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

Yi Huang Sung Kwan Lee Bernard Yeung Yujun Zhou

October 25, 2026

Abstract

We examine firm-level stock-market reactions to AI model releases and find that open-source releases generate more favorable valuation responses than proprietary releases. Relative to proprietary releases, open-source releases disproportionately benefit smaller, capital-constrained firms, consistent with their greater accessibility and lower user costs. Valuation gains after open-source releases concentrate among firms with greater exposure to AI-suitable tasks, particularly those involving human–AI augmentation; comparable patterns are absent after proprietary releases. Using input-output exposure measures, we further document positive valuation spillovers from high-AI-exposure firms to upstream and downstream partners in low-AI-exposure industries, but only following open-source releases. Taken together, our findings suggest that open-source AI broadens the expected economic benefits of technological progress by lowering adoption barriers and extending benefits across firms and interconnected industries.

Keywords: Open-source AI, Technology Diffusion, Financial Constraints, Supply-Chain Spillovers, Generative AI

JEL Classification: D85, G14, L17, L86, O31, O36

*Yi Huang: Bank for International Settlements (BIS) and CEPR, Yi.Huang@bis.org; Sung Kwan Lee: The Chinese University of Hong Kong, Shenzhen, leesungkwan@cuhk.edu.cn; Bernard Yeung: Asian Bureau of Finance and Economic Research, National University of Singapore Business School, and Southern University of Science and Technology, byeung@nus.edu.sg; Yujun Zhou: Tsinghua University, zhouyujun23@mails.tsinghua.edu.cn. We thank Pablo Hernández de Cos, Steve Davis, Jon Frost, Bengt Holmström, Chen Lin, Guangli Lu, Ye Luo, Fernando Perez-Cruz, Richard Portes, Hélène Rey, Damiano Sandri, Thomas Sargent, Hyun-Song Shin and Michael Spence for providing valuable feedback. We are also most grateful for helpful comments from workshop participants at National University of Singapore Business School, Georgetown University, and the Peterson Institute of International Economics. 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

Artificial intelligence (AI) is a general-purpose technology (GPT) with the potential to transform production across a broad range of sectors (Bresnahan and Trajtenberg, 1995; Eloundou et al., 2024). Yet, as with other GPTs, its economic impact depends not only on advances in the underlying technology but also on complementary innovations, organizational capabilities, and skills that enable widespread adoption (Brynjolfsson et al., 2019). Realizing AI's value requires more than access to powerful models: users must identify tasks where AI is effective and build the experience needed to integrate it into workflows and use its capabilities productively.

GPT development is evolutionary: technological advances interact with complementary innovations and adoption, sometimes reaching a tipping point when greater accessibility and lower user costs spur infrastructure development, applications, widespread adoption, and new business models. As firms learn to identify suitable tasks and integrate AI into workflows, deployment costs fall further, accelerating diffusion. The automobile illustrates this process. The first full-scale steam-powered road vehicle appeared in 1769, while the practical four-stroke internal-combustion engine emerged in 1876. More than a century later, the Ford Model T, introduced in 1908 and mass-produced using the moving assembly line from 1913, provided the turning point: it helped transform automobiles from luxury goods into mass-market products. By 1917, Model Ts accounted for over 40% of U.S. cars. Mass adoption then stimulated investment in roads, petrol stations, repair services, motels, trucking, and passenger transport, transforming production, transportation, commerce and society (Eli and Rhode, 2025). The recent emergence of open-source AI models may represent a comparable turning point.

McCulloch and Pitts (1943) laid the intellectual foundation for neural networks and deep learning. Then it took nearly 80 years to make generative AI a practical tool for ordinary

users.[1] OpenAI launched ChatGPT in November 2022, bringing powerful AI to the public through a natural-language interface, followed by proprietary models such as Anthropic's Claude (March 2023) and Google's Gemini (December 2023). These developers retain control over model weights and access terms, monetizing their capabilities through usage-based APIs, subscriptions, enterprise contracts, and premium services. Despite expanding access, inference costs and deployment requirements can still constrain adoption, particularly among cost-sensitive users and organizations.

Meanwhile, BLOOM and Meta's OPT broadened researchers' access to large language models in 2022. Meta released LLaMA to selected researchers in February 2023 and Llama 2 for research and commercial use under its community license in July. Subsequent releases of Qwen, Mistral, and DeepSeek in August, September, and November 2023, respectively, accelerated the growth of the open-weight ecosystem. By making trained weights available, these models enabled developers to fine-tune, distill, and combine existing models for specific tasks, languages, and industries without training foundation models from scratch (Shrestha et al., 2023; Yeverechyahu et al., 2024). Together with open-source tools and more accessible computing, this lowered experimentation costs and fostered a proliferation of derivative models. Licensing terms vary, however, and open weights do not necessarily mean models are fully open-source.

This development illustrates the economics of open-source technology sharing described by Lerner and Tirole (2005): firms can share core technology while capturing value through complementary products, services, and strategic advantages. Open-weight releases separate model development from inference provision, allowing independent firms to run shared mod-

[1] 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.

els and compete in inference APIs, cloud computing, fine-tuning, and customized enterprise deployment. Value can thus accrue not only to model developers but also to providers of applications, hardware, infrastructure, and complementary services. By enabling multiple firms to build on shared capabilities, open-source models can broaden participation, intensify competition in AI services, and accelerate diffusion.

Evidence from Demirer et al. (2025) and Demirer et al. (2026) is consistent with a substantial expansion in the supply of AI capabilities coinciding with the releases of popular open-source models. From late 2023 through 2025, the market for commercially available models and inference services grew rapidly, driven largely by open-source model releases; the number of releases reached millions. Competition has driven the price of intelligence down roughly a thousandfold, with open-source models costing about 90% less than comparable closed-source models with similar intelligence. These developments expanded the range of AI capabilities firms could obtain from competing suppliers, not just from the original model developers. Although open-source models are not perfect substitutes for proprietary models, they can substantially reduce the cost of accessing useful AI capabilities. Together, lower user costs and high accessibility suggest that open-source models may serve as a tipping point for broader AI adoption.

The distinct business models of proprietary and open-source AI may shape adoption patterns. Proprietary models can command a premium when their superior capabilities, reliability, or integrated services justify the cost. Open-source models, by contrast, can offer cost advantages, customization, and deployment flexibility. These differences may make open-source models particularly attractive to cost-sensitive users and organizations, while proprietary models remain attractive where performance and reliability are paramount. The two approaches may therefore serve overlapping but differentiated market segments.

Notably, the advance of open-source AI models may benefit all AI providers, including the proprietary AI modelers. On the one hand, open-source models intensify price compe-

tition and can substitute for proprietary offerings, albeit imperfectly. On the other hand, by lowering the cost of experimentation and deployment, open-source models broaden AI adoption, stimulate complementary innovation in applications and infrastructure, and build users' capabilities to integrate AI into their workflows. These developments can broaden demand for AI capabilities, including proprietary models that offer differentiated performance or integrated services.

Finally, as a general-purpose technology, AI can generate productivity spillovers along supply chains. AI-enabled suppliers may improve product quality, reduce costs, shorten delivery times, and provide more reliable inputs, benefiting downstream customers even if they do not adopt AI themselves. Conversely, AI-enabled customers may improve demand forecasting, specify orders more precisely, and coordinate production more effectively, benefiting suppliers through more predictable demand and better coordination, although meeting higher standards for quality, traceability, and delivery may entail additional costs. Such spillovers may follow only after open-source model releases lower deployment barriers, enabling more firms to adopt AI and intensifying competition that limits the ability of a few high-AI-adoption firms to collect rents from their AI-enabled capabilities. If these releases mark a turning point in AI diffusion, they may also herald broader productivity gains that propagate through supply chains and across the economy.

In summary, the above narratives suggest a few useful conjectures about the expected economic impact of open-source models. First, proprietary and open-source models have different impacts: the former favors well-endowed firms, while the latter's benefits are more inclusive, reaching smaller, more constrained firms. The drivers of the differentiation are accessibility and user costs. Second, firms with more AI-suitable tasks, particularly those with greater experience using AI, stand to benefit more; open-source models can also extend these benefits to smaller, resource-constrained firms in ways that proprietary models may not. Third, when open-source AI models are released, value will spill over from upstream

(downstream) high AI-exposed firms to downstream (upstream) low AI-exposed firms, but not from proprietary AI models.

However, doubts remain about these narratives: Widder et al. (2024) argue that labeling an AI model "open" does not necessarily make the AI system or the industry producing it genuinely open. Most open-source releases provide model weights, but they do not include all the resources needed to develop and deploy AI at scale; barriers remain in computing, access to training data, and deployment infrastructure (Azoulay et al., 2024; Korinek and Vipra, 2024). Thus, although open-source models can stimulate complementary innovation and expand the ecosystem of AI capabilities, applications, and deployment infrastructure, they do not necessarily decentralize control over critical resources or ensure effective competition. Moreover, wider access raises non-trivial safety and governance concerns because safeguards can be removed and released weights cannot be recalled (Kapoor et al., 2024; Bengio et al., 2026). Ultimately, whether open-source models translate lower access costs and complementary innovation into widespread productivity gains, including for resource-constrained users, depends on firms' deployment capabilities and the competitive conditions governing AI adoption.

We therefore empirically examine the expected economic impacts of proprietary and open-source AI model releases using an event-study framework. We use a novel dataset of 67 influential proprietary and open-source model releases spanning November 2022 to August 2025.[2] Because these model releases are very recent, we do not have a long enough firm-level panel to further validate investor expectations. Still, we use the longest possible period of firm-level EBITDA margins and operating costs to provide corroborative evidence. Proprietary AI model releases are sourced from all major global AI-related technology firms, including OpenAI, Anthropic, Amazon, Google, and ByteDance. We identify open-source AI model releases from Hugging Face, the leading global open-source software platform, by

[2]We provide details on the sample model selection in Section III.C, and the full list is in Appendix IA.1.

selecting the five most popular commercial AI models each month.[3] 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.

Our analysis reveals several key findings. First, as shown in Figure 1, 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: the lower-end CARs shift left by more than the upper-end CARs shift right. Likewise, firm-level CARs are positively associated with accessibility measures and negatively associated with user costs, suggesting investors expect accessibility and affordable user costs to drive value-creating AI adoption.

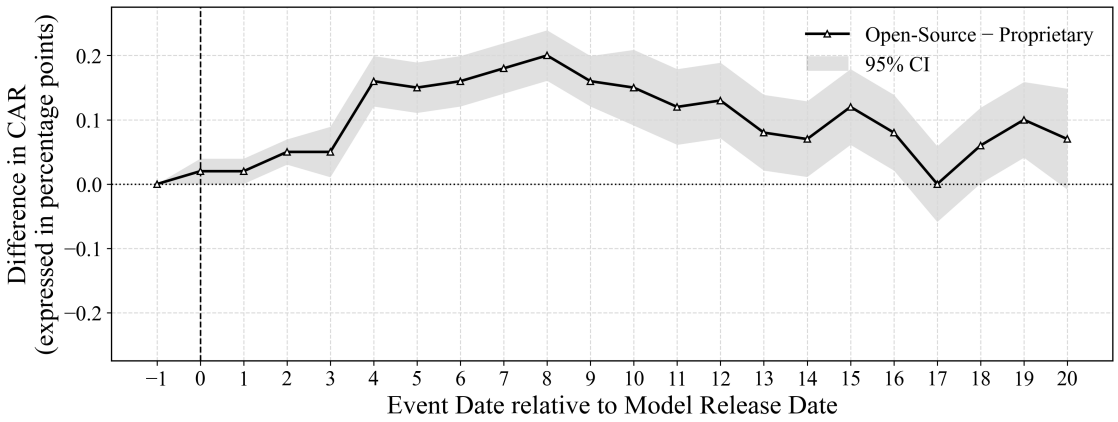


Figure 1. Market-Implied Valuation Gap: Open-Source vs. Proprietary
The figure illustrates the market-implied valuation gap between open-source and proprietary model releases. The solid line represents the cumulative difference in firm valuations following open-source versus proprietary model releases, while the shaded band indicates 95% confidence intervals.

Second, we find significant firm-level heterogeneity driven by adoption cost constraints. Consistent with the lower user costs and competitive supply of inference and API services

[3]We provide details on technical distinctions between proprietary and open-source AI models in Section II.

that open-source models make possible, smaller and financially constrained firms see significantly higher gains after open-source releases than after proprietary releases. By contrast, valuation gains following proprietary releases accrue mostly to firms with greater financial resources. These results suggest that open-source AI may lower barriers to access and participation in AI development and deployment, potentially broadening participation among resource-constrained firms.

Third, our analysis identifies firms' task suitability for AI innovation as a primary driver of these valuation gains. We proxy suitability with two complementary measures: a labor-based estimate of AI exposure (Eisfeldt et al., 2023) and a real-world measure of AI usage experience from 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.

To dig deeper, we decompose 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, for smaller firms, these augmentation-driven gains arise only after open-source releases; larger firms capture similar gains from proprietary models.

Fourth, the economic benefits of open-source model releases extend through supply chains, raising valuations even among firms in low-AI-exposure industries that are vertically linked to high-AI-exposure suppliers or buyers. Using upstream and downstream exposure measures constructed from the Organisation for Economic Co-operation and Development (OECD) input–output tables, we show that these vertical spillovers are specific to open-source releases. The lack of comparable effects after proprietary model releases suggests investors expect open-source AI to spur complementary AI adoption across interconnected firms, generating gains beyond those firms that directly adopt AI. These spillovers are consistent with improved procurement and interfirm collaboration. By contrast, investors appear

to expect firms in high-AI-exposure industries to capture the capability rents from proprietary AI, limiting the anticipated gains to vertically linked firms.

We subject our results to a battery of robustness checks. First, we confirm that our findings hold with unwinsorized returns and weighted least squares estimation, suggesting the results reflect significant market-wide valuation shifts rather than statistical noise or outlier influence. Second, to address model-type imbalance and timing bias, we use propensity score matching, entropy balancing, and year-month fixed effects. Third, we strengthen causal identification using a stacked difference-in-differences design and a suite of counterfactual estimators, including the fixed-effects counterfactual estimator, the interactive fixed-effects estimator, the matrix completion estimator, and the synthetic matching methodology, following Liu et al. (2024b) and Goldsmith-Pinkham and Lyu (2025). Fourth, we confirm the valuation effects by employing alternative measures of AI exposure constructed from 10-K filings and earnings call transcripts, as well as alternative outcome measures such as cumulative raw returns and Fama–French-adjusted abnormal returns.

Although our focus is on investors' expectations, we provide early, suggestive evidence that these expectations have begun to materialize in firms' operating outcomes. Using quarterly Compustat data from 2019 to 2026, we find that after open-source models became competitive with leading proprietary models following the release of DeepSeek-R1 in January 2025, financially constrained firms with greater actual AI usage experienced larger increases in EBITDA margins and larger declines in operating costs relative to sales than unconstrained firms. These patterns are tied to firms' realized, augmentation-oriented AI usage, measured by the Anthropic Economic Index, rather than to the potential exposure of their tasks, and they are not yet accompanied by higher sales growth or capital expenditures, consistent with low-cost adoption that operates in the short run through cost savings. These results suggest that the valuation gains we document reflect more than investor optimism, although the short post-release window cannot yet establish whether the operating gains

persist.

The remainder of this paper is organized as follows. Section II distinguishes proprietary from open-source AI models in their technical and economic characteristics and delineates our paper's contribution to the literature. Section III describes the data sources, variable construction, and empirical strategy. Section IV presents the empirical results on valuation effects and the distribution of benefit, while Section V discusses the differential impacts of the two AI release types across firms and industries and provides early evidence from firms' operating outcomes. Section VI shows robustness checks using alternative specifications and measures. Section VII concludes by summarizing the key findings and discussing their implications for future research.

II. Distinguishing AI Model Types & Contribution to Literature

A. Proprietary vs. Open-source AI Models

Our empirical work differentiates proprietary and open-source AI models.[4] 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 let firms adopt AI without the prohibitive capital costs of proprietary systems (Bommasani et al., 2021; Kapoor et al., 2024), enabling AI localization in which small and medium-sized businesses leverage local data and existing infrastructure to reduce cloud costs. Demirer et al. (2025) and Demirer et al. (2026) show a substantial expansion in the supply of AI capabilities, coinciding with the release of popular open-source models from late 2023 through 2025. As a result, competition from open-source releases may shift value creation from model builders to adopters

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

and complementors (Foster, 2025).

Building on Shrestha et al. (2023), Yeverechyahu et al. (2024), Liu et al. (2024a), Dai et al. (2024), Bommasani et al. (2021), and Kapoor et al. (2024), Table I presents these fundamental technical and economic distinctions: while proprietary models minimize initial setup complexity via APIs, open-source models offer enhanced customization and lower long-term marginal costs through code and weight accessibility. This structural shift, with lower adoption costs and greater accessibility, enables broader, economically relevant AI adoption, in contrast to high-market-power proprietary models. In summary, open-source AI releases may shift AI from exclusive experimentation to inclusive economic adoption.

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 has 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), Liu et al. (2024a), Dai et al. (2024), Bommasani et al. (2021), and Kapoor et al. (2024)

B. Contribution to the Literature

Our paper contributes to three strands of literature on the economics of technological change, unified by the question of how the dissemination of AI capabilities shapes their expected economic benefits.

First, we contribute to the literature on AI and firm value by shifting attention from

advances in model capabilities to the economics of access and diffusion. Brynjolfsson et al. (2019) emphasize that realizing the productivity gains from a general-purpose technology requires complementary investments and widespread adoption. Lerner and Tirole (2005) examine the incentives underlying open-source development. We connect these perspectives by examining how modes of AI dissemination shape user firms' expected gains. Open-source models can lower user costs, facilitate customization, and spur complementary innovation and competition among inference and API providers (Demirer et al., 2025). We find that open-source releases elicit more favorable investor reactions than proprietary releases, particularly among resource-constrained firms. Our results thus highlight the importance of diffusion, alongside frontier innovation, in determining the economic value of AI.

Second, we extend research on AI exposure and technology adoption by showing that firms' gains depend on how AI complements their production activities. Prior work examines how firms' exposure to AI shapes market valuations (Eisfeldt et al., 2023) and emphasizes the role of organizational capabilities in technology adoption (Felten et al., 2019; Hershbein and Kahn, 2018; Nagle, 2019). We distinguish potential AI suitability from realized AI usage and further separate augmentation-oriented from displacement-oriented applications, as in Chen and Wang (2025). Valuation gains after open-source releases concentrate among firms with greater exposure to AI-suitable tasks and higher realized AI usage, particularly where AI complements human work. These findings suggest that broader, competitively priced access creates opportunities, but realized value depends on firms' ability to integrate AI into existing activities.

Third, we show how openness shapes the distribution of expected AI gains across firms and their transmission through supply chains. Financial frictions can impede the adoption of productive technologies (Midrigan and Xu, 2014). Consistent with lower access costs easing these barriers, smaller, financially constrained firms fare relatively better after open-source than proprietary releases (see also Cai et al., 2025; Bertomeu et al., 2025). Moreover, open-

source releases generate positive valuation spillovers to upstream and downstream partners in less AI-exposed industries, whereas comparable spillovers are absent following proprietary releases. Building on research linking supply-chain relationships to innovation (Chu et al., 2019), our findings suggest that open access can extend AI's expected benefits beyond directly exposed firms to interconnected industries.

III. 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 use 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 quarterly control variables. We also use these fundamental variables 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 includes 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 the next subsection describes in greater detail.

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 return (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.

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

Because user costs and accessibility are expected to affect how Gen-AI model releases affect stock value, we quantify them directly. 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.[5] 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. First, we aggregate employee-level records to obtain the number of employees per job position within each firm each year, defining job positions 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.[6] We then calculate firm-level Gen-AI labor exposure, $AI\ Exposure^{raw}$, 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.[7] 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

[5] 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.

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

[7] 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.[8] 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

[8] 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.

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 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 purchases from industry k.

Similarly, for each focal industry with below-median industry AI exposure, we construct a measure of downstream-industry AI exposure 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). Their creators fully control these models, and they do not allow local deployment or transparency in implementation (e.g., GPT-3.5, Claude 2.0). *Open-source* models are those that publicly share code, model weights, and sometimes even training data (e.g., DeepSeek-R1, Gemma-2B), a feature that can alter user cost structures by enabling local optimization and potentially lowering compute, memory, or inference latency. This two-way classification allows us to isolate the effect of open-source AI releases, relative to proprietary releases, on firm valuations.

We then collect release events for influential proprietary and open-source Gen-AI models, along with their corresponding event dates, to build 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.[9] We manually collect the initial public release dates of major models from 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 use Hugging Face as the primary platform for disseminating them. For each model type, we record cumulative downloads from its release date through our data collection cutoff on August 27, 2025. Each month and year, we select the top five models by cumulative downloads to focus on the most widely adopted. 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.[10]

We then pool the proprietary and open-source model release dates into a single event timeline. When 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, illustrated in Appendix IA.1, comprises 67 unique event dates representing influential Gen-AI model releases that form the basis of our

[9]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).

[10]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.

event study analyses.

D. Empirical Strategy

We empirically assess the differential expected economic impacts of open-source and proprietary AI models by examining firm-level stock market reactions to model-release announcements. Our analysis rests on the premise that AI is a general-purpose technology whose economic impact depends on widespread adoption. By lowering access barriers, open-source models encourage entry by intermediaries and complementary service providers, intensify competition, and reduce user costs. Together, these forces expand accessibility and facilitate the broader diffusion and adoption of AI.

Our empirical strategy has two components. The first component appears in Section IV, where we show that open-source model releases elicit stronger positive stock-market reactions than proprietary releases, reflecting lower user costs and greater accessibility. Consistent with this interpretation, investors expect open-source releases to drive broader market diffusion, even among financially constrained firms. The second component appears in Section V and shows that lower adoption barriers let constrained firms unlock the latent value of their workforces, reflected in higher valuation gains for firms whose tasks are more amenable to Gen-AI adoption. Augmentation-oriented AI applications drive these gains. We further examine the conjecture that these anticipated efficiency gains generate broader economic effects through supply-chain spillovers from high-AI-exposure firms to vertically linked, low-AI-exposure firms, and that these spillovers arise only after open-source models are released.

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 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 in 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 specification enables us to estimate how stock market reactions to different types of AI model releases vary between firms with high and low Gen-AI exposure.

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

where $AI\ Exposure_{ig}^{raw}$ 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, $AI\ Exposure^{claude}$, $AI\ Exposure^{10k}$, $AI\ Exposure^{ec}$ sourced from the Anthropic Economic Index, 10-K filings, and earnings call transcripts, respectively.

To identify supply-chain spillovers, we restrict attention to focal firms in industries with below-median AI exposure. We examine whether their stock market reactions to AI model release events are shaped by exposure to upstream and downstream industries highly exposed to AI, distinguishing between 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 $AI\ Exposure_{jg}^{upstream}$ is the weighted sum of the AI exposure of upstream industries with above-median AI exposure, with weights equal to the focal industry's shares of intermediate purchases (Equation 3), $AI\ Exposure_{jg}^{downstream}$ is the weighted sum of the AI exposure of downstream industries with above-median AI exposure, with weights equal to the focal industry's shares of intermediate sales (Equation 4), and $AI\ Exposure_{jg}^{ind}$ denotes the AI exposure of 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 also adopt a stacked difference-in-differences (DID) design as an alternative empirical specification alongside 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.

$AI\ Exposure_{ig}$ 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.

IV. Empirical Results: Valuation Effects and Benefit Distribution

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 2 illustrates this progression by plotting the cumulative number of newly added generative AI models from November 2022 to August 2025. Consistent with Demirer et al. (2025), the patterns contrast sharply: proprietary models 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) grow exponentially, 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 3 reports user costs, measured as USD per one million tokens, for each Gen-AI model in our sample.[11] 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 and further broadening participation. Panel B shows that open-source models are significantly less expensive, while proprietary models remain relatively costly.

[11] 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.

Notably, the relationship between open and proprietary development is not one-sided. Proprietary models also serve as an important upstream source for open-source development. Conversely, while open-source 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 pattern is consistent with an innovation trajectory in which radical proprietary innovations reach a pivotal point, after which improvements and broader diffusion follow. Proprietary models can deliver major technological advances but often entail high user costs, concentrating adoption among resource-rich firms. Open-source AI, by contrast, advances through successive algorithmic and architectural improvements, while competition among inference and API providers accelerates cost reductions and expands access (see also Demirer et al., 2025). This trajectory resembles the takeoff phase of technological diffusion, in which cumulative improvements and decentralized innovation enable widespread adoption.[12] As a general-purpose technology, generative AI can reshape economic activities in both phases. However, as discussed in the Introduction, we expect the gains from proprietary innovations to accrue disproportionately to resource-rich adopters, whereas open-source AI facilitates broader adoption and more widely distributed economic benefits 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 reactions, and finally attribute the observed differences to heterogeneity in model user costs and accessibility.

[12] Appendix IA.2 reviews the history of earlier general-purpose technologies, including the steam engine, electricity, and the automobile.

Stock Price Reactions and AI Model Releases

Figure 1 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 1. 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 we use cumulative raw returns (CRR as in Equation 2) instead of cumulative abnormal returns.

To formally quantify the valuation gap and ensure firm fundamentals do not drive the results, we pool the model releases and regress the cumulative abnormal returns (CARs) 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 model 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 enable 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. The motivation is the conjecture that proprietary models benefit well-endowed firms, while open-source models also benefit constrained small firms. In Appendix IB.1, we find that while overall CAR variance is comparable across both release types, the distributions exhibit distinct skewness: CARs for proprietary releases have a slightly heavier left tail, while those for open-source releases have a slightly heavier right tail.[13]

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

[13]In particular, proprietary releases show slightly negative skewness: 4.42% of CARs are 1.65 standard deviations above 0, compared with 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.

expand right-tail growth options for AI-exposed firms, or mitigate their exposure to left-tail risk, or do both.

To distinguish between these conjectures, Table V reports firm-level quantile regressions of firms' cumulative abnormal returns, CAR[-1, +5], assessing whether open-source releases' effects 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 shows that gains in the left tail (5th percentile) from open-source releases come almost entirely from firms with high AI exposure, suggesting significant mitigation of extreme downside valuation risk. However, in the right tail (95th percentile), high-AI-exposure firms see no additional gains. For proprietary releases, firms with high AI exposure show no corresponding changes, except for 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 offers little protection against extreme negative events and provides benefits only 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 do not 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 open-source models' structural advantages: their openness enables broad adoption, and their lower user costs reduce barriers to entry. To test this conjecture, we examine whether accessibility and affordability drive 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, consistent with the view that open access and cost efficiency, features largely absent in proprietary models, are central to the value created by open-source models.

C. Open-Source Releases and the Spread of Adoption

In this subsection, we assess whether open-source AI releases produce stock price reactions consistent with widespread adoption, relative to proprietary releases. Because of their lower user cost and greater accessibility, open-source releases will stimulate inference and API services. The resulting expansion in intermediary markets will generate substantial 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 measure, 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 resource-constraint proxy. The coefficient on OS_g captures the valuation advantage of open-source over proprietary releases within each subsample. Panel A shows that this relative advantage is significantly larger for smaller firms than for larger firms.[14] Likewise, Panels B and C show that the open-source relative advantage is significantly larger for more capital-constrained firms, as measured by the SA and WW indices, than for less constrained firms.[15]

Taken together, these findings indicate that, relative to proprietary model releases, open-source model releases disproportionately benefit resource-constrained firms; the financial market appears to perceive these firms' adoption of open-source AI as value-enhancing. Perhaps because open-source models stimulate intermediary market services that lower entry barriers and user costs, they do not discriminate against resource-constrained firms that adopt them.

[14] Cai et al. (2025) and Bertomeu et al. (2025) demonstrate that smaller firms derive greater benefits from AI adoption in entrepreneurial activities and employment.

[15] Our point is that, relative to proprietary model releases, open-source releases disproportionately benefit more financially constrained firms. For robustness check, we use alternative measures of adoption cost constraints in Appendix IC.1. 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.

V. Impact on Firms and Industries

Another observation, comparing the coefficient on the open-source release dummy across Panels A-C of Table VII, is that, relative to proprietary releases, open-source releases are associated with more favorable valuation responses across firms, particularly among resource-constrained firms. In this section, we examine the leverage firms gain from aligning their task compositions with 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 AI's latent value for workforce augmentation, while gains from labor displacement are less certain, perhaps because of organizational inertia. Accordingly, open-source releases elicit larger valuation gains for firms whose core tasks are highly amenable to AI-driven augmentation. We then examine whether these expected gains have begun to appear in firms' operating outcomes, in the form of higher EBITDA margins and lower operating costs. We further examine whether these firm-level efficiency gains propagate more broadly through supply-chain spillovers.

A. Stock Price Responses and AI Suitability

To assess which firms benefit from open-source models, we examine heterogeneous valuation responses 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-source release indicator in Table VIII. Consistent with the idea that lower adoption costs increase investors' expectations of open-source AI benefits, Panel A examines whether firms

with work tasks more amenable to AI adoption experience significantly greater valuation gains following open-source model releases. From the user side, firms with a higher share of AI-relevant tasks stand to benefit more from AI adoption. Open-source model releases may signal amplified benefits by lowering adoption barriers and fostering competition among AI service providers and complementary technologies, leading investors to assign greater value to these releases. To better capture readiness and lower adjustment costs, 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.[16]

Over the window [-1, +5], the CAR is 1.05 percentage points for firms with a workforce characterized by high potential AI exposure, as shown in Panel A. The corresponding CAR is 2.08 percentage points higher 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 of AI exposure explains the cross-sectional variation in CARs following proprietary model releases, further underscoring the distinction between open-source and proprietary models.[17] 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.

The contrasting observations between open-source and proprietary AI releases in this subsection and the previous sections align with 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, 2025). In

[16] $AI\ Exposure^{raw}$ 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.

[17] In Appendix Section IC.2, 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.

contrast, open-source releases appear to yield broader gains by increasing expected investment opportunities. This interpretation is consistent with Björkegren (2026), who shows that open-source and proprietary AI model advancements have opposite effects on long-term interest rates, with open-source advancements raising long-term Treasury yields.

B. Augmentation- vs. Displacement-Oriented AI

While the preceding results show that valuation gains concentrate among firms with high AI exposure, the economic impact likely depends on the nature of that exposure. Following Chen and Wang (2025), we distinguish the effects of augmentation- vs. displacement-oriented AI innovation on firm value. We use 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 shows a 2.18-percentage-point increase in CAR[-1, +5]. In contrast, Panel B shows that firms whose AI use is more displacement-oriented see 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. This asymmetry is consistent with the idea that displacement entails cutting and creating new posts, which likely face 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 face less implementation friction 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.[18]

[18] In Appendix IC.3, 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.

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

For displacement-oriented applications, we do not observe significant negative effects, suggesting that disruption from startups, which tend to adopt AI more readily and flexibly, has not yet been large enough to generate measurable adverse impacts on listed firms. 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. Early Observations on Operating Outcomes

Our analyses thus far rely on stock-market reactions to capture investors' expectations about the economic implications of AI model releases. Whether these expectations translate into realized operating gains, however, remains an open question. As with other general-purpose technologies, realizing AI's productivity potential requires complementary investments in organizational change, worker retraining, and integrating new tools into existing processes; consequently, productivity gains may emerge only gradually in firms' operating performance and financial statements (Brynjolfsson et al., 2019). The growing competitiveness of open-source models relative to frontier proprietary systems since 2025 provides an opportunity to examine early evidence of whether firms that investors expected to benefit from these releases are beginning to realize operational gains.

To this end, we construct a quarterly panel of U.S. non-financial firms from Compustat for 2019Q1–2026Q2, including firms that subsequently delisted, and examine two operating outcomes: EBITDA margin and operating costs (cost of goods sold plus SG&A) relative

to sales. Because realized operating gains likely depend on how AI is used, rather than merely whether it is used, we construct a firm-level measure of augmentation-oriented AI usage. Specifically, we map occupation-level evidence on augmentation-oriented AI use from the Anthropic Economic Index to firms' workforce composition. We then examine the relationship between operating outcomes and this measure separately for financially constrained and unconstrained firms, classified using their SA and WW indices in 2021, before ChatGPT (GPT-3.5) was released in November 2022.

Appendix IB.2 shows that before 2025, constrained and unconstrained firms follow similar paths, with no response following the release of ChatGPT (GPT-3.5). After open-source models reached performance comparable to leading proprietary models with the release of DeepSeek-R1 in January 2025, constrained firms with higher augmentation-oriented usage show rising margins and declining operating costs. In contrast, unconstrained firms show at most a modest change. Table X confirms that these differences are statistically significant, at the 10% level using the SA index and at the 5% level or better using the WW index. Because sales growth and capital expenditures show no comparable change, the gains appear to operate through lower operating costs rather than expansion or new investment, consistent with AI adoption that requires little upfront capital.

These results reflect early observations consistent with expectations embedded in stock prices, rather than a direct test of the open-source mechanism. Because firms can access both open-source and proprietary models after 2025, the timing of the improvements is suggestive but does not isolate open-source adoption from AI adoption more broadly. Moreover, with only six quarters of post-period data (2025Q1–2026Q2) available, the persistence and longer-run productivity effects are uncertain. Subject to these qualifications, the results suggest that constrained firms whose workforces are concentrated in occupations where AI tends to augment human work have begun to show operating improvements, which is consistent with investors' expectations.

D. Spillover Effects Through the Supply Chain

Having taken a firm-level approach to examine how AI exposure affects firm valuations after AI model releases, we now turn to an industry-level approach by examining supply-chain channels. Specifically, we examine how AI exposure among upstream suppliers and downstream customers affects focal firms' valuations.

Table XI shows supportive regression results. Columns 3 and 4 show that firms in industries with low AI exposure see an increase in firm value after open-source releases, driven by supply-chain spillovers originating from suppliers and buyers with high AI exposure. This result holds even after controlling for the AI exposure of the firm's own industry, which typically exerts a larger direct effect. By contrast, Columns 1 and 2 show no supply-chain spillovers following proprietary releases.[19]

In Figure 6, we plot dynamic spillover effects through the supply chain over a 20-day horizon using the spillover coefficients from the specifications in Table XI. 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.

Our economic intuition is that open-source AI generates vertical spillovers to firms in low-AI-exposure industries through two complementary channels. Upstream spillovers arise when AI-enabled suppliers improve input quality, reduce costs, shorten delivery times, and provide more reliable inputs, raising downstream productivity and lowering coordination and integration costs. Downstream spillovers arise when AI-enabled customers improve demand forecasting, specify orders more precisely, and coordinate production more effectively, benefiting suppliers through more predictable demand, better capacity utilization, and lower

[19] In Appendix IC.4, we use FactSet Revere as an alternative source of network information 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.

inventory costs. These gains may, however, be offset by the costs of meeting higher standards for quality, traceability, and delivery, as well as by price pressure from more efficient buyers.

Positive net spillovers therefore may primarily follow open-source releases. By lowering barriers to AI deployment, open-source models enable broader adoption across firms, intensifying competition and limiting the concentration of productivity gains among a small number of AI-intensive firms. This diffusion lets efficiency gains propagate through supply chains, raising the value of firms in low-AI-exposure industries even when they do not adopt AI intensively themselves. By contrast, proprietary releases are less likely to generate comparable vertical spillovers: when access to advanced AI capabilities remains more restricted, productivity gains are more likely to accrue to adopting firms and model providers, with less diffusion to firms in indirectly exposed industries. It follows that proprietary releases are unlikely to produce systematic positive valuation spillovers along supply chains.

Taken together, we provide evidence that the open-source AI value creation propagates through supply chain networks to firms with low AI exposure via their economic ties to upstream and downstream industries with high AI exposure. Furthermore, releasing proprietary models does not generate spillover benefits across supply chains.

VI. Robustness Checks

In this section, we examine the validity and generalizability of our findings through four analytical lenses: economic magnitude, model type imbalance and timing bias, causal identification, and alternative measurement.

First, we evaluate the robustness of the economic magnitude and aggregate impact of AI model releases. Specifically, to test whether our findings are sensitive to extreme values, we report results without winsorization in Panel A of Table XII. They show that stock market reactions remain significantly negative for proprietary releases and significantly positive for open-source ones across the full distribution. To show that these valuation gains represent

meaningful wealth creation, we also use a Weighted Least Squares (WLS) approach in Panel B, weighting observations by each firm's market capitalization measured one trading day before the respective model release. The consistency of these results shows that valuation gains from open-source releases are economically significant, extend to the largest firms in the economy, and are not driven merely by idiosyncratic volatility in small-cap stocks.

Second, we address the sample imbalance between 41 proprietary and 26 open-source releases, as well as potential bias from announcement timing. In Panel A of Table XIII, we conduct propensity score matching (PSM) on firm-event observations. Specifically, we pool the firm-event observations associated with 41 proprietary Gen-AI model releases and 26 open-source Gen-AI model releases. Using the pooled sample, we then estimate a logit model following Bilicka (2019), where the dependent variable is the open-source release indicator, OS_g, equal to 1 if event g is an open-source Gen-AI model release and 0 if it is a proprietary Gen-AI model release. The explanatory variables are firm-level stock returns over the two to five trading days preceding the release of model g. Using the estimated propensity scores, we implement nearest-neighbor matching without replacement, with a caliper of 0.01 standard deviations of the logit of the estimated propensity score, to match each observation associated with an open-source AI model release to an observation associated with a proprietary AI model release with comparable pre-release return dynamics. This procedure improves comparability between firm-event observations in the open-source release sample and the proprietary release sample in pre-event stock-market performance. The estimates remain robust in the matched sample.

While propensity score matching addresses selection bias, it has limitations, such as discarding unmatched observations and sensitivity to the caliper choice. Therefore, we complement our analysis with the entropy balancing methodology following McMullin and Schonberger (2020), which reweights all firm-event observations associated with proprietary Gen-AI model releases so that the moments of their pre-event returns, over the two to five

trading days before each release, match those of firm-event observations associated with an open-source Gen-AI model release. Panel B of Table XIII consistently shows negative impacts on firm valuations for proprietary releases and positive impacts for open-source releases.

We also include year-month fixed effects to ensure our results reflect the specific nature of the model being released, not broad market cycles. Panel C of Table XIII shows consistent results. Collectively, these imbalance and timing-bias tests confirm that valuation gains are driven by the open-source nature of the releases, not announcement timing.

Third, as an alternative to the event-study approach, we employ a stacked triple difference-in-differences strategy specified in Equation 8 (Section III.D), with results presented in Table XIV. For this analysis, the unit of observation is the firm-day, spanning one trading day before through five trading days after the corresponding model release. By constructing event-specific firm-day panels, this framework addresses selection bias arising from endogenous firm characteristics. Specifically, it mitigates concerns that firms self-select into high AI-exposure categories based on unobserved growth trajectories or pre-event momentum. Panel A of Table XIV shows that, after the release of open-source models rather than proprietary models, firms with higher AI exposure experienced significantly improved firm valuations. While we build on the PSM and entropy-balancing methodologies detailed in Table XIII, we modify them here to explicitly balance high- and low-exposure firms on pre-event run-up returns. Reflecting this adjustment, Panel B confirms that these results are robust in the propensity score-matched (PSM) sample using high AI exposure firms as treated.[20] In Panel

[20]Following Bilicka (2019), we use propensity score matching to restrict the sample of control firms. For each event g, we classify firms into treated and control groups based on $AI\ Exposure_{ig}^{raw}$, a continuous Gen-AI exposure score from Eisfeldt et al. (2023). Treated firms are those with Gen-AI exposure above the sample median, for which $AI\ Exposure_{ig} = 1$, while control firms are those with Gen-AI exposure below the sample median, for which $AI\ Exposure_{ig} = 0$. We then estimate a logit model for each event in which $AI\ Exposure_{ig}$ is the dependent variable and firms' pre-event abnormal daily returns over trading days $t = -5, -4, -3$, and -2, where $t = 0$ denotes the release date of event g, are the predictors, which give us propensity scores. Note that we use $AI\ Exposure_{ig}$ as the dependent variable in this process, instead of OS_g, an open-source release indicator used in Panel A of Table XIII, as we intend to alleviate selection

C, we apply entropy balancing to achieve exact moment-matching for ex-ante stock returns across high- and low-AI-exposure firms within each event stack.[21] The results continue to support our main findings.

To further bolster the causal interpretation and address potential violations of the parallel-trends assumption, we supplement our stacked DID results with state-of-the-art counterfactual estimators. Recent work demonstrates that linear two-way fixed-effects (TWFE) estimators can be biased and 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 counterfactual estimators that accommodate richer unobserved time-varying heterogeneity: (i) the fixed-effects counterfactual estimator (FEct), (ii) the interactive fixed-effects counterfactual estimator (iFEct), (iii) the matrix completion estimator (MC), and (iv) synthetic matching.

Specifically, for each event, we define treated firms as those with $AI\ Exposure_{ig}^{raw}$ above the sample median for the corresponding year. For each treated observation, we predict the

bias across firm types. Using these propensity scores, we implement nearest-neighbor matching without replacement within a tight caliper, defined as 0.01 standard deviations of the logit of the propensity score. This procedure matches each treated firm with high Gen-AI exposure to the most comparable control firm with low Gen-AI exposure, based on pre-event stock-market performance. After matching, we construct an event-level panel that includes daily abnormal returns for treated firms and their matched control firms before and after the event. We then stack the event-level panels across all events and estimate the stacked difference-in-differences specification.

[21] In the entropy balancing, firms with high Gen-AI exposure constitute the treatment group, while firms with low Gen-AI exposure serve as controls. Note that we match the run-up return moments of firms with high versus low AI exposure in this process, as opposed to matching the run-up return moments of firms during proprietary versus open-source model releases as in Panel B of Table XIII, as we intend to alleviate selection into firm types. The balancing covariates are the same pre-event stock returns employed in the propensity-score matching exercise, measured over the two to five trading days before each release. Entropy balancing reweights the control firms so that the moments of the covariate distribution among controls exactly replicate those of the treatment firms.

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 a single dummy variable, OS_g, and report the estimation results in Table XV. 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), that improves on standard propensity score matching by constructing control firms whose pre-event abnormal return trajectories closely resemble those of treated firms. We use the 250- to 30-trading-day period before the first event announcement (the release of OpenAI's GPT-3.5) as the estimation window. For each firm in the treatment group with $AI\ Exposure_{ig}^{raw}$ above the sample median in the corresponding year, we estimate synthetic-control weights by minimizing the distance between the treated firm's daily abnormal returns and the weighted average daily abnormal returns of the donor pool over the estimation window. For each trading day after the event date, we construct the counterfactual outcome for each treated firm as the weighted average of its synthetic control's abnormal returns, using the pre-estimated weights. The results reported in Column 4 of Table XV further support the robustness of our baseline findings.

Fourth, we evaluate alternative measures of firms' exposure to Gen-AI, drawing on multiple sources and as reported in Section III.B, to confirm the consistency of market reactions to open-source model releases. Table XVI presents the AI exposure measure based on 10-K filings and earnings call transcripts, along with the results from Equation 6. All confirm that firms more deeply integrated with AI technologies show stronger positive market reac-

tions to open-source model releases. 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, to enhance robustness, we use alternative outcome variables instead of cumulative abnormal returns (CAR) in Panel A of Table XVII. 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 this case, the estimated beta may no longer reflect underlying market dynamics, potentially biasing abnormal return calculations. The results show that firms with higher AI exposure still see significant gains after the release of open-source models, even when using raw returns. The effect is economically meaningful, suggesting investors 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 XVII. 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.

VII. Conclusion

The economic impact of AI depends not only on advances in model capabilities, but also on who can access those capabilities and at what cost. Using stock-market reactions to 67 influential AI model releases from November 2022 to August 2025, we find more favorable valuation responses to open-source releases than to proprietary releases. These gains concentrate among firms with greater AI task suitability, particularly where AI augments rather than substitutes for human input. Relative to proprietary releases, open-source releases also disproportionately benefit smaller, financially constrained firms. These patterns are con-

sistent with open-source models lowering deployment barriers and with competition among complementary service providers broadening access to AI. Preliminary operating evidence is also consistent with these expectations, although results from a short window can only be suggestive.

Importantly, the anticipated benefits extend beyond firms that use AI intensively themselves. Open-source releases are associated with positive valuation spillovers to firms in low-AI-exposure industries linked to high-AI-exposure suppliers or customers; we find no comparable spillovers following proprietary releases. This contrast suggests that broader AI adoption can transmit efficiency gains through supply chains, extending expected benefits to firms that do not directly adopt AI intensively. The economic significance of openness thus lies not only in expanding adoption, but also in allowing its benefits to reach interconnected businesses.

These findings contribute to our understanding of how a general-purpose technology diffuses and distributes economic value. Open-source models can stimulate competition and complementary innovation in inference, APIs, and applications, potentially shifting some of the benefits of technological progress from model developers toward adopters and complementors. Yet access alone does not guarantee productivity gains: computing resources, workforce skills, and organizational adaptation remain important constraints. Moreover, stock-market reactions capture investors' expectations, not necessarily realized productivity growth. Our early operating evidence is consistent with firms beginning to realize some anticipated gains. Still, the short period since early 2025 does not establish persistence or isolate the effects of open-source adoption from broader AI adoption. Further investigation with a longer data window is warranted.

For economies seeking to benefit from AI without leading frontier-model development, open-source models may broaden participation opportunities by reducing access costs and supporting complementary innovation. Realizing this potential requires the skills, infras-

tructure, and organizational capabilities to turn access into productive use. Ultimately, our findings suggest that the economic consequences of AI depend not only on how far model capabilities advance, but also on how widely those capabilities diffuse and how their benefits are distributed across firms and industries. For policymakers, the challenge is therefore not simply to promote access to AI, but to build the conditions for access to translate into productive use and broadly shared economic gains.

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.

Azoulay, Pierre, Joshua L. Krieger, and Abhishek Nagaraj, 2024, Old moats for new models: Openness, control, and competition in generative AI, NBER Working Paper No. 32474.

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.

Bengio, Yoshua, Stephen Clare, Carina Prunkl, et al., 2026, International AI safety report 2026, arXiv:2602.21012.

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.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, 2019, Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics, in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, 23–57 (University of Chicago Press, Chicago).

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.

Chu, Yongqiang, Xuan Tian, and Wenyu Wang, 2019, Corporate innovation along the supply chain, *Management Science* 65, 2445–2945.

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, and Nadav Tadelis, 2026, The emerging market for intelligence: How firms buy and sell AI, *Journal of Economic Perspectives* 40, 23–46.

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.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, 2024, GPTs are GPTs: Labor market impact potential of LLMs, *Science* 384, 1306–1308.

Felten, Edward W., Manav Raj, and Robert Seamans, 2019, The occupational impact of artificial intelligence: Labor, skills, and polarization, Working paper, National Bureau of Economic Research.

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.

Hershbein, Brad, and Lisa B. Kahn, 2018, Do recessions accelerate routine-biased technological change? Evidence from vacancy postings, *American Economic Review* 108, 1737–1772.

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.

Korinek, Anton, and Jai Vipra, 2024, Concentrating intelligence: Scaling and market structure in artificial intelligence, NBER Working Paper No. 33139.

Lerner, Josh, and Jean Tirole, 2005, The economics of technology sharing: Open source and beyond, *Journal of Economic Perspectives* 19, 99–120.

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

Midrigan, Virgiliu, and Daniel Yi Xu, 2014, Finance and misallocation: Evidence from plant-level data, *American Economic Review* 104, 422–58.

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

Nagle, Frank, 2019, Open source software and firm productivity, *Management Science* 65, 1191–1215.

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.

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.

Widder, David Gray, Meredith Whittaker, and Sarah Myers West, 2024, Why 'open' AI systems are actually closed, and why this matters, *Nature* 635, 827–833.

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.

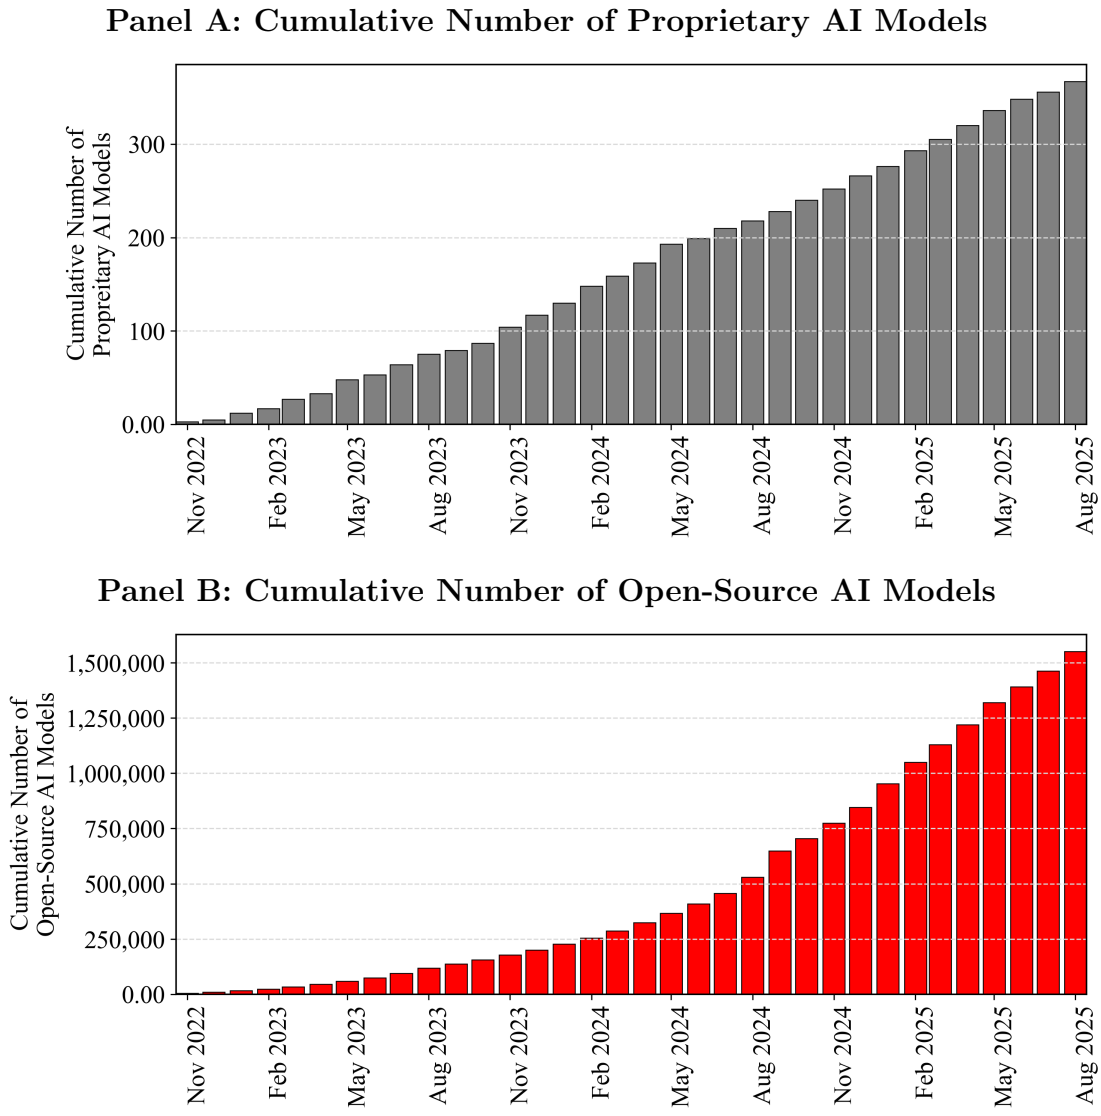


Figure 2. Cumulative Number of AI Models: Proprietary vs. Open-Source

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.

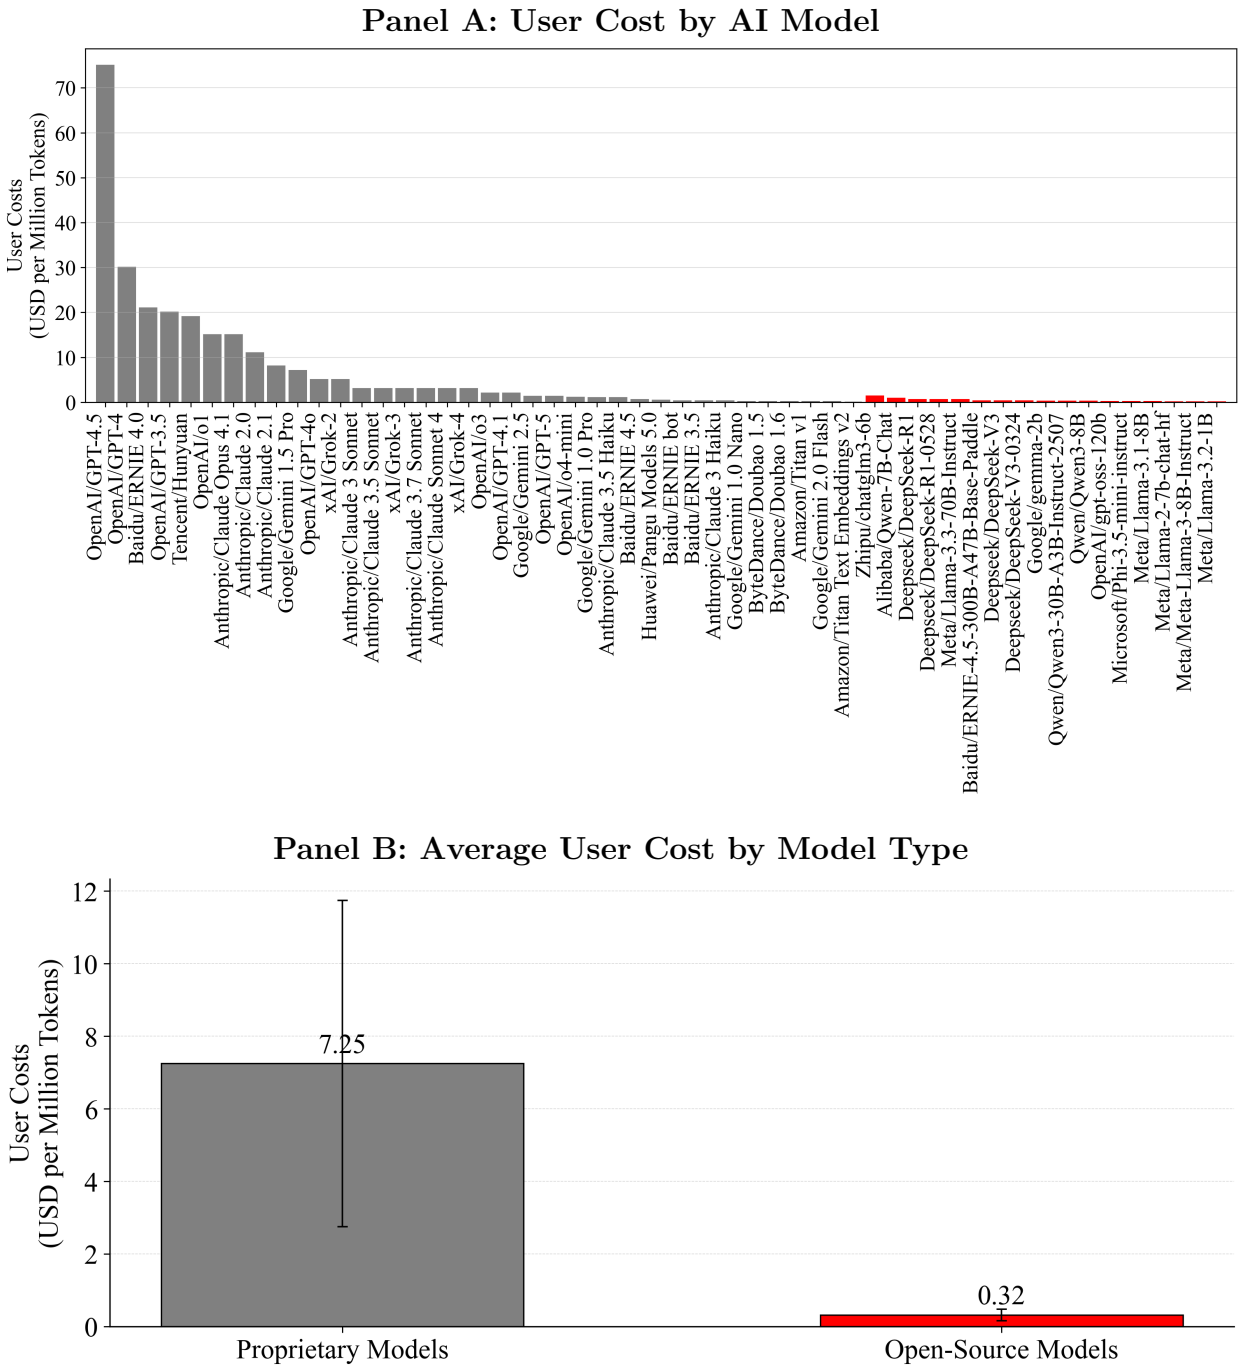
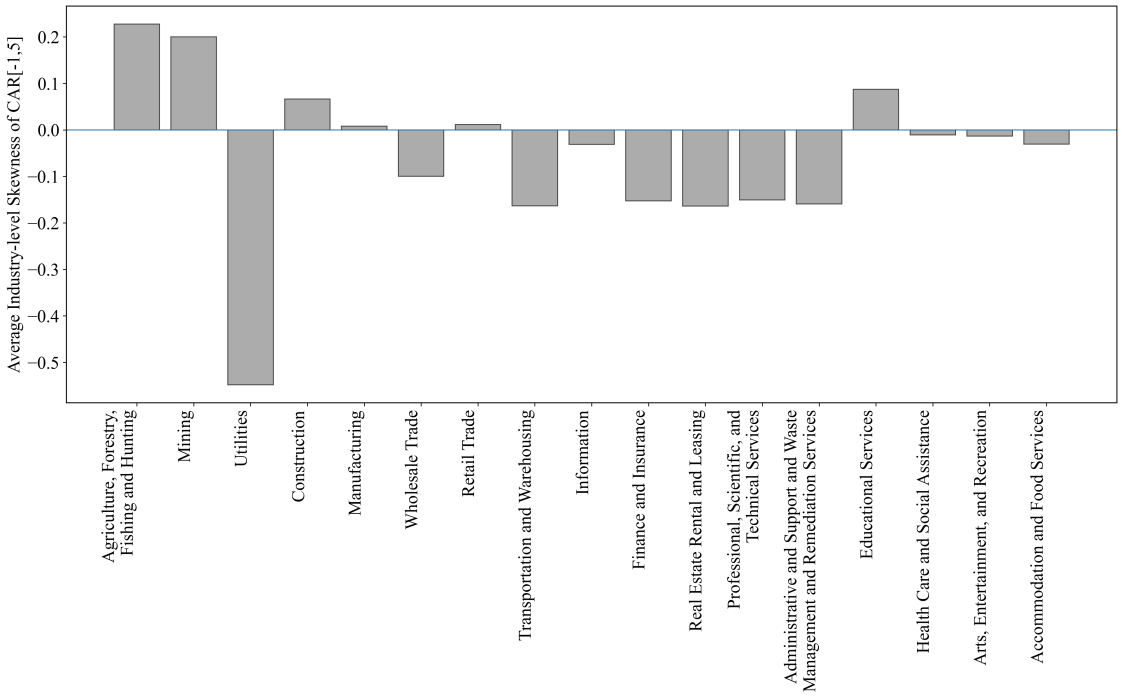


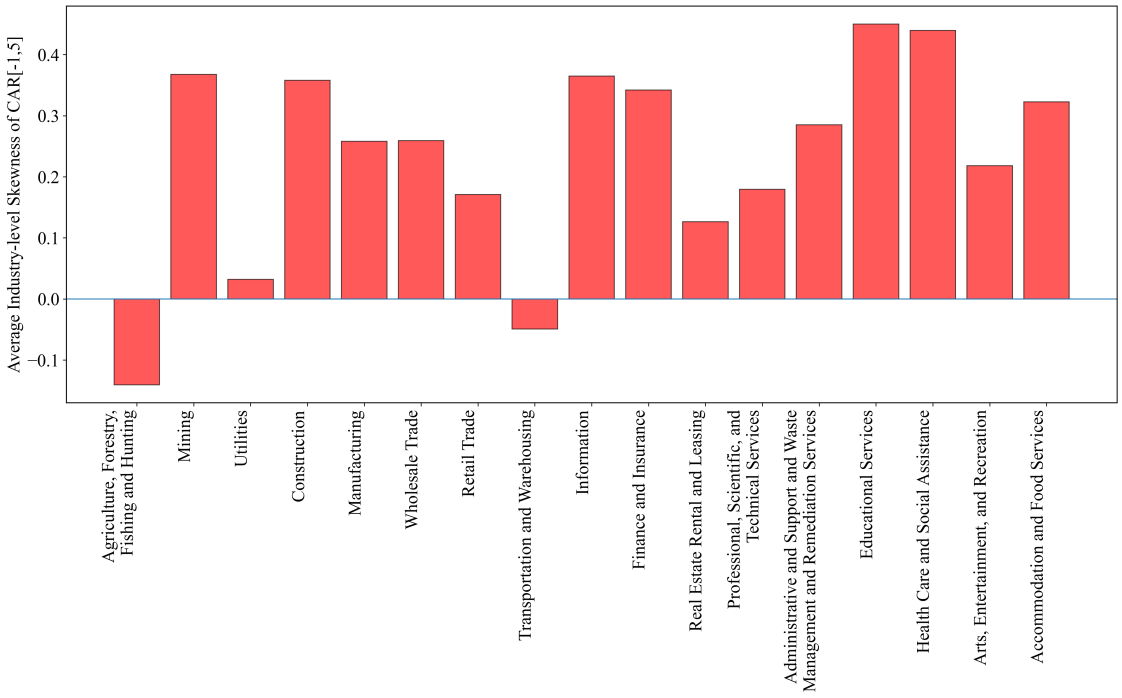
Figure 3. User Costs of AI Models: Proprietary vs. Open-Source

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.

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



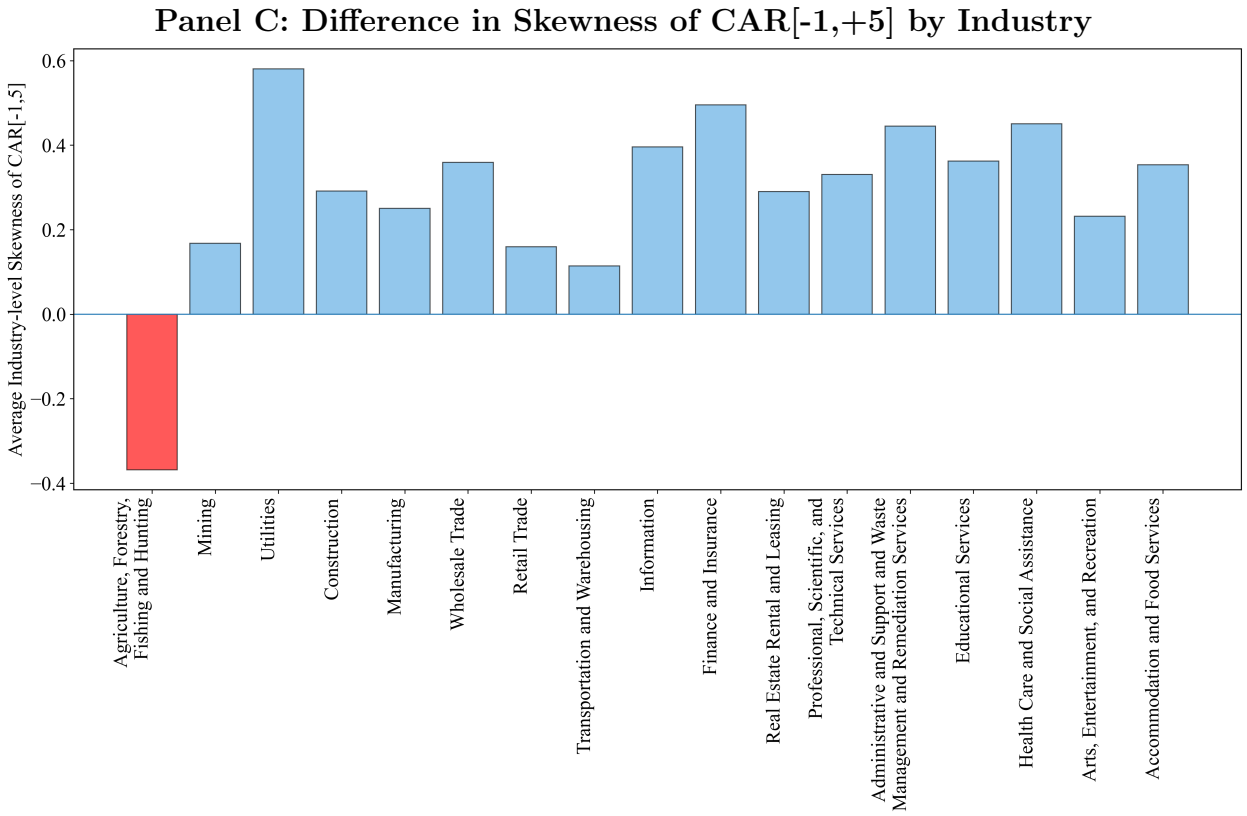


Figure 4. Heterogeneity in Skewness by Industry: Proprietary vs. Open-Source

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.

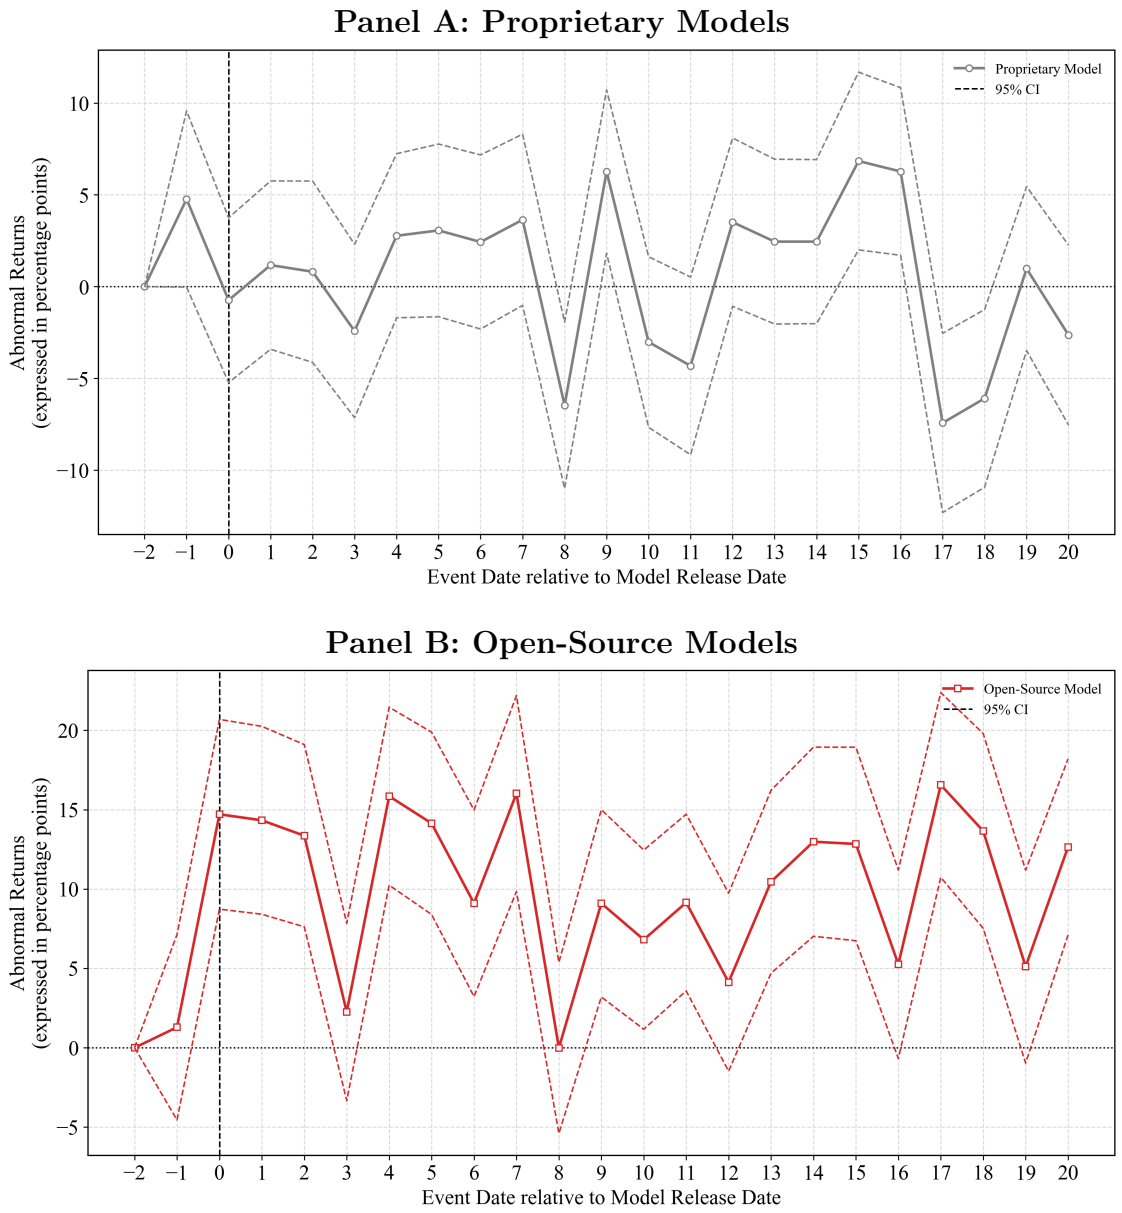
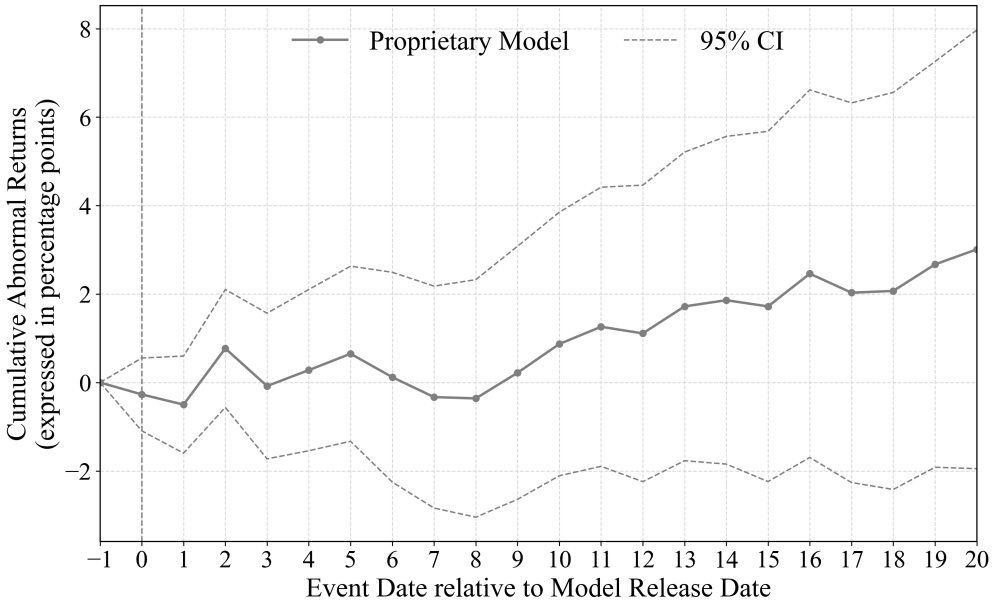


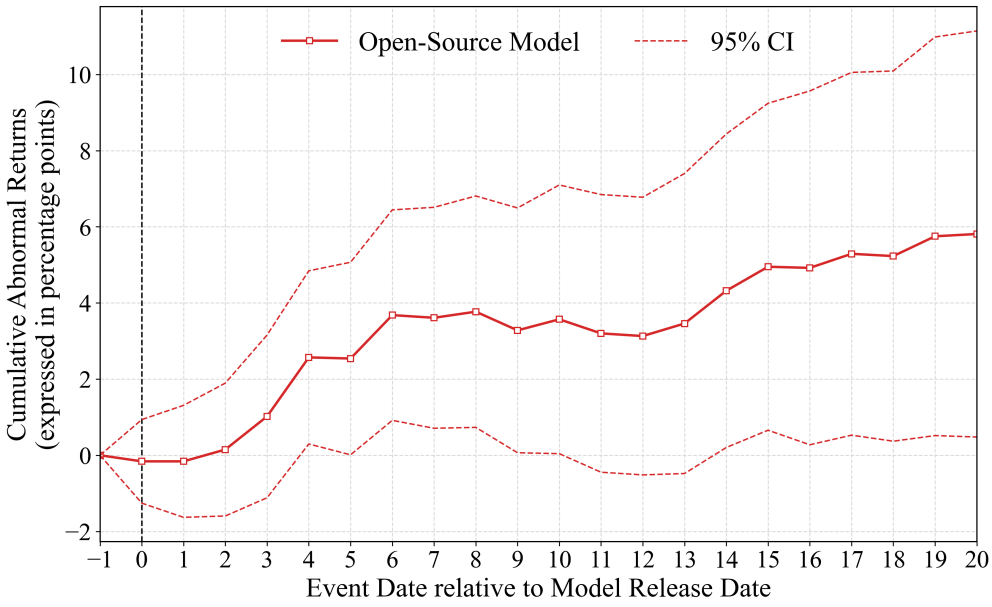
Figure 5. Heterogeneity in Abnormal Returns by AI Exposure

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.

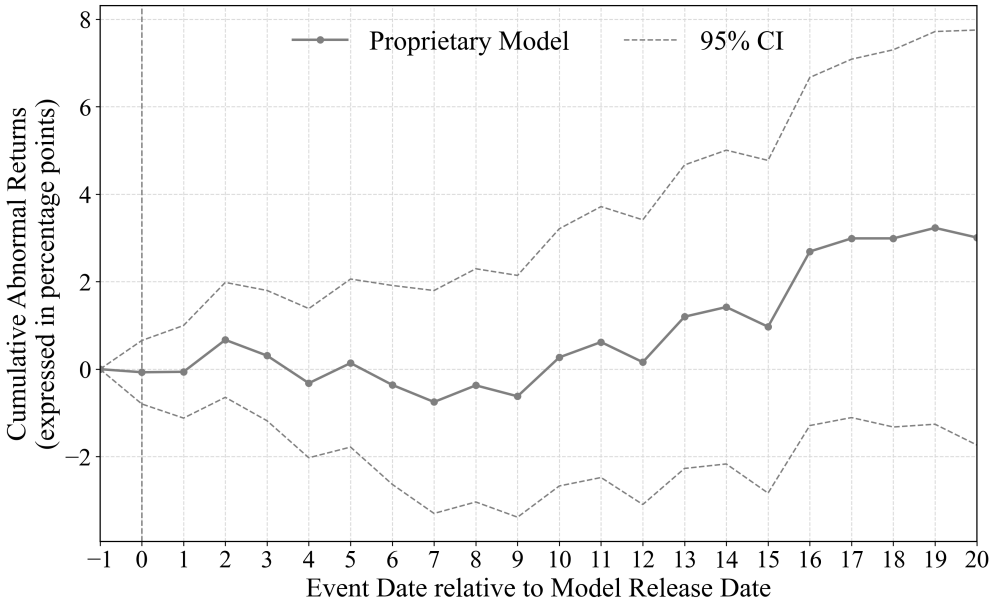
Panel A: High Upstream AI Exposure and Proprietary Releases



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

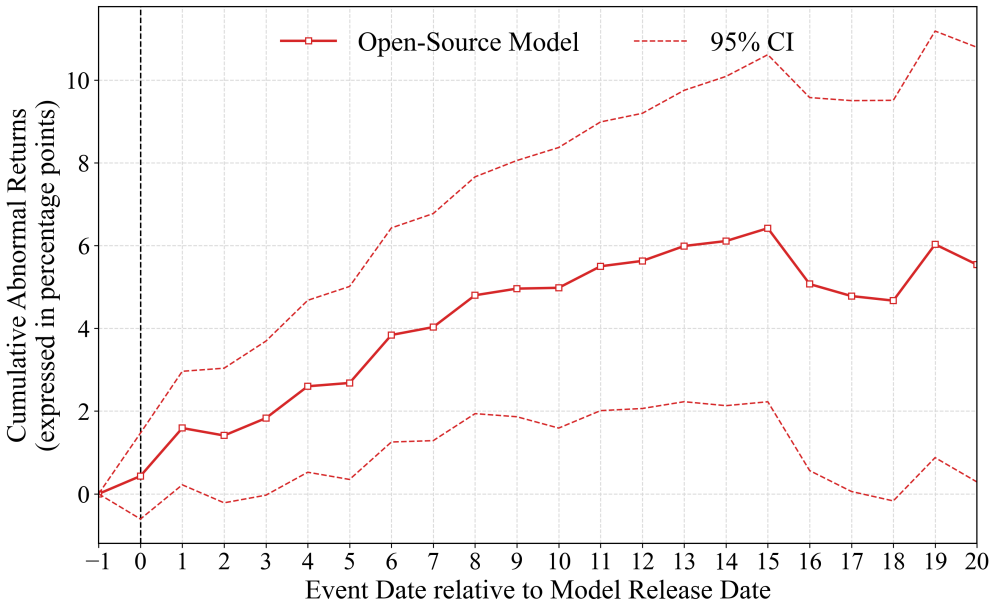


Figure 6. Dynamic Spillover Effects Through Supply Chain

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 XI. 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 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) $\text{CAR}[-1,+0]_{ig}$	(2) $\text{CAR}[-1,+1]_{ig}$	(3) $\text{CAR}[-1,+5]_{ig}$
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
AI Augmentation Usage and Operating Outcomes after DeepSeek-R1

This table reports triple-difference estimates of the change in operating outcomes following the release of DeepSeek-R1. We estimate

$$Y_{i,t} = \alpha_i + \delta_{j,t} + \beta\,(\textit{Post-R1}_t \times \textit{Usage}_i) + \theta\,(\textit{Post-R1}_t \times \textit{Usage}_i \times \textit{Constrained}_i) + \gamma\,(\textit{Post-R1}_t \times \textit{Constrained}_i) + \varepsilon_{i,t},$$

where $Y_{i,t}$ is the EBITDA margin (columns 1 and 3) or operating costs (cost of goods sold plus SG&A) relative to sales (columns 2 and 4) of firm i in quarter t. $\textit{Post-R1}_t$ equals one for 2025Q1–2026Q2 and zero otherwise. $\textit{Usage}_i$ is the firm's augmentation-oriented AI usage from the Anthropic Economic Index, standardized to mean zero and unit standard deviation. $\textit{Constrained}_i$ equals one if the firm's SA index (columns 1–2) or WW index (columns 3–4) in 2021 is at or above the sample median. α_i denotes firm fixed effects, which absorb $\textit{Usage}_i$ and $\textit{Usage}_i \times \textit{Constrained}_i$; $\delta_{j,t}$ denotes two-digit NAICS industry × quarter fixed effects, which absorb $\textit{Post-R1}_t$. θ measures the additional post-R1 change for constrained firms relative to unconstrained firms per one-standard-deviation increase in augmentation usage, and β measures the corresponding change for unconstrained firms. The sample covers U.S. non-financial, non-utility firms in Compustat, including firms that subsequently delisted, from 2022Q4 to 2026Q2, so the comparison period spans the release of ChatGPT (GPT-3.5) through the release of DeepSeek-R1. 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)
Constraint index	SA	SA	WW	WW
Dependent variable	EBITDA margin	Costs / sales	EBITDA margin	Costs / sales
$\textit{Post-R1}_t \times \textit{Usage}_i \times \textit{Constrained}_i$	0.0123*	-0.0121*	0.0149**	-0.0178***
	(0.0067)	(0.0064)	(0.0066)	(0.0066)
$\textit{Post-R1}_t \times \textit{Usage}_i$	0.0063*	-0.0057*	0.0009	0.0009
	(0.0032)	(0.0034)	(0.0034)	(0.0035)
$\textit{Post-R1}_t \times \textit{Constrained}_i$	-0.0013	-0.0019	0.0096	-0.0098
	(0.0066)	(0.0062)	(0.0065)	(0.0065)
Firm FE	Y	Y	Y	Y
NAICS2 × Quarter FE	Y	Y	Y	Y
Firms	2,038	1,937	1,859	1,769
Adj. R^2	0.825	0.821	0.816	0.809
Obs.	27,847	26,625	25,606	24,472

Table XI
Spillover Effects Through Supply Chain

This table presents how stock market reactions to proprietary and open-source AI model releases spill over through the supply chain. The sample is restricted to firms in industries with below-median AI exposure. The dependent variable is the cumulative abnormal return, expressed in percentage points, for firm i in industry j from one trading day before to five trading days after event g. Industries and input–output weights are based on the OECD Input–Output Table for the United States. $AI\ Exposure_{jg}^{upstream}$ is the weighted sum of the AI exposure of upstream industries with above-median AI exposure, with weights equal to the share of the focal industry's total intermediate input purchases sourced from each upstream industry (Equation 3). $AI\ Exposure_{jg}^{downstream}$ is the weighted sum of the AI exposure of downstream industries with above-median AI exposure, with weights equal to the share of the focal industry's total intermediate sales directed to each downstream industry (Equation 4). $AI\ Exposure_{jg}^{ind}$ is the average firm-level AI exposure in the focal industry. Columns 1 and 2 present results for proprietary AI model releases, and Columns 3 and 4 present results for open-source AI model releases. Columns 1 and 3 include upstream exposure, and Columns 2 and 4 include downstream exposure. 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	
	$CAR[-1,+5]_{ig}$			
$AI\ Exposure_{jg}^{upstream}$	0.652		2.586**	
	(1.006)		(1.289)	
$AI\ Exposure_{jg}^{downstream}$		0.138		2.708**
		(0.983)		(1.190)
$AI\ Exposure_{jg}^{ind}$	1.154	-0.660	2.797	1.041
	(4.105)	(3.494)	(4.718)	(3.928)
Constant	-2.810*	-2.081	-1.490	-0.772
	(1.584)	(1.346)	(1.848)	(1.496)
Controls		log(Assets), D/E, ROA		
Industry FE	Y	Y	Y	Y
Adj. R^2	0.018	0.018	0.011	0.011
Obs.	32,173	32,173	20,400	20,400

Table XII
Robustness: Economic Magnitude

This table presents robustness tests regarding the economic significance of the valuation effects to ensure the results are not driven by idiosyncratic volatility in small-cap stocks. Panel A reports results using cumulative returns without winsorization to capture the full distribution of market responses. Panel B reports results estimated via Weighted Least Squares (WLS); following Eisfeldt et al. (2023), observations are weighted by each firm's market capitalization measured one trading day prior to the respective model release (e.g., market capitalization on November 29, 2022, is used for OpenAI's GPT-3.5). 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. 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: No Winsorization

	(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: Weighted Least Squares

	(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

Table XIII
Robustness: Addressing Model Type Imbalance and Timing Bias

This table presents robustness tests addressing potential sample imbalance and timing bias between the 41 proprietary and 26 open-source model releases. Panel A reports market reactions based on a propensity-score-matched (PSM) sample. Using the sample pooling all firm-level observations associated with 41 proprietary Gen-AI model releases and 26 open-source Gen-AI model releases, we first estimate a logit model following Bilicka (2019) where the treatment indicator of open-source Gen-AI model releases, OS_g, is regressed on firms' pre-event stock returns over the two to five trading days preceding each release. Using a caliper of 0.01 standard deviations of the logit of the propensity score, we implement nearest-neighbor matching without replacement to match each firm observation associated with an open-source Gen-AI model release to a firm observation associated with a proprietary Gen-AI model release that exhibits comparable pre-release return dynamics. Panel B utilizes entropy balancing on the same pre-event return covariates following McMullin and Schonberger (2020). This procedure reweights observations in the sample of proprietary model releases such that the moments of their covariate distribution exactly replicate those of the sample of open-source model releases. Panel C includes year-month fixed effects to control for broader market cycles and seasonality. 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. 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: Propensity Score Matching

	(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: Entropy Balancing

	(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 C: Including Year-Month Fixed Effects

	(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

Table XIV
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). We construct the regression sample by stacking event-specific samples, where each event-specific sample consists of daily observations for all firms over the event window spanning one trading day before through five trading days after the corresponding model release. The dependent variable is the daily 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. Panel A reports baseline results using the full sample of firms. Panel B reports estimates based on a propensity-score-matched sample, where each treated firm with high AI exposure is matched to a control firm with low AI exposure based on pre-event abnormal stock return trajectories. See Footnote 20 for details on the propensity score matching procedure. Panel C reports results utilizing an entropy-balanced sample, where firms with low AI exposure are reweighted to ensure their ex-ante stock return distribution moments exactly replicate those of firms with high AI exposure within each event cohort. 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: DID Estimation

	AR_{itg}	
	(1)	(2)
$OS_g \times Post_{tg} \times AI\ Exposure_{ig}$	0.150***	0.119***
	(0.0300)	(0.0340)
$Post_{tg} \times AI\ Exposure_{ig}$	-0.0410**	-0.0720***
	(0.0190)	(0.0210)
Constant	0.222	0.040
	(1.557)	(1.514)
Controls	log(Assets), D/E, ROA	
Event-Firm FE	Y	Y
Event-Date FE	Y	Y
Event-Industry-Date FE	N	Y
Adj. R^2	0.0270	0.0490
Obs.	1,429,842	1,429,842

Panel B: DID Estimation with Propensity Score Matching		
	AR_{itg}	
	(1)	(2)
$OS_g \times Post_{tg} \times AI\ Exposure_{ig}$	0.114***	0.0970***
	(0.0310)	(0.0360)
$Post_{tg} \times AI\ Exposure_{ig}$	-0.0180	-0.0500**
	(0.0200)	(0.0230)
Constant	0.610	0.208
	(1.851)	(1.768)
Controls	log(Assets), D/E, ROA	
Event-Firm FE	Y	Y
Event-Date FE	Y	Y
Event-Industry-Date FE	N	Y
Adj. R^2	0.0260	0.0480
Obs.	1,222,844	1,222,830
Panel C: DID Estimation with Entropy Balancing		
	AR_{itg}	
	(1)	(2)
$OS_g \times Post_{tg} \times AI\ Exposure_{ig}$	0.150***	0.130***
	(0.031)	(0.036)
$Post_{tg} \times AI\ Exposure_{ig}$	-0.044**	-0.077***
	(0.020)	(0.022)
Constant	0.546	0.352
	(1.576)	(1.551)
Controls	log(Assets), D/E, ROA	
Event-Firm FE	Y	Y
Event-Date FE	Y	Y
Event-Industry-Date FE	N	Y
Adj. R^2	0.027	0.049
Obs.	1,429,602	1,429,602

Table XV
Robustness: Counterfactual Estimators

This table presents the effect of Gen-AI model releases on U.S. firm value, using counterfactual estimators. 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

Table XVI
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 XVII
Robustness: Alternative Outcome Variables

This table presents the effect of Gen-AI model releases on U.S. firm value, 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

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

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 and Vivarelli, 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 and Gurbaxani, 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 and Vivarelli, 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 1770s and 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 1879 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 U.S. homes electrified 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, introduced in 1908 and mass-produced on the moving assembly line from 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 1920s (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.

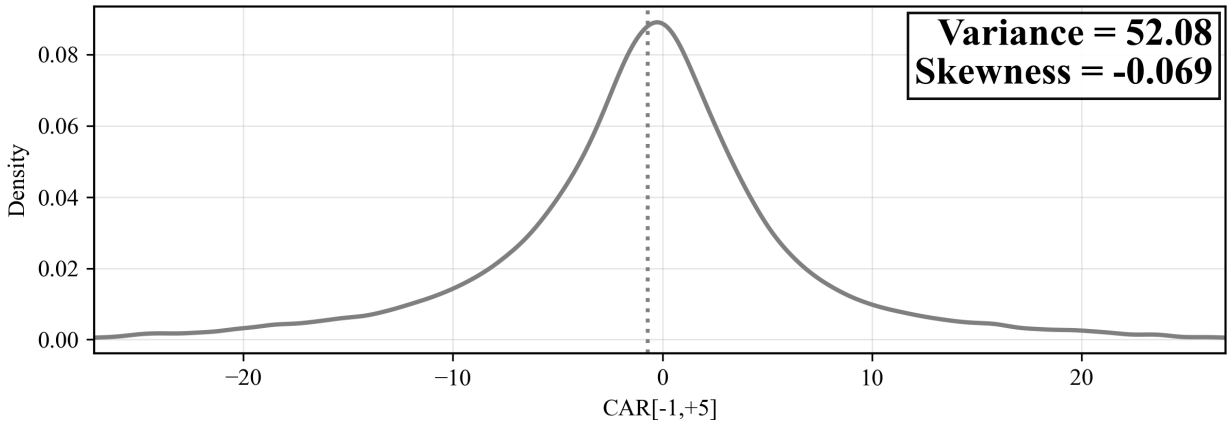
IA.3 AI Dictionary

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', 'con-

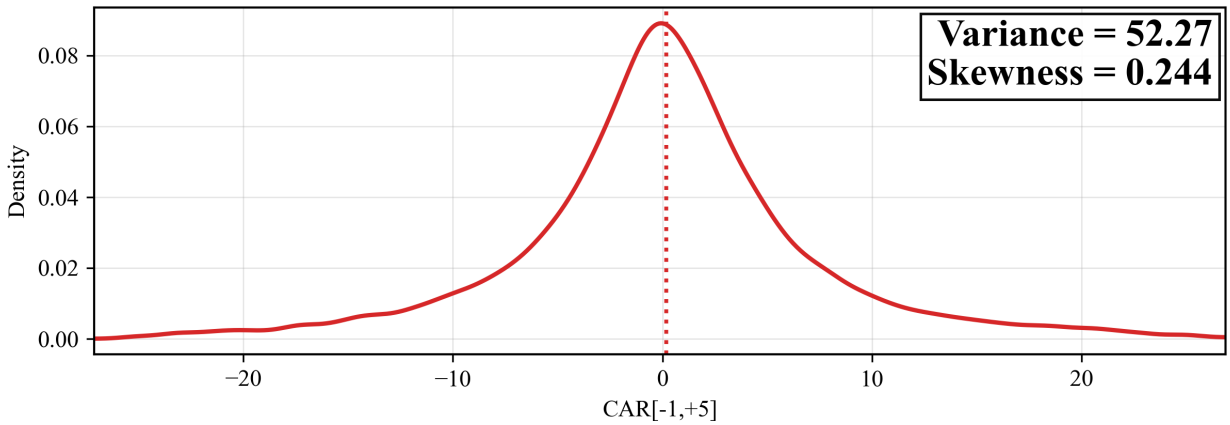
textualized embeddings', 'few-shot learning', 'zero-shot learning', 'generative adversarial network', 'self-supervised learning', 'mahout', 'opencv', 'libsvm', 'xgboost', 'word2vec', 'kernel methods', 'keras', 'support vector machines', 'latent semantic analysis', 'latent dirichlet allocation', 'knowledge 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

Panel A: Distribution of CAR[-1,+5] (Proprietary Models)

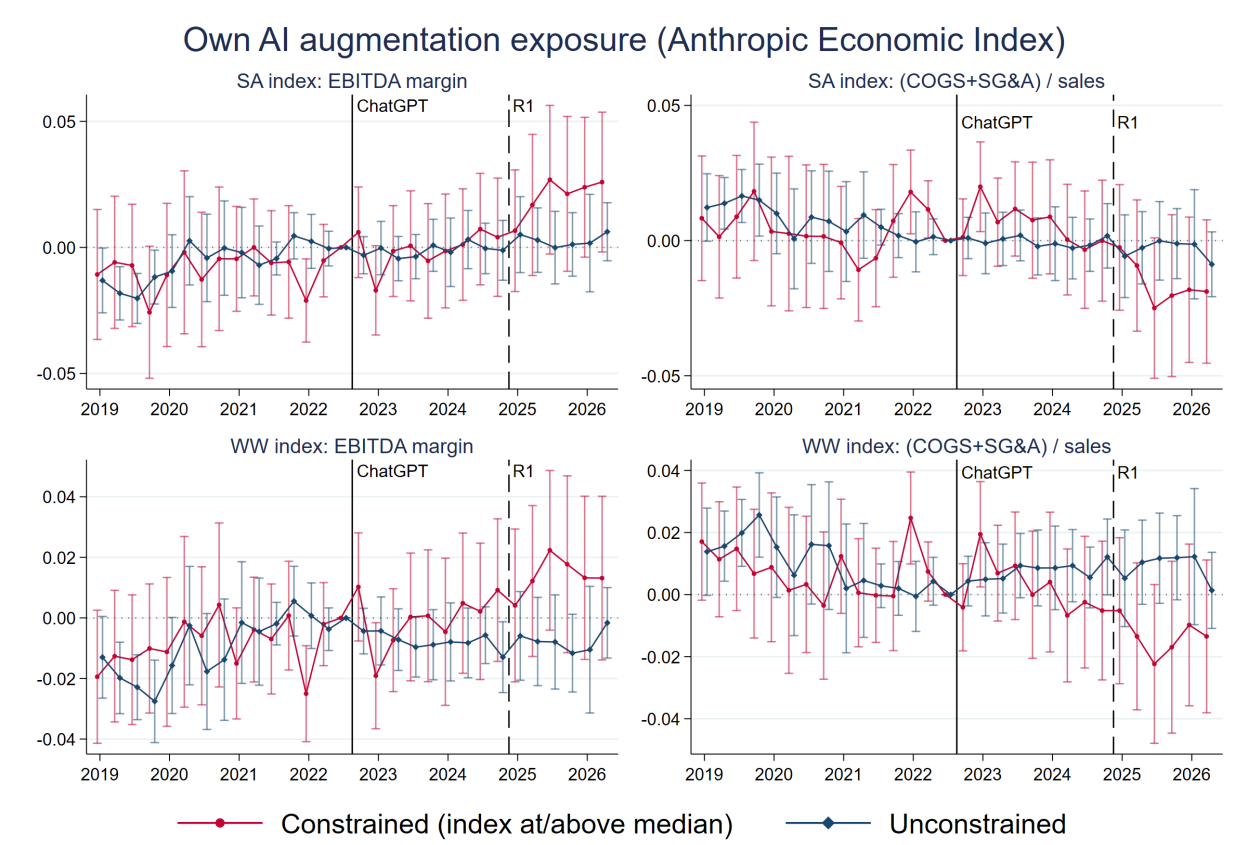


Panel B: Distribution of CAR[-1,+5] (Open-Source Models)



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%.

IB.2 AI Augmentation Usage and Operating Outcomes: Constrained vs. Unconstrained Firms



This figure plots the coefficients β_k from the following regression, estimated separately for financially constrained (red circles) and unconstrained (navy diamonds) firms:

$$Y_{i,t} = \alpha_i + \delta_{j,t} + \sum_{k \neq 2022Q3} \beta_k \times \mathbf{1}[t = k] \times Usage_i + \varepsilon_{i,t},$$

where $Y_{i,t}$ is the EBITDA margin (left) or operating costs (COGS plus SG&A) relative to sales (right), $\mathbf{1}[t = k]$ is an indicator for quarter k, and $Usage_i$ is the firm's augmentation-oriented AI usage from the Anthropic Economic Index, standardized to mean zero and unit standard deviation. α_i and $\delta_{j,t}$ denote firm and two-digit NAICS industry $\times$ quarter fixed effects. The reference quarter is 2022Q3. Firms are classified as constrained if their SA index (top row) or WW index (bottom row) in 2021 is at or above the sample median. Vertical lines mark the release of ChatGPT (GPT-3.5) in 2022Q4 (solid) and DeepSeek-R1 in 2025Q1 (dashed). Bars show 95% confidence intervals based on standard errors clustered at the firm level. The sample covers U.S. non-financial, non-utility firms in Compustat, including firms that subsequently delisted, from 2019Q1 to 2026Q2.

IC.1 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.2 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.3 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.4 Robustness: Spillover Effects Using FactSet Revere Data

This table presents how stock market reactions to proprietary and open-source AI model releases spill over through the supply chain. The sample includes only firms in industries with below-median industry-level AI exposure. Sales flows between suppliers and customer firms are identified based on FactSet Revere. 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. $AI\ Exposure_{jg}^{upstream}$ is the simple average AI exposure of the focal firms' suppliers. $AI\ Exposure_{jg}^{downstream}$ is the simple average AI exposure of the focal firms' customers. $AI\ Exposure_{jg}^{ind}$ 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 XI. 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	
	$CAR[-1,+5]_{ig}$			
$AI\ Exposure_{jg}^{upstream}$	-0.287*		0.541***	
	(0.168)		(0.189)	
$AI\ Exposure_{jg}^{downstream}$		0.161		0.342**
		(0.159)		(0.173)
$AI\ Exposure_{jg}^{ind}$	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