rapidly evolving technologies and startup exits". In: *Management Science* 69.2, pp. 940–967.

Chen, Mark, Wu, Qinxi, and Yang, Baozhong (2019). "How valuable is FinTech innovation?" In: *The Review of financial studies* 32.5, pp. 2062–2106.

D'Acunto, Francesco, Prabhala, Nagpurnanand, and Rossi, Alberto G (2019). "The promises and pitfalls of robo-advising". In: *The Review of Financial Studies* 32.5, pp. 1983–2020.

Eisfeldt, Andrea L, Schubert, Gregor, and Zhang, Miao Ben (2025). "Generative AI and firm values". In: *Forthcoming, Journal of Finance*.

Fang, Yue, Goyal, Vidhan, He, Yingxuan, and Zhang, Zilong (2025). "When AI Acquires Data: Strategic Complementarities in M&A". In: *Available at SSRN 6100206*.

Frésard, Laurent, Hoberg, Gerard, and Phillips, Gordon M (2020). "Innovation activities and integration through vertical acquisitions". In: *The Review of Financial Studies* 33.7, pp. 2937–2976.

Fuster, Andreas, Goldsmith-Pinkham, Paul, Ramadorai, Tarun, and Walther, Ansgar (2022). "Predictably unequal? The effects of machine learning on credit markets". In: *The Journal of Finance* 77.1, pp. 5–47.

Gofman, Michael and Jin, Zhao (2024). "Artificial intelligence, education, and entrepreneurship". In: *The Journal of Finance* 79.1, pp. 631–667.

Gupta, Abhinav, Qian, Franklin, Simintzi, Elena, and Sun, Yifan (2025). "Generative AI and Entrepreneurship". In: *Available at SSRN 5971474*.

Jansen, Mark, Nguyen, Hieu Quang, and Shams, Amin (2025). "Rise of the machines: The impact of automated underwriting". In: *Management Science* 71.2, pp. 955–975.

Jovanovic, Boyan and Rousseau, Peter L (2002). "The Q-theory of mergers". In: *American Economic Review* 92.2, pp. 198–204.

Lerner, Josh and Seru, Amit (2022). "The Use and Misuse of Patent Data: Issues for Finance and Beyond". In: *Review of Financial Studies* 35.6, pp. 2667–2704.

Ma, Wenting, Ouimet, Paige, and Simintzi, Elena (2025). "Mergers and acquisitions, technological change, and inequality". In: *Journal of Financial Economics* 172, p. 104136.

Maksimovic, Vojislav and Phillips, Gordon (2001). "The market for corporate assets: Who engages in mergers and asset sales and are there efficiency gains?" In: *The journal of finance* 56.6, pp. 2019–2065.

— (2002). "Do conglomerate firms allocate resources inefficiently across industries? Theory and evidence". In: *The Journal of Finance* 57.2, pp. 721–767.

Marchesi, Edoardo and Tang, Sixun (2025). "When AI Meets Entrepreneurship: Evidence from the Commercialization of ChatGPT". In: *Available at SSRN 5759264*.

Mitchell, Mark L and Mulherin, J Harold (1996). "The impact of industry shocks on takeover and restructuring activity". In: *Journal of financial economics* 41.2, pp. 193–229.

Pairolero, Nicholas A., Giczy, Anthony V., Torres, Guillermo, Islam Erana, Tasmia, Finlayson, Mark A., and Toole, Andrew A. (2025). "The Artificial Intelligence Patent Dataset (AIPD) 2023 Update". In: *Journal of Technology Transfer*.

Phillips, Gordon M and Zhdanov, Alexei (2013). "R&D and the incentives from merger and acquisition activity". In: *The Review of Financial Studies* 26.1, pp. 34–78.

Schoar, Antoinette (2002). "Effects of corporate diversification on productivity". In: *The Journal of Finance* 57.6, pp. 2379–2403.

Yang, Chih-Hai (2022). "How artificial intelligence technology affects productivity and employment: Firm-level evidence from Taiwan". In: *Research Policy* 51.6, p. 104536.

Figures and Tables

Figure 1: Quarterly M&A Activity by Acquirer and Target AI Status

This figure plots quarterly M&A activity by acquirer and target AI status. Panel A reports the total number of completed acquisitions. Panel B restricts attention to acquisitions of AI targets, while Panel C reports acquisitions of non-AI targets. Acquirers are classified as acquirers that are AI firms (either AI-producing, AI-adopting, or both) or non-AI firms, following the classification described in Section 3. Solid lines correspond to acquirers that are AI firms, while dotted lines correspond to non-AI acquirers. Targets are classified as AI or non-AI in a similar way. The figure illustrates the evolution of acquisition activity before and after the AI shock in 2022Q4. The vertical dashed line denotes the release of ChatGPT in 2022Q4, which serves as the event date in the subsequent difference-in-differences analysis.

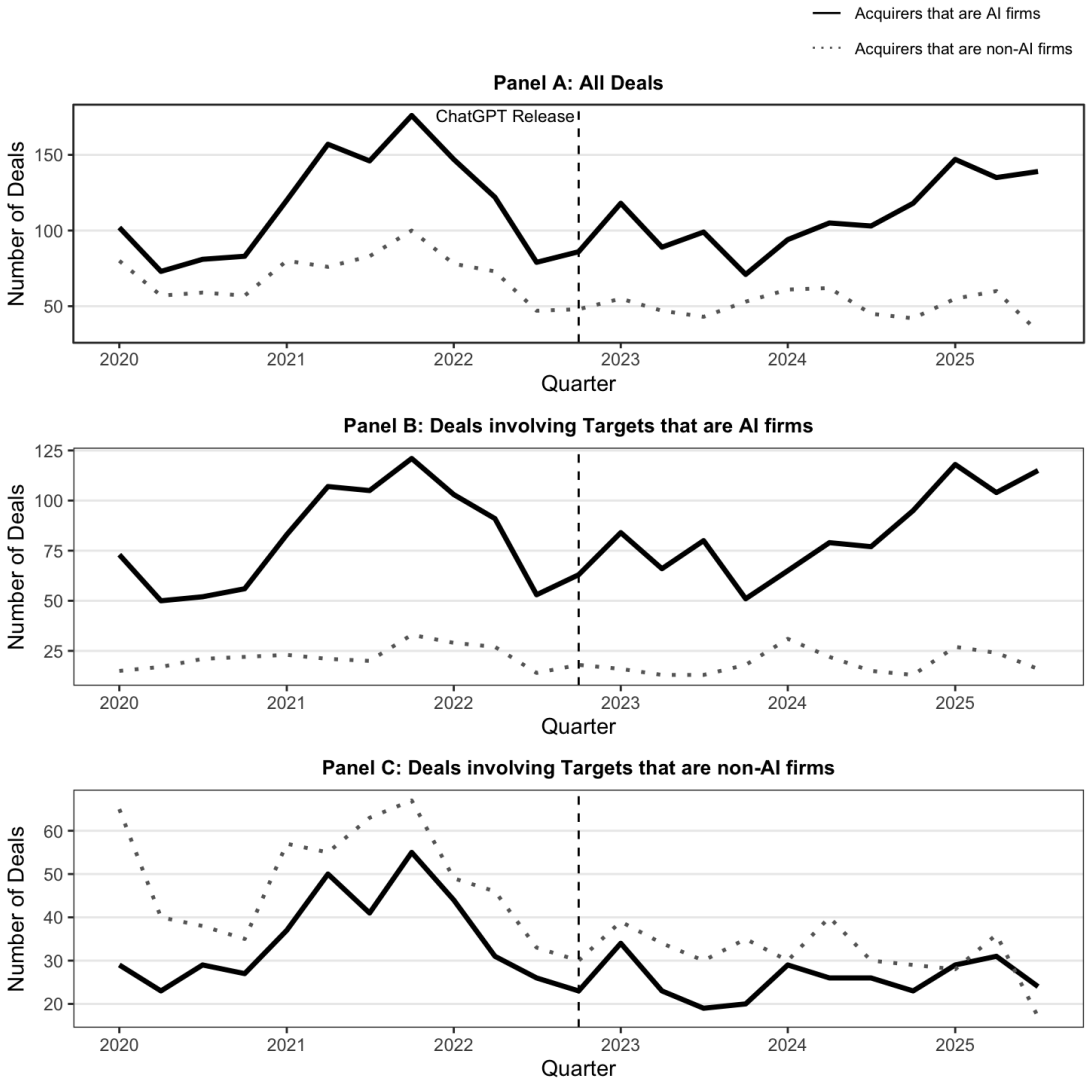


Figure 2: M&A Activity by Buyer and Target AI Type

This figure plots quarterly M&A activity by buyer AI type and target AI type. Panel A focuses on buyers that produce AI, while Panel B focuses on buyers that adopt AI as their input, following the classification in Section 3. Within each panel, solid lines denote acquisitions of AI-producing targets and dotted lines denote acquisitions of AI-adopting targets. The figure illustrates the evolution of acquisition patterns before and after the AI shock in 2022Q4. The vertical dashed line denotes the release of ChatGPT in 2022 Q4, which serves as the event date in the subsequent difference-in-differences analysis.

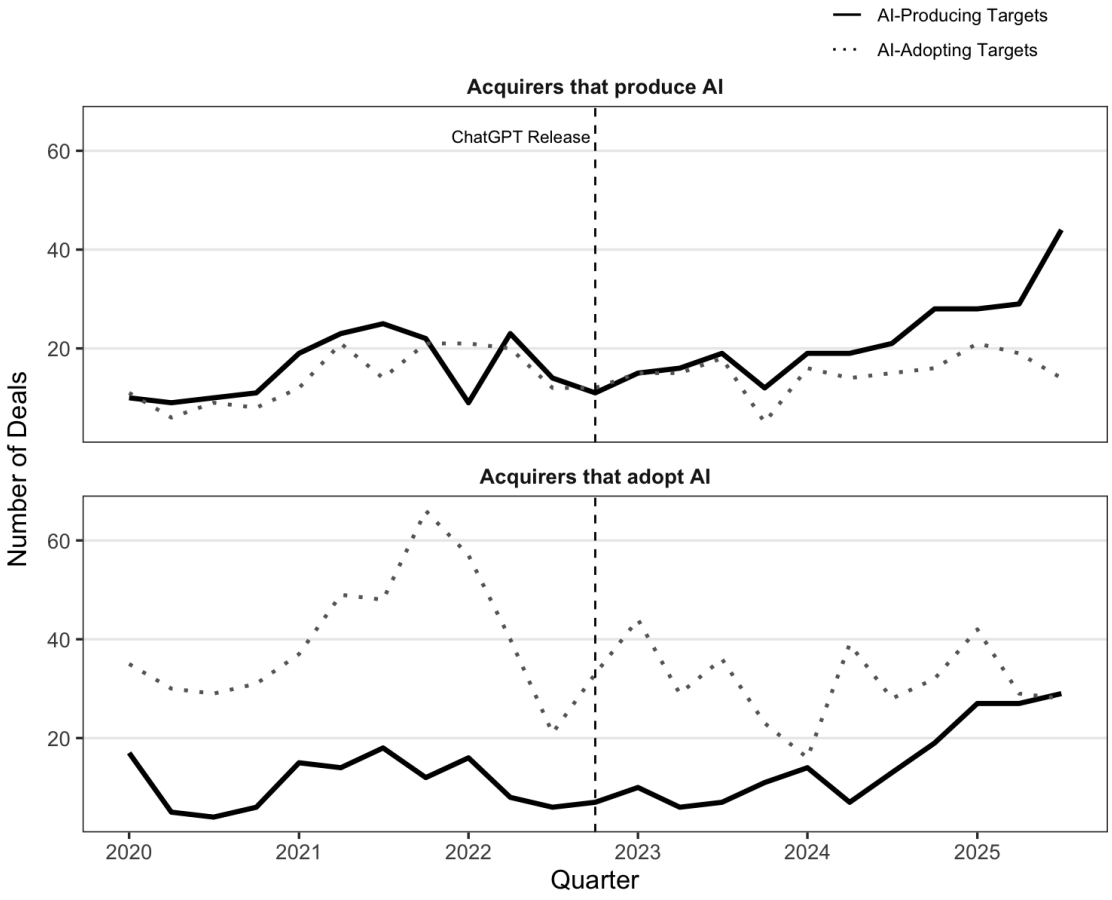


Figure 3: Dynamic Effects of AI on M&A Activity

This figure plots event-study estimates of acquisition activity following the AI shock in 2022Q4, using Equation (2), and the matched sample. The matched sample is constructed by pairing each acquirer that is an AI firm with the closest acquirer that is a non-AI acquirer based on pre-shock firm characteristics, using nearest-neighbor matching with replacement and exact matching on the acquirer's primary industry sector. Panel A reports the event-study estimates of the probability that an acquirer that is an AI firm completes any AI-related acquisitions, Panel B reports the estimates of the probability that an acquirer that is an AI firm acquires an AI-producing target, Panel C reports the estimates of the probability that an acquirer that is an AI firm acquires a young target, Panel D reports the estimates of the probability that an acquirer that is an AI firm acquires a small target, and Panel E reports the estimates of the probability that an acquirer that is an AI firm acquires a target that is both small and young. Coefficients are plotted relative to the quarter immediately preceding the AI shock ($k = -1$), which serves as the reference period. All specifications are estimated at the acquirer-quarter level and include acquirer fixed effects and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the acquirer level, and error bars represent 95% confidence intervals.

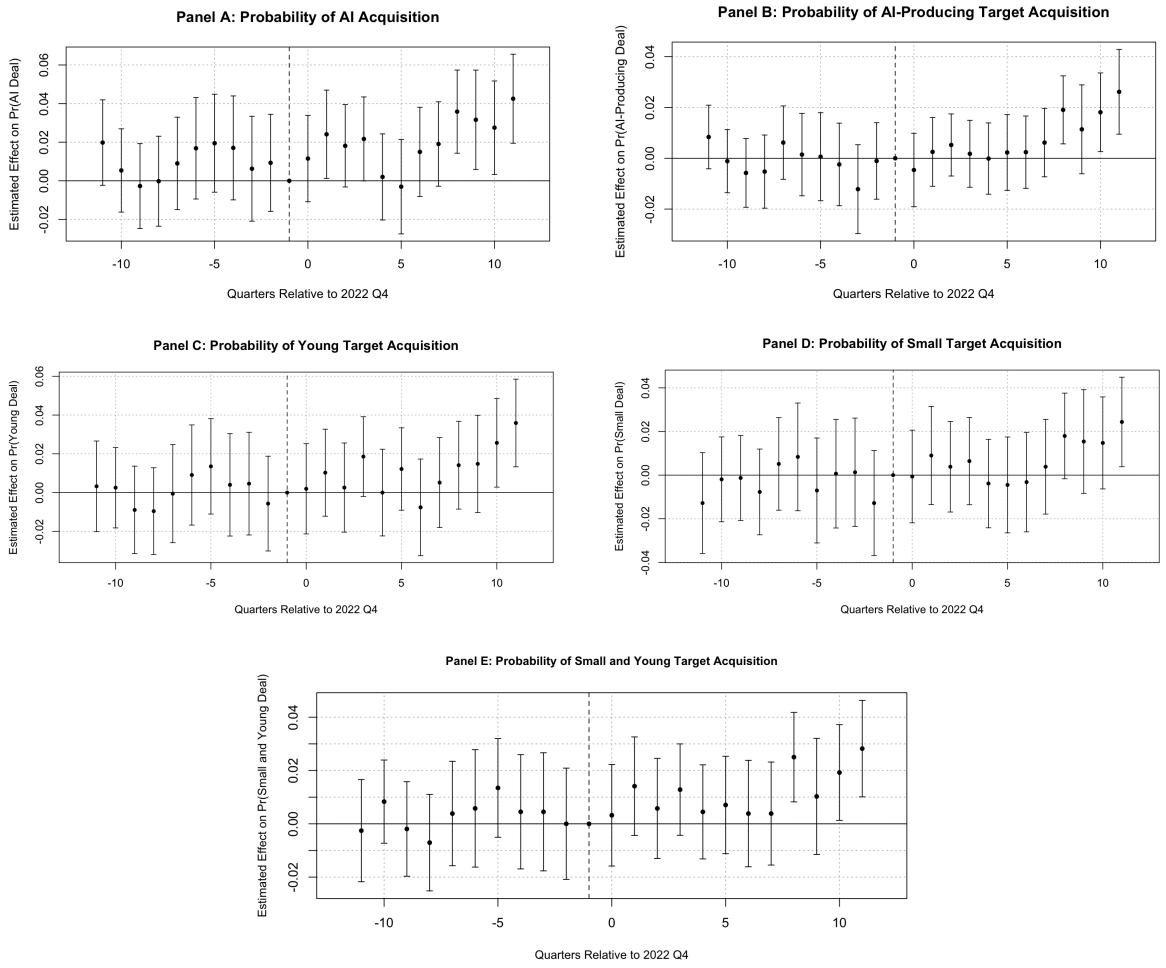


Figure 4: Dynamic Effects of AI on Acquisitions of AI-Producing Targets by Target Characteristics

This figure plots event-study estimates of acquisition activity following the AI shock in 2022Q4, using Equation (2), and the matched sample. The matched sample is constructed by pairing each acquirer that is an AI firm with the closest acquirer that is a non-AI acquirer based on pre-shock firm characteristics, using nearest-neighbor matching with replacement and exact matching on the acquirer's primary industry sector. Panel A reports the event-study estimates of the probability that an acquirer that is an AI firm acquires a small AI-producing target, Panel B reports the estimates of the probability that an acquirer that is an AI firm acquires a young AI-producing target, and Panel C reports the estimates of the probability that an acquirer that is an AI firm acquires a target that is simultaneously small, young, and AI-producing. Small and young targets are defined using the median values of target employment and target age, respectively. Coefficients are plotted relative to the quarter immediately preceding the AI shock ($k = -1$), which serves as the reference period. All specifications are estimated at the acquirer-quarter level and include acquirer fixed effects and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the acquirer level, and error bars represent 95% confidence intervals.

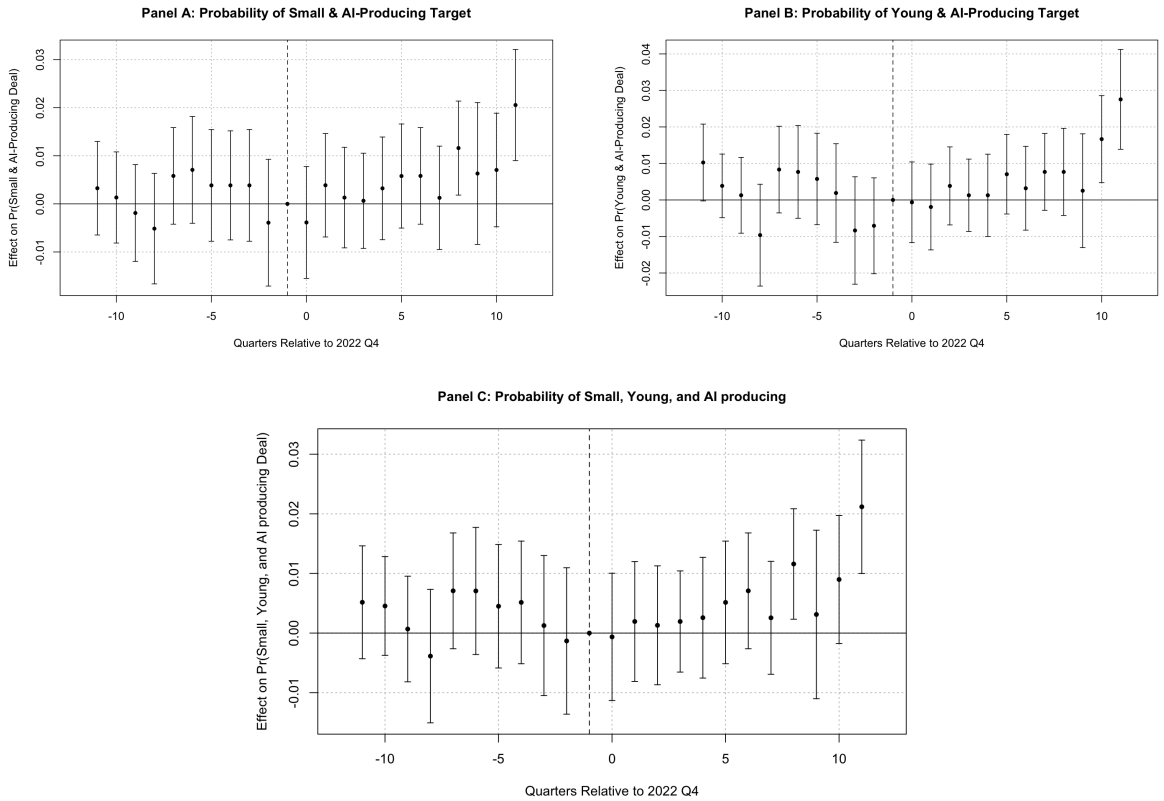


Table 1: Descriptive Statistics by Acquirer Type

This table reports descriptive statistics for firms classified by acquirer type. Firms are classified based on their AI exposure as AI producers, AI adopters, or non-AI firms using PitchBook verticals and emerging-space classifications, as described in Section 3. Panel A reports characteristics measured at the firm level, including firm age, public listing status, employment, revenue, and earnings. Panel B reports acquisition activity at the firm-quarter level, including the probability of conducting any deal and the probability of conducting a deal in which the target firm is an AI firm (AI deal), with a breakdown before and after the AI shock in 2022Q4. Panel C reports deal-level characteristics conditional on an acquisition occurring, including deal size and target AI status.

	Acquirers that are non-AI firms (Control)	**Acquirers that are AI firms** (Treated)
Panel A: Firm Characteristics (Unit: Firm)		
Number of Firms	993	1,617
Publicly Traded (%)	32.4%	14.0%
Firm Age (Years)	11.9	7.77
Employees (Mean)	5,261	2,433
Revenue (Mean $M)	3,716	2,284
EBIT (Mean $M)	1,237	1,829
Panel B: Acquisition Activity (Unit: Firm-Quarter)		
Number of Firm-Quarters	22,839	37,191
Pr(Any Deal)	5.85%	6.5%
Pre-GPT Period	6.85%	6.65%
Post-GPT Period	4.94%	6.36%
Pr(AI Deal)	2.02%	4.82%
Mean Deals per Quarter	0.061	0.070
Panel C: Deal Characteristics (Unit: Deal)		
Number of Deals	1,399	2,593
Mean Deal Size (Log)	4.33	4.38
Target is AI Company (%)	33.7%	72.7%
Target is AI Producer (%)	15.1%	28.3%

Table 2: M&A Deal Counts by Acquirers based on Target Type

This table reports the number of M&A deals by acquirer type and target type before and after the release of ChatGPT in 2022Q4. Acquirers are classified based on their AI exposure as AI producers, AI adopters, or non-AI firms using PitchBooks' information on verticals and emerging spaces described in Section 3. Targets are classified as AI targets or non-AI targets, with AI targets further divided into AI adopters and AI producers according to the same classification. The After Shock period refers to deals occurring in 2022Q4 and later, while the Before Shock period refers to deals occurring prior to 2022Q4.

Acquirer Type	**Target Type**	**After Shock**	**Before Shock**
Acquirers that are AI firms	AI Target (Total)	996	890
	AI Adopter	559	594
	AI Producer	437	296
	Non-AI Target	315	392
Acquirers that are non-AI firms	AI Target (Total)	224	247
	AI Adopter	116	144
	AI Producer	108	103
	Non-AI Target	381	547

Table 3: Balance Table: Unmatched vs. Matched Sample

This table reports the characteristics of treated and control groups before and after the matching procedure. Treated firms are acquirers that are AI firms (either AI producers, AI adopters, or both), while non-treated firms are non-AI firms. Firms are classified based on their AI exposure as AI producers, AI adopters, or non-AI firms using PitchBooks' information on verticals and emerging spaces as described in Section 3. The unmatched sample includes all treated and non-treated firms that satisfy the sample selection criteria. The matched sample is constructed by selecting, for each treated firm, the closest non-treated firm using nearest-neighbor matching based on Mahalanobis distance, with replacement, and exact matching on the firm's primary industry sector. We report mean values of four pre-shock firm characteristics: the total number of acquisitions, the log of total acquisition value, the log of employment, and the share of AI-related acquisitions. For each characteristic, the table reports the mean for treated and control firms, the difference in means, and the p-value.

	Unmatched Sample				**Matched Sample**			
Variable	Treated	Control	Diff.	p-value	Treated	Control	Diff.	p-value
Pre Total Count	0.797	0.808	-0.011	0.799	0.797	0.729	0.068	0.252
Pre Log Value	0.724	1.101	-0.377	0.000	0.724	0.667	0.056	0.545
Pre Log Employment	5.273	5.316	-0.042	0.614	5.273	5.208	0.065	0.485
Pre AI Share	0.375	0.170	0.204	0.000	0.375	0.365	0.009	0.694

Table 4: The Impact of AI on M&A Activity

This table reports difference-in-differences estimates of acquisition activity following the AI shock in 2022Q4, using Equation (1) and the matched sample. The matched sample pairs each AI-exposed acquirer with the closest non-AI acquirer based on pre-shock firm characteristics, using nearest-neighbor matching with replacement and exact matching on the acquirer's primary industry sector. The main explanatory variable is the interaction between an indicator for the post-shock period (2022Q4 and later) and an indicator for AI-exposed acquirers. Panel A reports outcomes for acquisitions of AI targets, Panel B for acquisitions of non-AI targets, and Panel C reports strategic reallocation toward AI targets, measured as differences between AI and non-AI activity. All specifications are estimated at the acquirer-quarter level and include acquirer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the acquirer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Acquisitions of AI Targets			
	Prob(AI Target Deal)	log(AI Target Count)	log(AI Target Deal Value)
Post × AI-exposed Acquirer	0.012***	0.008***	0.008
	(0.004)	(0.003)	(0.010)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel B: Acquisitions of Non-AI Targets			
	Prob(Non-AI Target Deal)	log(Non-AI Target Count)	log(Non-AI Target Deal Value)
Post × AI-exposed Acquirer	0.001	0.000	−0.009
	(0.004)	(0.003)	(0.010)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel C: Strategic Reallocation Toward AI			
	Share of AI	log Count(AI) − log Count(Non-AI)	log Value(AI) − log Value(Non-AI)
Post × AI-exposed Acquirer	0.012***	0.008**	0.016
	(0.004)	(0.004)	(0.013)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Table 5: The Impact of AI on M&A Activity across Target Types

This table reports difference-in-differences estimates of acquisition activity following the AI shock in 2022Q4, using Equation (1) and the matched sample. The matched sample pairs each AI-exposed acquirer with the closest non-AI acquirer based on pre-shock characteristics, using nearest-neighbor matching with replacement and exact matching on the acquirer's primary industry sector. The main explanatory variable is the interaction between an indicator for the post-shock period and an indicator for AI-exposed acquirers. Panels A and B focus on AI-producing and AI-adopting targets, respectively. Panel C reports outcomes for young targets, Panel D for small targets, and Panel E for targets that are both young and small. All specifications are estimated at the acquirer-quarter level and include acquirer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the acquirer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: AI-Producing Target M&A Outcomes			
	Prob(AI Producer Deal)	log(AI Producer Count)	log(AI Producer Deal Value)
Post × AI-exposed Acquirer	0.009***	0.006***	0.012*
	(0.003)	(0.002)	(0.006)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel B: AI-Adopting Target M&A Outcomes			
	Prob(AI Adopter Deal)	log(AI Adopter Count)	log(AI Adopter Deal Value)
Post × AI-exposed Acquirer	0.003	0.002	−0.005
	(0.003)	(0.002)	(0.008)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel C: Young Target M&A Outcomes			
	Prob(Young Deal)	log(Young Deal Count)	log(Young Deal Value)
Post × AI-exposed Acquirer	0.010**	0.008***	−0.009
	(0.004)	(0.003)	(0.008)
Num. Obs.	48,001	48,001	48,001
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel D: Small Target M&A Outcomes			
	Prob(Small Deal)	log(Small Deal Count)	log(Small Deal Value)
Post × AI-exposed Acquirer	0.009**	0.007***	−0.002
	(0.004)	(0.003)	(0.005)
Num. Obs.	47,955	47,955	47,955
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel E: Young & Small Target M&A Outcomes			
	Prob(Young & Small Deal)	log(Young & Small Count)	log(Young & Small Deal Value)
Post × AI-exposed Acquirer	0.009***	0.007***	−0.001
	(0.003)	(0.002)	(0.004)
Num. Obs.	47,955	47,955	47,955
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Table 6: AI Buyer Type and Acquisitions of AI-Producing and AI-Adopting Targets

This table reports difference-in-differences estimates of acquisition activity by buyer type following the release of ChatGPT. The main explanatory variable is the interaction between an indicator for the post-shock period and an indicator for the relevant AI buyer type. Panels A and B report acquisition outcomes for the subsample of AI-adopter buyers, while Panels C and D report acquisition outcomes for the subsample of AI-producer buyers. Panels A and C focus on acquisitions of AI-producing targets, whereas Panels B and D focus on acquisitions of AI-adopting targets. All specifications are estimated at the acquirer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the firm level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: AI-Adopter Buyers — AI-Producing Target Outcomes

	Prob(Producer Target)	log(Producer Count)	log(Producer Value)
Post × AI Adopter Buyer	0.007**	0.005***	0.010
	(0.003)	(0.002)	(0.007)
Num. Obs.	30,337	30,337	30,337
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Panel B: AI-Adopter Buyers — AI-Adopting Target Outcomes

	Prob(Adopter Target)	log(Adopter Count)	log(Adopter Value)
Post × AI Adopter Buyer	0.003	0.002	−0.007
	(0.004)	(0.003)	(0.011)
Num. Obs.	30,337	30,337	30,337
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Panel C: AI-Producer Buyers — AI-Producing Target Outcomes

	Prob(Producer Target)	log(Producer Count)	log(Producer Value)
Post × AI Producer Buyer	0.014**	0.010**	0.010
	(0.006)	(0.005)	(0.015)
Num. Obs.	9,154	9,154	9,154
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Panel D: AI-Producer Buyers — AI-Adopting Target Outcomes

	Prob(Adopter Target)	log(Adopter Count)	log(Adopter Value)
Post × AI Producer Buyer	−0.002	−0.002	0.006
	(0.005)	(0.004)	(0.015)
Num. Obs.	9,154	9,154	9,154
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Table 7: Pre-Shock Innovation Intensity and AI-Adopter Acquisition Responses

This table reports triple-difference estimates testing whether the post-ChatGPT acquisition response among AI-adopter buyers varies with pre-shock innovation intensity. The sample is restricted to patent-filing firms that are AI adopters but not AI producers or non-AI firms; AI-producing firms are excluded. The treatment group consists of AI-adopter-only buyers, while non-AI firms serve as the control group. Pre-shock innovation intensity is measured as $\log(1 + \text{Pre-Shock Patents})$, where pre-shock patents are cumulative patents filed before 2022Q4. Panel A compares acquisitions of AI and non-AI targets, while Panel B separates acquisitions of AI-producing and AI-adopting targets. All specifications are estimated at the buyer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the buyer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: AI and Non-AI Targets				
	Prob(AI Target)	log(AI Target Count)	Prob(Non-AI Target)	log(Non-AI Target Count)
Post × AI Adopter Buyer × log(1 + Pre-Shock Patents)	-0.017**	-0.010**	-0.005	-0.003
	(0.007)	(0.005)	(0.006)	(0.004)
Num. Obs.	11,868	11,868	11,868	11,868
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes
Sample	Patent-filing firms	Patent-filing firms	Patent-filing firms	Patent-filing firms
Panel B: AI-Producing and AI-Adopting Targets				
	Prob(AI Prod Target)	log(AI Prod Count)	Prob(AI Adopt Target)	log(AI Adopt Count)
Post × AI Adopter Buyer × log(1 + Pre-Shock Patents)	-0.007*	-0.003	-0.008	-0.007*
	(0.004)	(0.003)	(0.006)	(0.004)
Num. Obs.	11,868	11,868	11,868	11,868
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes
Sample	Patent-filing firms	Patent-filing firms	Patent-filing firms	Patent-filing firms

Table 8: AI Target Acquisitions and Patent Reallocation

This table reports triple-difference estimates of the effect of acquiring AI targets on patenting activity following the AI shock in 2022Q4. The sample is restricted to the 711 acquiring firms that can be matched to at least one patent in the USPTO data. Panel A examines acquisitions of AI-producing targets, while Panel B examines acquisitions of AI-adopting targets. The main explanatory variable is the triple interaction between an indicator for the post-shock period (2022Q4 and later), an indicator for AI buyers, and an indicator for buyers that acquire at least one AI target of the relevant type after the shock. The dependent variables are log(1+Non-AI Patents), log(1+AI Patents), and the difference between the two measures. All specifications are estimated at the buyer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the buyer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: AI Producer Acquisitions			
	log(Non-AI Patents)	log(AI Patents)	log(Non-AI) minus log(AI)
Post × AI Buyer × AI Producer Acquisition	0.538***	0.294***	0.243**
	(0.150)	(0.105)	(0.108)
Num. Obs.	16,353	16,353	16,353
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Sample	Patent-filing firms	Patent-filing firms	Patent-filing firms
Panel B: AI Adopter Acquisitions			
	log(Non-AI Patents)	log(AI Patents)	log(Non-AI) minus log(AI)
Post × AI Buyer × AI Adopter Acquisition	0.188	0.144	0.045
	(0.153)	(0.098)	(0.098)
Num. Obs.	16,353	16,353	16,353
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Sample	Patent-filing firms	Patent-filing firms	Patent-filing firms

Table 9: AI Producer Acquisitions and Patent Reallocation by Domain

This table reports triple-difference estimates examining how acquisitions of AI-producing targets affect the composition of subsequent innovation activity following the AI shock in 2022Q4. The sample is restricted to the 711 acquiring firms that can be matched to at least one patent in the USPTO data. The main explanatory variable is the triple interaction between an indicator for the post-shock period (2022Q4 and later), an indicator for AI buyers, and an indicator for buyers that acquire at least one AI-producing target after the shock. Panel A reports patenting outcomes across major non-AI patent domains, while Panel B reports outcomes across AI patent subtypes. All specifications are estimated at the buyer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the buyer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Non-AI Patent Domains					
	Human Necessities	Chemistry	Computing/Software	Business Methods	Physics
Post × AI Buyer × AI Producer Acquisition	0.064	0.052	0.339**	0.084	0.217***
	(0.065)	(0.070)	(0.134)	(0.057)	(0.082)
Num. Obs.	16,353	16,353	16,353	16,353	16,353
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes
Panel B: AI Patent Subtypes					
	ML/Neural Nets	Computer Vision	Speech/NLP	Language Processing	Pattern Recognition
Post × AI Buyer × AI Producer Acquisition	0.239***	0.042	0.027	0.077	0.071
	(0.088)	(0.066)	(0.044)	(0.050)	(0.047)
Num. Obs.	16,353	16,353	16,353	16,353	16,353
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes

Table 10: AI Adopter Acquisitions and Patent Reallocation by Domain

This table reports triple-difference estimates examining how acquisitions of AI-adopting targets affect the composition of subsequent innovation activity following the AI shock in 2022Q4. The sample is restricted to the 711 acquiring firms that can be matched to at least one patent in the USPTO data. The main explanatory variable is the triple interaction between an indicator for the post-shock period (2022Q4 and later), an indicator for AI buyers, and an indicator for buyers that acquire at least one AI-adopting target after the shock. Panel A reports patenting outcomes across major non-AI patent domains, while Panel B reports outcomes across AI patent subtypes. All specifications are estimated at the buyer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the buyer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Non-AI Patent Domains					
	Human Necessities	Chemistry	Computing/Software	Business Methods	Physics
Post × AI Buyer × AI Adopter Acquisition	0.030	-0.001	0.265**	0.050	0.135*
	(0.061)	(0.046)	(0.126)	(0.055)	(0.076)
Num. Obs.	16,353	16,353	16,353	16,353	16,353
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes
Panel B: AI Patent Subtypes					
	ML/Neural Nets	Computer Vision	Speech/NLP	Language Processing	Pattern Recognition
Post × AI Buyer × AI Adopter Acquisition	0.140*	-0.006	0.022	0.045	0.047
	(0.084)	(0.053)	(0.041)	(0.050)	(0.044)
Num. Obs.	16,353	16,353	16,353	16,353	16,353
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes

Table 11: Motives and Patent Domains of AI-Producing Target Acquisitions

This table reports difference-in-differences estimates examining the characteristics of AI-producing targets acquired by AI buyers following the AI shock in 2022Q4. The sample consists of buyer-quarter observations from the matched sample. The main explanatory variable is the interaction between an indicator for the post-shock period (2022Q4 and later) and an indicator for AI buyers. Panel A classifies AI-producing target acquisitions according to their economic relationship with the acquirer. A target is classified as a substitute if the buyer and target are identified as direct competitors in PitchBook's competitor network. Among non-competitor acquisitions, a target is classified as a complement if the Jaccard similarity of buyer and target vertical classifications is above the sample median and as unrelated otherwise. Panel B classifies AI-producing target acquisitions according to the technological domain of the target. To construct the mapping, we use target verticals, emerging spaces, keywords, and industry classifications to assign targets to broad patent-relevant domains, including Computing/Software, Human Necessities, Chemistry, and Physics. The dependent variables measure the probability of acquiring an AI-producing target belonging to the indicated category. All specifications are estimated at the buyer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the buyer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Economic Relationship of AI-Producing Targets

	Prob(Substitute)	Prob(Complement)	Prob(Unrelated)
Post × AI Buyer	0.001	0.003**	0.002
	(0.000)	(0.001)	(0.001)
Num. Obs.	60,030	60,030	60,030
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Panel B: Technological Domain of AI-Producing Targets

	Computing/ Software Targets	Human Necessities Targets	Chemistry Targets	Physics Targets
Post × AI Buyer	0.004***	0.001**	0.000	−0.000
	(0.001)	(0.000)	(0.000)	(0.000)
Num. Obs.	60,030	60,030	60,030	60,030
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes	Yes / Yes

Internet Appendix
for "Buy or Build? How AI Shapes Corporate Acquisitions and Innovation"

This Internet Appendix is organized as follows. Section A provides detailed information on the data construction and the classification of firms in our sample. Section B reports the results from the Section 6 and additional pre-trends results.

A. Data Construction and Classification of Firms

Appendix Table A.1: Illustrative PitchBook Category Descriptions and Classifications

This table presents two examples of PitchBook category descriptions and their resulting classifications. The *Artificial Intelligence and Machine Learning* vertical is classified as AI-producing because the development of AI and machine learning technologies constitutes the primary activity described. The *Construction Technology* vertical is classified as AI-adopting because AI-related technologies—such as construction robots and automated surveying drones—appear as enabling inputs within a sector whose primary focus is the construction industry. Descriptions are excerpted from PitchBook's publicly available category definitions.

Category	PitchBook Description (Excerpt)	Classification
Artificial Intelligence and Machine Learning	Artificial intelligence (AI) focuses on creating intelligent computers that can perceive their environment and make decisions to maximize the chances of reaching their goals. Machine learning (ML) is a subfield of AI that gives computers the ability to learn iteratively, improve predictive models and find insights from data without being explicitly programmed to do so. AI/ML applications include speech recognition, computer vision, robotic control and accelerating processes in the empirical sciences. Silicon Valley startup SambaNova Systems developed an advanced systems platform and hardware to power machine learning and data analytics via a system architecture that is flexible, scalable and optimized for dataflow.	AI-Producing
Construction Technology	Companies in this space develop hardware and software solutions that increase the efficiency and productivity of construction, building, and infrastructure operations. This includes project management software, construction robots, job site survey drones, building information modeling, augmented and virtual reality applications for construction, modular and pre-fabrication solutions, equipment marketplaces and risk management tools. PlanGrid's software gives real-time updates, blueprints, field reports and project documents enabling construction companies to simply and easily manage their projects.	AI-Adopting

Appendix Table A.2: List of Industry Verticals and Emerging Industry Spaces

This table reports the list of PitchBook verticals and emerging spaces used to classify firms' AI exposure, as described in Section 3.

Panel A: List of Verticals
AdTech, Advanced Manufacturing, AgTech, Artificial Intelligence, AudioTech, Augmented Reality, Autonomous Vehicles, B2B Payments, Beauty, Big Data, Cannabis, Carsharing, CleanTech, Climate Tech, CloudTech & DevOps, Construction, Cryptocurrency, Cybersecurity, Digital Health, E-Commerce, EdTech, Ephemeral Content, Esports, FemTech, FinTech, FoodTech, Gaming, HealthTech, HR Tech, Impact Investing, Industrials, Infrastructure, InsurTech, Internet of Things, Legal Tech, Life Sciences, LOHAS & Wellness, Manufacturing, Marketing Tech, Micro-Mobility, Mobile, Mobile Commerce, Mobility Tech, Mortgage Tech, Nanotechnology, Oil & Gas, Oncology, Pet Technology, Real Estate Tech, Restaurant Tech, Ridesharing, Robotics and Drones, SaaS, Space Technology, Supply Chain Tech, TMT, Virtual Reality, Wearables & Devices, 3D Printing.
Panel B: List of Emerging Spaces
3D Printed Buildings, AI in Foodtech, Autonomous Farming, Autonomous Mobile Robots, Autonomous Trucking, Blockchain Real Estate, Cannabis Breeding, Climate Risk Modeling, Commercial Space, Construction Robotics, Digital Freight Forwarding, eDiscovery Platforms, Election Tech, Electric Flight, Food Service Robotics, Ghost Kitchens, Indoor Mapping, Industrial Workforce Safety, In-Space Manufacturing, LPWAN, Microweather, Modular Construction, Natural Disaster Preparedness, Ocean Data Collection, Resilient PNT, Small Satellites, Smart Mirrors, Smart Packaging, Sports Tech, Supersonic Transport, Sustainable Packaging, Urban Planning Software, Virtual Events, Warehouse Management, 3D Printed Food, AI-Enhanced Surveillance, Air Taxis, Art Trading Platforms, Auto Commerce, Cannabis Beverages, Cashierless Checkout, Clean Meat, Connected Fitness, Electric Vehicle Charging Infrastructure, End-of-Life Planning, Hyperloop, Income Share Agreements, Insect-Based Proteins, Pandemic Tracking, Smart Clothing, Smart Home Security, Smart Locks, Smart Jewelry, Social Audio, Space Tourism, Sustainable Fashion, Sustainable Textiles, Batteryless IoT, Carbon Capture & Storage, Concentrated Solar Power, Electric Vehicle Battery Tech, Fusion Energy, Hydrogen Energy, Lithium Extraction, Long-Duration Energy Storage, Next-Generation Batteries, Renewable Chemicals, Smart Grid, Small Modular Reactors, Waste to Energy, Alternative Housing Finance, Banking as a Service, Conversational AI, Decentralized Finance (DeFi), Microinsurance, NFTs, Payment Facilitation, Real Estate Crowdfunding, Security Deposit Alternatives, Youth Banking, AI-Powered Drug Discovery, Anti-Aging, Assistive Tech, CRISPR Diagnostics, Fertility Tech, Gene Therapies, Medical Exoskeletons, Medical Robotics, Mental Health Tech, Nanomedicine, Neurotechnology, Psychedelics, Sleep Tech, VR Health, 4D Printing, AIOps, Blockchain Gaming, Cloud Gaming, Cloud Workload Protection, Cognitive Computing, Computational Biology, Contract Management, Crowdsourced Testing, DAOs, Data Center Infrastructure, Database Sharding, Digital Avatars, Digital Twins, DevSecOps, DNA Data Storage, Edge Computing, FinOps, Generative AI, Graph Databases, IoT Security, Lidar, Low Code / No Code, Next-Generation Cyber, Passwordless Authentication, Post-Quantum Cryptography, Quantum Computing, Robotic Process Automation, Security Orchestration, Service Mesh, Silicon Photonics, Swarm AI, TinyML, V2X, Cellular Agriculture, Desalination, Indoor Farming, Lithium Battery Recycling, Livestock Health, Mining Tech, Pollination Tech, Reforestation, Regenerative Agriculture, Smart Waste Management, Artificial General Intelligence, Counter-Unmanned Aircraft Systems, LLM Agents, AI Neoclouds, Synthetic Data, PayFac Enablement, Spatial Biology, Humanoid Robots, Longevity Tech, Trust & Safety, Deep Sea Mining.

Appendix Table A.3: Illustrative Examples of AI Firm Classification

Firms are classified using PitchBook's AI-related verticals and emerging-space tags. A firm is classified as an AI producer if it develops AI technologies or infrastructure. A firm is classified as an AI adopter if it primarily deploys AI within products, services, or business operations. The examples illustrate how the classification maps into well-known firms. Notably, firms such as Stryker and Thermo Fisher Scientific are not classified as AI firms because PitchBook does not assign them AI-related verticals or emerging-space tags.

Firm	Classification	Producer	Adopter	Evidence
Anthropic	AI Producer & Adopter	1	1	AI/ML vertical; AGI and Generative AI emerging spaces.
OpenAI	AI Producer & Adopter	1	1	AI/ML vertical; Generative AI emerging space.
Nvidia	AI Producer Only	1	0	Generative AI and autonomous vehicle simulation emerging spaces; AI processor technologies.
Salesforce	AI Adopter Only	0	1	SaaS, CloudTech, IoT, and marketing technology verticals; AI-enabled enterprise applications.
Stryker	Non-AI	0	0	Manufacturing vertical; no AI-related PitchBook verticals or emerging-space classifications.
Thermo Fisher Scientific	Non-AI	0	0	Life Sciences and Manufacturing verticals; no AI-related PitchBook verticals or emerging-space classifications.

Appendix Table A.4: Top Verticals of AI-Producing Targets Acquired by AI Buyers

This table reports the most common PitchBook vertical classifications among AI-producing targets acquired by AI buyers in our sample. AI buyers are firms classified as AI producers, AI adopters, or both based on PitchBook's verticals and emerging spaces. AI-producing targets are firms classified as developing AI technologies or AI-related infrastructure. The table reports the number of acquisitions associated with each target vertical and the corresponding share of all AI-producing target acquisitions. Because firms may be assigned multiple PitchBook verticals, counts need not sum to the total number of acquisitions.

Vertical	Count	Share
Artificial Intelligence & Machine Learning	407	0.321
SaaS	253	0.200
Big Data	134	0.106
TMT	48	0.038
FinTech	43	0.034
Mobile	35	0.028
Marketing Tech	31	0.024
Cybersecurity	30	0.024
HealthTech	29	0.023
Robotics and Drones	23	0.018
CloudTech & DevOps	22	0.017
Digital Health	14	0.011
Real Estate Technology	14	0.011
Supply Chain Tech	12	0.009
HR Tech	11	0.009
Internet of Things	11	0.009
Manufacturing	11	0.009

Appendix Table A.5: Top Emerging Spaces of AI-Producing Targets Acquired by AI Buyers

This table reports the most common PitchBook emerging-space classifications among AI-producing targets acquired by AI buyers in our sample. AI buyers are firms classified as AI producers, AI adopters, or both based on PitchBook's verticals and emerging spaces. AI-producing targets are firms classified as developing AI technologies or AI-related infrastructure. The table reports the number of acquisitions associated with each target emerging space and the corresponding share of all AI-producing target acquisitions. Because firms may be assigned multiple PitchBook emerging-space classifications, counts need not sum to the total number of acquisitions.

Emerging Space	Count	Share
Generative AI	53	0.434
LLM Agents	11	0.090
AI-powered Drug Discovery	6	0.049
AI Neoclouds	5	0.041
Digital Twins	4	0.033
AI-powered Code Completion	3	0.025
DevSecOps	3	0.025
Quantum Computing	3	0.025
Cloud Workload Protection	2	0.016
Construction Robots	2	0.016
Contract Management Automation	2	0.016
IoT Security	2	0.016
NFTs	2	0.016
Workstyle Analytics	2	0.016
AI-Enhanced Learning	1	0.008
AIOps	1	0.008
Alternative Home Financing	1	0.008
Art Trading Platforms	1	0.008
Artificial General Intelligence (AGI) Research	1	0.008
Auto Commerce	1	0.008
Autonomous Maritime Systems	1	0.008
Autonomous Trucking	1	0.008
Cognitive Computing	1	0.008
DAOs	1	0.008
Database Sharding	1	0.008
Digital Identity Passports	1	0.008
Electric Vehicle Charging Infrastructure	1	0.008
Humanoid Robotics	1	0.008
Lidar	1	0.008
Mining Tech	1	0.008
Robotic Foundation Models	1	0.008
Sleep Tech	1	0.008
Small Satellites	1	0.008
Smart Waste Management	1	0.008
Synthetic Data	1	0.008
Warehouse Robotics	1	0.008

Appendix Table A.6: Vertical Shift in AI-Producing Targets Acquired by AI Buyers

This table reports the change in the composition of AI-producing targets acquired by AI buyers before and after the release of ChatGPT in 2022Q4. AI buyers are firms classified as AI producers, AI adopters, or both based on PitchBook's verticals and emerging spaces. AI-producing targets are firms classified as developing AI technologies or AI-related infrastructure. For each vertical, the table reports its share among all AI-producing target acquisitions before the shock (Pre), after the shock (Post), and the corresponding change. Panel A reports the largest increases in share, while Panel B reports the largest decreases. Because firms may be assigned multiple PitchBook verticals, shares need not sum to one.

Panel A: Largest Increases			
Vertical	Pre	Post	Change
SaaS	0.114	0.200	0.086
Artificial Intelligence & Machine Learning	0.276	0.321	0.045
FinTech	0.020	0.034	0.014
HealthTech	0.011	0.023	0.012
Real Estate Technology	0.003	0.011	0.008
Digital Health	0.004	0.011	0.007
Cybersecurity	0.018	0.024	0.006
Oncology	0.001	0.006	0.005
Legal Tech	0.000	0.005	0.005
CloudTech & DevOps	0.013	0.017	0.004
Panel B: Largest Decreases			
TMT	0.148	0.038	-0.110
Big Data	0.126	0.106	-0.020
Industrials	0.020	0.005	-0.015
Mobility Tech	0.016	0.006	-0.010
Augmented Reality	0.012	0.002	-0.010
Marketing Tech	0.031	0.024	-0.007
Mobile	0.035	0.028	-0.007
Autonomous Cars	0.008	0.002	-0.006
Virtual Reality	0.008	0.002	-0.006
Internet of Things	0.013	0.009	-0.004

Appendix Table A.7: Technological Overlap Between AI-Producing Targets and AI Adopters

This table compares the technological composition of AI-producing targets acquired by AI adopters with the technological composition of the AI-adopter buyers themselves. Patent domains are constructed using broad technology classes derived from firms' patent portfolios. The *Target Share* column reports the fraction of patents associated with acquired AI-producing targets that fall within each patent domain, while the *Buyer Share* column reports the corresponding fraction for AI-adopter acquirers. The table provides evidence on the degree of technological overlap between AI-producing targets and the firms acquiring them.

Patent Domain	Target Share	Buyer Share
Computing / Software	63.2%	47.3%
Telecom / Networks	20.7%	28.6%
Business Methods	2.5%	8.6%
Physics (Other)	2.0%	3.7%
Chemistry	2.8%	3.5%
Human Necessities	1.5%	2.2%
Electricity (Other)	0.5%	1.4%
Wireless	0.0%	1.4%
Operations / Transport	0.1%	1.3%
Image Processing	6.1%	1.0%
Video / Display	0.5%	0.9%

Appendix Table A.8: Post-Shock Non-AI Patents Filed by AI-Adopting Firms: WIPO Technology Fields

This table reports the distribution of non-AI patents filed by AI-adopting firms after the release of ChatGPT in 2022Q4 across World Intellectual Property Organization (WIPO) technology fields. AI-adopting firms are defined using PitchBook's verticals and emerging-space classifications. Non-AI patents are patents that do not satisfy our AI patent classification criteria. Patent technology fields are assigned using the WIPO technology concordance, which maps patent classifications into broad technological domains. The table reports the number of patents and the corresponding share of all non-AI patents filed by AI-adopting firms in the post-shock period.

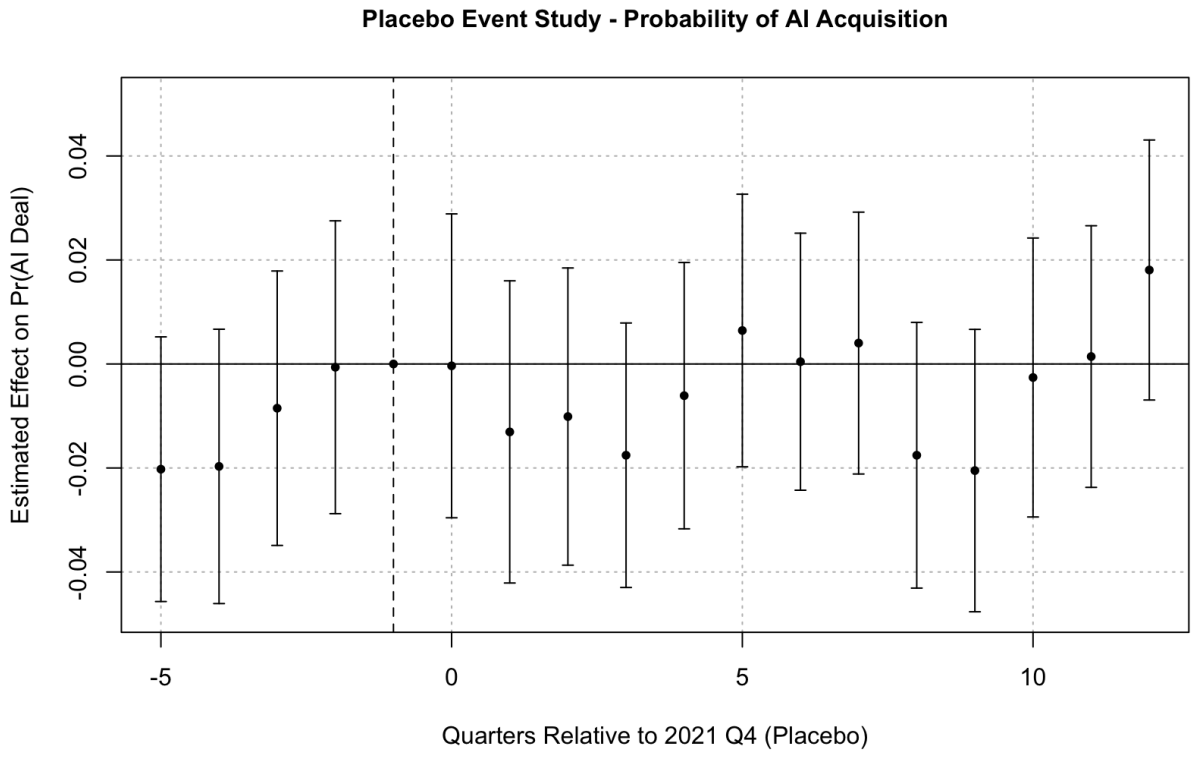
WIPO Technology Field	Count	Share (%)
Computer Technology	1,810	44.6
Digital Communication	1,210	29.8
IT Methods for Management	372	9.2
Medical Technology	96	2.4
Biotechnology	91	2.2
Control	80	2.0
Telecommunications	67	1.6
Audio-Visual Technology	57	1.4
Measurement	54	1.3
Basic Communication Processes	26	0.6
Handling	22	0.5
Chemical Engineering	20	0.5
Electrical Machinery, Apparatus, Energy	18	0.4
Materials, Metallurgy	18	0.4
Transport	17	0.4
Analysis of Biological Materials	13	0.3
Organic Fine Chemistry	11	0.3
Other Consumer Goods	10	0.2
Furniture, Games	9	0.2
Optics	8	0.2
Pharmaceuticals	8	0.2
Other Special Machines	7	0.2
Food Chemistry	6	0.1
Textile and Paper Machines	6	0.1
Semiconductors	5	0.1
Mechanical Elements	4	0.1
Basic Materials Chemistry	3	0.1
Environmental Technology	3	0.1
Civil Engineering	2	0.0
Micro-Structural and Nano-Technology	2	0.0
Thermal Processes and Apparatus	2	0.0
Engines, Pumps, Turbines	1	0.0
Machine Tools	1	0.0
Macromolecular Chemistry, Polymers	1	0.0
Surface Technology, Coating	1	0.0
Total Non-Missing Patents	4,061	100.0

B. Robustness Checks and Additional Pre-trends

B.1 Robustness Checks

Appendix Figure B.1: Placebo Event Study

This figure plots event-study estimates of acquisition activity using a placebo shock date of 2021Q4, one year prior to the release of ChatGPT. The matched sample is constructed by pairing each acquirer that is an AI firm with the closest acquirer that is a non-AI acquirer based on pre-shock firm characteristics, using nearest-neighbor matching with replacement and exact matching on the acquirer's primary industry sector. The figure is designed to assess whether AI and non-AI firms exhibited differential acquisition behavior prior to the actual AI shock or whether the main results are driven by broader macroeconomic or policy developments. Coefficients are plotted relative to the quarter immediately preceding the placebo shock ($k = -1$), which serves as the reference period. All specifications are estimated at the acquirer-quarter level and include acquirer fixed effects and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the acquirer level, and error bars represent 95% confidence intervals.



Appendix Table B.1: AI-Producing Target Age and Size

This table reports difference-in-differences estimates of acquisition activity following the AI shock in 2022Q4. The main explanatory variable is the interaction between an indicator for the post-shock period and an indicator for AI buyers. The dependent variables measure acquisitions of AI-producing targets that also satisfy target size and age criteria. Young targets are defined as firms below the median of target age, and small targets are defined as firms below the median of target employment. All specifications are estimated at the acquirer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the firm level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Small and AI-Producing Target M&A Outcomes			
	Prob(Small & AI-Prod Deal)	log(Small & AI-Prod Count)	log(Small & AI-Prod Value)
Post × AI Buyer	0.004**	0.003**	−0.004**
	(0.002)	(0.001)	(0.002)
Num. Obs.	47,725	47,725	47,725
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel B: Young and AI-Producing Target M&A Outcomes			
	Prob(Young & AI-Prod Deal)	log(Young & AI-Prod Count)	log(Young & AI-Prod Value)
Post × AI Buyer	0.005**	0.004**	0.004
	(0.002)	(0.001)	(0.004)
Num. Obs.	47,840	47,840	47,840
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel C: Young, Small, and AI-Producing Target M&A Outcomes			
	Prob(Young, Small & AI-Prod Deal)	log(Young, Small & AI-Prod Count)	log(Young, Small & AI-Prod Value)
Post × AI Buyer	0.003*	0.002*	−0.002
	(0.002)	(0.001)	(0.002)
Num. Obs.	47,748	47,748	47,748
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Appendix Table B.2: Target Age and Size

This table reports difference-in-differences estimates of acquisition activity following the AI shock in 2022Q4. The main explanatory variable is the interaction between an indicator for the post-shock period and an indicator for AI buyers. Young targets are defined as firms in the lowest tercile of target age, and small targets are defined as firms in the lowest tercile of target employment. All specifications are estimated at the acquirer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the firm level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Young Target M&A Outcomes			
	Prob(Young Deal)	log(Young Deal Count)	log(Young Deal Value)
Post × AI Buyer	0.011***	0.008***	0.001
	(0.003)	(0.002)	(0.005)
Num. Obs.	48,001	48,001	48,001
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel B: Small Target M&A Outcomes			
	Prob(Small Deal)	log(Small Deal Count)	log(Small Deal Value)
Post × AI Buyer	0.008***	0.006***	−0.002
	(0.003)	(0.002)	(0.003)
Num. Obs.	47,955	47,955	47,955
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel C: Young and Small Target M&A Outcomes			
	Prob(Young & Small Deal)	log(Young & Small Count)	log(Young & Small Value)
Post × AI Buyer	0.008***	0.006***	0.001
	(0.002)	(0.001)	(0.002)
Num. Obs.	47,955	47,955	47,955
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Appendix Table B.3: AI Target Acquisitions and Patent Reallocation (Expanded CPC)

This table reports triple-difference estimates of the effect of acquiring AI targets on patenting activity following the AI shock in 2022Q4. The sample is restricted to the 711 acquiring firms that can be matched to at least one patent in the USPTO data. Panel A examines acquisitions of AI-producing targets, while Panel B examines acquisitions of AI-adopting targets. The main explanatory variable is the triple interaction between an indicator for the post-shock period (2022Q4 and later), an indicator for AI buyers, and an indicator for buyers that acquire at least one AI target of the relevant type after the shock. AI patents are defined using the expanded CPC-based AI patent classification, which augments the baseline CPC definition by including CPC class G06F40 (Natural Language Processing). The dependent variables are log(1+Non-AI Patents), log(1+AI Patents), and the difference between the two measures. All specifications are estimated at the buyer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the buyer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: AI Producer Acquisitions			
	log(Non-AI Patents)	log(AI Patents)	log(Non-AI) minus log(AI)
Post × AI Buyer × AI Producer Acquisition	0.529***	0.312***	0.217**
	(0.150)	(0.106)	(0.107)
Num. Obs.	16,353	16,353	16,353
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
AI Patent Definition	CPC + G06F40	CPC + G06F40	CPC + G06F40
Sample	Patent-filing firms	Patent-filing firms	Patent-filing firms
Panel B: AI Adopter Acquisitions			
	log(Non-AI Patents)	log(AI Patents)	log(Non-AI) minus log(AI)
Post × AI Buyer × AI Adopter Acquisition	0.192	0.152	0.040
	(0.152)	(0.100)	(0.094)
Num. Obs.	16,353	16,353	16,353
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
AI Patent Definition	CPC + G06F40	CPC + G06F40	CPC + G06F40
Sample	Patent-filing firms	Patent-filing firms	Patent-filing firms

Appendix Table B.4: AI Target Acquisitions and Patent Reallocation (USPTO AIPD)

This table reports triple-difference estimates of the effect of acquiring AI targets on patenting activity following the AI shock in 2022Q4 using the USPTO Artificial Intelligence Patent Dataset (AIPD) Predict93 classification. The sample is restricted to the 711 acquiring firms that can be matched to at least one patent in the USPTO data. Panel A examines acquisitions of AI-producing targets, while Panel B examines acquisitions of AI-adopting targets. The main explanatory variable is the triple interaction between an indicator for the post-shock period (2022Q4 and later), an indicator for AI buyers, and an indicator for buyers that acquire at least one AI target of the relevant type after the shock. The dependent variables are log(1+Non-AI Patents), log(1+AI Patents), and the difference between the two measures. Because the USPTO AIPD classification is currently available only through 2023, the sample period is restricted through 2023Q4. All specifications are estimated at the buyer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the buyer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: AI Producer Acquisitions			
	log(Non-AI Patents)	log(AI Patents)	log(Non-AI) minus log(AI)
Post × AI Buyer × AI Producer Acquisition	0.557***	0.547***	0.010
	(0.196)	(0.175)	(0.104)
Num. Obs.	11,376	11,376	11,376
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
AI Patent Definition	USPTO AIPD Predict93	USPTO AIPD Predict93	USPTO AIPD Predict93
Sample Period	Through 2023Q4	Through 2023Q4	Through 2023Q4
Sample	Patent-filing firms	Patent-filing firms	Patent-filing firms
Panel B: AI Adopter Acquisitions			
	log(Non-AI Patents)	log(AI Patents)	log(Non-AI) minus log(AI)
Post × AI Buyer × AI Adopter Acquisition	0.300	0.321*	-0.021
	(0.187)	(0.166)	(0.087)
Num. Obs.	11,376	11,376	11,376
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
AI Patent Definition	USPTO AIPD Predict93	USPTO AIPD Predict93	USPTO AIPD Predict93
Sample Period	Through 2023Q4	Through 2023Q4	Through 2023Q4
Sample	Patent-filing firms	Patent-filing firms	Patent-filing firms

Appendix Table B.5: Alternative Matching Specification Excluding Pre-Shock AI Share

This table reports difference-in-differences estimates of acquisition activity following the AI shock in 2022Q4 using an alternative matching procedure. Unlike the baseline specification, the matching procedure does not include the pre-shock share of AI-related acquisitions as a matching covariate. Instead, AI buyers are matched to non-AI buyers based on firm size, firm age, pre-shock acquisition activity, and primary industry sector using nearest-neighbor matching with replacement and exact matching on sector. Panel A reports acquisition outcomes for AI-producing targets, Panel B reports acquisition outcomes for AI-adopting targets, and Panel C examines the relative acquisition of AI-producing versus AI-adopting targets following the release of ChatGPT. All specifications are estimated at the acquirer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the buyer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: AI-Producing Target M&A Outcomes			
	Prob(AI Producer Deal)	log(AI Producer Count)	log(AI Producer Deal Value)
Post × AI Buyer	0.007***	0.005***	0.007
	(0.002)	(0.002)	(0.006)
Num. Obs.	48,622	48,622	48,622
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel B: AI-Adopting Target M&A Outcomes			
	Prob(AI Adopter Deal)	log(AI Adopter Count)	log(AI Adopter Deal Value)
Post × AI Buyer	-0.000	-0.000	-0.010
	(0.003)	(0.002)	(0.007)
Num. Obs.	48,622	48,622	48,622
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel C: Relative Acquisitions of AI-Producing versus AI-Adopting Targets			
	Producer Share of AI Deals	log(Producer) − log(Adopter) Count	log(Producer) − log(Adopter) Value
Post × AI Buyer	0.006***	0.005*	0.017*
	(0.002)	(0.003)	(0.009)
Num. Obs.	48,622	48,622	48,622
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Appendix Table B.6: Controlling for Venture-Capital-Backed Targets

This table reports difference-in-differences estimates of acquisition activity following the AI shock in 2022Q4 after controlling for the share of each buyer's quarterly acquisitions involving venture-capital-backed targets. To construct this measure, we identify venture capital funds in PitchBook's fund database, trace their associated investors through fund-investor linkages, and then identify all target firms that previously received financing from those investors. A target is classified as VC-backed if it appears in at least one deal involving a venture capital investor. We then compute, for each buyer-quarter, the share of acquisitions involving VC-backed targets and include this measure as an additional control variable. This robustness check addresses the concern that the post-shock increase in AI-related acquisitions may reflect an expansion in the supply of venture-capital-backed AI startups rather than an increase in acquirer demand for AI capabilities. Panel A reports acquisition outcomes for AI targets, Panel B reports acquisition outcomes for non-AI targets, Panel C examines the relative acquisition of AI versus non-AI targets, and Panel D reports acquisition outcomes for AI-producing targets. All specifications are estimated at the acquirer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the firm level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: AI Target Outcomes			
	Prob(AI Deal)	log(AI Count)	log(AI Deal Value)
Post × AI Buyer	0.010***	0.007***	0.006
	(0.003)	(0.002)	(0.010)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel B: Non-AI Target Outcomes			
	Prob(Non-AI Deal)	log(Non-AI Count)	log(Non-AI Deal Value)
Post × AI Buyer	-0.000	-0.000	-0.010
	(0.004)	(0.003)	(0.010)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel C: Relative Acquisitions of AI versus Non-AI Targets			
	AI Share	log(AI) − log(Non-AI) Count	log(AI) − log(Non-AI) Value
Post × AI Buyer	0.010***	0.007*	0.015
	(0.003)	(0.004)	(0.013)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel D: AI-Producing Target Outcomes			
	Prob(Producer Target)	log(Producer Count)	log(Producer Value)
Post × AI Buyer	0.008***	0.006***	0.011*
	(0.002)	(0.002)	(0.006)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

Appendix Table B.7: Section 174 Exposure and AI-Producing Target Acquisitions

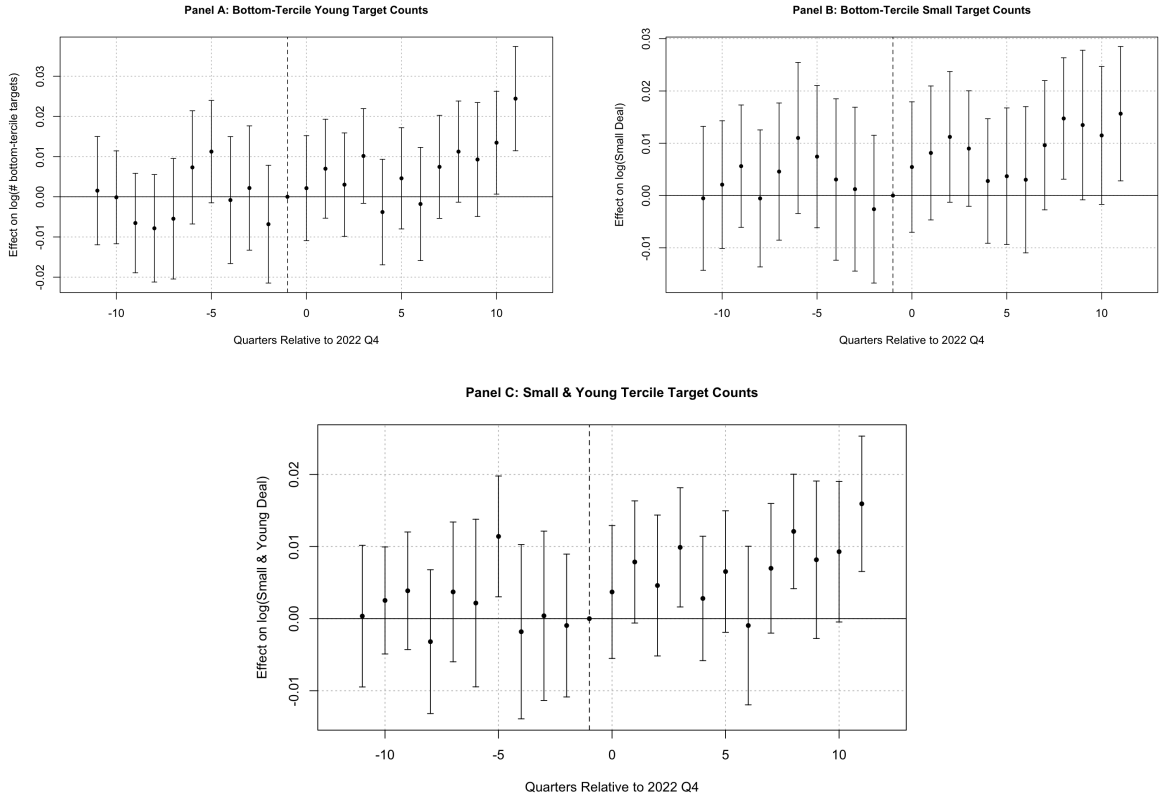
This table examines whether the increase in acquisitions of AI-producing targets following the release of ChatGPT can be explained by differential exposure to Section 174 of the Tax Cuts and Jobs Act. Beginning in 2022, Section 174 required firms to capitalize and amortize R&D expenditures, potentially affecting acquisition activity among R&D-intensive firms. We augment the baseline difference-in-differences specification with interactions between AI-buyer status, a post-2022Q1 indicator, and an industry-level measure of Section 174 exposure constructed using pre-2022 R&D intensity from Compustat firms. Panel A measures exposure at the PitchBook Industry Sector level using median pre-2022 R&D intensity. Panel B measures exposure at the more granular PitchBook Industry Group level using median pre-2022 R&D intensity. All specifications are estimated at the buyer-quarter level and include buyer and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the buyer level, and reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: PitchBook Industry Sector Exposure			
	Prob(AI Producer Deal)	log(AI Producer Count)	log(AI Producer Deal Value)
Post-ChatGPT × AI Buyer	0.012***	0.009***	0.013
	(0.004)	(0.003)	(0.009)
Post-174 × AI Buyer	-0.005	-0.004	0.003
	(0.005)	(0.003)	(0.011)
Post-174 × Section 174 Exposure	-0.002	-0.001	0.003
	(0.003)	(0.002)	(0.006)
Post-174 × AI Buyer × Section 174 Exposure	0.000	0.001	-0.005
	(0.004)	(0.003)	(0.009)
Num. Obs.	48,392	48,392	48,392
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes
Panel B: PitchBook Industry Group Exposure			
	Prob(AI Producer Deal)	log(AI Producer Count)	log(AI Producer Deal Value)
Post-ChatGPT × AI Buyer	0.012***	0.009***	0.013
	(0.004)	(0.003)	(0.009)
Post-174 × AI Buyer	-0.005	-0.004	0.004
	(0.004)	(0.003)	(0.010)
Post-174 × Section 174 Exposure	-0.005**	-0.003**	0.002
	(0.002)	(0.001)	(0.004)
Post-174 × AI Buyer × Section 174 Exposure	0.001	0.001	-0.017*
	(0.003)	(0.002)	(0.010)
Num. Obs.	48,300	48,300	48,300
Buyer FE / Quarter FE	Yes / Yes	Yes / Yes	Yes / Yes

B.2 Additional Pre-trends

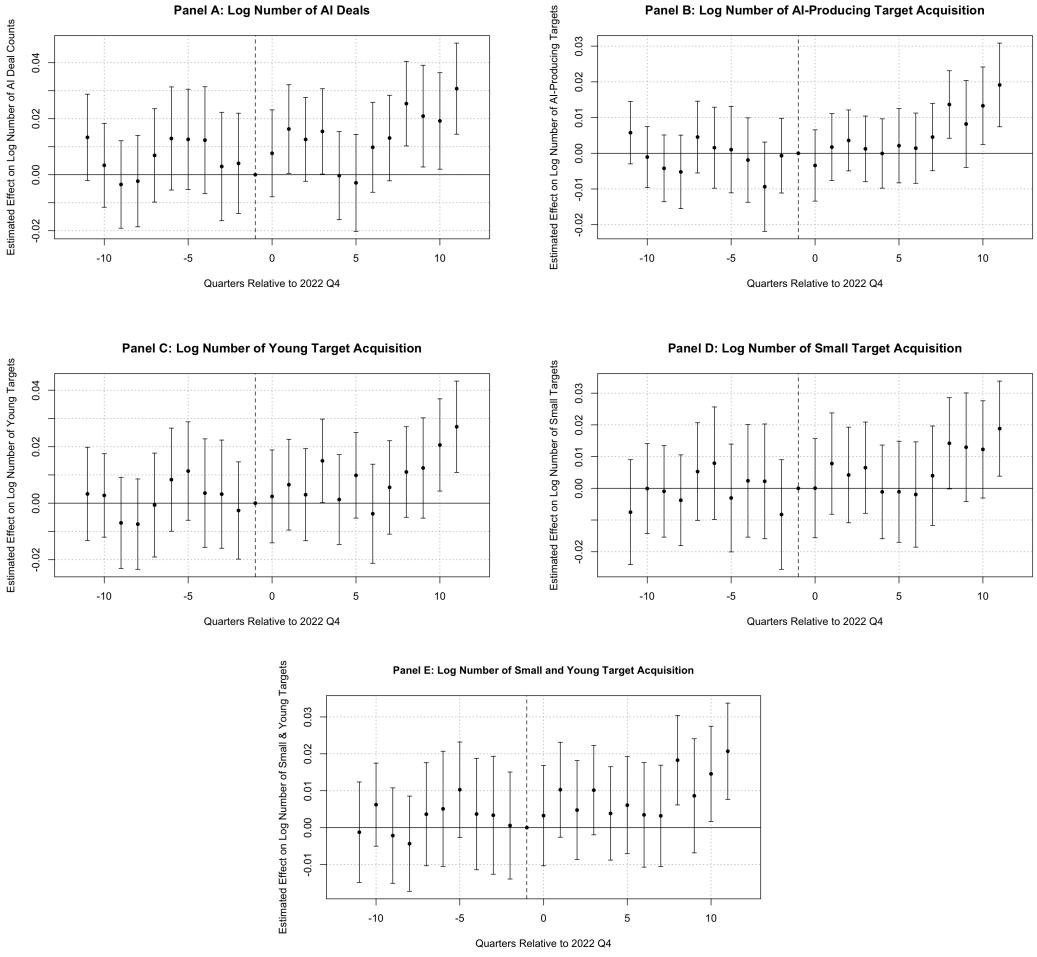
Appendix Figure B.3: Dynamic Effects of AI on Target Acquisition Counts: Tercile Cutoffs

This figure plots event-study estimates of acquisition activity following the AI shock in 2022Q4, using Equation (2), and the matched sample. The matched sample is constructed by pairing each acquirer that is an AI firm with the closest acquirer that is a non-AI acquirer based on pre-shock firm characteristics, using nearest-neighbor matching with replacement and exact matching on the acquirer's primary industry sector. Panel A reports the event-study estimates of the log number of young target acquisitions, Panel B reports the estimates of the log number of small target acquisitions, and Panel C reports the estimates of the log number of target acquisitions that are both young and small. Young targets are defined as firms in the lowest tercile of target age, and small targets are defined as firms in the lowest tercile of target employment. The dependent variable in each specification is $\log(1 + \text{Target Count})$. Coefficients are plotted relative to the quarter immediately preceding the AI shock ($k = -1$), which serves as the reference period. All specifications are estimated at the acquirer-quarter level and include acquirer fixed effects and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the acquirer level, and error bars represent 95% confidence intervals.



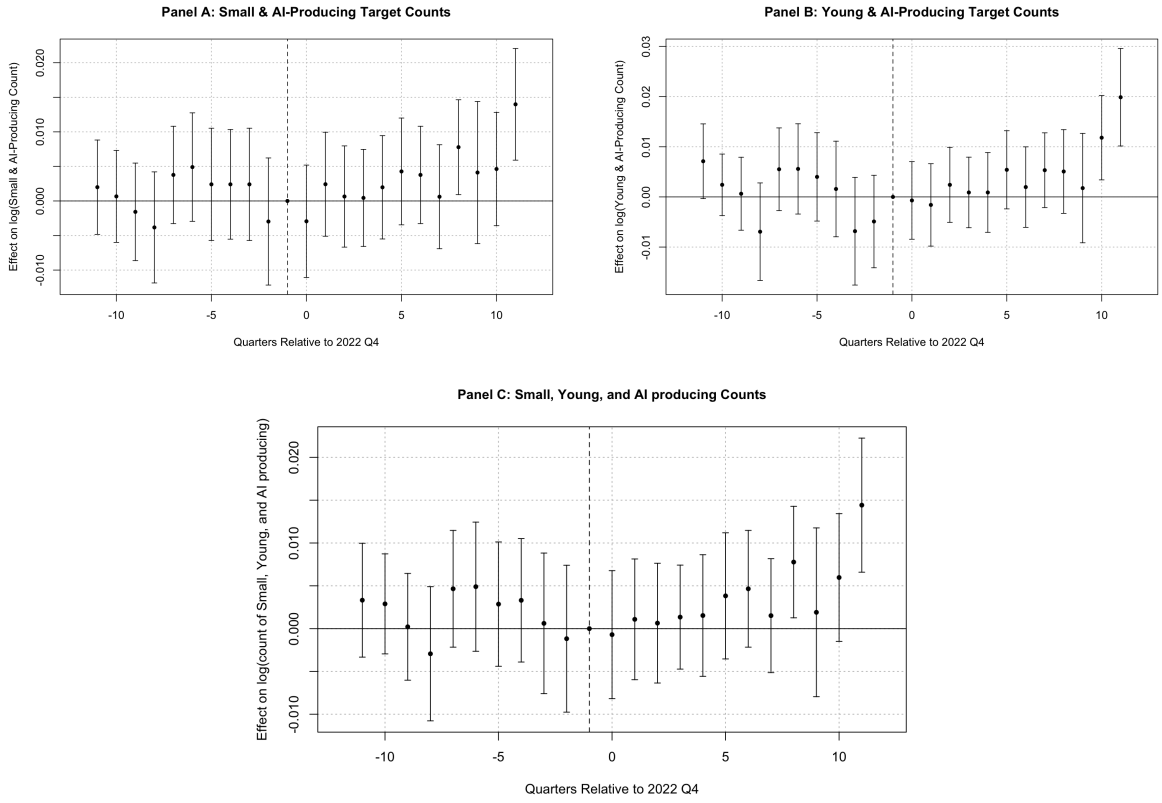
Appendix Figure B.2: Dynamic Effects of AI on Target Acquisition Counts: Baseline Definitions

This figure plots event-study estimates of acquisition activity following the AI shock in 2022Q4, using Equation (2) and the matched sample. The matched sample is constructed by pairing each acquirer that is an AI firm with the closest acquirer that is a non-AI acquirer based on pre-shock firm characteristics, using nearest-neighbor matching with replacement and exact matching on the acquirer's primary industry sector. Panel A reports event-study estimates for the log number of AI target acquisitions, Panel B reports estimates for the log number of AI-producing target acquisitions, Panel C reports estimates for the log number of young AI target acquisitions, Panel D reports estimates for the log number of small AI target acquisitions, and Panel E reports estimates for the log number of AI target acquisitions that are both small and young. Young targets are defined as firms whose age is below the sample median at the time of acquisition, and small targets are defined as firms whose employment is below the sample median at the time of acquisition. The dependent variable in each specification is $\log(1 + \text{Target Count})$. Coefficients are plotted relative to the quarter immediately preceding the AI shock ($k = -1$), which serves as the reference period and is denoted by the vertical dashed line. All specifications are estimated at the acquirer-quarter level and include acquirer fixed effects and quarter fixed effects. Standard errors are clustered at the acquirer level, and error bars represent 95% confidence intervals.



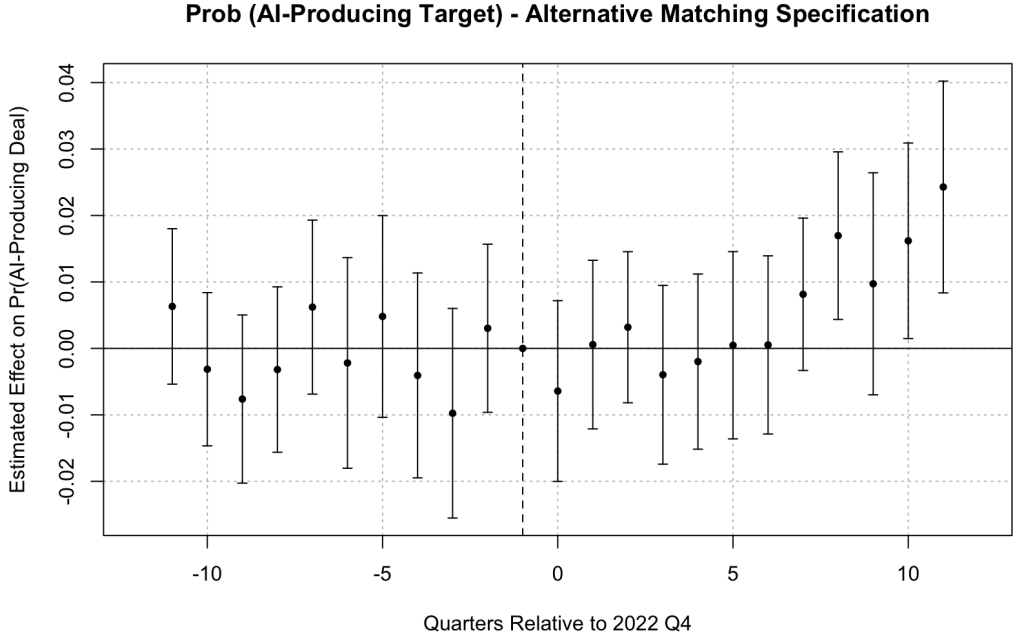
Appendix Figure B.4: Dynamic Effects of AI on Acquisitions of AI-Producing Targets by Target Characteristics: Counts

This figure plots event-study estimates of acquisition activity following the AI shock in 2022Q4, using Equation (2), and the matched sample. The matched sample is constructed by pairing each acquirer that is an AI firm with the closest acquirer that is a non-AI acquirer based on pre-shock firm characteristics, using nearest-neighbor matching with replacement and exact matching on the acquirer's primary industry sector. Panel A reports the event-study estimates of the log number of small AI-producing target acquisitions, Panel B reports the estimates of the log number of young AI-producing target acquisitions, and Panel C reports the estimates of the log number of target acquisitions that are simultaneously small, young, and AI-producing. Small and young targets are defined using the median values of target employment and target age, respectively. The dependent variable in each specification is $\log(1 + \text{Target Count})$. Coefficients are plotted relative to the quarter immediately preceding the AI shock ($k = -1$), which serves as the reference period. All specifications are estimated at the acquirer-quarter level and include acquirer fixed effects and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the acquirer level, and error bars represent 95% confidence intervals.



Appendix Figure B.5: Dynamic Effects of AI on AI-Producing Target Acquisitions: Alternative Matching Specification

This figure plots event-study estimates of AI-producing target acquisition activity following the AI shock in 2022Q4, using Equation (2), and an alternative matched sample. Unlike the baseline specification, the matching procedure does not include the pre-shock share of AI-related acquisitions as a matching covariate. Instead, AI buyers are matched to non-AI buyers based on firm size, firm age, pre-shock acquisition activity, and primary industry sector using nearest-neighbor matching with replacement and exact matching on sector. The figure reports the event-study estimates of the probability that an acquirer that is an AI firm acquires an AI-producing target. Coefficients are plotted relative to the quarter immediately preceding the AI shock ($k = -1$), which serves as the reference period. All specifications are estimated at the acquirer-quarter level and include acquirer fixed effects and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the acquirer level, and error bars represent 95% confidence intervals.



Appendix Figure B.6: Dynamic Effects of AI Acquisition Activity Controlling for VC-Backed Targets

This figure plots event-study estimates of acquisition activity following the AI shock in 2022Q4 after controlling for the share of each buyer's quarterly acquisitions involving venture-capital-backed targets. To construct this measure, we identify venture capital funds in PitchBook's fund database, trace their associated investors through fund-investor linkages, and then identify all target firms that previously received financing from those investors. A target is classified as VC-backed if it appears in at least one deal involving a venture capital investor. We then compute, for each buyer-quarter, the share of acquisitions involving VC-backed targets and include this measure as an additional control variable. This robustness check addresses the concern that the post-shock increase in AI-related acquisitions may reflect an expansion in the supply of venture-capital-backed AI startups rather than an increase in acquirer demand for AI capabilities. Panel A reports the event-study estimates of the probability that an AI buyer acquires any AI target, Panel B reports the estimates of the probability that an AI buyer acquires an AI-producing target, Panel C reports the estimates for the logarithm of AI-target acquisition counts, and Panel D reports the estimates for the logarithm of AI-producing target acquisition counts. Coefficients are plotted relative to the quarter immediately preceding the AI shock ($k = -1$), which serves as the reference period. All specifications are estimated at the acquirer-quarter level and include acquirer fixed effects and quarter fixed effects. Standard errors are heteroskedasticity robust and clustered at the acquirer level, and error bars represent 95% confidence intervals.

