Economic Impacts of Open-Source AI Models: Firm-Level Evidence *

Yi Huang Sung Kwan Lee Bernard Yeung Tao Zhang Yujun Zhou

March 18, 2026

Abstract

We examine firm-level stock market reactions to AI model releases and find that open-source model releases are associated with significant valuation gains, whereas proprietary model releases exhibit negligible valuation effects. These positive effects are widespread across firms, particularly benefiting smaller and capital-constrained firms due to the lower user costs and greater accessibility of open-source models. The valuation effects vary significantly across firms, with the largest gains observed among those with a higher share of AI-relevant tasks, particularly augmentation-intensive ones. Furthermore, using an input-output exposure measure, we document positive valuation spillovers from high-AI exposure firms to low-AI-exposure firms in both upstream and downstream markets, which are associated only with open-source releases, not with proprietary ones. Taken together, our findings highlight how open-source releases reduce adoption costs and entry barriers, facilitating broader AI diffusion and enhancing firm value.

Keywords: Open-source AI, Cost Efficiency, Process Innovation, Gen-AI
JEL Classification: D85, G14, L17, L86, O31, O36

*Yi Huang: Bank for International Settlements (BIS) and CEPR. Email: Yi.Huang@bis.org; Sung Kwan Lee: The Chinese University of Hong Kong, Shenzhen. Email: leesungkwan@cuhk.edu.cn; Bernard Yeung: Asian Bureau of Finance and Economic Research, Southern University of Science and Technology, and National University of Singapore Business School. byeung@nus.edu.sg; Tao Zhang: Bank for International Settlements (BIS). Tao.Zhang@bis.org; Yujun Zhou: School of Economics and Management, Tsinghua University. zhouyujun23@mails.tsinghua.edu.cn. We thank Pablo Hernández de Cos, Hélène Rey, Bengt Holmström, Richard Portes, Damiano Sandri, Jon Frost, Chen Lin, Ye Luo and Fernando Perez-Cruz for providing valuable feedback. The views expressed in this publication are those of the authors and do not necessarily reflect the views of the BIS or its member central banks.

I. Introduction

Recent advancements in artificial intelligence (AI), particularly in generative AI models such as ChatGPT, Llama, Claude, and DeepSeek, have led to a significant shift in the innovation landscape. This shift marks a transition from proprietary models dominated by a small number of large firms to an ecosystem where open-source model releases play a growing role. Open-source models are typically more accessible and cost-effective, offering substantial economic benefits by lowering adoption barriers and accelerating the diffusion of AI technologies.[1] However, the firm-level economic impacts of open-source AI models remain underexplored. While their effects on macroeconomic variables, such as bond yields, have been studied (e.g., Björkegren, 2026), it remains unclear to what extent they enhance market valuations, which types of firms benefit the most, and what micro-level transformations they may drive.

This paper examines the economic impacts of open-source AI models compared to proprietary ones, using firm-level evidence from stock market reactions to model release announcements. The underlying intuition is that while proprietary innovation often delivers radical breakthroughs that push technological frontiers, open-source innovation enables more inclusive access and facilitates numerous economically relevant incremental advancements. These two approaches are not mutually exclusive but rather complementary and sequential, each deserving careful study. Proprietary AI innovations typically raise expectations of dramatic performance improvements but impose high entry barriers, such as elevated user costs, which lead to an uneven outcome benefiting resource-rich companies and expose other firms to downside risks. In contrast, open-source innovations significantly lower adoption barriers, lead to expectations of positive gains, and promote widespread diffusion (Bommasani et al., 2021; Kapoor et al., 2024).

Consistent with this distinction, Andrews and Farboodi (2025) show that announcements of proprietary AI models lower long-term Treasury yields, while Björkegren (2026) finds that open-source AI news has the opposite effect. These findings align with the interpretation that

[1] *The Economist* (2024), "Why open-source AI models are good for the world"

proprietary AI news increases perceived future wealth volatility and precautionary savings, whereas open-source AI news boosts expectations of widespread adoption, higher productivity, and increased future investment. We extend the distinction between open-source and proprietary releases to the firm level, highlighting their nuanced impacts, and contrasting them with prior research focused on proprietary shocks (e.g., Eisfeldt et al., 2023; Bertomeu et al., 2025).

To empirically capture the distinct effects of these two AI model release types, we employ an event-study framework to analyze a novel dataset of 67 influential model releases covering the period from November 2022 to August 2025. Proprietary AI model releases are sourced from all major global AI-related technology firms, including OpenAI, Anthropic, Amazon, Google, and ByteDance. Open-source AI model releases are identified from Hugging Face, the leading global open-source software platform, by selecting the five most popular commercial AI models each month. We further validate the greater accessibility and affordability of open-source models by demonstrating that they are more likely to release model weights and code, utilize fewer parameters, and offer lower per-token user costs.

We explore different economic impacts of these models through multiple dimensions. First, we quantify the magnitude and distributional effects of stock market reactions to the two types of model releases. Second, we examine the relationship between adoption barriers and firm valuations and provide evidence that open-source releases lead investors to expect more inclusive adoption. Third, we explore the differential impacts of the two release types across firms and industries. Details on the sample model selection are in Section II.C, and the full list of models is in Appendix IA.1. The technical and economic distinctions between proprietary and open-source models are in Table I.[2]

Our analysis reveals several key findings. First, open-source model releases are associated with significantly larger cumulative abnormal returns (CARs) than proprietary model releases. Furthermore, open-source releases are associated with a positively skewed CAR distribution at both the stock market and industry levels, shifting lower-end returns to the right. By contrast, proprietary model releases are associated with negatively skewed returns,

[2]Also, we conduct a battery of robustness checks, including alternative specifications and measures.

with lower-end CARs shifting left and upper-end CARs shifting right. Likewise, cross-model variation in CARs is positively associated with accessibility measures and negatively associated with user costs, suggesting that open-source AI models drive adoption by making innovation more accessible and affordable.

Second, we find significant firm-level heterogeneity driven by adoption cost constraints. Prior work such as Demirer et al. (2025) finds that open-source models are about 90% cheaper than comparable closed-source models of the same intelligence quality. Consistent with this cost-efficiency mechanism, while valuation gains are broad-based, smaller and financially constrained firms experience significantly higher gains exclusively upon the release of open-source models. By contrast, valuation gains following proprietary releases accrue mostly to firms with greater financial resources. These results suggest that open-source AI may enhance economic inclusiveness by lowering barriers to access and participation in AI development and deployment, thereby enabling less advantaged firms to compete more effectively.

Third, our analysis identifies firms' task suitability for AI innovation as a primary driver of these valuation gains. We proxy suitability using two complementary measures: a labor-based estimate of AI exposure (Eisfeldt et al., 2023) and a real-world AI usage measure via the *Anthropic Economic Index*. High-suitability firms experience significantly larger gains from open-source releases but exhibit no substantial reaction to proprietary ones. In contrast, low-suitability firms show no significant gains from either paradigm.

Decomposing suitability into augmentation-driven vs. automation-driven exposures, we find valuation gains are driven primarily by augmentation-intensive tasks, with little evidence of significant market response to automation-oriented applications. Again, these augmentation-driven gains, for smaller firms, arise only after open-source releases; larger firms capture similar gains from proprietary models as well.

Fourth, we find that the economic benefits of open-source releases further propagate through supply chains, lifting valuations even for firms in low-AI-exposure industries that maintain vertical links to high-AI-exposure partners. Utilizing weighted measures of upstream and downstream exposure from Organization for Economic Co-operation and De-

velopment (OECD) input–output tables, we demonstrate that these vertical spillovers are unique to open-source releases. The absence of comparable effects for proprietary models indicates that open-source AI elevates market expectations of vertical spillovers by enabling coordinated gains from AI adoption among interconnected firms within a production network. The spillovers likely stem from improved procurement and collaborative operations, indicating that the economic impacts of open-source AI extend beyond the directly adopting firm.

The remainder of this paper is organized as follows. Section II describes the data sources, variable construction, and empirical strategy. Section III presents the empirical results, while Section IV discusses the differential impacts of the two AI release types across firms and industries. Section V shows robustness checks using alternative specifications and measures. The paper concludes in Section VI by summarizing the key findings and discussing their implications for future research.

II. Data, Measurement, and Empirical Strategy

This section describes the data sources and the construction of key variables, presents the event list of proprietary and open-source AI model releases, and details the event-study design for measuring the stock-market response to AI model releases.

A. Data

For stock return data, we use the Compustat North America Security Price dataset as of August 2025. For firm-level financial variables, we rely on the Compustat North America Security Fundamentals data. We extract the natural logarithm of total assets, return on assets (ROA), and the debt-to-equity ratio as control variables, measured at the quarterly frequency. These fundamental variables are also used to filter out non-stock securities, such as Exchange Traded Funds (ETFs) and mutual funds, from the Security Price data. After applying these filters, the final sample consists of 5,656 unique U.S. listed firms.

For occupational employment data, we utilize Revelio Labs, which provides detailed information on employees' job titles, positions, Standard Occupational Classification (SOC) codes, and employer affiliations, sourced primarily from LinkedIn. We use these data to

construct measures of occupational exposure to Gen-AI, which are described in greater detail in the next subsection.

Data on 10-K filings for U.S. publicly listed companies are sourced from the U.S. Securities and Exchange Commission (SEC)'s Electronic Data Gathering, Analysis, and Retrieval (EDGAR) database, while earnings call transcripts are obtained from Capital IQ. We use these textual data to construct an alternative measure of firm-level AI exposure, which is explained in the next subsection. We also source the Anthropic Economic Index from Claude to construct firm-level AI exposure in an alternative way.

B. Measures of Outcome Variables and AI Exposure

Our main outcome variable is the cumulative abnormal returns (CAR) of firm i for each event g over the window [-1,k], where day 0 is the announcement date of the Gen-AI model releases.

$$CAR_{ig}[-1,k] = [\Pi_{t=-1}^{k}(1+R_{i,t})-1] - \beta_{1,i,g} * [\Pi_{t=-1}^{k}(1+R_{m,t})-1] \tag{1}$$

where $R_{i,t}$ is the daily raw return for stock i on date t, $R_{m,t}$ is the corresponding stock market index m's return on date t, and $\beta_{1,i,g}$ of each stock i is estimated using daily stock returns over the one-year period preceding each event g as in the event study in Wagner et al. (2018).

Furthermore, given the possibility of a "regime shift" from the releases of Gen-AI models, which are events that could alter the underlying risk profile of firms, we acknowledge that traditional market betas may no longer remain stable. In such cases, a stock's beta could shift to a different structural value, potentially biasing market-adjusted return estimates. To ensure robustness against such structural changes in risk dynamics, we alternatively use cumulative raw returns (CRR) in our event study analysis. This approach avoids reliance on beta stability and provides a more conservative estimate of abnormal performance.

$$CRR_{ig}[-1,k] = \Pi_{t=-1}^{k}(1+R_{i,t})-1 \tag{2}$$

Given that user costs and accessibility are expected to affect the impact of Gen-AI model releases on stock value, we directly quantify them. We measure accessibility as the average of three indicator variables capturing (a) whether a model releases its parameter weights, (b) whether a model releases its training code, and (c) whether it has fewer than 3 billion parameters. We measure user cost as the dollar amount charged per million tokens.[3] Then we standardize these variables to construct *Accessibility* and *User Costs*.

To examine heterogeneous effects of open-source AI models across firms with varying labor-market exposure to Gen-AI, we construct a firm-level exposure measure using Revelio Labs' occupational employment data.[4] First, we aggregate employee-level records to obtain the number of employees per job position within each firm for each year, where job positions are defined at the 6-digit SOC occupation level. Next, we match each job position to its corresponding Gen-AI exposure score from Eisfeldt et al. (2023), which quantifies the share of tasks in a given occupation that Gen-AI can enhance to improve productivity.[5] Firm-level Gen-AI labor exposure, $AI\ Exposure^{raw}$, is then calculated as the weighted average of these occupation-level exposure scores, with weights proportional to the number of employees in each position. We use $AI\ Exposure^{raw}$ as a proxy for a firm's potential suitability for deploying Gen-AI applications.

A limitation of the measure proposed by Eisfeldt et al. (2023), $AI\ Exposure^{raw}$, is that it captures potential Gen-AI exposure, based on ChatGPT's assessment of task suitability, rather than actual adoption. To further verify whether workforce usage drives firm valuation, we construct an alternative metric, $AI\ Exposure^{claude}$, using the Anthropic Economic Index. Unlike the theoretical inferences of the prior measure, the Anthropic Economic Index is empirically grounded in millions of real-world user interactions with Claude.[6] We calculate this measure as the weighted average of occupation-level exposures, using employee counts per position as weights. This approach allows us to capture the realized Gen-AI exposure of

[3] We used the Wayback Machine to collect user cost data from the official websites of AI model publishers, open-source platform disclosures, and news media reports. Consequently, the information reflects the costs at the time of each model's release or within the news coverage immediately following those dates.

[4] See Appendix IA.2 for background information about Gen-AI.

[5] This task-based exposure methodology is also used in related work, including Auer et al. (2024).

[6] See: https://www.anthropic.com/news/the-anthropic-economic-index

the firm's workforce, and we use $AI\ Exposure^{claude}$ as another proxy for a firm's suitability for deploying Gen-AI applications.

The Anthropic Economic Index has two sub-indices: an Augmentation Index and a Displacement Index. The Augmentation Index measures the extent to which Gen-AI is used for creative and productivity-enhancing activities such as idea generation and problem solving. In contrast, the Displacement Index captures the extent to which Gen-AI is used for automation and efficiency-oriented tasks such as translation, writing, and document processing. Building on these measures, we construct the occupation-weighted firm-level AI exposure indices that separately capture AI-driven task augmentation and displacement: *AI Augmentation* and *AI Displacement*.

As a robustness check for AI exposure, we adopt an alternative approach of analyzing 10-K filings and directly identifying firms that explicitly mention AI-related keywords, following Lu et al. (2024) and Bertomeu et al. (2025). Specifically, we use a predefined dictionary of AI-related terms (see Appendix IA.3) to screen Items 1, 1A, and 7 of each firm's 10-K filing.[7] We extract sentences containing at least one AI-related keyword and manually verify whether they pertain to genuine AI-related business activities. We then construct a firm-level dummy variable, $AI\ Exposure^{10k}$, which equals 1 if the firm's operations are substantively linked to AI and the keyword is mentioned. Similarly, we use the presentation and all Q&A responses from the earnings call transcripts to construct a firm-level AI exposure measure, $AI\ Exposure^{ec}$. This method provides both historical (10-K) and forward-looking (earnings calls) measures of AI exposure.

Beyond the firm-level exposure to AI, we also construct AI exposure along supply chains. As the first step, we aggregate firm-level $AI\ Exposure^{raw}$ to industry-level $AI\ Exposure^{ind}$. Then, we use the U.S. domestic input-output table from the OECD in 2022 and construct industry weights based on the shares of a focal industry's purchases from upstream suppliers and sales to downstream customer industries. In particular, for each focal industry

[7] In Item 1, a firm describes its core operations, which include its technology adoption. In Item 1A, a firm discusses the risks it faces, including technological and competitive changes. Item 7 is the Management Discussion and Analysis – management's explanation of a firm's financial results and conditions.

with below-median industry AI exposure, we construct a measure of upstream-industry AI exposure that is based on upstream industries with above-median AI exposure as follows.

$$AI\ Exposure_j^{\text{upstream}} = \sum_{j \neq k} w_{jk}\, AI\ Exposure_k^{\text{ind}} \tag{3}$$

where for each industry j, w_{jk} represents the fraction of purchase from industry k.

Similarly, for each focal industry with below-median industry AI exposure, we construct a measure of down-industry AI exposure that is based on downstream industries with above-median AI exposure as follows.

$$AI\ Exposure_j^{\text{downstream}} = \sum_{j \neq k} w_{jk}\, AI\ Exposure_k^{\text{ind}} \tag{4}$$

where for each industry j, w_{jk} represents the fraction of sales to industry k.

C. Key Event Dates

First, we classify Gen-AI models into two categories: *proprietary* and *open-source*. *Proprietary* models are those for which the code, weights, and training data are not publicly released and whose use, modification, and redistribution are restricted by the developer—typically accessible only via an Application Programming Interface (API). These models are fully controlled by their creators and do not permit local deployment or transparency in implementation (e.g., GPT-3.5, Claude 2.0). *Open-source* models are those for which code, model weights, and/or training data are publicly shared and often deliver material gains in computational efficiency by reducing compute, memory usage, energy consumption, or latency during training or inference (e.g., DeepSeek-R1, Gemma-2B). This two-way classification allows us to isolate the effect of open-source AI releases, relative to proprietary releases, on firm valuations.

Then, we collect release events for influential proprietary and open-source Gen-AI models, along with their corresponding event dates, to construct a comprehensive timeline of major advancements.

For proprietary Gen-AI models, we focus on leading firms in the United States and China,

the two primary centers of AI innovation.[8] We hand-collect the initial public release dates of major models developed by five prominent U.S. companies (OpenAI, Anthropic, Amazon, xAI, and Google) and five leading Chinese firms (Baidu, Huawei, Tencent, ByteDance, and Moonshot AI). This effort yields a sample of 41 proprietary model releases from November 2022 to August 2025.

To identify influential open-source models, we leverage Hugging Face as the primary platform for disseminating open-source models. For each model type, we record cumulative downloads from its release date through our data collection cutoff on August 27, 2025. Within each month and year, we select the top five models by cumulative downloads to focus on those with the broadest adoption. To further ensure prominence, we restrict the sample to models released by widely recognized commercial developers such as Meta, Google, and OpenAI, thereby focusing on high-impact, credible open-source innovations. This process yields 26 open-source Gen-AI model release events over the same sample period.[9]

We then pool the proprietary and open-source model release dates into a single event timeline. In cases where multiple models were released on the same date, we retain only the model with the greatest expected influence, prioritizing those from more prominent developers or with higher anticipated reach. The final dataset is illustrated in Appendix IA.1 and comprises 67 unique event dates, representing influential Gen-AI model releases that serve as the basis for our event study analyses.

D. Empirical Strategy

We plan to empirically assess the economic impact of open-source AI by analyzing firm-level stock market reactions to model release announcements. Our study rests on the premise that while generative AI is a general-purpose technology, open-source models—unlike proprietary

[8] While notable models have emerged from other jurisdictions, the United States and China remain the two dominant centers of AI model development by a wide margin. See Maslej et al. (2025) and Rishabh and Shreeti (2026).

[9] Following Trajtenberg (1990), who established that patent quality is measured by future citations, we posit that the quality of a newly released open-source model is similarly captured by its future number of downloads. We recognize that future downloads could be correlated with contemporaneous model adoptions. However, this does not raise endogeneity concerns, because our dependent variables are driven by market expectations of model quality at the time of release, not by actual post-release usage or adoptions.

ones— enable more inclusive access and widespread incremental innovation.

Our empirical strategy has two components. The first component is presented in Section III, where we show that open-source model releases elicit more positive stock-market reactions than proprietary releases, reflecting lower user costs and greater accessibility. Consistent with this interpretation, investors expect open-source releases to facilitate more inclusive adoption, even among financially constrained firms. The second component is reported in Section IV, which shows that lower adoption barriers allow firms to unlock the latent value of their workforces, reflected in higher valuation gains for firms whose tasks are highly amenable to Gen-AI adoption. These gains are driven by augmentation-oriented AI applications. We further examine whether these anticipated efficiency gains generate broader economic effects through supply-chain spillovers.

To conduct these analyses, we extract for each model release date, a_g, the daily stock returns for all listed firms, i, in the U.S. over a time window spanning one trading day before and k trading days after a_g, $CAR[-1,+k]_{ig}$. We pool these event-specific samples across all model releases to yield the final dataset used in the regressions Equation 5.

$$CAR[-1,+k]_{ig} = \beta_0 + \beta_1 OS_g + \alpha X_i + FE + \varepsilon_{ig} \tag{5}$$

where i denotes the firm and g denotes the event. The explanatory variable, OS_g, is a dummy variable that equals 1 if event g is the release of an open-source Gen-AI model and 0 otherwise. By design, β_0 and β_1 capture the valuation effects of releasing proprietary and open-source models, respectively. The control variables X_i include firm size, leverage ratio, and ROA. FE represents firm fixed effects.

To examine heterogeneous effects by AI exposure, we interact the model release dummy with the degree of Gen-AI exposure. This classification enables us to estimate the differential stock market reactions between firms with higher versus lower exposure to Gen-AI.

$$\begin{aligned} CAR[-1,+k]_{ig} = \beta_0 + \beta_1 OS_g \times AI\ Exposure_{ig}^{raw} \\ + \beta_2 AI\ Exposure_{ig}^{raw} + \beta_3 OS_g + \alpha X_i + FE + \varepsilon_{ig} \end{aligned} \tag{6}$$

where $AIExposure_{ig}^{raw}$ is a weighted measure of Gen-AI exposure of firm i for a given event g, following Eisfeldt et al. (2023), and FE represents industry (2-digit NAICS) fixed effects. Alternatively, we use other measures of AI exposure, $AIExposure^{Claude}$, $AIExposure^{10-K}$, $AIExposure^{ec}$ sourced from Anthropic Economic Index, 10-K filing, and earnings call transcripts, respectively.

To assess spillover effects of AI exposure along supply chains, we examine stock market reactions of focal firms with varying upstream and downstream AI exposure around AI model release events, separately for proprietary and open-source models. In particular, we run the following specification:

$$\begin{aligned} CAR[-1,+k]_{ijg} = \beta_0 &+ \beta_1 AI\ Exposure_{jg}^{upstream} \\ &+ \beta_2 AI\ Exposure_{jg}^{downstream} \\ &+ \beta_3 AI\ Exposure_{jg}^{ind} \\ &+ \alpha X_i + FE + \varepsilon_{ijg} \end{aligned} \tag{7}$$

where $AIExposure_{jg}^{upstream}$ is a weighted measure of $AIExposure_{jg}^{ind}$ using industry-level purchase as weights, $AIExposure_{jg}^{downstream}$ is a weighted measure of $AIExposure_{jg}^{ind}$ using industry-level sales as weights, and $AIExposure_{jg}^{ind}$ denotes industry-level AI exposure for the focal industry j. FE represents industry (2-digit NAICS) fixed effects.

To mitigate confounders that may simultaneously influence Gen-AI model releases and stock market reactions, we further adopt a stacked difference-in-differences (DID) design as an alternative empirical specification in addition to Equation 6. Specifically, for event g, we use daily abnormal returns (AR) as our primary measure, rather than cumulative abnormal returns (CAR) over an event window. Then we estimate the following empirical model to estimate the impact on firm value across the nature of released Gen-AI models, following the insight of stacked DID:

$$\begin{aligned} AR_{itg} = \beta_0 + \beta_1 OS_g \times Post_{tg} \times AI\ Exposure_{ig} \\ + \beta_2 Post_{tg} \times AI\ Exposure_{ig} + \alpha X_i + \theta_{gi} + \gamma_{gt} + \varepsilon_{itg} \end{aligned} \tag{8}$$

where i denotes the firm, t denotes the date, and g denotes the event (Gen-AI model release). The outcome variable, AR_{itg}, is the market-adjusted return in percentage points for firm i on date t during the [-1, +5] time window around event g. The dummy variable, $Post_{tg}$, equals 1 if date t is on or after the event date a_g around event g, and 0 otherwise. $AIExposure_{ig}$ is a dummy variable that equals 1 if AI Exposure$_{ig}^{raw}$, a continuous Gen-AI exposure score from Eisfeldt et al. (2023), is above the sample median in the year of the given event g, and 0 otherwise. Control variables X_i include firm size, leverage ratio, and ROA. θ_{gi} represents the event-firm fixed effects. γ_{gt} represents the event-date fixed effects.

III. Empirical Results: Valuation Effects and Inclusiveness

A. The Evolution of AI Models

Before analyzing how proprietary and open-source model releases affect firm valuations, we first outline the evolution of the AI model development landscape. Figure 1 illustrates this progression by plotting the cumulative number of newly added generative AI models from November 2022 to August 2025. The trajectories reveal a stark contrast: proprietary models, sourced from Epoch AI (Panel A), show a steady linear increase from 0 to fewer than 400, while open-source models on Hugging Face (Panel B) exhibit dramatic exponential growth, surging from 0 to over 1.5 million. This explosive expansion signals a structural shift from a closed, vendor-controlled innovation regime to an open ecosystem where frontier capabilities diffuse at an accelerating rate.

Against this backdrop, Panel A of Figure 2 reports user costs, measured as USD per one million tokens, for each Gen-AI model in our sample.[10] After 2022, intensified competition coincides with a sharp decline in user costs. Over the same period, the supply of model-development talent expanded, with universities producing more AI-capable human capital,

[10]We measure cost based on the expense incurred when users access AI models through the publisher's API services. While deploying open-source models locally may result in lower marginal monetary costs compared to using paid API services, we rely on published API pricing to ensure a consistent comparison between proprietary and open-source models, as the costs of local deployment for proprietary models are not readily quantifiable. Additionally, user-cost data are missing for some models, and because models were released at different times, these comparisons reflect price differences over time rather than a contemporaneous cross-section.

further broadening participation. Panel B shows that open-source models are significantly less expensive, while proprietary models remain relatively costly.

Notably, the relationship between open and proprietary development is not one-sided. Proprietary models constitute an important upstream source of open-source development. Conversely, while open-weight models primarily advance open-source ecosystems, they may also indirectly benefit proprietary models: developers can learn from and integrate techniques derived from open releases. This interplay underscores the need for a nuanced measure of accessibility; beyond open weights alone, we incorporate open-source training code and parameter size into a broader accessibility index, which we detail in the next subsection.

This evolutionary pattern is consistent with the literature on a canonical innovation trajectory from disruptive proprietary innovation to process innovation and diffusion, characterized by declining adoption barriers and broader utilization. Proprietary models typically deliver radical advances but impose high user costs, limiting adoption to a narrow set of firms. In contrast, open-source AI progresses through rapid, successive releases, enabled by algorithmic and architectural improvements that accelerate cost reductions and expand accessibility. The second trajectory resembles the takeoff phase of technological diffusion, where cumulative improvements and decentralized collaboration drive rapid and broad adoption. As a general-purpose technology, Gen-AI can reshape economic activities in both phases. However, our intuition is that the gains from proprietary AI innovations are likely concentrated among a narrow set of resource-rich adopters, potentially disadvantaging others. In contrast, open-source AI, by systematically improving accessibility and affordability, supports more inclusive adoption across firms and industries.

B. Valuation Effects of AI Model Releases

Building on the preceding discussion, we analyze the firm-level valuation effects of AI model releases. We begin by contrasting stock price reactions to proprietary and open-source releases, then examine the distributions of these two sets of stock price reactions, and finally attribute the observed differences to heterogeneity in model user costs and accessibility.

Stock Price Reactions and AI Model Releases

Figure 3 shows cumulative abnormal returns (CAR as in Equation 1), around model release dates for open-source and proprietary models. The data reveal a stronger market response to open-source releases, with the return differential widening over the following two to three trading weeks.

We estimate the firm-level valuation impact of AI model releases using an event-study framework, as in Equation 2. Table II presents the abnormal stock market responses by model type. Panel A shows that, over the event window of [-1, +5] days, the average firm experiences a statistically significant CAR of -0.69 percentage points for proprietary model releases, compared to a statistically significant +0.17 percentage points for open-source releases. Panel B demonstrates a similar asymmetric pattern when cumulative raw returns (CRR as in Equation 1) are used instead of cumulative abnormal returns.[11]

As our event sample contains more proprietary releases than open-source releases, we address this imbalance through three robustness checks. First, in Appendix IC.2, we weight observations by market capitalization measured one trading day prior to each event, following Eisfeldt et al. (2023), rather than assigning equal weights. Second, in Appendix IC.3, we employ propensity score matching to balance firm-level covariates across release types, following Bilicka (2019). Third, in Appendix IC.4, we apply entropy balancing to ensure pre-event covariate distributions are comparable, following McMullin and Schonberger (2020). Across all specifications, our core result holds: proprietary releases are associated with negative stock market reactions, while open source releases elicit positive reactions.

In addition, in order to alleviate concerns that the estimates are confounded by contemporaneous macroeconomic news or systematic differences in release timing relative to the AI development cycle, we include year-month fixed effects as an alternative specification in Appendix IC.5. The results consistently show that stock market reactions are significantly negative for proprietary model releases, whereas they remain positive for open-source

[11] Instead of winsorizing cumulative abnormal returns in the baseline specification, we also examine the results without winsorization in Appendix IC.1. This robustness check confirms that our core findings are not driven by extreme values: stock market reactions remain significantly negative to proprietary AI model releases and significantly positive to open-source releases.

releases.

To further assess robustness, we pool the cumulative abnormal returns (CARs) and regress them on an open-source release dummy and controls, as in Equation 5. Table III shows that, compared to proprietary models, open-source AI releases are associated with significantly higher stock market valuations. On average, open-source releases generate a CAR that is 0.23 percentage points higher over the 2-day window, increasing to 0.35 percentage points and 0.87 percentage points over the 3-day and 7-day horizons, respectively. The impact of open-source releases remains persistently positive across time horizons. In contrast, the market reaction to proprietary releases is initially negative over the 3-day horizon and becomes insignificant thereafter. Overall, the sustained positive returns following open-source releases are consistent with the view that open, collaborative models facilitate faster innovation, broader scrutiny, and community-driven improvements, enhancing perceived model robustness and long-term value.

The Distribution of Stock Price Reactions and Model Releases

In this subsection, we investigate whether proprietary versus open-source model releases are associated with any differences in the distributions of stock price reactions. In Appendix IB.1, we find that while the overall variance in CARs is comparable across both release types, the distributions exhibit distinct skewness. In particular, the proprietary releases are associated with slightly negative skewness: 4.42% of CARs are 1.65 standard deviations above 0, compared to 6.11% below 0. In contrast, the open-source releases exhibit significantly positive skewness: 5.45% of CARs are 1.65 standard deviations above 0, compared to 4.90% below 0. Thus, CARs for proprietary releases are slightly heavier on the left tail, while those for open-source releases are slightly heavier on the right tail.

We then plot skewness at the 2-digit NAICS industry level in Figure 4, where the contrasting patterns become more evident. Panel A shows consistently negative skewness across industries following proprietary releases, while Panel B demonstrates consistently positive skewness following open-source releases. Panel C highlights differences in skewness in CAR[-1, +5] between open-source and proprietary releases across industries, with the gap consis-

tently positive, except in agriculture.

We formally test this heterogeneity in distributional moments in Table IV, where industry-level skewness is computed for each model release. Panel A of Column 1 confirms that open-source releases are associated with significantly greater positive skewness in CARs. Column 2 shows that this result holds for cumulative raw returns CRRs as well.

To explore the drivers of this result, we examine how the skewness of industry-level returns varies with an industry-level AI exposure measure, constructed by aggregating the firm-level metric from Eisfeldt et al. (2023). Panel B of Table IV presents regression estimates of the effects of open-source model releases stratified by firms' levels of AI exposure. It consistently shows that the positive skewness in CARs from open-source releases is concentrated in industries with high AI exposure. This result suggests the open-source models either create superstar potential for AI-exposed firms, protect them from left-tail risk, or do both.

To distinguish between these conjectures, Table V reports the results of firm-level quantile regressions of firms' cumulative abnormal returns, CAR[-1, +5], assessing whether the effects of open-source releases differ across the distribution of stock returns. Panels A and B present unconditional quantile regression results. The constant terms in Panel A show that proprietary AI releases have statistically significant negative effects on lower-ranked CARs and positive effects on higher-ranked CARs, resulting in minimal net change in overall skewness. In contrast, open-source model releases generally increase CARs across the distribution, producing an overall rightward shift.

Panel B further reveals that the left tail (5th percentile) gains associated with open-source releases are driven almost entirely by firms with high AI exposure, suggesting a strong protective effect against extreme negative valuation outcomes. However, in the right tail (95th percentile), high-AI-exposure firms do not experience additional gains. For proprietary releases, firms with high AI exposure show no corresponding changes in all but a significant left-tail loss.

Panels C and D, which employ conditional quantile regressions, yield similar findings:

open-source model releases raise firm valuations overall, with disproportionately large benefits observed in the lower tail of the return distribution.

Taken together, these results indicate that open-source model releases increase the valuation of AI-exposed firms, mitigating their downside risk. For firms without AI exposure, open-source adoption does not offer strong protection against extreme negative events and offers benefits in the middle of the return distribution. The asymmetric protection of the left tail for AI-exposed firms primarily drives the more positive aggregate skewness observed with open-source releases. In contrast, proprietary model releases fail to mitigate downside risk for AI-exposed firms.

User Costs and Accessibility Drive the Difference in Stock Price Reactions

This asymmetric distributional effect and the gaps in stock price reactions likely stem from the structural advantages of open-source models: their openness enables broad adoption, and their lower user costs reduce barriers to entry. To validate this conjecture, we run more direct validation tests examining accessibility and affordability as key drivers of open-source model impact, with results reported in Table VI.

Column 1 uses a standardized accessibility index—measured as the average of three binary indicators for whether the model provides open-source weights, provides open-source training code, and has a parameter count below 3 billion—and shows that more accessible models are associated with significantly higher firm valuations over the event window [-1, +5]. Column 2 uses the standardized user cost of AI models, originally measured in USD per million tokens, as a proxy for affordability, and shows that cheaper model releases are associated with higher firm valuations over the same event window. In economic terms, a one-standard-deviation reduction in user costs (11.542 USD per million tokens) is associated with a 0.313 percentage point increase in CAR[-1, +5]. Specifically, a 1 USD decrease per million tokens corresponds to a 0.027 percentage point gain in CAR[-1, +5]. Collectively, stock market reactions are consistently positive for AI models that are both widely available and economically viable, confirming that open access and cost efficiency are central to the process innovation enabled by open-source models, and notably absent in their proprietary

counterparts.

C. Open-Source Releases and Inclusiveness

In this subsection, we assess whether open-source AI releases produce stock price reactions consistent with inclusive adoption, relative to proprietary releases. Open-source releases, due to their lower user cost and higher accessibility, will generate substantially larger valuation gains for resource-constrained firms, whereas proprietary releases may yield muted or even negative reactions.

We employ a comprehensive set of firm-level proxies to capture these constraints, including firm asset size and two established measures of financial constraints from the corporate finance literature: the Size-Age (SA) index (Hadlock and Pierce, 2010), which incorporates firm size and age, and the Whited-Wu (WW) index (Whited and Wu, 2006), which incorporates cash flow, dividend payment status, leverage, firm size, industry sales growth, and firm sales growth. For each of these measures, we split firms into terciles and compare the top and bottom terciles by estimating the impact of open-source releases on CAR[-1, +5].

We then estimate the differential impact of open-source model releases across levels of resource constraints in Table VII using top- and bottom-tercile subsample analysis for each variable proxying for resource constraints. Panel A shows that while smaller firms experience negative stock returns upon proprietary model releases, valuation improvements follow open-source model releases. These findings corroborate the idea that open-source AI models lower critical access and cost barriers, particularly for resource-constrained businesses, aligning with Cai et al. (2025) and Bertomeu et al. (2025), who demonstrate that smaller firms derive greater benefits from AI adoption in entrepreneurial activities and employment.

Likewise, Panels B and C show that, relative to proprietary model releases, open-source releases lead to larger valuation gains for more financially constrained firms. By contrast, following proprietary model releases, valuation gains accrue primarily to less financially constrained firms.[12]

[12]Our point is that, relative to proprietary model releases, open-source releases lead to larger valuation gains for more financially constrained firms. Nevertheless, all three panels show that the stock price reactions to the model releases are greater for less financially constrained firms. For robustness check, we use alternative

Taken together, these findings highlight that open-source AI operates as a powerful, inclusive catch-up technology: it lowers entry barriers, enabling firms with limited financial resources to access frontier AI tools without bearing the full cost of innovation.

IV. Impact on Firms and Industries

The preceding section shows that open-source model releases generate broad-based valuation gains, with particularly strong effects for resource-constrained firms that benefit from lower adoption barriers. In this section, we examine the leverage firms obtain from the alignment between their task compositions and the suitability of those tasks for Gen-AI adoption. We argue that the reduced adoption frictions inherent in open-source models enable firms to more readily unlock the latent value of AI for workforce augmentation, while gains from labor displacement are less certain, perhaps due to organizational inertia. Accordingly, open-source releases elicit larger valuation gains for firms whose core tasks are highly amenable to AI-driven augmentation. We further examine whether these firm-level efficiency gains propagate more broadly through supply-chain spillovers.

A. Stock Price Responses and AI Suitability

To further assess which firms benefit from open-source models, we examine heterogeneous valuation responses across firms by assessing whether their tasks are well-suited to AI adoption.

First, we plot a time-series of the gap in abnormal returns between firms with high versus low AI exposure in Figure 5. Panel A shows that this gap is muted around the release of proprietary models. In contrast, Panel B shows that a clear gap emerges around the release of open-source models. This visual evidence suggests that firms with greater AI exposure benefit more from open-source model releases.

We then formally estimate this heterogeneity by interacting AI exposure with the open-

measures of adoption cost constraints in Appendix IC.6. These include employment size as a proxy for firm size; debt-to-asset ratio, interest expense scaled by revenue, and interest expense scaled by debt (following Frank and Shen (2016)) as proxies for higher costs of capital; and operating net cash flow and net cash flow over total assets as proxies for the ability to finance AI adoption. All these attempts provide evidence consistent with those reported in Table VII.

source release indicator in Table VIII. Consistent with the idea that lower adoption costs increase investors' expected impact of open-source AI, Panel A examines whether firms with workforces potentially more suited to AI adoption experience significantly greater valuation gains following open-source model releases. From the user-side perspective, a higher share of AI-relevant tasks facilitates the redeployment of labor toward AI-enabled workflows, reducing internal adjustment costs and strengthening the market response. In Panel B, we use $AI\ Exposure^{claude}$, a measure of realized AI usage intensity derived from the *Anthropic Economic Index*, as an alternative robustness check for AI exposure.[13]

We find that, over the window [-1, +5], the CAR is 1.05 percentage points for firms with a workforce characterized by high potential Gen-AI exposure, as shown in Panel A. Similarly, the corresponding CAR is 2.08 percentage points for firms with a workforce characterized by high realized AI usage intensity, as shown in Panel B. These results suggest that valuation gains accrue when firms are not only well-positioned to adopt AI but also actively using it, consistent with open-source availability lowering implementation frictions and thereby facilitating adoption. Interestingly, neither measure has any explanatory power for the CARs in the case of proprietary releases, further underscoring the distinction between open-source and proprietary models.[14]

These results indicate that open-source AI models are more effective than proprietary models in enabling firms to realize the potential benefits of AI innovation. Put differently, open-source models substantially amplify the realization of the potential benefits initially created by proprietary AI models.

The contrasting observations between open-source and proprietary AI releases in the current subsection and the previous sections align with the expectations that proprietary AI model releases impose net competitive pressures on most firms rather than generating broad-based benefits, potentially increasing precautionary savings (Andrews and Farboodi,

[13] $AI\ Exposure^{raw}$ and $AI\ Exposure^{claude}$ are highly correlated, with a correlation coefficient of 0.6. Given this level of multicollinearity, we refrain from including both variables in the same regression equation.

[14] In Appendix Section IC.7, we re-estimate the regressions under firm fixed-effect specifications and find robust results that firms with high potential Gen-AI exposure and those with high realized Gen-AI usage intensity experienced significant valuation gains following open-source releases.

2025). In contrast, open-source releases appear to produce more inclusive gains by increasing expected investment opportunities, which, in turn, raise long-term Treasury yields. These findings are consistent with Björkegren (2026), who shows that open and closed AI advancements have distinctly opposite effects on long-term interest rates.

B. Augmentation vs. Displacement-oriented AI

While the preceding results establish that valuation gains are concentrated among firms with high AI exposure, the economic impact likely depends on the specific nature of that exposure. Following Chen and Wang (2025), we distinguish the effects of augmentation- vs. displacement-oriented AI innovation on firm value. We leverage two sub-indices from the Anthropic Economic Index to distinguish between augmentation-oriented and displacement-oriented use. Panel A of Table IX shows that valuation gains around open-source releases are concentrated among firms whose AI use is more augmentation-oriented. Specifically, Column 3 indicates a 2.18 percentage point increase in CAR[-1, +5]. In contrast, Panel B shows that firms whose AI-usage is more displacement-oriented experience little change in valuation. These results suggest that while open-source releases affect firm value, the market primarily rewards productivity-enhancing augmentation applications rather than labor-substituting automation. The observation asymmetry is consistent with the idea that displacement entails cutting and creating new posts, which are likely subject to organizational inertia, whereas workforces accept augmentation more readily. Indeed, the literature on skill-biased and task-based technological change emphasizes that complementarities with existing labor tend to encounter lower implementation frictions than substitutionary adjustments (Autor et al., 2003; Acemoglu et al., 2016).

Notably, the standalone AI-augmentation variable is insignificant, suggesting that proprietary releases are less effective than open-source releases at generating positive expectations for augmentation-oriented AI innovation.[15]

We next partition the sample by asset size in Panel C. Column 1 shows that small firms experience the positive effect of augmentation-oriented AI applications only with open-source

[15]In Appendix IC.8, we re-estimate the regressions under firm fixed-effect specifications and find robust results that augmentation-oriented firms gained firm value following open-source releases.

releases. In contrast, large firms experience gains from augmentation-oriented adoption following both types of releases. This pattern suggests that the lower adoption costs of open-source models broaden access to AI-driven gains, particularly for resource-constrained firms.

For displacement-oriented applications, we do not observe significant negative effects, indicating that disruption within the startup ecosystem is insufficient to generate measurable adverse impacts. Overall, the results support a complementarity perspective: open-source models create the most value where firms possess the workforce structure and organizational capacity to convert reduced adoption costs into productivity-enhancing augmentation.

C. Spillover Effects Through the Supply Chain

Having taken a firm-level approach to examine how the degree of AI exposure affects firm valuations following the release of open-source models, we now turn to an industry-level approach by examining supply-chain channels. Specifically, we investigate how AI exposure among upstream suppliers and downstream customers affects the valuations of focal firms.

In Panel A of Table X, we find that firms in industries with low AI exposure experience an increase in firm value following open-source releases, driven by spillovers. This result holds even after controlling for the AI exposure of the firm's own industry, which typically exerts a larger direct effect. In Panel A, we find no evidence of supply chain spillovers following proprietary releases. On the contrary, we observe significant positive spillovers from both upstream and downstream industries with high AI exposure following open-source releases.[16]

In Figure 6, we plot dynamic spillover effects through the supply chain over a 20-day horizon. Again, we find insignificant spillovers following proprietary releases. In contrast, firms with low AI exposure experience persistent valuation gains following open-source releases from high-AI-exposure firms in both upstream and downstream industries.

We conjecture that firms in low-AI-exposure industries benefit from upstream spillovers,

[16] In Appendix IC.9, we use market capitalization as the weight for constructing industry-level AI exposure, and we find robust results of upstream and downstream spillovers. Furthermore, in Appendix IC.10, we use FactSet Revere to identify inter-firm relationships and compute new upstream and downstream exposure measures. We find robust evidence of positive spillovers from upstream suppliers and downstream customers in industries with high AI exposure to focal firms in industries with low AI exposure.

as they would gain access to more efficient production inputs and face reduced integration frictions. This eliminates the need for costly investments in standardizing or customizing inventories and products. Simultaneously, these firms benefit from downstream spillovers. High-AI-exposure customers would drive market expansion and enable superior demand forecasting, enabling low-AI-exposure firms to reduce inventory-holding costs. Although downstream firms may leverage their enhanced efficiency to negotiate lower purchase prices, the gains from increased sales volume and operational stability outweigh this negative price effect.

In Panels B and C of Table X, we split the sample by firm size and financial constraints to examine heterogeneous spillover effects following open-source releases. We find that larger firms and firms with financial resources do not receive any positive spillovers. In contrast, smaller firms and firms with higher financial constraints drive these benefits. Unlike large incumbents, these firms typically lack the resources to maintain extensive internal departments for quality control, procurement, or logistics. Consequently, upstream AI adoption would serve as a substitute for these missing capabilities, enabling small firms with financial constraints to obtain higher-quality inputs without incurring the associated monitoring costs. Furthermore, small firms may suffer less from organizational inertia, enabling them to adapt more quickly to reconfigured supply chains.

Similarly, financially constrained firms often lack internal market research capabilities. When downstream customers utilize AI to forecast demand accurately and share production schedules, they would effectively externalize this planning capability to the focal firm. While large firms possess the financial slack to weather demand volatility, small firms and financially constrained firms are highly vulnerable to inventory shocks. Therefore, the demand stabilization provided by downstream AI would offer a critical liquidity benefit for the focal firm.

In Panels D and E, we conduct analogous analyses of stock market reactions to proprietary releases, stratified by firm size and financial constraints. In contrast to open-source releases, where we observed distinct spillover benefits for financially constrained and smaller

firms, the market reaction to proprietary releases exhibits a fundamentally different pattern: the coefficients for upstream and downstream industry-level AI exposure are statistically insignificant across all specifications. This suggests that, unlike open-source technology, which diffuses rapidly due to low adoption barriers, proprietary models with licensing costs and closed ecosystems do not generate immediate, value-relevant spillovers for supply chain partners.

Taken together, we provide evidence that the value creation of open-source AI propagates through supply chain networks, accruing to smaller and more financially constrained firms with low AI exposure through their economic ties to upstream and downstream industries with high AI exposure. Larger and less financially constrained firms with low AI exposure do not receive these spillover benefits. Furthermore, there is no spillover benefit across supply chains from the release of proprietary models.

V. Robustness Checks

In this robustness section, we examine the validity and generalizability of the findings through multiple analytical lenses. First, as an alternative specification to the event study approach, we employ a stacked difference-in-differences strategy specified in Equation 8, with results presented in Table XI. Columns 1 and 2 indicate that, following the release of open-source models, as opposed to proprietary models, firms with higher AI exposure experienced significantly improved firm valuations. Columns 3 and 4 confirm that these results are robust in the propensity score-matched (PSM) sample.[17]

In addition, we evaluate alternative measures of firms' exposure to Gen-AI, drawing on multiple sources, to confirm the consistency of market reactions to open-source model releases. Table XII presents the AI exposure measure based on 10-K filings and earnings call transcripts, along with the results from Equation 6. All of them confirm that firms more

[17]Following Bilicka (2019), we use propensity score matching to limit the sample of control firms. Specifically, we estimate a logit model for each event g with the treatment indicator $AIExposure_{ig}$ on the left-hand side and pre-event outcome variables (i.e., abnormal returns from trading days t=-5 to t=-2 before the model release) as the predictors on the right-hand side, which gives us propensity scores. We use a nearest-neighborhood matching strategy within a small caliper radius (i.e., 0.01 standard deviation of the logit of the propensity scores) without replacement, which, for each treated firm with higher exposure to Gen-AI, finds the closest comparable firm with lower exposure to Gen-AI.

deeply integrated with AI technologies show stronger positive market reactions to the release of open-source models. The robustness of this finding underscores the market's recognition of open-source innovation as a key driver of value creation in AI-intensive firms.

Furthermore, we use alternative outcome variables instead of cumulative abnormal returns (CAR) to enhance robustness in Panel A of Table XIII. We use cumulative raw returns (CRR) as an alternative outcome variable, as the stock market may have entered a new regime following the emergence of AI models. In such a scenario, the estimated beta may no longer accurately reflect the underlying market dynamics, potentially biasing abnormal return calculations. It shows that the significant gains for firms with higher AI exposure following the release of open-source models persist even when raw returns are used. The magnitude of the effect is economically meaningful, suggesting that investors positively revalued firms well-positioned to leverage these technological advancements.

Moreover, we account for systematic risks beyond market exposure when examining the impact of AI model releases. Specifically, we use abnormal returns estimated from the Fama-French three-factor model as an alternative outcome variable in Panel B of Table XIII. It again shows that firms better positioned to adopt AI innovation benefited from the release of open-source models. Hence, the positive stock market reactions are robust across different return measures.

Recent work demonstrates that linear two-way fixed-effects (TWFE) estimators can be misleading in settings with staggered treatment adoption and heterogeneous dynamic effects, because the TWFE coefficient is a non-convex weighted average of cohort- and horizon-specific treatment effects. Our stacked difference-in-differences design mitigates these weighting issues by aligning firms in event time and restricting the estimation sample to a narrow window around each event, but identification still relies on a parallel-trends assumption for untreated potential outcomes. Following Liu et al. (2024b), we therefore complement our baseline estimates with three counterfactual estimators that accommodate richer unobserved time-varying heterogeneity: (i) the fixed-effects counterfactual estimator (FEct), (ii) the interactive fixed-effects counterfactual estimator (iFEct), and (iii) the matrix completion

estimator (MC), and (iv) synthetic matching.

Specifically, for each event, we define treated firms as those with $AIExposure_{ig}^{raw}$ above the sample median for the corresponding year. For each treated observation, we predict the counterfactual outcome, denoted as *Counterfactual* AR_{itg}, and then compute the individual treatment effect as the difference between the actual and counterfactual outcomes. This procedure yields a panel of individual treatment effects for all firms in the treatment groups over the periods following the open-source model releases. Using this panel, we regress the individual treatment effects on one dummy variable, OS_g, and report the estimation results in Table XIV. Column 1 presents the results based on the FEct estimator, Column 2 presents the results based on the iFEct estimator, and Column 3 presents the results based on the MC estimator. The results indicate that, regardless of the counterfactual estimation method employed, our baseline findings remain robust.

Additionally, we implement a synthetic matching methodology, following Goldsmith-Pinkham and Lyu (2025), which has advantages over standard propensity score matching, because it compares firms that are similar in the behavior of their pre-event abnormal returns. We use the period from 250 trading days to 30 trading days before the announcement date of the first event (the release of OpenAI's GPT-3.5) as the estimation window. For each firm in the treatment group with $AIExposure_{ig}^{raw}$ above the sample median for the corresponding year, we obtain the weights for firms in the synthetic control group by minimizing the discrepancy between the treated firm's daily abnormal returns and the weighted average of the synthetic control group's daily abnormal returns over the estimation window. For each trading day after the event date, we calculate the counterfactual outcome for each treated firm as the weighted average of the abnormal returns of firms in the synthetic control group, using the pre-estimated weights. The results reported in Column 4 of Table XIV further support the robustness of our baseline findings.

VI. Conclusion

This paper examines the economic impacts of open-source AI model releases on firm value. Using an event study of major AI model announcements from 2022 to 2025, we find that

open-source releases generate positive stock market reactions, while proprietary releases elicit negligible responses for most firms. Notably, these effects exhibit significant heterogeneity. In particular, valuation gains are pronounced for smaller, financially constrained firms, whereas these firms exhibit muted value changes following proprietary AI model releases. This pattern supports the view that open-source AI models lower deployment barriers, facilitating broader access and adoption.

Using labor-based measures of AI exposure, we find that valuation gains are concentrated in firms with greater AI suitability and higher realized AI usage, particularly when AI is used to augment rather than replace human input. These findings highlight the critical need for firms to align their human capital strategies with AI-related innovations. By integrating AI to complement human skills—particularly in creativity, analytical thinking, and technical expertise—firms can enhance workforce capabilities, enhance innovation, and maintain a competitive edge in an AI-driven economy.

Finally, we observe positive valuation spillovers from open-source model releases, benefiting firms in low-AI-exposure industries through AI adoption in both upstream and downstream industries. Absent these spillovers, such firms would experience negative valuation effects. These benefits are particularly pronounced for smaller, financially constrained firms. In contrast, proprietary model releases are associated with negative valuation responses for low-exposure firms. These findings highlight the broader and more diffuse impacts of open-source AI models.

Overall, this study highlights investor expectations that open-source AI promotes innovation, operational efficiency, and productivity, as well as expansive attempts to develop and deploy AI at the firm level. It also holds significant implications for developing economies, which often lack resources for capital-intensive AI advancements, but support many small-scale businesses focused on AI applications. These economies can benefit from the accessibility and affordability of open-source AI, provided they overcome bottlenecks such as adoption costs and entry barriers, and further promote technological diffusion through workforce training or trade with high-AI-exposure countries that host upstream and downstream industries.

References

Acemoglu, Daron, Simon Johnson, Amir Kermani, James Kwak, and Todd Mitton, 2016, The value of connections in turbulent times: Evidence from the United States, *Journal of Financial Economics* 121, 368–391.

Alexopoulos, M., and J Cohen, 2016, The medium is the measure: Technical change and employment 1909—1949, *Review of economics and statistics* 98, 792–810.

Andrews, Isaiah, and Maryam Farboodi, 2025, Do markets believe in transformative AI?

Auer, Raphael, David Köpfer, and Josef Sveda, 2024, The rise of generative artificial intelligence: Modelling effects on the US labour market.

Autor, David H., Frank Levy, and Richard J. Murnane, 2003, The skill content of recent technological change: An empirical exploration get access arrow, *The Quarterly Journal of Economics* 118, 1279–1333.

Bekar, C., K. Carlaw, and R Lipsey, 2018, General purpose technologies in theory, application and controversy: a review, *Journal of Evolutionary Economics* 28, 1005–1033.

Bertomeu, Jeremy, Yupeng Lin, Yibin Liu, and Zhenghui Ni, 2025, The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy, *Journal of Accounting and Economics* 1, 101782.

Bilicka, K. A., 2019, Comparing UK tax returns of foreign multinationals to matched domestic firms., *American Economic Review* 109, 2921–2953.

Björkegren, Daniel, 2026, Market beliefs about open vs. closed AI.

Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, et al., 2021, On the opportunities and risks of foundation models, arXiv:2108.07258.

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

Brynjolfsson, E., and L. M Hitt, 2000, Beyond computation: Information technology, organizational transformation and business performance., *Journal of Economic Perspectives* 14, 23–48.

Brynjolfsson, E., D. Rock, and C Syverson, 2021, The productivity J-curve: How intangibles complement general purpose technologies, *American Economic Journal: Macroeconomics* 13, 333–372.

Cai, Jeff Junhui, Xian Gu, Liugang Sheng, Linda Zhao Mengjia Xia, and Wu Zhu, 2025, AI as "co-founder": Generative AI for entrepreneurship, Working Paper.

Carlaw, K. I., and R. G. Lipsey, 2002, Externalities, technological complementarities and sustained economic growth, *Research Policy* 31, 1305–1315.

Chen, Mark A., and Joanna Wang, 2025, Displacement or augmentation? The effects of AI innovation on workforce dynamics and firm value, Working Paper.

Cheng, D., H. Crucini, M. J.and Oh, and H Yilmazkuday, 2025, Early 20th century american exceptionalism: Production, trade and diffusion of the automobile, *Journal of International Economics* 153.

Chun, H., J. W. Kim, and R Morck, 2011, Varying heterogeneity among us firms: facts and implications., *Review of economics and statistics* 93, 1034–1052.

Chun, H., and D Shin, 2025, Beyond automation: The multifaceted impact of advanced digital technologies on employment dynamics, *Research Policy* 54, 105321.

Dai, Damai, Chengqi Deng, Chenggang Zhao, R.X. Xu, Huazuo Gao, Deli Chen, et al., 2024, Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models, arXiv:2401.06066.

David, Paul A., and Julie Ann Bunn, 1988, The economics of gateway technologies and network evolution: Lessons from electricity supply history, *Information Economics and Policy* 3, 165–202.

Demirer, Mert, Andrey Fradkin, Nadav Tadelis, and Sida Peng, 2025, The emerging market for intelligence: Pricing, supply, and demand for LLMs, NBER Working paper.

Eisfeldt, Andrea L., Gregor Schubert, Bledi Taska, and Miao Ben Zhang, 2023, Generative AI and firm values, NBER Working Paper.

Eli, Hausman J. K., S. and, and P. W. Rhode, 2025, The model T, *The Journal of Economic History* 85, 110–151.

Eli, Shari, Joshua K. Hausman, and Paul W. Rhode, 2022, Transportation revolution: The car in the 1920s, *AEA Papers and Proceedings* 112, 219–23.

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux, 2009, Unconditional quantile regressions, *Econometrica* 77, 953–973.

Foster, Christopher, 2025, Openness in AI and downstream governance: A global value chain approach, arXiv:2509.10220.

Frank, Murray Z., and Tao Shen, 2016, Investment and the weighted average cost of capital, *Journal of Financial Economics* 119, 300–315.

Giglio, Stefano, Bryan Kelly, and Seth Pruitt, 2016, Systemic risk and the macroeconomy: An empirical evaluation, *Journal of Financial Economics* 119, 457–471.

Goldfarb, A., B. Taska, and F Teodoridis, 2023, Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings., *Research Policy* 52, 104653.

Goldsmith-Pinkham, Paul, and Tianshu Lyu, 2025, Causal inference in financial event studies, Working Paper.

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

Kapoor, Sayash, Rishi Bommasani, Kevin Klyman, Shayne Longpre, Ashwin Ramaswami, Peter Cihon, et al., 2024, On the societal impact of open foundation models, arXiv:2403.07918.

Kassem, D., 2024, Does electrification cause industrial development? Grid expansion and firm turnover in Indonesia, *Journal of Development Economics* 167.

Lewis, J., and E. Severnini, 2020, Short-and long-run impacts of rural electrification: Evidence from the historical rollout of the US power grid., *Journal of Development Economics* 143.

Liu, Aixin, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, et al., 2024a, Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model, arXiv:2405.04434.

Liu, Licheng, Ye Wang, and Yiqing Xu, 2024b, A practical guide to counterfactual estimators for causal inference with time-series cross-sectional data, *American Journal of Political Science* 68, 160–176.

Lo, S. T., and Sutthiphisal, 2008, Crossover inventions and knowledge diffusion of general purpose technologies? Evidence from the electrical technology, *National Bureau of Economic Research* No. w14043.

Lu, Yao, Gordon M. Phillips, and Jia Yang, 2024, The impact of cloud computing and AI on industry dynamics and concentration.

Machado, José A.F., and J.M.C. Santos Silva, 2019, Quantiles via moments, *Journal of Econometrics* 213, 145–173, Annals: In Honor of Roger Koenker.

Maslej, Nestor, Loredana Fattorini, Raymond Perrault, et al., 2025, Artificial intelligence index report 2025, Technical report, Stanford University, Human-Centered Artificial Intelligence (HAI).

McCulloch, Warren .S., and Walter H. Pitts, 1943, A logical calculus of the ideas immanent in nervous activity, *Bulletin of Mathematical Biology* 5, 115–133.

McMullin, Jeff L., and Bryce Schonberger, 2020, Entropy-balanced accruals, *Review of Accounting Studies* 25, 84–119.

McNerney, James, J. Doyne Farmer, and Jessika E. Trancik, 2011, Historical costs of coal-fired electricity and implications for the future, *Energy Policy* 39.

Melville, Kraemer K. Gurbaxani V., N., 2004, Information technology and organizational performance: An integrative model of it business value, *MIS quarterly* 283–322.

Mo, Hongwei, and Shumiao Ouyang, 2025, AI in financial economics, *Journal of Chinese Economic and Business Studies* 23, 509–587.

Montobbio, Staccioli J. Virgillito M. E. Vivarelli M., F., 2024, The empirics of technology, employment and occupations: Lessons learned and challenges ahead., *Journal of Economic Surveys* 38, 1622–1655.

Petralia, S, 2020, Mapping general purpose technologies with patent data, *Research Policy* 49, 104013.

Rios Avila, Fernando, 2019, Recentered influence functions in stata: Methods for analyzing the determinants of poverty and inequality.

Rishabh, Kumar, and Vatsala Shreeti, 2026, The geography of AI firms, BIS Working Paper.

Rosenberg, N., and M Trajtenberg, 2004, A general-purpose technology at work: The corliss steam engine in the late-nineteenth-century United States., *The Journal of Economic History* 64, 61–99.

Seely, B. E, 1987, *Building the American Highway System: Engineers as Policy Makers*.

Shrestha, Yash Raj, Georg von Krogh, and Stefan Feuerriegel, 2023, Building open-source AI, *Nature Computational Science* 3, 908–911.

Spear, Brian, 2008, James watt: The steam engine and the commercialization of patents, *World Patent Information* 30, 53–58.

Tanriverdi, H, 2006, Performance effects of information technology synergies in multibusiness firms, *MIS quarterly* 57–77.

Trajtenberg, Manuel, 1990, A penny for your quotes: Patent citations and the value of innovations, *RAND Journal of Economics* 21, 172–187.

Turing, Alan M., 1950, Computing machinery and intelligence, *Mind* 59, 433–460.

U.S. Bureau of the Census, 1922, Central electric light and power stations in the United States.

Wagner, Alexander F., Richard J. Zeckhauser, and Alexandre Ziegler, 2018, Company stock price reactions to the 2016 election shock: Trump, taxes, and trade, *Journal of Financial Economics* 130, 428–451.

Whited, Toni M., and Guojun Wu, 2006, Financial constraints risk, *The Review of Financial Studies* 19, 531–559.

Woolf, Arthur G., 1987, The residential adoption of electricity in early twentieth-century, *The Energy Journal* 8, 19–30.

Yeverechyahu, Doron, Raveesh Mayya, and Gal Oestreicher-Singer, 2024, The impact of large language models on open-source innovation: Evidence from github copilot, arXiv:2409.08379.

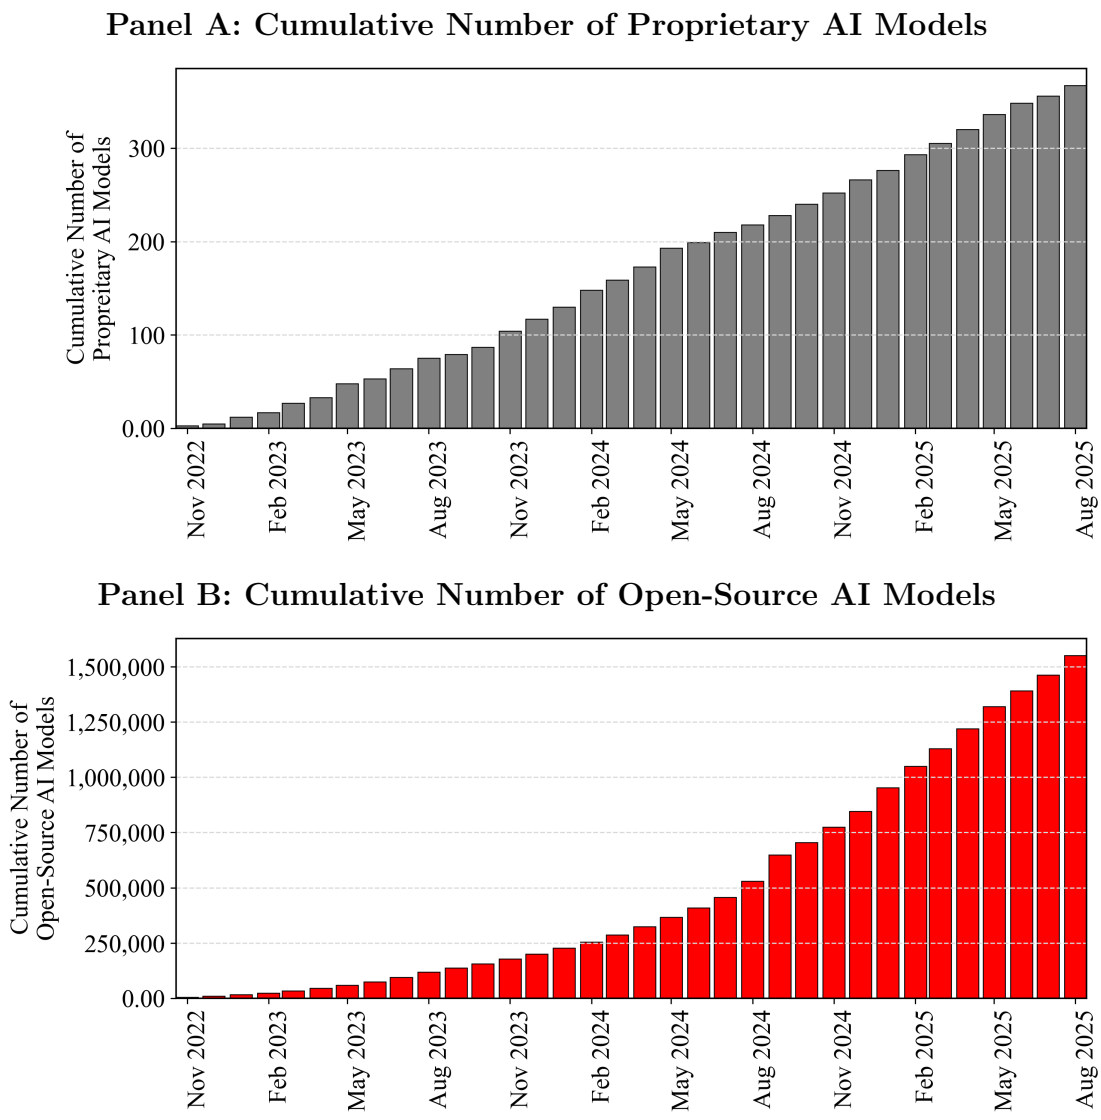


Figure 1. Cumulative Number of AI Models: Proprietary vs. Open-Source

This figure illustrates the time-series of the cumulative number of AI models for text generation from November 2022 to August 2025. Panel A shows the cumulative number of newly added proprietary AI models accessible via API. Data on proprietary models sourced from Epoch AI and manually verified. Panel B shows the cumulative number of open-source AI models released on Hugging Face, the primary platform for disseminating open-source models.

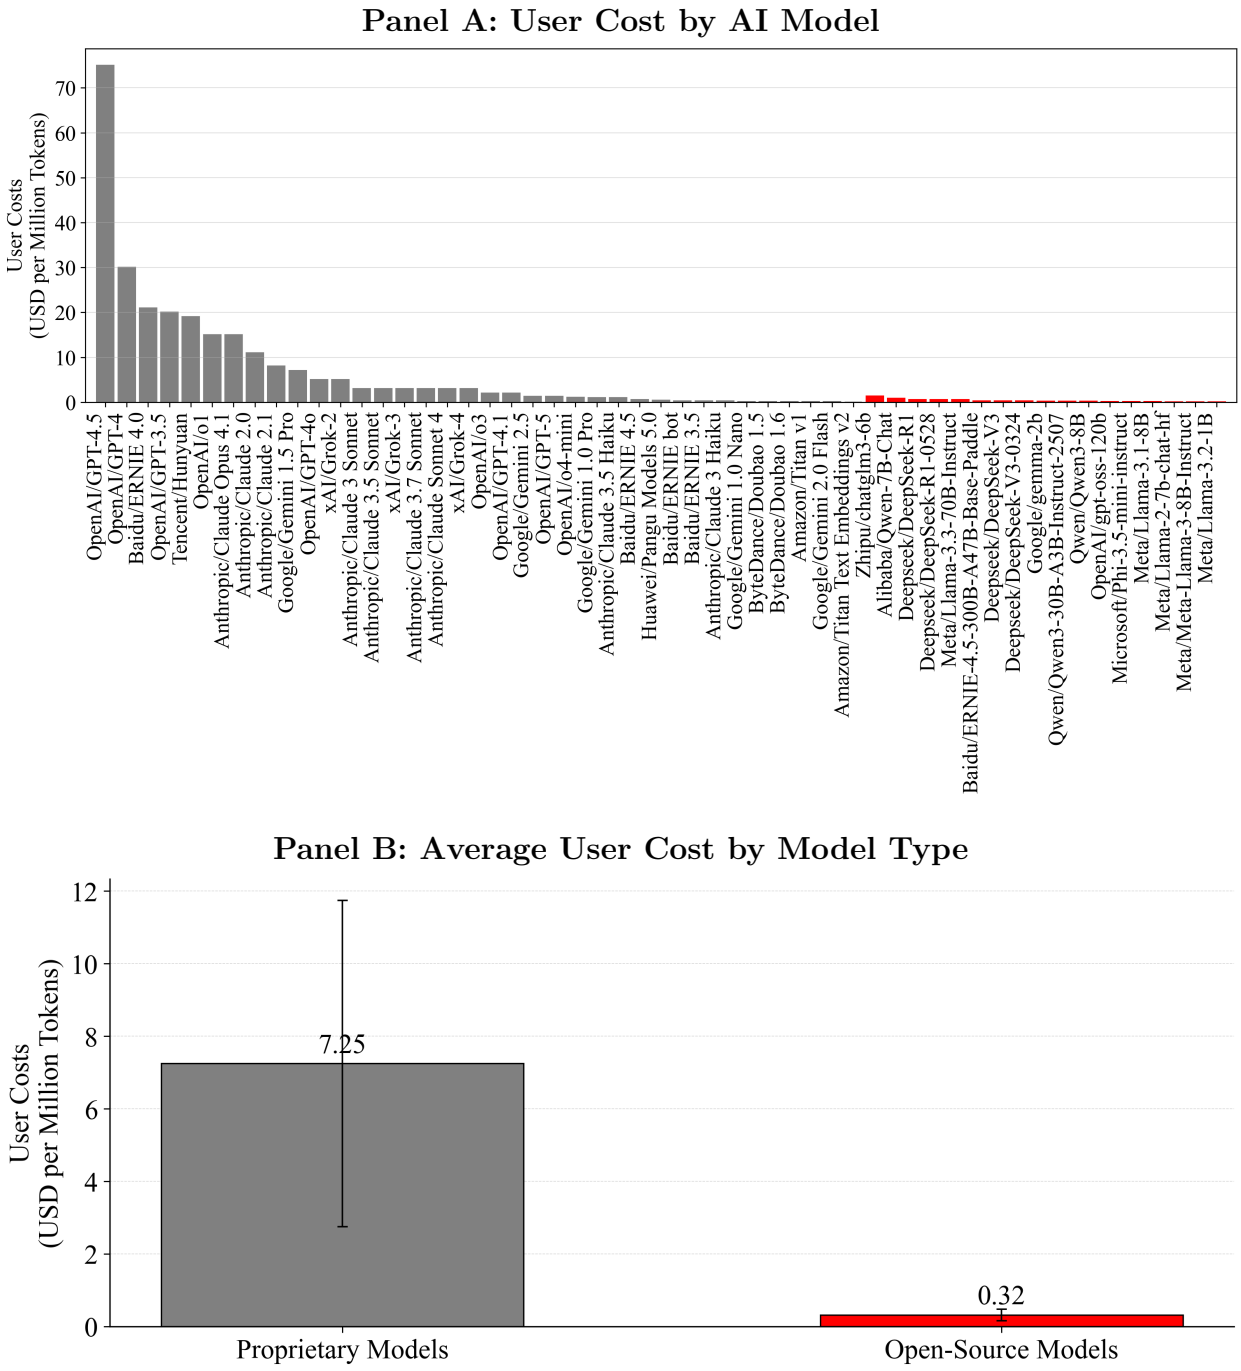


Figure 2. User Costs of AI Models: Proprietary vs. Open-Source

This figure presents the average user costs (USD per million tokens) for Gen-AI models listed in Appendix IA.1. Panel A displays the user cost for each individual model, whereas Panel B shows the average user cost by model type. Gray bars represent proprietary models, and red bars represent open-source models. Confidence intervals are shown at the 5% significance level. Note that pricing data are unavailable for some models, as they are not yet commercially available.

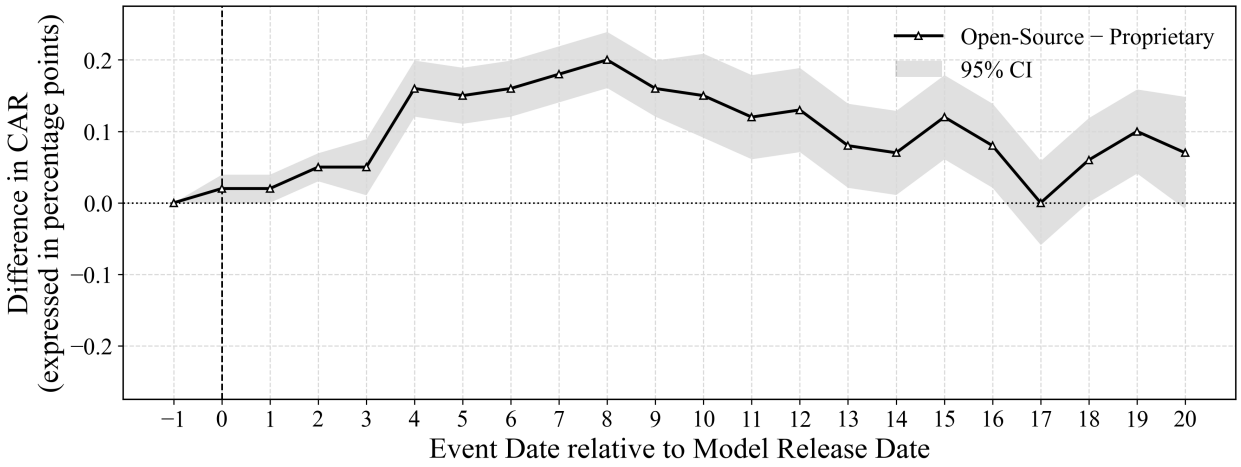
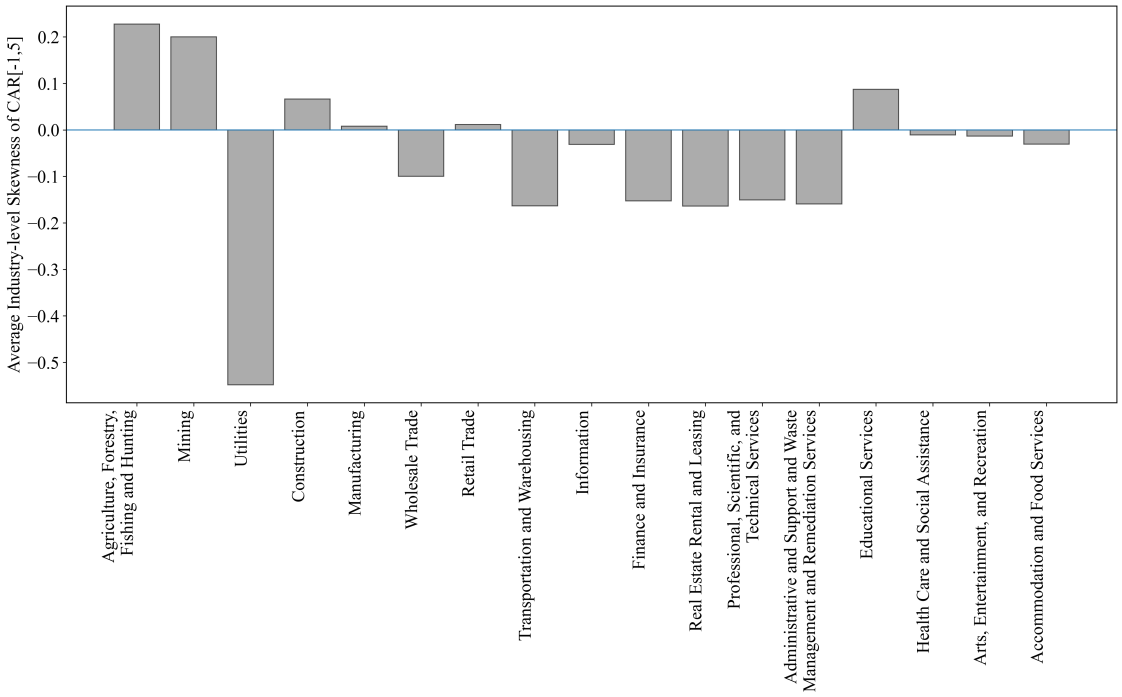


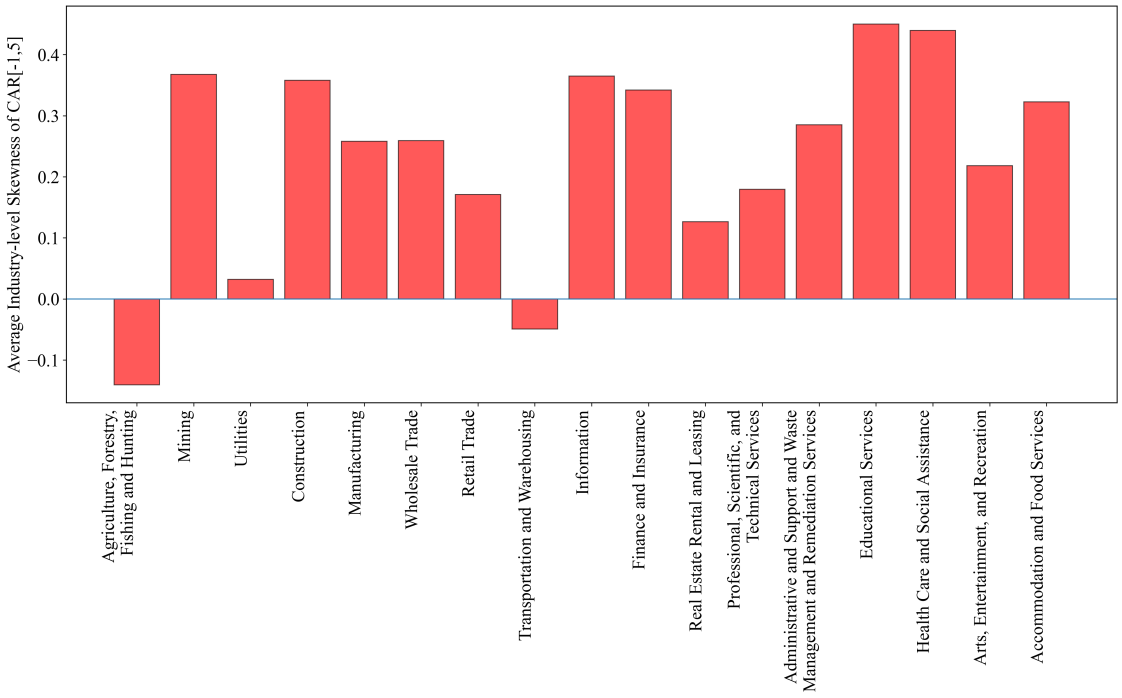
Figure 3. Market-Implied Valuation Gap: Open-Source vs. Proprietary

The figure illustrates the market-implied valuation gap between open-source and proprietary model releases. The solid line represents the cumulative difference in firm valuations following open-source versus proprietary model releases, while the shaded band indicates 95% confidence intervals.

Panel A: Skewness of CAR[-1,+5] by Industry for Proprietary Models



Panel B: Skewness of CAR[-1,+5] by Industry for Open-Source Models



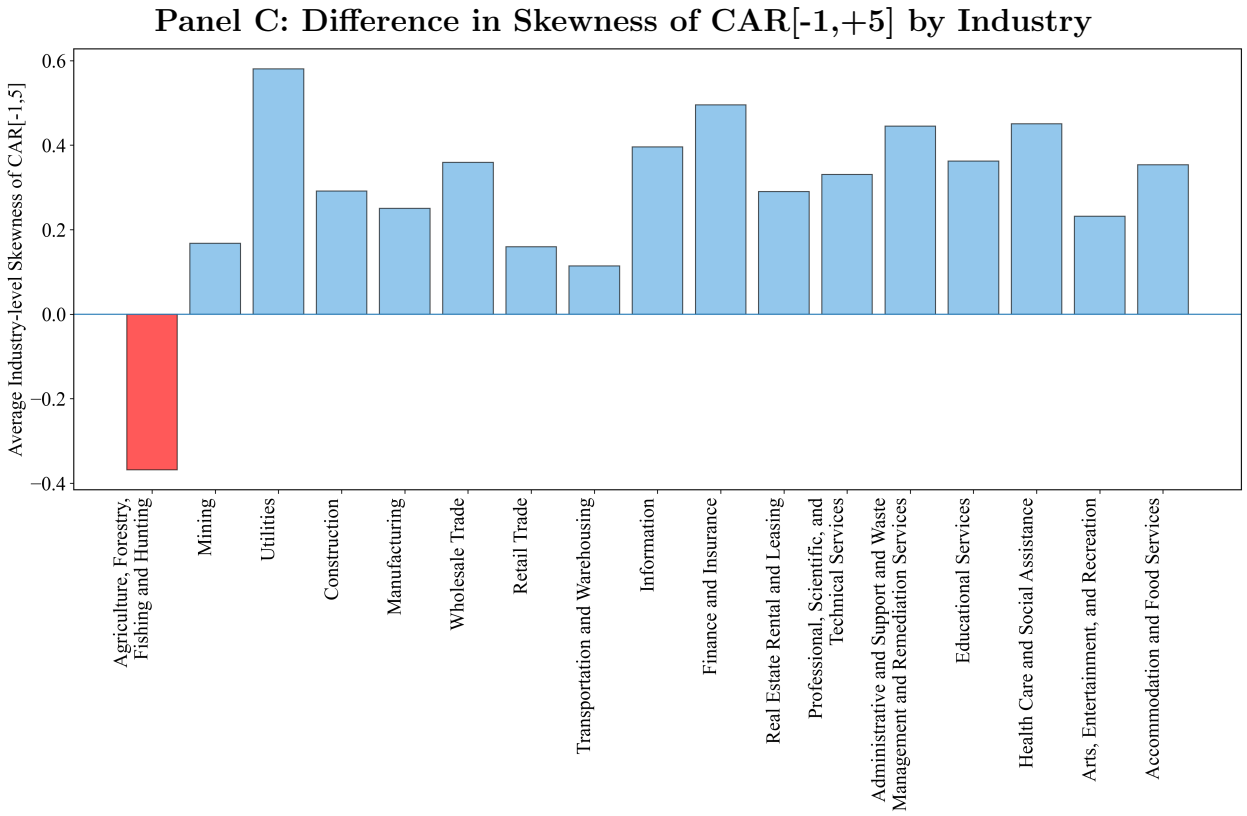


Figure 4. Heterogeneity in Skewness by Industry: Proprietary vs. Open-Source

This figure presents the skewness of cumulative abnormal returns over the event window [–1, 5] across industries, classified by two-digit NAICS codes. Panel A shows skewness around proprietary model releases. Panel B shows skewness around open-source model releases. Panel C shows the difference in skewness between open-source and proprietary models.

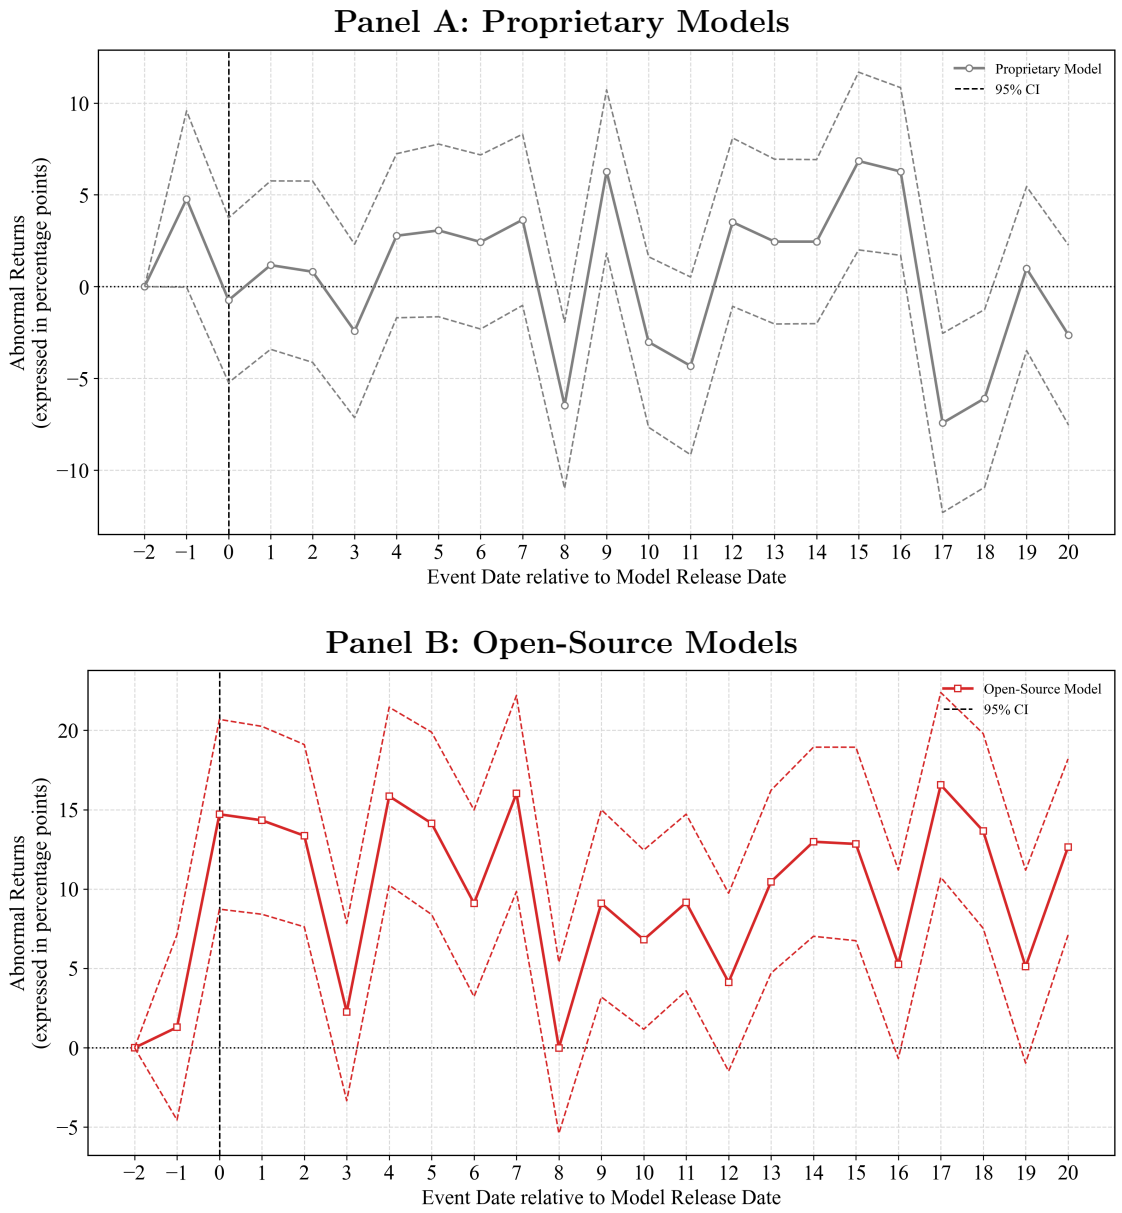
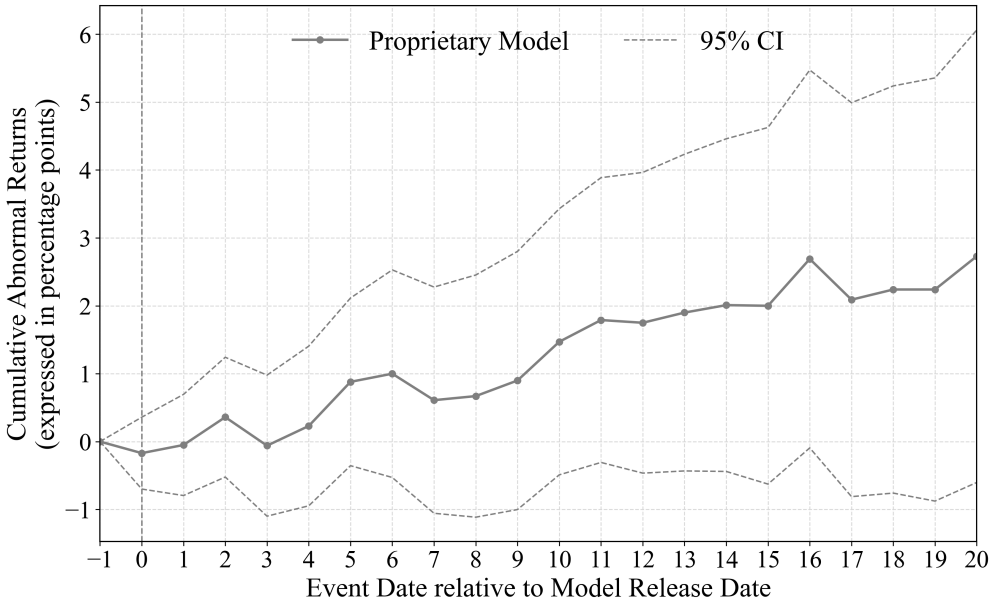


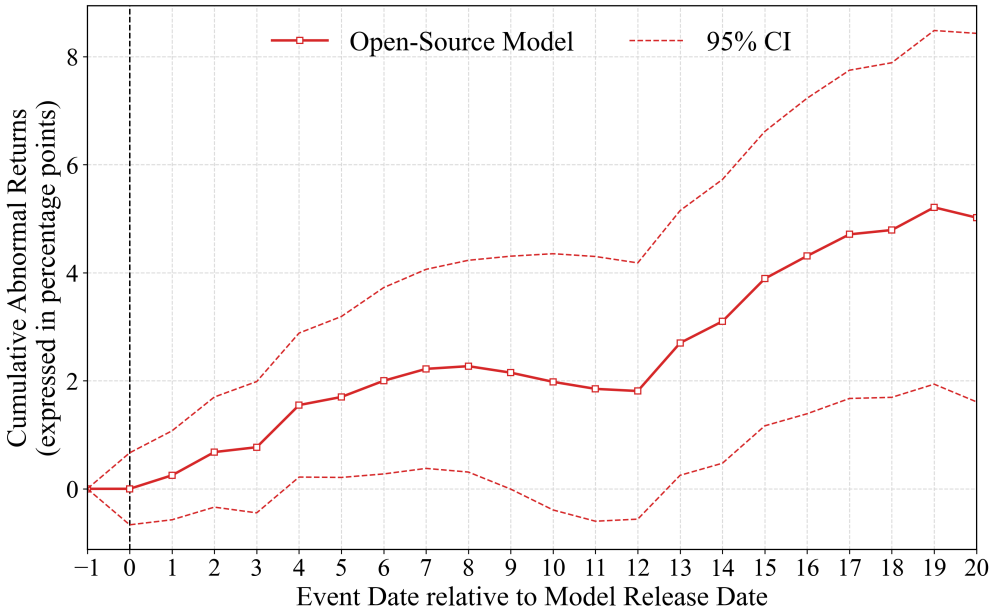
Figure 5. Heterogeneity in Abnormal Returns by AI Exposure

This figure plots the difference in daily abnormal returns over [-1,20] event window, expressed in percentage points, between firms with AI exposure above and below the sample median over the event window surrounding model releases. The sample is restricted to the period from two trading days before to twenty trading days after the release date, with the second pre-release trading day serving as the baseline. Panel A restricts the sample to proprietary model releases, and Panel B restricts the sample to open-source model releases. AI exposure is measured by AI Exposure$_{ig}^{raw}$, a continuous Gen-AI exposure score from Eisfeldt et al. (2023), which quantifies the share of tasks in a given occupation that Gen-AI can enhance to improve productivity.

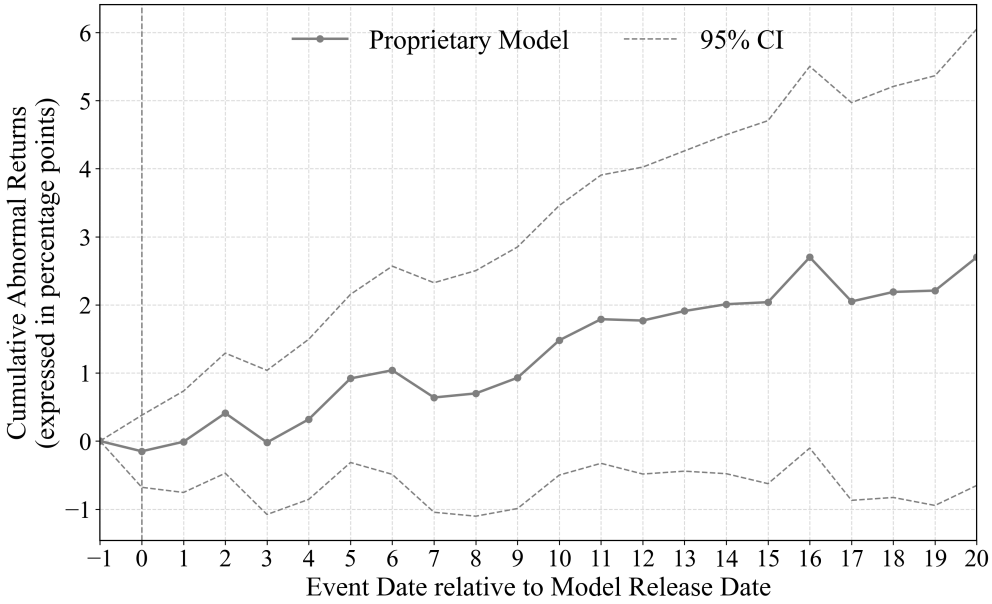
Panel A: High Upstream AI Exposure and Proprietary Releases



Panel B: High Upstream AI Exposure and Open-Source Releases



Panel C: High Downstream AI Exposure and Proprietary Releases



Panel D: High Downstream AI Exposure and Open-Source Releases

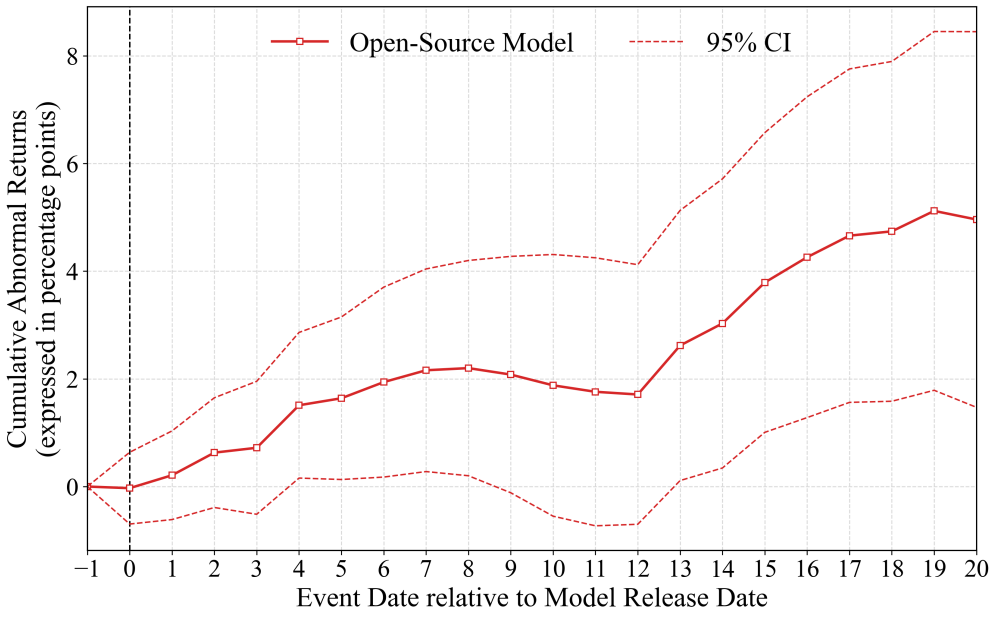


Figure 6. Dynamic Spillover Effects Through Supply Chain

Note: This figure illustrates how stock market reactions to proprietary and open-source AI model releases spill over through the supply chain. The empirical specifications and construction of upstream and downstream AI exposure are consistent with those used in Table X. Panels A and B plot responses to upstream AI exposure, while Panels C and D plot responses to downstream AI exposure. Confidence intervals are indicated at the 5% significance level.

Table I
Open-Source vs. Proprietary AI Models: Technical and Economic Distinctions

	Open-Source Models	Proprietary Models
Accessibility	Model parameters, source code, and/or data information released; runnable locally under license.	Access via API or closed product terms.
Costs	Lower marginal inference costs at scale when self-hosted; requires compute and engineering setup.	Usage-based pricing (e.g., per token/request), which potentially brings higher long-term costs.
Customization	Fine-tuning and architectural modifications are possible.	Customization is constrained (prompting and provider-approved fine-tuning).

Sources: Shrestha et al. (2023), Yeverechyahu et al. (2024), Bommasani et al. (2021), and Kapoor et al. (2024)

Table II
Market Response by Model Type: Open-Source vs. Proprietary

This table reports stock market reactions to the release of proprietary and open-source AI models. In Panel A, the dependent variable is the cumulative abnormal return (in percentage points) for firm i from one trading day before to five trading days after event g. In Panel B, the dependent variable is the cumulative raw return (in percentage points) for firm i over the same event window. Column 1 presents results for the subsample of proprietary AI models. Column 2 presents results for the subsample of open-source AI models. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Cumulative Abnormal Return

	(1)	$CAR[-1,+5]_{ig}$	(2)
	Proprietary AI Models		Open-Source AI Models
Constant	-0.693***		0.174***
	(0.000)		(0.000)
Controls	N		N
Firm FE	Y		Y
Adj. R^2	0.021		0.010
Obs.	125,505		79,264

Panel B: Cumulative Raw Return

	(1)	$CRR[-1,+5]_{ig}$	(2)
	Proprietary AI Models		Open-Source AI Models
Constant	-0.102***		0.975***
	(0.000)		(0.000)
Controls	N		N
Firm FE	Y		Y
Adj. R^2	0.013		0.008
Obs.	125,505		79,264

Table III
Market Response to Open-Source vs. Proprietary AI Models

This table presents the effect of Gen-AI model releases on U.S. firm value (Equation 5), using stock market reactions to the releases listed in Appendix IA.1. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i from one trading day before to k trading days after event g. The dummy variable, OS_g, equals 1 if event g is the release of an open-source Gen-AI model, and 0 otherwise. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) $CAR[-1,+0]_{ig}$	(2) $CAR[-1,+1]_{ig}$	(3) $CAR[-1,+5]_{ig}$
OS_g	0.230***	0.349***	0.867***
	(0.0140)	(0.0180)	(0.0310)
Constant	-0.508	-0.984*	1.069
	(0.393)	(0.582)	(1.119)
Controls		log(Assets), D/E, ROA	
Firm FE		Y	
Adj. R^2	0.00300	0.00600	0.0210
Obs.	204,332	204,332	204,332

Table IV
Skewness in Market Response to Open-Source vs. Proprietary AI Models

This table presents how open-source Gen-AI model releases affect the distributional moments of stock returns at the event-industry level. The dependent variables are moments of cumulative abnormal returns and cumulative raw returns computed for industry j around event g, including (i) $CAR\ Skewness_{jg}$, skewness of cumulative abnormal returns, and (ii) $CRR\ Skewness_{jg}$, skewness of cumulative raw returns. The dummy variable, OS_g, equals 1 if event g is the release of an open-source Gen-AI model, and 0 otherwise. $AI\ Exposure_j$ is a continuous measure of Gen-AI exposure, which is constructed by aggregating the firm-level score from Eisfeldt et al. (2023) to the 2-digit NAICS industry level. Standard errors are two-way clustered at the industry and event levels and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: The Effects of Open-Source AI Models on Return Skewness

	(1) CAR Skewness$_{jg}$	(2) CRR Skewness$_{jg}$
OS_g	0.156***	0.129***
	(0.031)	(0.031)
Constant	0.010	0.053*
	(0.024)	(0.025)
Adj. R^2	0.006	0.004
Obs.	1,139	1,139

Panel B: Return Skewness Responses Across AI Exposure

	(1) CAR Skewness$_{jg}$	(2) CRR Skewness$_{jg}$
OS_g	-0.219***	-0.246***
	(0.070)	(0.070)
AI Exposure$_j$	-0.542**	-0.639**
	(0.204)	(0.287)
$OS_g \times$ AI Exposure$_j$	1.026***	1.024***
	(0.161)	(0.165)
Constant	0.209***	0.287***
	(0.039)	(0.053)
Adj. R^2	0.005	0.003
Obs.	1,139	1,139

Table V
Market Response Across Return Quantiles: Open-Source vs. Proprietary

This table reports results from quantile regressions of firms' cumulative abnormal returns over the event window from one trading day before to five trading days after each model release, CAR[-1,+5]ig. We report results from unconditional quantile regressions in Panels A and B, following Rios Avila (2019), and results from conditional quantile regressions in Panels C and D, following Machado and Santos Silva (2019). Columns correspond to different quantiles of the stock return distribution. The dummy variable, OS_g, equals 1 if event g is the release of an open-source Gen-AI model, and 0 otherwise. $AI\ Exposure_{ig}^{raw}$ is a continuous Gen-AI exposure score from Eisfeldt et al. (2023), which quantifies the share of tasks in a given occupation that can be enhanced by Gen-AI to improve productivity. Quantiles include the 5th, 20th, 50th, 80th, and 95th percentiles, following Giglio et al. (2016). All columns include industry (2-digit NAICS) fixed effects. Standard errors are bootstrapped following Firpo et al. (2009) and are reported in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: The Effects of Open-Source AI Models (Unconditional)

	(1)	(2)	(3)	(4)	(5)
			CAR[-1,+5]$_{ig}$		
Quantile:	0.05	0.20	0.50	0.80	0.95
OS_g	1.325***	0.795***	0.564***	0.829***	1.364***
	(0.101)	(0.038)	(0.024)	(0.043)	(0.131)
Constant	-21.066***	-9.369***	-0.998***	6.767***	20.114***
	(0.275)	(0.108)	(0.048)	(0.081)	(0.304)
Controls			log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y	Y
Adj. R^2	0.051	0.065	0.010	0.021	0.028
Obs.	204,332	204,332	204,332	204,332	204,332

Panel B: The Effects Across AI Exposure (Unconditional)

	(1)	(2)	(3)	(4)	(5)
			CAR[-1,+5]$_{ig}$		
Quantile:	0.05	0.20	0.50	0.80	0.95
OS_g	-0.406	0.289	0.710***	0.480*	0.788
	(0.555)	(0.229)	(0.155)	(0.258)	(0.589)
$OS_g \times AI\ Exposure_{ig}^{raw}$	4.518***	1.320**	-0.381	0.910	1.504
	(1.515)	(0.561)	(0.405)	(0.669)	(1.478)
$AI\ Exposure_{ig}^{raw}$	-2.249**	0.374	0.304	-0.741	1.265
	(1.117)	(0.477)	(0.314)	(0.462)	(1.290)
Constant	-20.194***	-9.531***	-1.117***	7.059***	19.591***
	(0.518)	(0.214)	(0.137)	(0.213)	(0.612)
Controls			log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y	Y
Adj. R^2	0.051	0.065	0.010	0.021	0.028
Obs.	204,332	204,332	204,332	204,332	204,332

Panel C: The Effects of Open-Source AI Models (Conditional)					
	(1)	(2)	(3)	(4)	(5)
			$CAR[-1,+5]_{ig}$		
Quantile:	0.05	0.20	0.50	0.80	0.95
OS_g	0.782***	0.829***	0.865***	0.901***	0.953***
	(0.056)	(0.035)	(0.029)	(0.036)	(0.062)
Constant	-19.392***	-8.973***	-1.233***	6.706***	18.155***
	(0.132)	(0.099)	(0.067)	(0.103)	(0.160)
Controls			log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y	Y
R^2	0.134	0.068	0.006	0.021	0.093
Obs.	204,332	204,332	204,332	204,332	204,332
Panel D: The Effects Across AI Exposure (Conditional)					
	(1)	(2)	(3)	(4)	(5)
			$CAR[-1,+5]_{ig}$		
Quantile:	0.05	0.20	0.50	0.80	0.95
OS_g	0.172	0.337*	0.460**	0.585**	0.766**
	(0.303)	(0.192)	(0.203)	(0.263)	(0.372)
$OS_g \times AI\ Exposure_{ig}^{raw}$	1.585**	1.282***	1.056*	0.825	0.492
	(0.774)	(0.486)	(0.542)	(0.665)	(0.978)
$AI\ Exposure_{ig}^{raw}$	-0.313	-0.125	0.014	0.157	0.363
	(0.563)	(0.334)	(0.332)	(0.431)	(0.573)
Constant	-19.280***	-8.931***	-1.247***	6.640***	17.998***
	(0.277)	(0.143)	(0.153)	(0.190)	(0.241)
Controls			log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y	Y
R^2	0.134	0.068	0.006	0.021	0.093
Obs.	204,332	204,332	204,332	204,332	204,332

Table VI
Validation Tests for Accessibility and Affordability

This table presents validation test results of accessibility and affordability as key factors of open-source AI models. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i from one trading day before to five trading days after event g. The explanatory variables are (i) $Accessibility_g$, measured as the average of three binary indicators for whether the model provides open-source weights, provides open-source training code, and has a parameter count below 3 billion, and (ii) $User\ Costs_g$, measured as USD per million tokens. Both $Accessibility_g$ and $User\ Costs_g$ are standardized to have a mean of 0 and a standard deviation of 1. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	$CAR[-1,+5]_{ig}$	
	(1)	(2)
$Accessibility_g$	0.358***	
	(0.0150)	
$User\ Costs_g$		-0.313***
		(0.0178)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.0200	0.0204
Obs.	204,332	164,248

Table VII
Heterogeneity Across Firms: Adoption Constraints

This table reports the heterogeneous effect of Gen-AI model releases on U.S. firm value by firm size and resource constraints. Moderator variables are constructed in the sample baseline period (i.e., the quarter of the first model-release event or, if the firm enters later, the first quarter in which the firm appears in the sample). We conduct subsample regressions for the highest and lowest terciles of the following moderators: total assets in Panel A, the SA index (Hadlock and Pierce, 2010) in Panel B, and the WW index (Whited and Wu, 2006) in Panel C. For both the SA and WW indices, firms in the top tercile are more constrained. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i from one trading day before to five trading days after event g. The dummy variable, OS_g, equals 1 if event g is the release of an open-source Gen-AI model, and 0 otherwise. Standard errors are clustered at the firm level and written in parentheses. In each Panel, we report tests of differences in the coefficient on OS_g across columns based on seemingly unrelated estimation. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Heterogeneity by Total Assets

	$CAR[-1,+5]_{ig}$	
	Bottom tercile	Top tercile
	(1)	(2)
OS_g	1.241***	0.467***
	(0.070)	(0.035)
Constant	-2.035**	2.681
	(0.882)	(3.144)
Diff. in Coef. on OS_g: (1) - (2)	0.774***	
	(0.078)	
Coefficient on OS_g + Constant	-0.794	3.148
	(0.884)	(3.146)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.023	0.015
Obs.	65,442	71,632

Panel B: Heterogeneity by SA Index

	$CAR[-1,+5]_{ig}$	
	Bottom tercile	Top tercile (higher constraints)
	(1)	(2)
OS_g	0.565***	1.153***
	(0.037)	(0.068)
Constant	7.314***	-2.187**
	(2.706)	(0.990)
Diff. in Coef. on OS_g: (1) - (2)		-0.588***
		(0.077)
Coefficient on OS_g + Constant	7.879***	-1.034
	(2.706)	(0.992)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.015	0.022
Obs.	68,323	68,066

Panel C: Heterogeneity by WW Index

	$CAR[-1,+5]_{ig}$	
	Bottom tercile	Top tercile (higher constraints)
	(1)	(2)
OS_g	0.458***	1.135***
	(0.036)	(0.058)
Constant	3.573	-0.707
	(3.931)	(1.074)
Diff. in Coef. on OS_g: (1) - (2)		-0.677***
		(0.068)
Coefficient on OS_g + Constant	4.031	0.428
	(3.934)	(1.075)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.015	0.022
Obs.	60,918	84,974

Table VIII
Heterogeneity Across Firms: Labor-Based AI Exposure

This table illustrates how firm value responds to model releases, depending on exposure to AI. (Equation 6). The dependent variable is the cumulative abnormal return in percentage points for firm i surrounding the event g. OS_g is an indicator variable equal to 1 if event g is the release of an open-source model. In Panel A, $AI\ Exposure_{ig}^{raw}$ is a continuous Gen-AI exposure score from Eisfeldt et al. (2023). In Panel B, $AI\ Exposure_{ig}^{claude}$ is constructed as the employment-weighted average of occupation-level exposures from the Anthropic Economic Index. All columns include industry (2-digit NAICS) fixed effects. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Workforce with Potential AI Exposure

	(1) $CAR[-1,+0]_{ig}$	(2) $CAR[-1,+1]_{ig}$	(3) $CAR[-1,+5]_{ig}$
$OS_g \times AI\ Exposure_{ig}^{raw}$	0.387**	0.822***	1.050**
	(0.190)	(0.266)	(0.458)
$AI\ Exposure_{ig}^{raw}$	-0.170	-0.236	0.018
	(0.146)	(0.201)	(0.396)
OS_g	0.081	0.034	0.463***
	(0.072)	(0.103)	(0.177)
Constant	-0.077	-0.140	-1.045***
	(0.066)	(0.092)	(0.183)
Controls		log(Assets), D/E, ROA	
Industry FE	Y	Y	Y
Adj. R^2	0.00300	0.00600	0.0210
Obs.	204,332	204,332	204,332

Panel B: Workforce with Realized AI Usage Intensity

	(1) $CAR[-1,+0]_{ig}$	(2) $CAR[-1,+1]_{ig}$	(3) $CAR[-1,+5]_{ig}$
$OS_g \times AI\ Exposure_{ig}^{claude}$	0.447	0.751*	2.084***
	(0.363)	(0.440)	(0.797)
$AI\ Exposure_{ig}^{claude}$	0.199	0.036	0.414
	(0.245)	(0.323)	(0.633)
OS_g	0.116	0.158	0.337*
	(0.091)	(0.112)	(0.201)
Constant	-0.190***	-0.235**	-1.115***
	(0.072)	(0.095)	(0.189)
Controls		log(Assets), D/E, ROA	
Industry FE	Y	Y	Y
Adj. R^2	0.00200	0.00400	0.0110
Obs.	200,788	200,788	200,788

Table IX
Heterogeneity Across Firms: AI Augmentation vs. Displacement

This table presents how the valuation effects of model releases vary across firms with different exposure to AI displacement versus AI augmentation. The sample includes stock market reactions to all open-source and proprietary model releases. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i from one trading day before to k trading days after event g. The dummy variable, OS_g, equals 1 if event g is the release of an open-source Gen-AI model, and 0 otherwise. $AI\ Augmentation_{ig}$ measures a firm's exposure to AI augmentation using a sub-index of the Anthropic Economic Index, which quantifies the share of tasks in a given occupation that can be augmented by Gen-AI. $AI\ Displacement_{ig}$ measures exposure to AI displacement using another sub-index of the Anthropic Economic Index, which quantifies the share of tasks in a given occupation that can be automated by Gen-AI. Panel C shows the heterogeneity analysis across firm size, which is measured by total assets in the baseline period (prior to the release date of the first model in our event set). All columns include industry (2-digit NAICS) fixed effects. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Exposure to AI Augmentation

	(1) $CAR[-1,+0]_{ig}$	(2) $CAR[-1,+1]_{ig}$	(3) $CAR[-1,+5]_{ig}$
$OS_g \times AI\ Augmentation_{ig}$	0.417	0.744**	2.175***
	(0.302)	(0.372)	(0.663)
$AI\ Augmentation_{ig}$	0.133	0.072	0.350
	(0.213)	(0.278)	(0.554)
OS_g	0.097	0.113	0.178
	(0.094)	(0.118)	(0.209)
Constant	-0.182**	-0.251**	-1.127***
	(0.077)	(0.100)	(0.202)
Controls		log(Assets), D/E, ROA	
Industry FE	Y	Y	Y
Adj. R^2	0.002	0.004	0.011
Obs.	200,788	200,788	200,788

Panel B: Exposure to AI Displacement

	(1) $CAR[-1,+0]_{ig}$	(2) $CAR[-1,+1]_{ig}$	(3) $CAR[-1,+5]_{ig}$
$OS_g \times AI\ Displacement_{ig}$	-0.233	-0.345	-0.556
	(0.354)	(0.454)	(0.793)
$AI\ Displacement_{ig}$	0.148	-0.303	0.701
	(0.235)	(0.313)	(0.616)
OS_g	0.284***	0.429***	0.997***
	(0.084)	(0.108)	(0.189)
Constant	-0.169**	-0.145	-1.160***
	(0.066)	(0.089)	(0.175)
Controls		log(Assets), D/E, ROA	
Industry FE	Y	Y	Y
Adj. R^2	0.002	0.004	0.011
Obs.	200,788	200,788	200,788

Panel C: Exposure to AI Augmentation and Displacement: Small vs. Large

	(1)	(2)	(3)	(4)
	$CAR[-1,+5]_{ig}$			
	Small (Bottom tercile)	Large (Top tercile)	Small (Bottom tercile)	Large (Top tercile)
$OS_g \times AI\ Augmentation_{ig}$	2.783*	-0.124		
	(1.498)	(0.756)		
$AI\ Augmentation_{ig}$	1.095	0.985*		
	(1.255)	(0.547)		
$OS_g \times AI\ Displacement_{ig}$			-1.248	0.407
			(0.987)	(0.527)
$AI\ Displacement_{ig}$			-0.098	0.392
			(0.681)	(0.349)
OS_g	0.340	0.487**	0.685***	0.124
	(0.487)	(0.238)	(0.238)	(0.124)
Constant	-2.704***	-0.337	-0.892***	0.345***
	(0.472)	(0.254)	(0.214)	(0.124)
Controls		log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y
Adj. R^2	0.009	0.010	0.003	0.005
Obs.	66,262	68,181	66,262	68,181

Table X
Spillover Effects Through Supply Chain

This table presents how stock market reactions to open-source model releases spill over through the supply chain. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i in industry j from one trading day before to k trading days after event g. The dummy variable, OS_g, equals 1 if event g is the release of an open-source AI model, and 0 otherwise. $AI\ Exposure_{ig}^{upstream}$ denotes the AI exposure of the upstream industry, computed as a weighted average of each upstream industry's AI exposure, with weights equal to the share of the focal industry's total intermediate input purchases sourced from that upstream industry. $AI\ Exposure_{ig}^{downstream}$ denotes the AI exposure of the downstream industry, computed as a weighted average of each downstream industry's AI exposure, with weights equal to the share of the focal industry's total intermediate supplies directed to that downstream industry. $AI\ Exposure_{ig}^{ind}$ is the simple average of AI exposure of the focal industry. The industry codes used to construct industry-level AI exposure are from the OECD Domestic Input–Output Table for the US. In each panel, the regression sample is restricted to firms in industries with AI exposure below the 50th percentile, while $AI\ Exposure_{ig}^{upstream}$ and $AI\ Exposure_{ig}^{downstream}$ are constructed using only industries with AI exposure above the 50th percentile. Panel A presents subsample analyses by event type (proprietary versus open-source AI models). Panels B and C examine heterogeneity in stock market reactions to open-source AI model releases, stratified by firm asset size and the SA index (Hadlock and Pierce, 2010), respectively. Panels D and E present analogous analyses for proprietary AI model releases, using the same firm characteristics. Firms with SA index values above the sample median are classified as more financially constrained. All columns include industry (2-digit NAICS) fixed effects. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Spillovers by Model Type: Open-Source vs. Proprietary

	(1)	(2)	(3)	(4)
	Proprietary AI Models		Open-Source AI Models	
	$CAR[-1,+5]_{ig}$			
$AI\ Exposure_{ig}^{upstream}$	0.870		1.740**	
	(0.631)		(0.754)	
$AI\ Exposure_{ig}^{downstream}$		0.905		1.672**
		(0.637)		(0.764)
$AI\ Exposure_{ig}^{ind}$	5.992	6.059	8.754**	8.183*
	(3.870)	(3.820)	(4.370)	(4.361)
Constant	-4.052***	-4.078***	-3.341*	-3.107*
	(1.519)	(1.496)	(1.727)	(1.719)
Controls	log(Assets), D/E, ROA			
Industry FE	Y	Y	Y	Y
Adj. R^2	0.009	0.009	0.005	0.005
Obs.	32,173	32,173	20,400	20,400

Panel B: Stock Market Reactions to Open-Source Models by Total Assets				
	(1)	(2)	(3)	(4)
			$CAR[-1,+5]_{ig}$	
	below median		above median	
$AI\ Exposure_{ig}^{upstream}$	3.528***		-0.801	
	(1.141)		(0.841)	
$AI\ Exposure_{ig}^{downstream}$		3.597***		-0.940
		(1.160)		(0.843)
$AI\ Exposure_{ig}^{ind}$	20.576***	20.520***	-3.487	-4.162
	(6.907)	(6.904)	(5.327)	(5.244)
Constant	-9.505***	-9.487***	2.130	2.386
	(2.716)	(2.715)	(2.017)	(1.975)
Controls		log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y
Adj. R^2	0.011	0.011	0.009	0.009
Obs.	9,771	9,771	10,629	10,629
Panel C: Stock Market Reactions to Open-Source Models by SA Index				
	(1)	(2)	(3)	(4)
			$CAR[-1,+5]_{ig}$	
	below median		above median (higher constraints)	
$AI\ Exposure_{ig}^{upstream}$	-0.657		3.543***	
	(0.773)		(1.148)	
$AI\ Exposure_{ig}^{downstream}$		-0.739		3.475***
		(0.779)		(1.165)
$AIExposure_{ig}^{ind}$	-2.448	-2.799	16.280**	15.499**
	(4.674)	(4.600)	(6.601)	(6.596)
Constant	2.013	2.153	-6.709***	-6.388**
	(1.818)	(1.783)	(2.590)	(2.583)
Controls		log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y
Adj. R^2	0.009	0.009	0.008	0.008
Obs.	10,334	10,334	9,961	9,961

Panel D: Stock Market Reactions to Proprietary Models by Total Assets

	(1)	(2)	(3)	(4)
			$CAR[-1,+5]_{ig}$	
	below median		above median	
$AI\ Exposure_{ig}^{upstream}$	1.089		0.662	
	(0.981)		(0.642)	
$AI\ Exposure_{ig}^{downstream}$		1.145		0.721
		(0.998)		(0.649)
$AI\ Exposure_{ig}^{ind}$	11.525*	11.677*	1.980	2.222
	(6.103)	(6.071)	(4.347)	(4.277)
Constant	-6.544***	-6.607***	-2.258	-2.347
	(2.455)	(2.441)	(1.618)	(1.582)
Controls		log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y
Adj. R^2	0.009	0.009	0.004	0.004
Obs.	15,362	15,362	16,811	16,811

Panel E: Stock Market Reactions to Proprietary Models by SA Index

	(1)	(2)	(3)	(4)
			$CAR[-1,+5]_{ig}$	
	below median		above median (higher constraints)	
$AI\ Exposure_{ig}^{upstream}$	1.067		0.951	
	(0.701)		(0.950)	
$AI\ Exposure_{ig}^{downstream}$		1.116		0.975
		(0.709)		(0.961)
$AI\ Exposure_{ig}^{ind}$	5.967	6.089	6.962	6.960
	(5.023)	(4.959)	(5.458)	(5.386)
Constant	-3.140*	-3.184*	-4.612**	-4.611**
	(1.903)	(1.874)	(2.151)	(2.118)
Controls		log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y
Adj. R^2	0.003	0.003	0.011	0.011
Obs.	16,345	16,345	15,634	15,634

Table XI
Robustness: Stacked DID Specification

This table presents the impact of AI model releases on firm value across U.S. firms with varying levels of AI exposure, using a difference-in-differences design (as specified in Equation 8). The dependent variable is the abnormal return, expressed in percentage points, for firm i on date t for event g. The dummy variable, OS_g, equals 1 if event g is the release of an open-source Gen-AI model, and 0 otherwise. $AI\ Exposure_{ig}$ is a dummy variable equal to 1 if $AI\ Exposure_{ig}^{raw}$, a continuous Gen-AI exposure score from Eisfeldt et al. (2023), is above the sample median for the corresponding year, and 0 otherwise. $Post_{tg}$ is a dummy variable equal to 1 if date t is observed after the model g is released and 0 otherwise. Columns 1 and 2 report results using the full sample. Columns 3 and 4 report results based on a propensity-score-matched subsample. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	AR_{itg}			
	Full Sample		Matched Sample	
	(1)	(2)	(3)	(4)
$OS_g \times Post_{tg} \times AI\ Exposure_{ig}$	0.150***	0.119***	0.114***	0.0970***
	(0.0300)	(0.0340)	(0.0310)	(0.0360)
$Post_{tg} \times AI\ Exposure_{ig}$	-0.0410**	-0.0720***	-0.0180	-0.0500**
	(0.0190)	(0.0210)	(0.0200)	(0.0230)
Constant	0.222	0.040	0.610	0.208
	(1.557)	(1.514)	(1.851)	(1.768)
Controls		log(Assets), D/E, ROA		
Event-Firm FE	Y	Y	Y	Y
Event-Date FE	Y	Y	Y	Y
Event-Industry-Date FE	N	Y	N	Y
Adj. R^2	0.0270	0.0490	0.0260	0.0480
Obs.	1,429,842	1,429,842	1,222,844	1,222,830

Table XII
Robustness: Alternative Measures of AI Exposure

This table presents the impact of AI model releases on firm value across U.S. firms with varying levels of AI exposure. Alternative measures of AI exposure are defined as follows: (i) $AI\ Exposure_{ig}^{10k}$, based on 10-K filings (constructed by identifying firms that explicitly mention AI-related keywords listed in Appendix IA.3 in Items 1, 1A, and 7, following Lu et al. (2024) and Bertomeu et al. (2025)); and (ii) $AI\ Exposure_{ig}^{ec}$, based on earnings-call transcripts (constructed by identifying AI-related keywords in the presentation and all Q&A responses). Panel A uses the 10-K measure, and Panel B uses the earnings-call measure. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: AI Exposure based on 10-K Filings

	(1) $CAR[-1,+0]_{ig}$	(2) $CAR[-1,+1]_{ig}$	(3) $CAR[-1,+5]_{ig}$	(4) $CAR[-1,+20]_{ig}$
$OS_g \times AI\ Exposure_{ig}^{10k}$	0.0780**	0.0550	0.192**	0.443***
	(0.0360)	(0.0460)	(0.0790)	(0.118)
OS_g	0.215***	0.338***	0.831***	0.639***
	(0.0150)	(0.0190)	(0.0350)	(0.0510)
Constant	-0.510	-0.985*	1.065	5.892**
	(0.393)	(0.582)	(1.119)	(2.897)
Controls		log(Assets), D/E, ROA		
Firm FE	Y	Y	Y	Y
Adj. R^2	0.00300	0.00600	0.0210	0.0610
Obs.	204,332	204,332	204,332	204,332

Panel B: AI Exposure based on Earnings Calls

	(1) $CAR[-1,+0]_{ig}$	(2) $CAR[-1,+1]_{ig}$	(3) $CAR[-1,+5]_{ig}$	(4) $CAR[-1,+20]_{ig}$
$OS_g \times AI\ Exposure_{ig}^{ec}$	0.0590*	0.0170	-0.0530	0.299***
	(0.0350)	(0.0430)	(0.0790)	(0.115)
OS_g	0.220***	0.346***	0.876***	0.675***
	(0.0150)	(0.0200)	(0.0340)	(0.0510)
Constant	-0.508	-0.984*	1.069	5.902**
	(0.393)	(0.582)	(1.119)	(2.897)
Controls		log(Assets), D/E, ROA		
Firm FE	Y	Y	Y	Y
Adj. R^2	0.00300	0.00600	0.0210	0.0610
Obs.	204,332	204,332	204,332	204,332

Table XIII
Robustness: Alternative Outcome Variables

This table presents the impact of AI model releases on firm value across U.S. firms with varying levels of AI exposure, using alternative outcome variables. Panel A uses cumulative raw returns, $CRR[-1,+k]_{ig}$, as the outcome variable. Panel B uses cumulative abnormal returns after taking into account for the Fama-French three factors, $CAR[-1,+k]_{ig}^{FF3}$, as the outcome variable. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Cumulative Raw Returns

	(1) CRR[-1,+0]$_{ig}$	(2) CRR[-1,+1]$_{ig}$	(3) CRR[-1,+5]$_{ig}$	(4) CRR[-1,+20]$_{ig}$
OS_g	0.0820***	0.411***	1.08***	1.26***
	(0.0140)	(0.0180)	(0.0300)	(0.0470)
Constant	-0.266	-1.089*	1.075	6.070**
	(0.434)	(0.636)	(1.221)	(3.033)
Controls		log(Assets), D/E, ROA		
Firm FE	Y	Y	Y	Y
Adj. R^2	0.00100	0.00500	0.0190	0.0530
Obs.	204,332	204,332	204,332	204,332

Panel B: Cumulative Abnormal Returns Adjusted by Fama-French 3 Factors

	(1) CAR$[-1,+0]_{ig}^{FF3}$	(2) CAR$[-1,+1]_{ig}^{FF3}$	(3) CAR$[-1,+5]_{ig}^{FF3}$	(4) CAR$[-1,+20]_{ig}^{FF3}$
OS_g	-0.00500	0.00200	0.0800***	0.453***
	(0.0130)	(0.0170)	(0.0290)	(0.0440)
Constant	-0.715*	-0.940*	1.015	5.497*
	(0.386)	(0.555)	(1.112)	(2.812)
Controls		log(Assets), D/E, ROA		
Firm FE	Y	Y	Y	Y
Adj. R^2	0.00200	0.00500	0.0190	0.0610
Obs.	204,332	204,332	204,332	204,332

Table XIV
Robustness: Counterfactual Estimators

This table presents the impact on firm value across firms with varying AI exposure in the U.S. following the disclosure of AI models. The regression sample includes all treated periods. The dependent variable is the difference between realized abnormal returns (AR_{itg}) and counterfactual abnormal returns (*Counterfactual* AR_{itg}). Counterfactual abnormal returns are estimated using the fixed-effects counterfactual estimator (FEct) in Column 1, the interactive fixed-effects estimator (iFEct) in Column 2, the matrix completion estimator (MC) in Column 3, and the synthetic matching methodology in Column 4, following Liu et al. (2024b) and Acemoglu et al. (2016). The dummy variable, OS_g, equals 1 if event g is the release of an open-source Gen-AI model, and 0 otherwise. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	AR_{itg} − Counterfactual AR_{itg}			
	FEct	iFEct	MC	Synthetic Matching
	(1)	(2)	(3)	(4)
OS_g	0.0900***	0.0560***	0.0570***	0.0230***
	(0.0100)	(0.00900)	(0.00900)	(0.00700)
Controls		log(Assets), D/E, ROA		
Firm FE	Y	Y	Y	Y
Date FE	Y	Y	Y	Y
Adj. R^2	0.00300	0.00200	0.00200	0.00300
Obs.	672,223	783,973	783,973	656,561

APPENDIX

IA.1 Gen-AI Model Release Dates

Event ID	Date	Company and Model Name	Open-Source
1	2022-11-30	OpenAI/GPT-3.5	
2	2023-01-26	Facebook/opt-iml-max-1.3b	✓
3	2023-03-14	OpenAI/GPT-4	
4	2023-03-16	Baidu/ERNIE bot	
5	2023-06-24	Zhipu/chatglm2-6b	✓
6	2023-06-27	Baidu/ERNIE 3.5	
7	2023-07-07	Huawei/Pangu Models 3.0	
8	2023-07-11	Anthropic/Claude 2.0	
9	2023-07-13	Meta/Llama-2-7b-chat-hf	✓
10	2023-08-03	Alibaba/Qwen-7B-Chat	✓
11	2023-09-06	Tencent/Hunyuan	
12	2023-09-28	Amazon/Titan v1	
13	2023-10-09	Moonshot AI/Kimi	
14	2023-10-17	Baidu/ERNIE 4.0	
15	2023-10-25	Zhipu/chatglm3-6b	✓
16	2023-11-04	xAI/Grok-1	
17	2023-11-14	Intel/neural-chat-7b-v3-1	✓
18	2023-11-21	Anthropic/Claude 2.1	
19	2023-11-29	Bigcode/Starcoder2-3b	✓
20	2023-12-06	Google/Gemini 1.0 Nano	
21	2023-12-13	Google/Gemini 1.0 Pro	
22	2024-01-22	Alibaba/Qwen1.5-0.5B	✓
23	2024-02-08	Google/gemma-2b	✓
24	2024-03-04	Anthropic/Claude 3 Sonnet	
25	2024-03-14	Anthropic/Claude 3 Haiku	

26	2024-04-09	Google/Gemini 1.5 Pro	
27	2024-04-12	Apple/OpenELM-1_1B-Instruct	✓
28	2024-04-17	Meta/Meta-Llama-3-8B-Instruct	✓
29	2024-04-30	Amazon/Titan Text Embeddings v2	
30	2024-05-13	OpenAI/GPT-4o	
31	2024-05-15	xAI/Grok-1.5	
32	2024-05-31	Alibaba/Qwen2-0.5B	✓
33	2024-06-20	Anthropic/Claude 3.5 Sonnet	
34	2024-06-21	Huawei/Pangu Models 5.0	
35	2024-07-14	Meta/Llama-3.1-8B	✓
36	2024-08-06	Amazon/Titan Image Generator G1 v2	
37	2024-08-13	xAI/Grok-2	
38	2024-08-16	Microsoft/Phi-3.5-mini-instruct	✓
39	2024-09-12	OpenAI/o1	
40	2024-09-18	Meta/Llama-3.2-1B	✓
41	2024-10-22	Anthropic/Claude 3.5 Haiku	
42	2024-11-26	Meta/Llama-3.3-70B-Instruct	✓
43	2024-12-11	Google/Gemini 2.0 Flash	
44	2024-12-25	Deepseek/DeepSeek-V3	✓
45	2025-01-20	Deepseek/DeepSeek-R1	✓
46	2025-01-22	ByteDance/Doubao 1.5	
47	2025-01-31	OpenAI/o3	
48	2025-02-17	xAI/Grok-3	
49	2025-02-24	Anthropic/Claude 3.7 Sonnet	
50	2025-02-27	OpenAI/GPT-4.5	
51	2025-03-10	Google/gemma-3-1b-it	✓
52	2025-03-16	Baidu/ERNIE 4.5	
53	2025-03-24	Deepseek/DeepSeek-V3-0324	✓
54	2025-03-25	Google/Gemini 2.5	
55	2025-04-14	OpenAI/GPT-4.1	
56	2025-04-16	OpenAI/o4-mini	
57	2025-04-27	Qwen/Qwen3-8B	✓
58	2025-05-22	Anthropic/Claude Sonnet 4	
59	2025-05-28	Deepseek/DeepSeek-R1-0528	✓
60	2025-06-11	ByteDance/Doubao 1.6	

61	2025-06-19	Google/t5gemma-b-b-prefixlm	✓
62	2025-06-28	Baidu/ERNIE-4.5-300B-A47B-Base-Paddle	✓
63	2025-07-09	xAI/Grok-4	
64	2025-07-28	Qwen/Qwen3-30B-A3B-Instruct-2507	✓
65	2025-08-04	OpenAI/gpt-oss-120b	✓
66	2025-08-05	Anthropic/Claude Opus 4.1	
67	2025-08-07	OpenAI/GPT-5	

IA.2 Background Information about General-Purpose Technologies and AI Innovation

A. The Transformative Role of General-Purpose Technologies

The literature on general-purpose technologies (GPTs) highlights their transformative role in driving economic and industrial change through multifaceted impacts and critical enabling conditions. A central theme is the distribution of benefits: early adopters initially secure temporary competitive advantages, though these diminish as adoption becomes widespread (Chun et al., 2011). At the firm level, GPTs are linked to tangible performance improvements, such as higher operating income (Petralia, 2020), while their macroeconomic effects include stimulating economic growth, creating new investment opportunities, and enhancing the efficiency of existing technologies (Carlaw and Lipsey, 2002; Bekar et al., 2018; Lewis and Severnini, 2020; Kassem, 2024). Notably, GPTs are powerful drivers of productivity, with studies showing they increase total factor productivity (Brynjolfsson et al., 2021; Alexopoulos and Cohen, 2016).

Employment dynamics under GPT adoption reflect both creation and displacement: while net job growth is observed, particularly in high-tech and knowledge-intensive sectors (Chun and Shin, 2025), a side effect is the reduction of employment in traditional or routine-task sectors, as innovative industries expand (Montobbio, 2024). Critical to unlocking these benefits are complementary investments in human capital and infrastructure, which enable firms and economies to fully leverage GPT potential (Lo and Sutthiphisal, 2008; Tanriverdi, 2006; Melville, 2004; Brynjolfsson and Hitt, 2000).

Historically, GPTs have included the steam engine, electric motors, semiconductors (Bresnahan and Trajtenberg, 1995), and automobiles (Alexopoulos and Cohen, 2016). Modern examples span industrial robots (Montobbio, 2024), software (Brynjolfsson et al., 2021), cloud computing and the Internet of Things (Chun and Shin, 2025), and broader information technology. Emerging GPTs today encompass advanced AI systems, including logic- and symbolism-based models, pattern-connecting technologies such as large language models (LLMs), and embodied AIs such as robots, with natural language processing (NLP) and machine learning standing out for their wide application scope (Goldfarb et al., 2023). Together, these technologies underscore the enduring role of GPTs as catalysts for structural economic change.

B. History of GPTs

A historical review of transformative general-purpose technologies reveals critical patterns in their adoption and societal impact, offering parallels to the evolution of modern AI. The steam engine, a cornerstone of the Industrial Revolution (1780s–1830s), illustrates this trajectory: James Watt's efficiency improvements in the 1780s dramatically reduced fuel consumption, making steam cost-competitive with waterwheels. The expiration of Watt's patents in 1800 then opened the market to new firms, driving competition that lowered prices and accelerated innovation between 1800 and 1830 (Spear, 2008). Concurrently, the expansion of coal mining and distribution networks in the 1800s built the necessary fuel infrastructure, enabling steam power to operate beyond localized sites. By the 1820s–1830s, steam locomotives and steamships revolutionized transportation, facilitating large-volume shipments and spurring urbanization (Rosenberg and Trajtenberg, 2004).

The second Industrial Revolution (1880s–1920s) saw electricity emerge as another transformative GPT. Thomas Edison's carbon filament lightbulb in 1878 and Nikola Tesla's alternating current (AC) system in 1888, which enabled long-distance transmission (David and Bunn, 1988), laid foundational technological groundwork. As plant construction costs fell (McNerney et al., 2011), the number of power plants surged to 5,000 by 1922 (U.S. Bureau of the Census, 1922), and electricity became the primary energy source, with 68% of the U.S. population gaining access by the end of the 1920s (Woolf, 1987).

The automobile, a defining GPT of the early 20th century (1910s–1930s), followed a similar path. Henry Ford's Model T assembly-line production, introduced in 1908 and scaled in 1913 (Eli and Rhode, 2025), drastically reduced costs: average wholesale prices of U.S. automobiles fell from 1,227 USD in 1913 to 659 in 1917 (Cheng et al., 2025). Infrastructure investments, including the 1916 Federal Aid Road Act (Seely, 1987), paved the way for the expansion of quality highways and enabled widespread adoption, culminating in an explosion of car ownership in the 1930s (Eli et al., 2022).

Across these examples, two prerequisites for GPT takeoff emerge as universal: accessible infrastructure and affordable user costs. These historical parallels underscore the importance of these factors in determining when a technology transitions from innovation to transformative societal impact.

C. History of Artificial Intelligence Innovation

Artificial intelligence was first formalized by McCulloch and Pitts (1943) and later conceptualized by Turing (1950), with the term "*artificial intelligence*" coined at the 1956 Dartmouth Conference. These foundational ideas remained largely theoretical until a major inflection point in the early 2010s, when advances in deep learning, the availability of large-scale data, and GPU computing (e.g., AlexNet, 2012) enabled rapid performance gains. Since then, AI has progressed from automating narrow tasks to performing high-order cognitive functions and coordinating complex workflows, emerging as a general-purpose technology.

Currently, the AI innovation paradigm has also shifted.[18] Early development emphasized

[18]See Mo and Ouyang (2025) for a review of finance papers on the impact of AI on financial markets and firms.

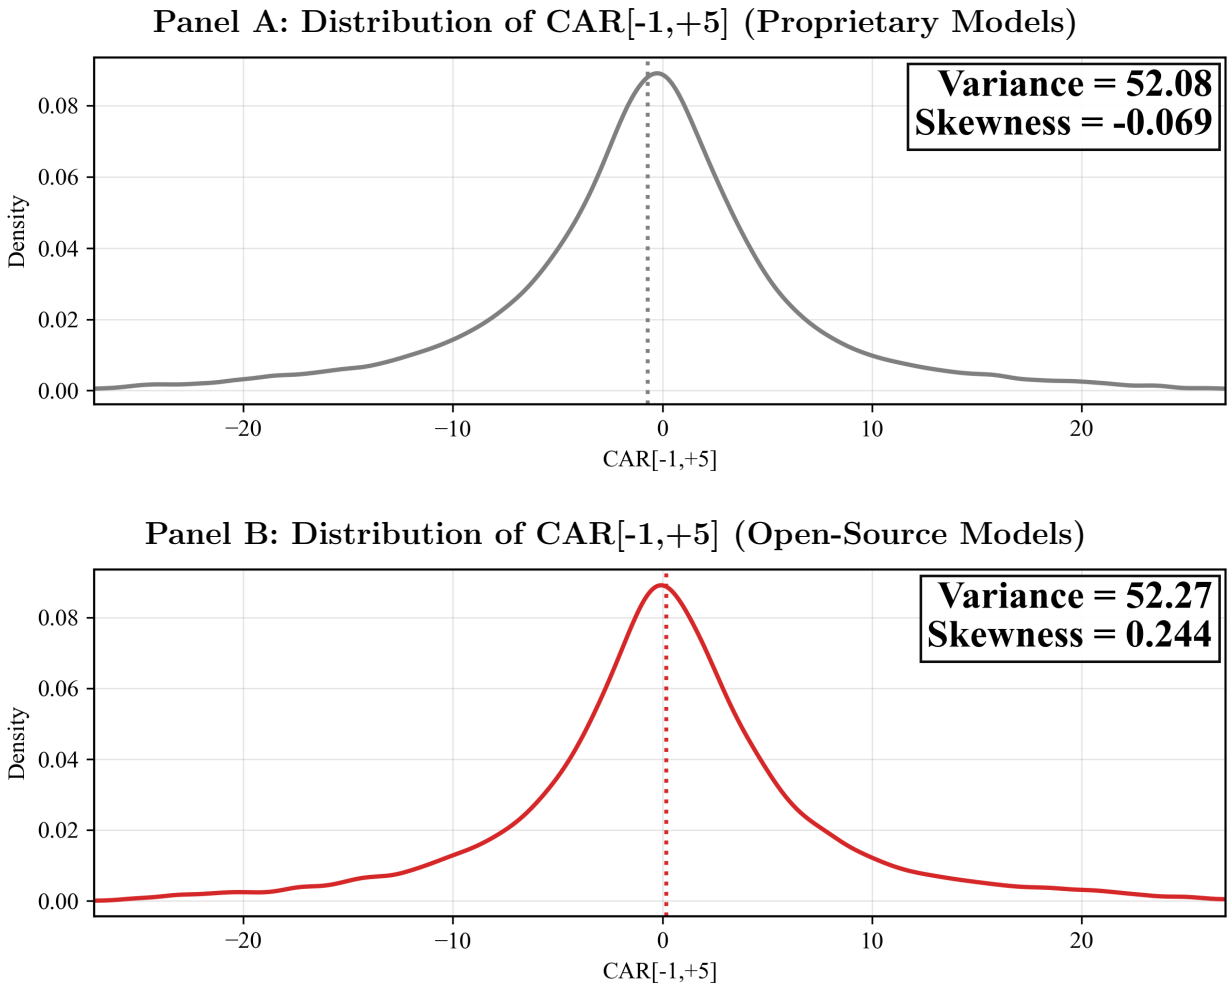
proprietary, product-centric models that prioritize performance at high computational and licensing costs (e.g., GPT-4). More recently, innovation has shifted toward open, process-oriented models that reduce user costs and adoption barriers (e.g., open-source LLMs). The current era of open-source collaboration is characterized not only by broader participation but also by intensified competition. Open-source AI enhances accessibility and collaboration across diverse contributors (Shrestha et al., 2023) and augments collaborative innovation processes (Yeverechyahu et al., 2024).

By treating AI models as shared digital infrastructure, the open-source ecosystem prioritizes computational sustainability and cost compression. Algorithmic advances, such as Multi-Head Latent Attention (MLA), which reduces cache requirements by 93% (Liu et al., 2024a), and Mixture of Experts (MoE), which enhances parameter efficiency (Dai et al., 2024), have drastically lowered the capital required for adoption. These innovations enable firms to adopt AI without the prohibitive capital requirements of proprietary systems (Bommasani et al., 2021; Kapoor et al., 2024), facilitating AI localization in which small and medium-sized businesses leverage local data and existing infrastructure to reduce cloud costs. As a result, competition induced by open-source releases may reallocate value creation from model builders to adopters and complementors (Foster, 2025).

IA.3 AI Dictionary

Following Lu, Phillips, and Yang (2024) and Bertomeu, Lin, Liu, and Ni (2025), we construct a dictionary for AI, containing the following words: 'ai', 'intelligence', 'artificial intelligence', 'algorithms', 'machine learning', 'deep learning', 'face recognition', 'neural networks', 'computer vision', 'natural language processing', 'automation', 'gpt', 'generative artificial intelligence', 'speech recognition', 'pattern recognition', 'supervised learning', 'unsupervised learning', 'generative adversarial network', 'transformer', 'language model', 'text generation', 'embeddings', 'reinforcement learning', 'backpropagation', 'generative model', 'bert', 'long short-term memory', 'cross-entropy', 'fine-tuning', 'generative pretrained transformer', 'recurrent neural network', 'tokenization', 'attention mechanism', 'transfer learning', 'contextualized embeddings', 'few-shot learning', 'zero-shot learning', 'generative adversarial network', 'self-supervised learning', 'mahout', 'opencv', 'libsvm', 'xgboost', 'word2vec', 'kernel methods', 'keras', 'support vector machines', 'latent semantic analysis', 'latent dirichlet allocation', 'knowledge graph', 'multi-modal learning', 'feature extraction', 'multi-task learning'.

IB.1 Heterogeneity in Distribution of CAR[-1,+5] by Model Types: Open-Source vs. Proprietary



This figure presents the probability distribution of firm-level cumulative abnormal returns over an event window spanning one trading day before to five trading days after model release. Panel A displays results for proprietary models, while Panel B shows those for open-source models. The top of each panel reports the respective distribution's variance and skewness. In Panel A, the share of CARs that exceed 0+1.65σ is 4.42%, and the share below 0-1.65σ is 6.11%. In Panel B, the share of CARs that exceed 0+1.65σ is 5.45%, and the share below 0-1.65σ is 4.90%.

IC.1 Robustness: Cumulative Stock Returns without Winsorization

This table reports stock market reactions to the release of proprietary and open-source AI models, using cumulative stock returns without winsorization. In Panel A, the dependent variable is the cumulative abnormal return (in percentage points) for firm i from one trading day before to five trading days after event g. In Panel B, the dependent variable is the cumulative raw return (in percentage points) for firm i over the same event window. In Panel A, the dependent variable is the cumulative abnormal return (in percentage points) for firm i from one trading day before to five trading days after event g. In Panel B, the dependent variable is the cumulative raw return (in percentage points) for firm i over the same event window. Column 1 presents results for the subsample of proprietary AI models. Column 2 presents results for the subsample of open-source AI models. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Cumulative Abnormal Return by Model Type

	(1)	(2)
	$CAR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.687***	0.183***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
Adj. R^2	0.022	0.012
Obs.	125,505	79,264

Panel B: Cumulative Raw Return by Model Type

	(1)	(2)
	$CRR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.096***	0.982***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
Adj. R^2	0.014	0.010
Obs.	125,505	79,264

IC.2 Robustness: Weighted Least Squares

This table reports stock market reactions to the release of proprietary and open-source AI models, estimated using weighted least squares. Following Eisfeldt et al. (2023), we weight each stock by its market capitalization measured one trading day prior to the model release (e.g., market capitalization on November 29, 2022, is used for OpenAI's GPT-3.5). In Panel A, the dependent variable is the cumulative abnormal return (in percentage points) for firm i from one trading day before to five trading days after event g. In Panel B, the dependent variable is the cumulative raw return (in percentage points) for firm i over the same event window. Column 1 presents results for the subsample of proprietary AI models. Column 2 presents results for the subsample of open-source AI models. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Cumulative Abnormal Return by Model Type

	(1)	(2)
	$CAR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.680***	0.004***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
R^2	0.012	0.018
Obs.	125,505	79,264

Panel B: Cumulative Raw Return by Model Type

	(1)	(2)
	$CRR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.293***	0.835***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
Adj. R^2	0.012	0.018
Obs.	125,505	79,264

IC.3 Robustness: Propensity Score Matching

This table reports stock market reactions to the release of proprietary and open-source AI models, based on a propensity-score-matched subsample. Following Bilicka (2019), we obtain propensity scores by estimating a logit model with the open-source indicator on the left-hand side and the stock returns measured from five trading days to two trading days prior to the model release. We then implement nearest-neighbor matching within a tight caliper, without replacement. In Panel A, the dependent variable is the cumulative abnormal return (in percentage points) for firm i from one trading day before to five trading days after event g. In Panel B, the dependent variable is the cumulative raw return (in percentage points) for firm i over the same event window. Column 1 presents results for the subsample of proprietary AI models. Column 2 presents results for the subsample of open-source AI models. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Cumulative Abnormal Return by Model Type

	(1)	(2)
	$CAR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.674***	0.173***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
Adj. R^2	0.020	0.010
Obs.	79,227	79,245

Panel B: Cumulative Raw Return by Model Type

	(1)	(2)
	$CRR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.093***	0.974***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
Adj. R^2	0.011	0.008
Obs.	79,227	79,245

IC.4 Robustness: Entropy Balancing

This table reports stock market reactions to the release of proprietary and open-source AI models, based on entropy balancing as in McMullin and Schonberger (2020). The covariates used to achieve balance in entropy balancing include stock returns measured from five trading days to two trading days prior to the model release. In Panel A, the dependent variable is the cumulative abnormal return (in percentage points) for firm i from one trading day before to five trading days after event g. In Panel B, the dependent variable is the cumulative raw return (in percentage points) for firm i over the same event window. Column 1 presents results for the subsample of proprietary AI models. Column 2 presents results for the subsample of open-source AI models. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Cumulative Abnormal Return by Model Type

	(1)	(2)
	$CAR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.687***	0.173***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
Adj. R^2	0.022	0.010
Obs.	125,485	79,245

Panel B: Cumulative Raw Return by Model Type

	(1)	(2)
	$CRR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.097***	0.974***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
Adj. R^2	0.014	0.008
Obs.	125,485	79,245

IC.5 Robustness: Including Year-Month Fixed Effects

This table reports stock market reactions to the release of proprietary and open-source AI models, controlling for year-month fixed effects. In Panel A, the dependent variable is the cumulative abnormal return (in percentage points) for firm i from one trading day before to five trading days after event g. In Panel B, the dependent variable is the cumulative raw return (in percentage points) for firm i over the same event window. Column 1 presents results for the subsample of proprietary AI models. Column 2 presents results for the subsample of open-source AI models. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Cumulative Abnormal Return by Model Type

	(1)	(2)
	$CAR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.693***	0.174***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
Year-Month FE	Y	Y
Adj. R^2	0.061	0.068
Obs.	125,505	79,264

Panel B: Cumulative Raw Return by Model Type

	(1)	(2)
	$CRR[-1,+5]_{ig}$	
	Proprietary AI Models	Open-Source AI Models
Constant	-0.102***	0.975***
	(0.000)	(0.000)
Controls	N	N
Firm FE	Y	Y
Year-Month FE	Y	Y
Adj. R^2	0.145	0.121
Obs.	125,505	79,264

IC.6 Robustness: Firms with Low vs. High Adoption Constraints

This table reports the heterogeneous effect of Gen-AI model releases on U.S. firm value by firm size and resource constraints. Moderator variables are constructed in the sample baseline period (i.e., the quarter of the first model-release event or, if the firm enters later, the first quarter in which the firm appears in the sample). We conduct subsample regressions for the highest and lowest terciles of the following moderators: employment size in Panel A, debt-to-asset ratio in Panel B, interest expense scaled by revenue in Panel C, interest expense scaled by debt, following Frank and Shen (2016) in Panel D, operating net cash flow in Panel E, and net cash flow over total assets in Panel F. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i from one trading day before to five trading days after event g. The dummy variable, OS_g, equals 1 if event g is the release of an open-source Gen-AI model, and 0 otherwise. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Heterogeneity by Employment Size

	$CAR[-1,+5]_{ig}$	
	Bottom tercile	Top tercile
	(1)	(2)
OS_g	0.926***	0.601***
	(0.064)	(0.039)
Constant	-1.884	11.107***
	(1.239)	(2.486)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.027	0.011
Obs.	66,126	71,808

Panel B: Heterogeneity by Debt-to-Asset Ratio

	$CAR[-1,+5]_{ig}$	
	Bottom tercile	Top tercile
	(1)	(2)
OS_g	0.852***	0.967***
	(0.058)	(0.061)
Constant	3.843**	-2.298
	(1.678)	(1.829)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.016	0.025
Obs.	59,352	61,690

Panel C: Heterogeneity by Interest Expense Scaled by Revenue

	$CAR[-1,+5]_{ig}$	
	Bottom tercile	Top tercile
	(1)	(2)
OS_g	0.895***	0.971***
	(0.060)	(0.051)
Constant	3.604**	-1.663
	(1.421)	(1.628)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.017	0.026
Obs.	58,746	86,001

Panel D: Heterogeneity by Interest Expense Scaled by Debt

	$CAR[-1,+5]_{ig}$	
	Bottom tercile	Top tercile
	(1)	(2)
OS_g	0.827***	1.007***
	(0.056)	(0.061)
Constant	1.579	-3.401
	(1.602)	(2.114)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.015	0.027
Obs.	60,992	60,985

Panel E: Heterogeneity by Operating Net Cash Flow

	$CAR[-1,+5]_{ig}$	
	Bottom tercile	Top tercile
	(1)	(2)
OS_g	1.365***	0.388***
	(0.067)	(0.035)
Constant	-1.323	7.568***
	(1.126)	(2.392)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.021	0.013
Obs.	66,227	71,086

Panel F: Heterogeneity by Net Cash Flow over Total Assets

	$CAR[-1,+5]_{ig}$	
	Bottom tercile	Top tercile
	(1)	(2)
OS_g	1.167***	0.732***
	(0.061)	(0.053)
Constant	-0.340	-0.435
	(1.208)	(2.508)
Controls	log(Assets), D/E, ROA	
Firm FE	Y	Y
Adj. R^2	0.024	0.019
Obs.	66,876	69,671

IC.7 Robustness: Heterogeneity Across Firms' AI Exposure

This table illustrates how firm value responds to model releases, depending on exposure to AI. (Equation 6). The dependent variable is cumulative abnormal return in percentage points for firm i surrounding the event g. OS_g is an indicator variable equal to 1 if event g is the release of an open-source model. In Panel A, $AI\ Exposure_{ig}^{raw}$ is a continuous AI exposure score from Eisfeldt et al. (2023). In Panel B, $AI\ Exposure_{ig}^{claude}$ is constructed as the employment-weighted average of occupation-level exposures from the Anthropic Economic Index. All columns include firm fixed effects. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Workforce with Potential AI Exposure

	(1) $CAR[-1,+0]_{ig}$	(2) $CAR[-1,+1]_{ig}$	(3) $CAR[-1,+5]_{ig}$
$OS_g \times AI\ Exposure_{ig}^{raw}$	0.403**	0.842***	1.08**
	(0.190)	(0.266)	(0.458)
OS_g	0.0760	0.0260	0.455**
	(0.0720)	(0.103)	(0.177)
Constant	-0.508	-0.984*	1.069
	(0.392)	(0.582)	(1.119)
Controls		log(Assets), D/E, ROA	
Firm FE	Y	Y	Y
Adj. R^2	0.00300	0.00600	0.0210
Obs.	204,332	204,332	204,332

Panel B: Workforce with Realized AI Usage Intensity

	(1) $CAR[-1,+0]_{ig}$	(2) $CAR[-1,+1]_{ig}$	(3) $CAR[-1,+5]_{ig}$
$OS_g \times AI\ Exposure_{ig}^{claude}$	0.449	0.748*	2.07***
	(0.363)	(0.441)	(0.798)
OS_g	0.115	0.158	0.341*
	(0.0910)	(0.112)	(0.201)
Constant	-0.485	-1.109*	0.875
	(0.412)	(0.612)	(1.166)
Controls		log(Assets), D/E, ROA	
Firm FE	Y	Y	Y
Adj. R^2	0.00300	0.00600	0.0210
Obs.	200,788	200,788	200,788

IC.8 Robustness: AI Augmentation vs. Displacement

This table presents how the valuation effects of model releases vary with firms' exposure to AI displacement versus AI augmentation. The sample includes stock market reactions to all open-source and proprietary model releases. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i from one trading day before to k trading days after event g. The dummy variable, OS_g, equals 1 if event g is the release of an open-source AI model, and 0 otherwise. *AI Augmentation*$_{ig}$ measures a firm's exposure to AI augmentation using a sub-index of the Anthropic Economic Index, which quantifies the share of tasks in a given occupation that can be augmented by Gen-AI. *AI Displacement*$_{ig}$ measures exposure to AI displacement using another sub-index of the Anthropic Economic Index, which quantifies the share of tasks in a given occupation that can be automated by Gen-AI. All columns include firm fixed effects. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Exposure to AI Augmentation

	(1) CAR[-1,+0]$_{ig}$	(2) CAR[-1,+1]$_{ig}$	(3) CAR[-1,+5]$_{ig}$
$OS_g \times$ *AI Augmentation*$_{ig}$	0.419	0.744**	2.17***
	(0.302)	(0.372)	(0.664)
OS_g	0.0970	0.113	0.180
	(0.0940)	(0.118)	(0.209)
Constant	-0.485	-1.110*	0.874
	(0.412)	(0.612)	(1.166)
Controls		log(Assets), D/E, ROA	
Firm FE	Y	Y	Y
Adj. R^2	0.00300	0.00600	0.0210
Obs.	200,788	200,788	200,788

Panel B: Exposure to AI Displacement

	(1) CAR[-1,+0]$_{ig}$	(2) CAR[-1,+1]$_{ig}$	(3) CAR[-1,+5]$_{ig}$
$OS_g \times$ *AI Displacement*$_{ig}$	-0.229	-0.339	-0.539
	(0.353)	(0.454)	(0.793)
OS_g	0.283***	0.428***	0.994***
	(0.0830)	(0.108)	(0.189)
Constant	-0.485	-1.109*	0.875
	(0.412)	(0.612)	(1.166)
Controls		log(Assets), D/E, ROA	
Firm FE	Y	Y	Y
Adj. R^2	0.00300	0.00600	0.0210
Obs.	200,788	200,788	200,788

IC.9 Robustness: Spillover Effects Using Value-Weighted Industry AI Exposure

This table reports a robustness check for Table X using an industry-level AI exposure measure constructed as the value-weighted average of firms' AI exposure within each industry. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i in industry j from one trading day before to k trading days after event g. The dummy variable, OS_g, equals 1 if event g is the release of an open-source AI model, and 0 otherwise. $AI\ Exposure_{ig}^{upstream}$ denotes the AI exposure of the upstream industry, computed as a weighted average of the AI exposure of each upstream industry, with weights equal to the share of the focal industry's total intermediate input purchases sourced from that upstream industry. $AI\ Exposure_{ig}^{downstream}$ denotes the AI exposure of the downstream industry, computed as a weighted average of the AI exposure of each downstream industry, with weights equal to the share of the focal industry's total intermediate supplies directed to that downstream industry. $AI\ Exposure_{ig}^{ind}$ is the weighted average AI exposure of the focal industry, where weights equal each firm's market capitalization divided by the total market capitalization of all firms in the focal industry. The industry codes used to construct industry-level AI exposure are from the OECD Domestic Input–Output Table for the US. The regression sample is restricted to firms in industries with AI exposure below the 50th percentile, while $AI\ Exposure_{ig}^{upstream}$ and $AI\ Exposure_{ig}^{downstream}$ are constructed using only industries with AI exposure above the 50th percentile. All columns include industry (2-digit NAICS) fixed effects. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)
	Proprietary AI Models		Open-Source AI Models	
	$CAR[-1,+5]_{ig}$			
$AI\ Exposure_{ig}^{upstream}$	0.080		2.617***	
	(0.346)		(0.536)	
$AI\ Exposure_{ig}^{downstream}$		0.078		2.840***
		(0.363)		(0.567)
$AI\ Exposure_{ig}^{ind}$	-4.057	-4.057	13.336***	13.752***
	(2.592)	(2.602)	(3.004)	(3.033)
Constant	-0.459	-0.457	-4.526***	-4.685***
	(0.892)	(0.897)	(1.083)	(1.094)
Controls		log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y
Adj. R^2	0.012	0.012	0.010	0.010
Obs.	32,173	32,173	20,400	20,400

IC.10 Robustness: Spillover Effects Through Supply Chain

This table presents how stock market reactions to open-source model releases spill over through the supply chain. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i from one trading day before to k trading days after event g. The dummy variable, OS_g, equals 1 if event g is the release of an open-source AI model, and 0 otherwise. AI exposure is a continuous Gen-AI exposure score from Eisfeldt et al. (2023), which quantifies the share of tasks in a given occupation that can be enhanced by Gen-AI to improve productivity. $AI\ Exposure_{ig}^{upstream}$ is the simple average AI exposure of the focal firms' suppliers. $AI\ Exposure_{ig}^{downstream}$ is the simple average AI exposure of the focal firms' customers. Sales flows between suppliers and customer firms are identified based on FactSet Revere. $AI\ Exposure_{ig}^{ind}$ is the simple average of AI exposure of the focal industry. The industry codes used to construct the AI exposure of the focal industry are from the OECD Domestic Input–Output Table for the US, to ensure consistency with Table X. The regression sample is restricted to firms with AI exposure below the 50th percentile, while $AI\ Exposure_{ig}^{upstream}$ and $AI\ Exposure_{ig}^{downstream}$ are constructed using only firms with AI exposure above the 50th percentile. Standard errors are clustered at the firm level and written in parentheses. Significance level: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)
	Proprietary AI Models		Open-Source AI Models	
	$CAR[-1,+5]_{ig}$			
$AI\ Exposure_{ig}^{upstream}$	-0.287*		0.541***	
	(0.168)		(0.189)	
$AI\ Exposure_{ig}^{downstream}$		0.161		0.342**
		(0.159)		(0.173)
$AI\ Exposure_{ig}^{ind}$	3.731	3.779	4.603*	4.566*
	(2.392)	(2.385)	(2.392)	(2.400)
Constant	-3.045***	-3.094***	-1.601*	-1.597*
	(0.829)	(0.825)	(0.847)	(0.851)
Controls		log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y
Adj. R^2	0.015	0.015	0.008	0.008
Obs.	57,277	57,277	36,165	36,165