AI Product Displacement Risk [*†]

Brett Green[‡] Renping Li[§] Dimitris Papanikolaou[¶]

June 17, 2026

Preliminary draft.

Abstract

Generative AI can raise firm value by lowering labor costs but destroy it by displacing the firm's products. We construct AI Product Displacement Risk (AI-PDR), a firm-level importance-weighted measure of how digitally deliverable a firm's products are. At the January 2026 launch of Cowork Plugins, AI agents designed to replace specific incumbent SaaS products, high AI-PDR firms exhibit large negative CARs, with top-quintile firms underperforming bottom-quintile firms by approximately 10 percentage points over the $[-1, +10]$ trading-day window. The AI-PDR effect is larger for firms whose products overlap with the launched plugins. Both supply- and demand-side shields attenuate the repricing, with supply-side shields providing the larger attenuation. Analyst target prices decline significantly for high-AI-PDR firms while near-term earnings and revenue forecasts do not, consistent with long-horizon repricing. Similar effects do not appear at earlier AI model releases that did not target specific products.

[*]Results and interpretations are subject to revision.

[†]The authors have no sources of research support, financial relationships, or other potential conflicts of interest to disclose.

[‡]Olin Business School, Washington University in St. Louis, b.green@wustl.edu

[§]Freeman School of Business, Tulane University, rli17@tulane.edu

[¶]Kellogg School of Management, Northwestern University and NBER, d-papanikolaou@kellogg.northwestern.edu

1 Introduction

New technologies create value by displacing old ones. Since Schumpeter (1942), economists have recognized that creative destruction reallocates value across firms: incumbents whose products are rendered obsolete lose, while new entrants and adaptable firms gain (Pástor and Veronesi, 2009; Gârleanu et al., 2012; Kogan et al., 2017). Generative AI may drive this displacement at unprecedented scale because it produces digital output at near-zero marginal cost, creating close substitutes for many information goods and services.

A growing literature examines how AI displaces labor (Acemoglu et al., 2022; Eloundou et al., 2024; Hampole et al., 2025; Webb, 2020) and how AI adoption raises firm value (Babina et al., 2024; Eisfeldt et al., 2025), but this work studies which workers can be replaced or which firms adopt AI. None of it measures the risk that a firm's *product itself* faces direct substitution. The distinction matters for equity valuation: a firm whose workers are replaceable by AI may benefit through lower costs, but a firm whose *product* is replaceable faces revenue erosion regardless of how efficiently it operates internally. Labor displacement can be good news for shareholders, but product displacement is not. Which firms are most vulnerable to this product-level threat, and whether the equity market prices this risk, remains an open question.

We introduce *AI Product Displacement Risk* (AI-PDR) to measure this product-level threat. Using a large language model, we extract individual product and service lines from the business description (Item 1) of each firm's 10-K filing and score each product on its *Digital Deliverability*, the upper bound on how much of the product could in principle be produced and delivered in purely digital form. AI-PDR is the firm-level importance-weighted average of these product-level Digital Deliverability scores, capturing the substitutability frontier of the firm's product mix.

AI-PDR is conceptually distinct from the Eisfeldt et al. (2025) (ESZT) labor-exposure measure. Their index captures which firms' *workers* are replaceable by AI, while AI-PDR captures which firms' *products* are replaceable. The two correlate at $\rho = 0.38$. AI-PDR ranks Communication Services and Information Technology highest, and Materials, Energy, and Consumer Staples lowest, among GICS sectors. High-AI-PDR firms include SaaS vendors such as Salesforce, Adobe, and Workday, financial and legal data providers such as S&P Global, Thomson Reuters, and RELX, IT consulting firms such as Accenture, and online education firms such as Duolingo and Chegg, all in the top decile of AI-PDR.

We study equity market reactions to four major AI releases: ChatGPT (November 2022), DeepSeek R1 (January 2025), Claude Opus 4.5 (November 2025), and Cowork Plugins (January 2026). The first three released general-purpose models that demonstrated new capabil-

ities without targeting specific product markets. The fourth, Cowork Plugins, was launched by Anthropic on January 30, 2026 and differs in kind. It offered 11 open-source AI agent plugins covering productivity, sales, customer support, legal, finance, data analysis, and other enterprise functions, all accessible through a $20-per-month subscription that undercut incumbent per-seat SaaS fees by an order of magnitude. Within days, software and professional-services stocks shed approximately $285 billion in market capitalization globally (CNN Business, 2026), with firms whose core products overlapped with these agents experiencing the sharpest declines.

Our central finding is that the equity market reprices product displacement risk sharply at events that introduce targeted product substitution and not at general capability demonstrations. Two of the three capability-demonstration events (ChatGPT and Claude Opus 4.5) produce small, insignificant AI-PDR coefficients under the baseline GICS 6-digit industry fixed effects, and DeepSeek R1 produces a marginally significant positive coefficient, attributable to a sector-level AI tailwind. Cowork Plugins, by contrast, produces a large negative AI-PDR coefficient of -0.040 ($t = -6.68$), implying that top-quintile firms underperform bottom-quintile firms by approximately 10 percentage points (p.p.) over the $[-1, +10]$ event window. The contrast across event types is consistent with a mechanism in which capability demonstrations and targeted product-substitution events resolve different forms of uncertainty for incumbents. Capability demonstrations leave the path from raw technology to actual product competition uncertain. A targeted substitution event collapses that uncertainty by eroding identifiable incumbent revenue streams.

Within-event variation sharpens the interpretation. Splitting the sample by a Cowork Overlap score that measures how closely each firm's products compete with the launched plugins reveals that the AI-PDR effect is stronger among high-overlap firms. At the same time, the effect is not limited to SaaS firms. It holds in the zero-Cowork-Overlap subsample and in the subsample of firms whose primary products are not classified as SaaS. The effect does not vary with firm size, R&D intensity, or financial constraints, consistent with broad-based repricing.

We next investigate what attenuates the AI-PDR repricing at Cowork Plugins. We score each product on four competitive-shield dimensions, two supply-side (legal and regulatory requirements, production moats) and two demand-side (customer trust, customer embeddedness), and interact the composite Shield with AI-PDR. A one-standard-deviation (one-SD) higher AI-PDR is associated with about 3.7 p.p. lower CAR. A one-SD higher supply-side shield attenuates this by about 37 percent, while a one-SD higher demand-side shield attenuates it by about 22 percent. Both classes therefore attenuate the repricing, with supply-side shields delivering the larger attenuation.

The repricing is also highly persistent. Beyond the baseline $[-1, +10]$ window, we extend the window end out to $[-1, +90]$, 90 trading days after the launch, and the AI-PDR effect does not reverse at any horizon. The persistence is itself shielded. Splitting firms at the median Shield score, high-Shield firms drift back toward zero over the post-event horizon, while low-Shield firms show no sign of reversal. The other three AI releases show no comparable pattern.

The Cowork Plugins result is robust along several dimensions. We address two specific confounders by controlling for pre-event interest-rate sensitivity and by excluding firms that reported earnings within the event window. The outcome is robust to alternative event-window lengths and to Fama-French five-factor-adjusted abnormal returns. The specification is robust to alternative industry-FE granularities from no-FE through GICS 8-digit, to cumulative addition of firm-level controls, and to adding ESZT as a regressor. It is also robust to reconstruction using GPT-3.5, which has a training data cutoff of September 2021 that predates all of our events, alleviating concerns about look-ahead bias from the LLM scorer.

Analyst forecast revisions corroborate the event-study findings. After Cowork Plugins, consensus target prices for high-AI-PDR firms fall (coefficient $t = -2.67$), with a magnitude in line with the Cowork Plugins CAR effect, while near-term EPS and revenue forecasts are essentially unchanged. The pattern, valuations adjust but near-term earnings forecasts do not, is consistent with displacement risk operating on long-horizon cash flows rather than current revenue streams. The shield effect from the event study carries over to the analyst evidence. Supply-side shields moderate the target-price revisions more than demand-side shields, so markets and analysts both treat supply-side barriers as more durable protection.

As a complementary test, we construct an AI Input Intensity measure that captures whether scaling AI requires more units of a firm's product. The event-study window around Cowork Plugins is not the ideal test for this measure because the high demand for compute, memory, and data-center infrastructure had already been priced in by 2024 through multi-year compute contracts signed by hyperscalers. We therefore look at the long horizon, where value-weighted buy-and-hold returns from January 2022 to May 2026 are 279 percent for firms in the top half of the non-zero AI Input Intensity distribution, 87 percent for firms in the bottom half, and 35 percent for firms with no AI-input classification.

We contribute to three literatures. The first studies AI's displacement effects. That literature has focused almost entirely on labor. Acemoglu and Restrepo (2019) provide the theoretical framework distinguishing task displacement from task creation. Acemoglu and Restrepo (2020) and Acemoglu et al. (2022) document employment effects of robots and AI. Felten et al. (2021), Felten et al. (2023), Webb (2020), and Eloundou et al. (2024) construct occupation-level exposure indices, while Hampole et al. (2025) study firm-level employment

effects. At the macro level, Acemoglu (2025) and Aghion et al. (2019) model equilibrium growth effects of automation, and Bick et al. (2026) document the rapid pace of generative AI adoption. We add the product dimension, showing that the demand-side threat to revenue can dominate the supply-side benefit to costs in equity valuations.

The second literature examines AI and firm value. The closest antecedent is Eisfeldt et al. (2025), who construct a labor-exposure measure and show that firms with more AI-exposed workers earned positive abnormal returns around the ChatGPT launch, consistent with expected cost savings. Our measure captures the opposite cross-section. We identify firms whose products are *vulnerable to* AI rather than firms that benefit from adopting it. Babina et al. (2024) find that AI-investing firms experience higher sales growth. Eisfeldt and Schubert (2025) survey the broader AI-and-finance literature, emphasizing labor-exposure and adoption channels. Andrews and Farboodi (2025) use bond yields to show that market beliefs about AI update in discrete jumps around major releases but do not examine cross-sectional equity returns. We identify *which* firms are repriced and show that the pattern depends on the *type* of release.

Finally, we connect to the broader technological displacement literature. Kogan and Papanikolaou (2014) show that innovation reallocates value across firms, and Kakhbod et al. (2024) develop a text-based measure of creative destruction from patents. We identify the product channel of displacement in a setting where the threat is measurable and the market response is observable. Methodologically, our use of 10-K business descriptions to characterize products builds on Hoberg and Phillips (2016); we replace their bag-of-words similarity approach with structured LLM extraction and scoring along economically motivated dimensions of displacement vulnerability.

The remainder of the paper proceeds as follows. Section 2 describes the data and measure construction. Section 3 presents the event-study results. Section 4 examines analyst forecast revisions. Section 5 reports an AI Input Intensity test and an extension to 26 AI release events. Section 6 concludes.

2 Data and Measure Construction

2.1 AI-PDR Measure

We construct AI-PDR from business descriptions (Item 1) of 10-K annual filings retrieved from SEC EDGAR for fiscal years 2018–2026. The construction proceeds in two stages (product extraction and product scoring), both implemented with GPT-4o mini at temperature zero with structured JSON output schemas.

Step 1: Product Extraction. For each filing, we provide the Item 1 text to the LLM with the prompt reproduced in Listing 1.

Listing 1: Product extraction prompt

```
SYSTEM:
You are an information extraction engine. Output must be valid
JSON following the schema exactly. No markdown, no commentary.

USER:
Extract the distinct products and services this company sells,
from the 10-K Item 1 (Business Description) below.

Rules:
1. Each entry is a distinct product/service a customer would
   recognize as something the firm sells.
2. Merge aggressively. Combine variants of the same business
   line into one entry. Examples:
   - All staffing types -> "Staffing Services"
   - All loan types -> "Commercial Lending"
   - All annuity types -> "Annuity Products"
   Do NOT list internal capabilities unless sold externally.
3. Description: 1-2 sentences grounded in the text. Focus on
   HOW the product is produced and delivered (e.g., requires
   physical manufacturing? financial capital? human workers?
   digital only?), not just what it is called.
4. Importance: "high" if clearly a primary revenue driver;
   "medium" if meaningful but not dominant; "low" if ancillary.
5. Evidence: 1-3 short quotes (<=20 words each) from the text
   that support inclusion and importance. Must be near-exact
   substrings.
```

Each extraction returns a list of products with names, descriptions, importance ratings, and supporting evidence quotes. Requiring textual evidence forces the LLM to ground each product entry in specific filing language. The extraction step yields approximately 339,000 product entries from roughly 10,100 unique firms across all filing years.

Step 2: Product Scoring. We score each extracted product on Digital Deliverability. The product-level prompt asks the LLM whether the core value of the product could be delivered as a purely digital output. Responses are integer scores on a 0–10 scale with anchors at every two points, reproduced in Listing 2. The JSON schema orders the reasoning field before the score field, requiring a chain-of-thought justification before each numerical rating. Each batch of 15 products mixes products from different firms to prevent same-firm anchoring. The prompt is AI-agnostic by design, asking only about digital deliverability to limit look-ahead bias from the LLM scorer.

Listing 2: Digital Deliverability scoring prompt

```
Score each product/service below on its DIGITAL DELIVERABILITY.

DIGITAL DELIVERABILITY (0-10 integer):
   Could the core value of this product/service be delivered
   as a purely digital output (information, software, or
   digital content), without physical goods, in-person
   service, or transfer of financial capital?
   Each level below describes a category; examples in
   parentheses are illustrative, not exhaustive. Score by
   matching the CATEGORY, not the specific examples.
   0 = physical product (e.g., steel, crude oil, manufactured
       drugs)
   2 = primarily physical or capital-based with minor digital
       wrapper (e.g., medical device + companion app, bank
       loan with online application, insurance policy
       serviced via portal)
   4 = mixed delivery; significant physical and significant
       digital components, neither dominant (e.g., hardware
       bundled with required firmware/service, on-site
       installation with ongoing software access)
   6 = mostly digital output but requires real-time human
       expertise (e.g., custom IT consulting, brand strategy,
       legal advisory)
   8 = digital output, some human oversight (e.g., analytics
       reports, curated data feeds, managed SaaS)
   10 = entirely digital, mass-producible commodity content
       (e.g., stock photography, royalty-free audio, digital
```

```
        templates)
    Higher = EASIER to deliver digitally

IMPORTANT SCORING RULES:
1. Write reasoning BEFORE the score. Think step-by-step
    about the form in which the customer receives the value,
    then commit to a number.
2. Do NOT default to 5. There is no anchor at 5.
3. Use the FULL 0-10 range. Odd scores are valid when a
    product falls between anchors.
4. Differentiate products within the same batch.
5. Base your score on the product DESCRIPTION, not just the
    name.
```

We also score each product on four competitive-shield dimensions (legal and regulatory requirements, production moats, customer trust, and customer embeddedness) that capture market-structure frictions attenuating the displacement threat. Each shield dimension is elicited in a separate API call to prevent cross-dimension contamination. The full prompts for the four shield dimensions are reproduced in Listing A1.

Aggregation. Digital Deliverability is a product property, the substitutability frontier set by whether the output could in principle be produced and delivered in purely digital form. The four shield dimensions describe market-structure frictions that determine how much of that frontier converts into realized displacement. We treat AI-PDR as the primary threat measure and the Shield composite as a moderator.

We set the product-level AI-PDR score equal to the Digital Deliverability rating, $\text{AI-PDR}_p = \text{Digital}_p$. The firm-level measure is an importance-weighted average across products, where weights reflect the LLM's assessment of each product's revenue importance to the firm (high = 3, medium = 2, low = 1):

$$\text{AI-PDR}_{i,t} = \frac{\sum_{p=1}^{P_i} \omega_p \cdot \text{Digital}_p}{\sum_{p=1}^{P_i} \omega_p},$$

with ω_p the importance weight of product p. The four shield scores (Legal and Regulatory Requirements, Production Moats, Customer Trust, and Customer Embeddedness) are aggregated to the firm level by the same importance-weighted averaging rule. The firm-level Shield composite used in moderation tests is the simple average of the four shield sub-dimension

scores. We winsorize every firm-level score at the 1st and 99th percentiles within each fiscal year. For the baseline specification we use event-year-specific scores, taking AI-PDR$_{i,t}$ from each firm's most recent 10-K filing before each event (FY 2021 for ChatGPT, FY 2024 for DeepSeek R1 and Claude Opus 4.5, and FY 2025 for Cowork Plugins). Section 3.3 reports a robustness check using the firm-level average across FY 2021–2026.

Measure validation. Figure 1 plots the distribution of firm-level AI-PDR by GICS sector. Communication Services and Information Technology rank highest, while Materials, Energy, and Consumer Staples rank lowest. Sectors whose outputs are primarily digital score highest, while those producing physical goods score lowest. Within every sector there is substantial variation. Within Health Care, for example, firms range from near zero (drug manufacturers) to well above the sample median (health IT platforms), confirming that AI-PDR captures within-industry heterogeneity rather than sector membership.

Figure 2 displays word clouds of product descriptions for each scoring dimension. High-scoring products on digital deliverability feature "insights," "video," "organizations," "employee," and "data analytics." Low-scoring products feature "theatres," "golf," "outdoor," "movie," "natural gas," and "dining." Table A4 lists the ten highest- and lowest-scoring firms. High-AI-PDR firms include enterprise data and developer software (Palantir, Adobe, Atlassian, Snowflake), social platforms (Reddit, Roblox), cybersecurity and cloud infrastructure (Zscaler, Cloudflare, Datadog), and advertising technology (AppLovin). Low-AI-PDR firms cluster in metals and mining (Freeport-McMoRan, Newmont, Southern Copper), building materials (Martin Marietta, Amrize, Smurfit Westrock), packaged food and tobacco (Tyson, Kraft Heinz, Altria), and aerospace components (Howmet), all firms whose products require physical extraction, manufacturing, or processing.

Table A5 reports correlations between AI-PDR and observables. The correlation with ESZT is 0.38, positive but far from collinear (Figure A1), confirming that product displacement and labor exposure capture distinct dimensions of AI risk. AI-PDR correlates negatively with tangibility (-0.35) and is essentially uncorrelated with R&D intensity (-0.02). The raw correlation with firm size is small (0.08), and the within-sector relation is modest. The correlation with profit margins (0.04) is near zero, so AI-PDR does not proxy for profitability.

Because AI-PDR derives from product descriptions that change slowly, it should be stable over time. Figure A2 confirms this. Adjacent-year Pearson correlations range from 0.92 to 0.94, and non-adjacent years (e.g., 2021 versus 2026) have correlations above 0.86. This stability confirms that AI-PDR captures durable product characteristics rather than filing-year idiosyncrasies.

2.2 Other Data and Summary Statistics

Daily stock returns through 2025 are from the CRSP Stock Files. For 2026, we supplement CRSP with daily returns from Yahoo Finance, validated against CRSP on the overlapping month of December 2025. We merge firms with Compustat via the CRSP-Compustat Merged (CCM) link table. We obtain firm-level generative-AI labor exposure scores from Eisfeldt et al. (2025). Control variables follow the standard Fama–French timing convention: for each firm, we use Compustat data from fiscal year $t-1$ to construct log(ME), Tobin's Q, ROA, and asset growth, matched to returns from July of year t through June of year $t+1$. Table 1 reports summary statistics and Table A1 provides definitions for all variables used in the analysis.

3 Results

3.1 Events

We study equity market reactions to four major generative-AI releases, namely ChatGPT (November 30, 2022), DeepSeek R1 (January 20, 2025), Claude Opus 4.5 (November 24, 2025), and Cowork Plugins (January 30, 2026).[1] We classify each release based on whether it targets the replacement of specific categories of incumbent products or services. At the first three events, the launching lab released a general-purpose model as a capability demonstration. Cowork Plugins, by contrast, made packaged AI agents available for purchase that target specific incumbent digitally-deliverable products.

Because our tests rely heavily on the Cowork Plugins event, we document its institutional features in this subsection. Anthropic released Claude Cowork as a desktop-based agentic AI tool on January 12, 2026. On January 30, Anthropic added plugin support and open-sourced 11 role-specific plugins on GitHub. The plugins covered productivity, enterprise search, marketing, sales, finance, data analysis, legal, customer support, product management, and biology research (TechCrunch, 2026). Each plugin bundles domain knowledge, tool connectors, and sub-agents, enabling the agent to function as a department-specific workflow automation tool. For example, the legal plugin automates contract review, NDA triage, compliance workflows, and legal briefings. The finance plugin handles month-end close sequencing, journal entry preparation, account reconciliation, and SOX 404 control testing. The sales plugin performs prospect research, CRM automation, pipeline review, and competitive intelligence.

[1] Section 5.2 extends the analysis to 26 AI release events from late 2022 through January 2026.

The pricing structure differed sharply from incumbent SaaS products. A Claude Pro subscription providing access to all plugins costs $20 per user per month, while a Claude Team seat costs between $25 and $125 per user per month, compared with enterprise SaaS subscriptions that can run thousands of dollars per seat annually for comparable functionality. For example, Salesforce Enterprise Edition lists at $165 per user per month (Salesforce, 2026), and Adobe Creative Cloud for Teams runs $80 per user per month (Adobe, 2026).

Market reactions appeared within days. By February 3–4, approximately $285 billion in market capitalization had been erased from software stocks globally (CNN Business, 2026). The iShares Expanded Tech-Software Sector ETF (IGV) fell 5.7% on February 4, its largest single-day loss since April 2025 (Fast Company, 2026). SaaS leaders took large declines. Salesforce, ServiceNow, and Adobe each declined roughly 7% (Fast Company, 2026). The core concern was what analysts termed "seat compression." Once AI agents perform the operational tasks, the number of paid human users for a SaaS platform shrinks, allowing businesses to reduce their SaaS license counts and threatening the per-seat pricing model that underpins most enterprise SaaS valuations.

The repricing was not confined to SaaS. It extended to digitally-deliverable products outside the standard SaaS definition. Among financial and legal data providers, Thomson Reuters posted its largest single-day decline on record at roughly 16%, RELX fell about 14%, and LegalZoom dropped about 20% (Fast Company, 2026). S&P Global declined 11.9% cumulatively from January 30 to February 4. IT consulting firms fell sharply, with Accenture down 7.5% over the same window and Indian IT services firms Infosys, TCS, and Wipro down by comparable amounts (CNBC, 2026). Online education firms Duolingo and Chegg fell 16.2% and 8.9%, respectively, over the same window. Freelancer marketplaces that intermediate the same digitally-deliverable knowledge work the plugins automate also fell, with Upwork down 18.9% and Fiverr down 7.2% (Habib-Valdhorn, 2026).

Table A6 documents the timeline of Anthropic product announcements surrounding our event window, which establishes that January 30, 2026 is the appropriate event date. The Cowork desktop agent previewed on January 12 was a general-purpose tool without domain-specific functionality. The Opus 4.5 release on November 24 was a model-capability upgrade. The Opus 4.6 release on February 5 was a further model upgrade that postdates the event window. The January 30 plugin launch is qualitatively different: it is the first date on which Anthropic released packaged, role-specific workflow automation tools that directly compete with identifiable categories of incumbent digitally deliverable products. Consistent with this dating, commentators attributed the February 3–4 selloff to the plugin launch rather than to the earlier Cowork preview or subsequent model releases (Fast Company, 2026; CNBC, 2026).

Figure 3 reports quintile-sorted CARs over the [−1, +10] window. For Cowork Plugins, CARs decline monotonically from Q1 to Q5, and the spread is an order of magnitude larger than at any of the three capability releases.

3.2 Cross-Sectional Regressions

3.2.1 Event Study Design

For each event, we compute CAPM-adjusted cumulative abnormal returns (CARs) over a $[-1, +10]$ trading-day window:

$$\mathrm{AR}_{i,t} = (r_{i,t} - r_{f,t}) - \hat{\alpha}_i - \hat{\beta}_i \, (r_{m,t} - r_{f,t}) \,, \qquad \mathrm{CAR}_i = \sum_{t=-1}^{10} \mathrm{AR}_{i,t},$$

where $r_{i,t}$ is firm i's return on day t relative to the event, $r_{f,t}$ is the daily risk-free rate, $r_{m,t}$ is the CRSP value-weighted market return, and $(\hat{\alpha}_i, \hat{\beta}_i)$ are firm i's pre-event intercept and market beta from OLS of $(r_{i,t} - r_{f,t})$ on $(r_{m,t} - r_{f,t})$ over $[-252, -11]$ trading days. We then estimate cross-sectional regressions:

$$\mathrm{CAR}_i = \delta_j + \beta_1 \cdot \text{AI-PDR}_i + \gamma' \mathbf{X}_i + \varepsilon_i, \tag{1}$$

where δ_j are GICS 6-digit industry fixed effects (with industries having fewer than five firms dropped) and $\mathbf{X}_i$ includes log(ME), Tobin's Q, ROA, and asset growth. All right-hand-side variables are winsorized at the 1st and 99th percentiles within each event sample. The estimation universe is firms with market capitalization above $100 million and Compustat total assets above $100 million. We report t-statistics from standard errors clustered at the GICS 6-digit industry level in parentheses throughout. Table 3 reports coefficient sensitivity across alternative FE granularities (GICS 2-digit through 8-digit). The appendix reports a complementary specification (Table A16) that includes ESZT as a regressor.

3.2.2 Baseline Results

Table 2 reports the baseline cross-sectional regression of CAPM-adjusted $[-1, +10]$ CARs on AI-PDR, separately for each of the four events. The specification includes GICS 6-digit industry fixed effects and controls for log(ME), Tobin's Q, ROA, and asset growth, all winsorized at the 1st and 99th percentiles. AI-PDR is standardized to mean zero and unit variance within each estimation sample. At Cowork Plugins, the AI-PDR coefficient is -0.0395 ($t = -6.68$), significant at the 1% level. The three capability-demonstration releases are statistically null or marginal under the same specification, ChatGPT Launch

yields 0.0067 ($t = 1.33$), DeepSeek R1 yields 0.0096 ($t = 1.78$), and Claude Opus 4.5 yields -0.0052 ($t = -1.24$). Because the standardized AI-PDR spread from quintile 1 to quintile 5 is about 2.6 sample standard deviations, the Cowork coefficient implies a Q5–Q1 CAR spread of roughly 10 p.p. over $[-1, +10]$ within narrowly defined industries.

Table 3 shows that this contrast survives across five alternative industry fixed-effect choices, from no industry FE through GICS 8-digit sub-industry. The Cowork Plugins coefficient remains negative and significant at the 1% level in every column, ranging from -0.0482 ($t = -8.72$) with no industry FE through -0.0335 ($t = -5.52$) at the GICS 8-digit sub-industry level. In contrast, the ChatGPT coefficient is marginal at GICS 2-digit (0.0071, $t = 1.82$) and 4-digit (0.0070, $t = 1.72$), and insignificant once industry FE is tightened to GICS 6-digit or finer. The DeepSeek R1 coefficient is significant at no-FE through GICS 4-digit (range 0.011–0.017) and only marginal at GICS 6-digit (0.0096, $t = 1.78$). These positive coefficients are consistent with a broad AI tailwind at the two releases moving high-AI-PDR sectors such as Communication Services and Information Technology, although their magnitude is small relative to the negative effect at the Cowork Plugins release. The Claude Opus 4.5 coefficient is significant at the GICS 2-digit (-0.0077, $t = -2.23$) and 4-digit (-0.0077, $t = -2.36$) levels but null once industry FE is tightened. One reading is that by then the market had partially incorporated AI product displacement risk, since Claude Opus 4.5's improved performance had made automated agent workflows a more credible threat.

Table 4 reports the AI-PDR coefficient under GICS 6-digit fixed effects across four event windows using CAPM-adjusted CARs. For Cowork Plugins, the coefficient is negative and significant at every horizon, growing from -0.0174 ($t = -4.72$) at $[-1, +3]$ to -0.0395 ($t = -6.68$) at $[-1, +10]$, moderating to -0.0316 ($t = -4.49$) at $[-1, +20]$, and stabilizing at -0.0204 ($t = -2.57$) at $[-1, +30]$. The 30-day persistence without reversal is consistent with permanent repricing rather than transient overreaction.

Figure 4 extends this scan out to 90 trading days after the Cowork Plugins event, plotting the AI-PDR coefficient from the cross-sectional regression of the $[-1, +d]$ CAR on AI-PDR for window ends d from 0 through 90 trading days. The displacement effect is highly persistent. The coefficient remains negative at every horizon and shows no tendency to revert toward zero by day +90. Figure 5 repeats the scan for the three capability releases and finds no comparable pattern, the one exception being Claude Opus 4.5, whose post-event window runs into the Cowork Plugins event at +45 trading days and therefore absorbs the same repricing at longer horizons. Overall, the effect is highly persistent and does not reflect transient market sentiment.

The divergence reflects how the two event types resolve different forms of uncertainty. General-purpose capability releases resolve uncertainty about what AI *can* do, but the path

from demonstrated capability to actual product competition involves uncertain intermediate steps such as packaging, distribution, overcoming switching costs, and satisfying enterprise procurement requirements. Each step attenuates the present-value threat to incumbents, who retain the option to adapt. A targeted substitution event collapses this chain of uncertainty. Markets observe direct entry into identifiable incumbent revenue streams. The repricing reflects the removal of this uncertainty discount. This interpretation extends Andrews and Farboodi (2025), who find that aggregate market beliefs about AI update in discrete jumps around releases; our evidence shows that the *cross-sectional* displacement effect depends on *what type* of release occurs.

3.2.3 Moderation by Competitive Shields

Table 5 extends the specification by adding a *Shield* composite, the simple average of four firm-level defense scores (legal and regulatory requirements, production moats, customer trust, and customer embeddedness), together with its interaction with AI-PDR. Figure A3 reports word clouds of product descriptions for the top and bottom deciles of each shield sub-dimension, confirming that the four constructs map cleanly to the underlying 10-K language. For example, products in the top decile of the legal and regulatory shield feature words such as "broadband," "satellite," "aircraft," "nuclear," and "defense," categories in which licensing, certification, or regulatory approval gates entry, while the bottom decile features "apparel," "furniture," "footwear," "golf," and "printing," consumer goods with minimal regulatory friction.

For Cowork Plugins at the baseline $[-1, +10]$ window, $\hat{\beta}_{\text{AI-PDR}\times\text{Shield}} = 0.0118$ ($t = 2.85$). This implies that while a one-SD higher AI-PDR is associated with about 3.7 p.p. lower CAR, a one-SD higher composite Shield attenuates the effect by about 32 percent. Panel B of Table 5 decomposes the four sub-dimensions into supply-side defenses and demand-side defenses. The former comprises legal and regulatory requirements and production moats, which capture barriers a new entrant would face producing and selling the same output, while the latter comprises customer trust and customer embeddedness, which capture frictions a customer would face switching to an alternative. The supply-side interaction $\hat{\beta}_{\text{AI-PDR}\times\text{Supply}}$ is 0.0136 ($t = 3.45$) at the $[-1, +10]$ window. The demand-side interaction $\hat{\beta}_{\text{AI-PDR}\times\text{Demand}}$ is 0.0082 ($t = 2.23$) at the same window. Both supply-side and demand-side defenses therefore moderate the displacement effect, with the supply-side coefficient roughly 66% larger in magnitude than the demand-side coefficient. Entering each sub-dimension separately at the $[-1, +10]$ window (Table A13) sharpens the picture. The supply-side moderation is concentrated in production moats (0.0143, $t = 3.06$), while the legal and regulatory dimension

alone is insignificant (0.0078, $t = 1.53$). On the demand side, customer trust is significant at the 5% level (0.0091, $t = 2.17$) and customer embeddedness is marginal (0.0055, $t = 1.84$).

The pattern is consistent with a mechanism in which entrant-side production assets and customer-side behavior respond differently to AI capability advances. Production moats such as proprietary data, fabrication assets, and network effects govern whether an AI entrant can produce the substitute at all, so a high-AI-PDR firm sitting behind such a barrier is partially protected even after AI capability advances. Customer trust and customer embeddedness rest on customer behavior that is itself responsive to the substitute. Trust thresholds fall as AI-produced output accumulates a track record, and switching costs erode as AI-native alternatives mature their interfaces and integrations. Markets price both classes of defense, with greater weight on the production-side barriers that are less responsive to the entrant's own progress.

The main effect of Shield is, by contrast, theoretically ambiguous. With AI-PDR and Shield each standardized and demeaned, $\hat{\beta}_{\text{Shield}}$ is the marginal effect of Shield evaluated at the sample mean of AI-PDR, which equals the average marginal effect of Shield across the AI-PDR distribution. Theory predicts that shields should protect firms only where the product channel is active. A non-digital firm has nothing for a shield to protect against, so the slope at low AI-PDR is approximately zero, while the slope at high AI-PDR is positive. The reported $\hat{\beta}_{\text{Shield}}$ is an average of these two pieces, mixed with any non-AI channels correlated with Shield. A positive, null, or modestly negative $\hat{\beta}_{\text{Shield}}$ is therefore consistent with the theory. Only the sign of the interaction is sharply predicted. In Panel A at Cowork Plugins, $\hat{\beta}_{\text{Shield}} = -0.0099$ ($t = -1.49$), modestly negative and statistically indistinguishable from zero.

The window scan of Figure 4, which plots the AI-PDR coefficient at window ends out to 90 trading days after the Cowork Plugins event, shows that this moderation strengthens at longer horizons. Splitting firms at the median Shield score, the AI-PDR slope for high-Shield firms drifts back toward zero over the post-event horizon, while the slope for low-Shield firms stays large and negative throughout, showing no sign of reversal.

3.2.4 Effect Heterogeneity

Because the main result relies on a single product-substitution event, we exploit within-event variation to discipline the interpretation. We construct a Cowork Overlap score measuring how closely each firm's products compete with the Cowork Plugins launched on January 30, 2026. The plugin definitions are from the Anthropic launch announcement.[2] An LLM scores

[2] `https://github.com/anthropics/knowledge-work-plugins/blob/main/README.md`, version committed January 30, 2026.

each extracted product on a 0–10 scale against each of the 10 plugin bundles independently (Listing A2). The product-level Cowork Overlap is the maximum across the 10 plugin scores, and the firm-level Cowork Overlap is the importance-weighted average across products. The firm-level distribution is right-skewed. About 78% of firms in the Cowork Plugins regression sample have zero overlap. Cowork Overlap correlates with AI-PDR at a Pearson $\rho = 0.57$.

Table 6 reports the AI-PDR slope estimated separately on two Cowork Overlap subsamples, zero-overlap firms (about 78% of the sample) and positive-overlap firms, across five industry fixed-effect granularities. The AI-PDR slope is negative and significant in every cell, and is consistently larger in absolute value for positive-overlap firms than for zero-overlap firms. At the baseline GICS 6-digit FE, the slope is -0.0266 ($t = -4.97$) on the zero-overlap subsample and -0.0358 ($t = -3.46$) on the positive-overlap subsample, with the gap widening further at coarser FE granularities. That positive-overlap firms bear a substantially larger AI-PDR penalty than zero-overlap firms is the cleanest evidence that the cross-sectional return effect operates through the product-displacement channel.

Table A12 interacts AI-PDR with five firm characteristics, namely size, R&D intensity, and three financial-constraint indices (Hadlock and Pierce, 2010; Whited and Wu, 2006; Kaplan and Zingales, 1997). All interaction terms are insignificant. The displacement effect is broad-based rather than concentrated among small, R&D-intensive, or financially constrained firms.

3.2.5 Is the Effect Limited to SaaS?

The displacement effect extends beyond firms GICS classifies as technology or communications. Table A7 lists the highest-AI-PDR firms after excluding GICS sectors 45 (Information Technology) and 50 (Communication Services). The remaining high-PDR firms span consumer discretionary, financials, industrials, and health care. GICS classifies firms by primary revenue source rather than product-delivery method, so firms outside the technology sector can still sell digitally-delivered products. Table A8 re-estimates the Cowork Plugins regression on the restricted sample. At the GICS 6-digit FE baseline, the AI-PDR coefficient is -0.0381 ($t = -6.83$) excluding Information Technology and -0.0363 ($t = -6.90$) excluding both Information Technology and Communication Services, confirming that the displacement effect originates broadly across sectors rather than from technology firms alone.

Even after excluding the two technology sectors, many of the surviving high-AI-PDR firms still sell software-as-a-service products, such as Veeva Systems and Doximity in Health Care, or Verisk Analytics and Broadridge in Industrials. To test whether the effect is limited to SaaS firms, we ask an LLM to flag each firm as SaaS or not. A firm is SaaS-classified

if any of its top-importance product lines in the most recent year qualifies as Software-as-a-Service under a cloud-delivery, multi-tenant, recurring-revenue definition. Table A9 illustrates the classification with random examples of SaaS and non-SaaS firms, together with the five non-SaaS firms holding the highest AI-PDR (Shutterstock, Getty Images, Exodus Movement, Circle Internet Group, and MongoDB). Figure A4 reports word clouds of product descriptions for the two groups, confirming that the classification is reflected in the underlying 10-K language. Table A10 re-estimates the Cowork Plugins regression separately on SaaS-classified and non-SaaS subsamples. The non-SaaS subsample carries an AI-PDR coefficient of -0.0278 ($t = -5.10$) at the GICS 6-digit FE, significant at the 1% level, showing that the displacement effect extends beyond SaaS revenue models to other digitally-delivered products.

3.3 Robustness

We subject the main finding to robustness tests organized into three groups, namely identification, measure construction, and regression specification. All robustness tables appear in the Internet Appendix.

3.3.1 Identification

Interest-rate sensitivity. Table A14 adds each firm's interest-rate sensitivity, estimated as the coefficient on Δy_{10} (the daily change in the 10-year Treasury yield) from a two-factor pre-event regression over $[-252, -11]$ trading days. For Cowork Plugins at GICS 6-digit FE, the AI-PDR coefficient is virtually unchanged once rate sensitivity is added as a control (-0.0392, $t = -6.62$ versus the -0.0395, $t = -6.68$ baseline), indicating that duration exposure is not the driver.

Coincident earnings announcements. The Cowork Plugins event window overlaps with the late-January earnings season. Table A11 excludes firms whose Compustat earnings-announcement date falls within the $[-1, +10]$ window, reducing N from 2,803 to 1,906 at the GICS 6-digit baseline. The AI-PDR coefficient remains negative and significant at the 1% level, at -0.0382 ($t = -5.29$) versus the -0.0395 ($t = -6.68$) baseline.

3.3.2 Measure Construction

A second set of tests verifies that the results are driven by underlying product characteristics rather than the specific scoring choices or filing year used to construct AI-PDR.

Alternative AI-PDR specifications. The baseline uses event-year-specific AI-PDR (the most recent filing year before each event). As a robustness check, we replace it with the within-firm mean of AI-PDR across the firm's fiscal-year 2021–2026 10-K filings, applying the same averaged measure at all four events. Table A15 reports the results. For Cowork Plugins at GICS 6-digit FE, the AI-PDR coefficient is -0.0469 ($t = -6.77$), slightly larger in magnitude than the -0.0395 baseline, indicating that AI-PDR captures a stable firm-level product characteristic rather than a one-year aberration.

Alternative scorer validation. To verify that the AI-PDR scores are not artifacts of the specific scoring model, we rescore a random 1,000-product sample drawn from the full product universe with a fixed seed, using GPT-4.1 (a more capable LLM than the GPT-4o-mini baseline), and compare the scores at the product level. Table A2 reports the cross-scorer agreement across the five LLM-scored dimensions. The two scorers agree at Spearman $\rho = 0.86$ on Digital Deliverability and ρ between 0.63 and 0.85 on the four shield sub-dimensions, with 42 to 68 percent of products receiving integer scores within ± 1 on the 0–10 scale. The high cross-scorer agreement indicates that AI-PDR captures stable product characteristics rather than scorer-specific noise.

Chronologically consistent scorer. A separate concern is that the scoring LLM's training data may encode which products AI will displace, so that digital deliverability ratings reflect learned associations rather than inherent product characteristics. He et al. (2025a,b) formalize this concern and develop chronologically consistent LLMs with strict knowledge cutoffs. Their ChronoGPT models at 1.5 billion parameters lack the instruction-following capacity needed for our scoring rubric. We therefore re-score all products using GPT-3.5-Turbo-0125, whose training data has a September 2021 cutoff that predates the ChatGPT launch (November 2022), our first event. Table A20 reports the results. The all-years GPT-3.5 AI-PDR correlates with the baseline at a Pearson $\rho = 0.94$, indicating broad agreement on which products are digitally deliverable despite a three-year gap in training data. For Cowork Plugins at GICS 6-digit FE, the AI-PDR coefficient is -0.0321 ($t = -5.45$), replicating the main finding. The consistency across training cutoffs confirms that AI-PDR captures durable product characteristics rather than ex-post knowledge of AI displacement events.

3.3.3 Specification and Controls

Adding the ESZT control. Table A16 adds ESZT as a control in a regression weighted by market capitalization at $t-1$ with NAICS 3-digit industry fixed effects, following the

specification of Eisfeldt et al. (2025). At ChatGPT, the ESZT coefficient is positive and highly significant (0.0168, $t = 2.99$), consistent with their original finding that labor-exposed firms outperformed on news of generative AI. At DeepSeek R1 and Claude Opus 4.5, ESZT is statistically indistinguishable from zero or marginally negative. For Cowork Plugins, the AI-PDR coefficient is -0.0328 ($t = -2.69$) and the ESZT coefficient is -0.0146 ($t = -1.83$). AI-PDR and ESZT therefore capture distinct cross-sections at every event, and the central AI-PDR result survives controlling for labor exposure.

Fama–French five-factor adjustment. Table A17 replaces CAPM-adjusted CARs with Fama–French five-factor adjusted CARs. For Cowork Plugins at GICS 6-digit FE, the AI-PDR coefficient is -0.0330 ($t = -5.71$). The coefficient is attenuated relative to the CAPM baseline because the five-factor model absorbs some cross-sectional return variation correlated with AI-PDR, but the effect persists across event windows. Table A18 reports the FF5-adjusted Cowork Plugins coefficient at -0.0103 ($t = -3.27$) at $[-1, +3]$, -0.0287 ($t = -3.45$) at $[-1, +20]$, and -0.0174 ($t = -2.01$) at $[-1, +30]$, confirming persistent repricing.

Additional controls. Table A19 cumulatively adds firm characteristics, namely total assets, tangibility, intangibles ratio, profit margin, leverage and cash, momentum and volatility, firm age, and R&D intensity. The AI-PDR coefficient ranges from -0.0343 to -0.0395 across the broad-sample specifications ($N \geq 2{,}258$) and remains highly significant (t between -5.83 and -6.68) throughout. Adding R&D intensity last reduces the sample to 1,217 firms, where the point estimate grows in magnitude to -0.0441 ($t = -4.39$). The stability across specifications indicates that AI-PDR does not proxy for standard firm characteristics.

4 Analyst Forecast Revisions

The event study shows that the equity market reprices high-AI-PDR firms at Cowork Plugins. We next ask whether sell-side analyst forecasts move with AI-PDR in a way consistent with the displacement story. Under that story, the repricing should appear in long-horizon valuations such as target prices, which embed expectations of cash flows over the indefinite future. It should not appear in near-term earnings and revenue forecasts, which are anchored on existing contracts, backlog, and order pipelines that have not yet been disrupted. We test this with a cross-sectional regression of scaled forecast revisions on standardized AI-PDR.

$$y_i = \alpha_{j(i)} + \beta \, \text{AI-PDR}_i + \theta' X_i + \varepsilon_i,$$

where α_j are industry fixed effects, X_i are the firm controls used in the event study, and y_i is the scaled revision (post$-$pre)/pre for EPS, revenue, and target price. For long-term growth, y_i is the level change post $-$ pre since it is already in p.p. Table 7 reports both GICS 2-digit sector and GICS 6-digit industry fixed effects side by side. The pre-event and post-event observations are both IBES mean consensus forecasts, for December 2025 and March 2026 respectively. We restrict to firms with December fiscal year-ends and match forecasts to the same target fiscal year (2026, ending December 31, 2026). We drop EPS observations with pre-event EPS below \$0.50 (which includes all negative-EPS firms) to avoid sign ambiguity from scaling by a near-zero or negative denominator. Table A3 reports summary statistics for the analyst forecast sample.

Panel A of Table 7 reports results for four forecast outcomes, namely consensus mean EPS, revenue, long-term growth (LTG), and target price. The AI-PDR coefficient is essentially zero for EPS ($\hat{\beta} = 0.010$, $t = 1.25$ at GICS 6-digit FE) and, if anything, marginally positive for revenue ($\hat{\beta} = 0.009$, $t = 2.08$), the opposite sign from displacement, so analysts did not revise near-term earnings or revenue downward for high-AI-PDR firms. The LTG coefficient is statistically insignificant ($\hat{\beta} = 0.455$, $t = 0.53$) on a much smaller sample.[3] Target prices show a robust negative revision, with $\hat{\beta} = -0.026$ ($t = -2.67$) at GICS 6-digit FE and $\hat{\beta} = -0.060$ ($t = -6.01$) at GICS 2-digit FE. A one-SD higher AI-PDR is associated with roughly a 3 to 6 percent larger downward revision in the target price. The magnitude lines up with the event-study CAR of about -4 p.p. per one-SD AI-PDR at Cowork Plugins (Table 2).

The contrast between target prices and near-term earnings is the central result of this section. The same sell-side analysts produce both forecasts using the same information set in the same revision window, yet only the long-horizon valuation moves with AI-PDR. The pattern is consistent with displacement risk operating on cash flows over a long horizon. Signed contracts and customer switching costs anchor near-term EPS and revenue, while target prices reflect expected longer-horizon profitability that displacement risk erodes.

Panel B tests whether Shields moderate the target-price revision using the same interaction structure as Table 5. At GICS 2-digit FE, the composite-Shield specification yields $\hat{\beta}_{\text{AI-PDR}} = -0.047$ ($t = -4.52$) and $\hat{\beta}_{\text{AI-PDR}\times\text{Shield}} = 0.019$ ($t = 3.24$). When the composite is split into supply-side and demand-side components, the supply-side interaction is the larger of the two and the only one that is statistically significant, with $\hat{\beta}_{\text{AI-PDR}\times\text{Supply}} = 0.023$ ($t = 4.56$) versus $\hat{\beta}_{\text{AI-PDR}\times\text{Demand}} = 0.008$ ($t = 1.34$). The supply-side coefficient is roughly

[3]LTG coverage is sparse. Only about 30% of firms in the IBES sample have a long-term growth estimate. Analysts produce LTG less consistently than EPS or revenue, and the forecasts are unanchored to a precise horizon and revised infrequently.

three times as large as the demand-side coefficient, mirroring the event-study finding (Table 5, Panel B), so markets and sell-side analysts agree on which class of shield provides more durable protection against AI substitution.

5 Additional Tests

5.1 AI Input Intensity

AI-PDR captures the threat that AI substitutes for a firm's output. By contrast, scaling AI training and inference requires more units of some firms' products. This input channel is conceptually distinct from AI-PDR. We construct AI Input Intensity using a taxonomy classifier (Listing A3) that assigns each extracted product to one of 23 AI-input categories grouped into four tiers (AI core, AI-adjacent infrastructure, semiconductor manufacturing, and general inputs), or to NONE if no category fits. Firm-level AI Input Intensity is the importance-weighted fraction of the firm's products classified into any T1–T4 category.

Table A21 verifies face validity. We show the top three firms by market capitalization within each category, which include both the major AI suppliers (e.g., NVIDIA, AMD, Broadcom, Alphabet, Microsoft) and specialized names in each specific category (e.g., CoreWeave for GPU-dedicated cloud, IonQ for quantum-as-a-service compute, Vertiv for AI-specific cooling, Digital Realty for data-center colocation). Application-layer firms that brand themselves as "AI-powered" but do not supply physical compute capacity are classified as `NONE`, consistent with the construct capturing upstream inputs to AI rather than downstream AI consumption.

Table A22 adds AI Input Intensity to the event-study specification across all four events and four event windows, with two FE levels (GICS 6-digit and GICS 2-digit) and two estimators (unweighted OLS and market-cap-weighted WLS). The picture is consistent across the four combinations within each event. At the two capability-demo events, ChatGPT Launch and Claude Opus 4.5, the AI Input Intensity coefficient is small and statistically insignificant.

DeepSeek R1 is materially different. At $[-1, +10]$, AI Input Intensity is significantly negative across both FE levels under OLS. Moving from 0 to 100 percent share of products classified into AI-input categories is associated with roughly a 9 p.p. lower $[-1, +10]$ CAR. Under WLS the GICS-2 estimate stays significant while the GICS-6 estimate weakens. AI Input Intensity is significantly associated with lower CARs at the $[-1, +30]$ horizon across all four specifications. The negative coefficient is consistent with an AI-cost-deflation narrative. More efficient reasoning models reduce the units of compute, memory, networking,

and physical infrastructure required to deliver a unit of AI output, eroding pricing power for incumbent suppliers. This mirrors "DeepSeek Monday" (January 27, 2025), on which NVIDIA's stock fell roughly 17 percent and lost about $600 billion in market capitalization, the largest single-day market-capitalization loss for any U.S.-listed stock on record.

Cowork Plugins shows a different pattern. At $[-1, +3]$ the AI Input Intensity coefficient is negative and significant, and turns positive and significant by $[-1, +30]$. Moving from 0 to 100 percent share of products classified into AI-input categories is associated with roughly a 6 p.p. lower CAR over the $[-1, +3]$ horizon and an 11 p.p. higher CAR over the $[-1, +30]$ horizon. One explanation is that, before Cowork Plugins, the market was pricing in a scenario where high-AI-PDR firms would each train their own AI models in-house, sustaining strong AI-infrastructure demand. The Cowork Plugins release prompted a revision toward a winner-take-all situation in which only a few AI leaders train models, implying lower per-firm AI-input demand and a negative repricing of AI-infrastructure suppliers. By day $[-1, +30]$ the market appears to recognize that even with a small number of frontier model trainers, agentic workloads imply higher inference demand than the pre-event consensus, leaving the AI-infrastructure cross-section with a positive abnormal return.

The event-study design works for AI-PDR because the Cowork Plugins event resolves uncertainty about whether AI can displace digitally-deliverable products, and AI-PDR is the cross-section that should reprice when this uncertainty resolves. For AI Input Intensity the analogous uncertainty had largely been resolved by 2024, when the market had absorbed that scaling AI requires large quantities of compute infrastructure, and the hyperscalers had publicly committed tens of billions in multi-year compute contracts. A short event window therefore picks up little incremental information about the AI Input Intensity cross-section.

Figure A5 reports value-weighted buy-and-hold cumulative returns for three portfolios sorted on AI Input Intensity, over January 3, 2022 through May 21, 2026. Firms in the top half of the non-zero AI Input Intensity distribution returned 278.6% over the period, firms in the bottom half returned 86.9%, and firms with no AI-input classification returned 34.9%. The monotonic ordering and the wide spread between the high and rest portfolios indicate that the AI Input Intensity measure captures a large and persistent cross-sectional return spread.

5.2 Other AI Releases

To test the AI-PDR result outside the four events in our headline event study, we re-estimate the Table 2 cross-sectional regression at 26 AI release events from late 2022 through January 2026. The 26-event set covers every flagship model release from OpenAI, Google, Anthropic,

xAI, and DeepSeek over this period, plus four releases that explicitly target product substitution (ChatGPT Plugins, March 2023; Apple Intelligence, June 2024; Claude 3.5 Computer Use, October 2024; and OpenAI Operator, January 2025). We classify each event into Category A (capability demonstration) or Category B (targeted product substitution) using the launching lab's official blog post and three contemporaneous news articles per event.

Figure A6 plots the AI-PDR coefficient with 95% confidence intervals at every event, across three industry-FE granularities. The pattern matches the headline finding. Cowork Plugins, the rightmost event, is the only release for which the AI-PDR coefficient is large, negative, and significant across all three FE levels. Other events produce coefficients that are small in magnitude and inconsistent across FE granularities. The four pre-Cowork Category B events also fail to produce a Cowork-style response. Two factors plausibly explain the difference. First, those releases targeted consumer-oriented workflows such as calendar management and flight booking rather than the enterprise software products that constitute most of the listed-equity universe. Second, the level of tool use and reliability needed to credibly substitute for incumbent enterprise products through multi-step workflows was not yet present around those four releases. Both conditions had changed by January 2026. Cowork Plugins target enterprise SaaS products, and the underlying agent infrastructure is now reliable enough to be a credible substitute.

6 Conclusion

Prior literature has documented that generative AI can raise firm value by lowering labor costs. This paper shows that AI can also destroy firm value by displacing products, and that equity markets distinguish between the two channels. We introduce AI Product Displacement Risk (AI-PDR) to measure this demand-side threat and find that the cross-sectional displacement effect strengthens as AI advances from capability demonstration to targeted product substitution. Digital deliverability is the binding determinant. Whether a product can be produced and delivered in purely digital form separates the firms that bear the repricing from those that do not.

The results raise a strategic question for affected firms. They face two broad paths. They can race (invest in AI-augmented versions of existing products and pivot to AI-native delivery) or retreat (harvest remaining revenue and return cash to shareholders). The optimal response depends on whether a firm holds complementary assets that survive the product transition, including brand, distribution networks, customer relationships, and proprietary data. Firms with strong complementary assets may integrate AI into their product delivery and preserve their market position, whereas pure-play information-goods producers without

such assets face acute long-run revenue risk. Whether high-AI-PDR firms race or retreat in the quarters following AI release events is a natural direction for future research.

A separate question is the structure of the displacing industry. Displacement could arrive from a small number of large AI providers (Anthropic, OpenAI, Google) substituting directly for incumbent software products, or from a wave of AI-native startups building vertical applications on top of foundation models. The welfare and policy implications differ. Displacement by a few dominant platforms raises concentration concerns and may transfer rents from incumbent SaaS firms to AI-infrastructure providers, whereas displacement by a competitive fringe of startups is more likely to pass cost savings through to consumers. Which market structure emerges will shape whether product displacement represents a net welfare gain or a redistribution of market power.

Several limitations point toward future work. AI-PDR is a static measure based on current 10-K filings and does not capture firms' ability to adapt or pivot, so it measures exposure rather than realized outcomes. The event study identifies market reactions, not realized displacement. The latter will take years to materialize in revenue and earnings data. Examining real effects such as changes in investment, hiring, and R&D at high-AI-PDR firms would test whether the repricing translates into resource reallocation. Our strongest evidence comes from a single product-substitution event, and the generalizability of the escalation pattern depends on whether future AI product launches follow a similar structure. Finally, AI-PDR captures only the demand side and does not account for second-order effects such as firms that benefit from the displacement of their competitors or the potential for displaced firms to become acquisition targets.

A complete accounting of AI's effect on firm value would integrate both channels (labor cost savings and product displacement), which may offset or reinforce each other within the same firm. AI currently threatens primarily digital products and services. As it extends into physical goods and regulated services, the frontier of displacement risk will expand. The conceptual distinction between labor displacement and product displacement applies wherever a new technology creates close substitutes for existing goods, making the framework portable beyond the generative AI setting studied here.

Acemoglu, Daron, "The Simple Macroeconomics of AI," *Economic Policy*, 2025, *40* (121), 13–58.

_ **and Pascual Restrepo**, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

_ **and** _ , "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

_ **, David Autor, Jonathon Hazell, and Pascual Restrepo**, "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.

Adobe, "Creative Cloud All Apps for Teams Pricing," adobe.com/creativecloud/business/teams, accessed May 2026 2026.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones, "Artificial Intelligence and Economic Growth," in "The Economics of Artificial Intelligence: An Agenda," University of Chicago Press, 2019, pp. 237–290.

Andrews, Isaiah and Maryam Farboodi, "Do Markets Believe in Transformative AI?," *NBER Working Paper No. 34243*, 2025.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, 2024, *151*, 103745.

Bick, Alexander, Adam Blandin, and David Deming, "The Rapid Adoption of Generative AI," *Management Science*, 2026.

CNBC, "AI Fears Pummel Software Stocks: Is It "Illogical" Panic or a SaaS Apocalypse?," CNBC February 6 2026.

CNN Business, "Anthropic's New AI Tool Sends Shudders Through Software Stocks," CNN February 4 2026.

Eisfeldt, Andrea L. and Gregor Schubert, "Generative AI and Finance," *Annual Review of Financial Economics*, 2025, *17*, 363–393.

_ **,** _ **, Miao Ben Zhang, and Bledi Taska**, "Generative AI and Firm Values," *Journal of Finance, forthcoming*, 2025.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: Labor Market Impact Potential of LLMs," *Science*, 2024, *384* (6702), 1306–1308.

Fast Company, "Why One Anthropic Update Wiped Billions Off Software Stocks," Fast Company February 5 2026.

Felten, Edward, Manav Raj, and Robert Seamans, "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal*, 2021, *42* (12), 2195–2217.

_ **,** _ **, and** _ , "How Will Language Modelers Like ChatGPT Affect Occupations and Industries?," *Working Paper*, 2023.

Gârleanu, Nicolae, Stavros Panageas, and Jianfeng Yu, "Technological Growth and Asset

Pricing," *Journal of Finance*, 2012, *67* (4), 1265–1292.

Habib-Valdhorn, Shiri, "Surviving the Software Storm," Globes May 26 2026.

Hadlock, Charles J. and Joshua R. Pierce, "New Evidence on Measuring Financial Constraints: Moving Beyond the KZ Index," *Review of Financial Studies*, 2010, *23* (5), 1909–1940.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller, "Artificial Intelligence and the Labor Market," *NBER Working Paper No. 33509*, 2025.

He, Songrun, Linying Lv, Asaf Manela, and Jimmy Wu, "Chronologically Consistent Large Language Models," Working Paper 2025.

_ , _ , _ , and _ , "Instruction Tuning Chronologically Consistent Language Models," Working Paper 2025.

Hoberg, Gerard and Gordon Phillips, "Text-Based Network Industries and Endogenous Product Differentiation," *Journal of Political Economy*, 2016, *124* (5), 1423–1465.

Kakhbod, Ali, Leonid Kogan, Peiyao Li, and Dimitris Papanikolaou, "Measuring Creative Destruction," *Working Paper*, 2024.

Kaplan, Steven N. and Luigi Zingales, "Do Investment-Cash Flow Sensitivities Provide Useful Measures of Financing Constraints?," *Quarterly Journal of Economics*, 1997, *112* (1), 169–215.

Kogan, Leonid and Dimitris Papanikolaou, "Growth Opportunities, Technology Shocks, and Asset Prices," *Journal of Finance*, 2014, *69* (2), 675–718.

_ , _ , Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," *Quarterly Journal of Economics*, 2017, *132* (2), 665–712.

Salesforce, "Sales Cloud Pricing," salesforce.com/sales/pricing, accessed May 2026 2026.

Schumpeter, Joseph A., *Capitalism, Socialism and Democracy*, New York: Harper and Brothers, 1942.

TechCrunch, "Anthropic Brings Agentic Plug-ins to Cowork," TechCrunch January 30 2026.

Ľuboš Pástor and Pietro Veronesi, "Technological Revolutions and Stock Prices," *American Economic Review*, 2009, *99* (4), 1451–1483.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," *Working Paper*, 2020.

Whited, Toni M. and Guojun Wu, "Financial Constraints Risk," *Review of Financial Studies*, 2006, *19* (2), 531–559.

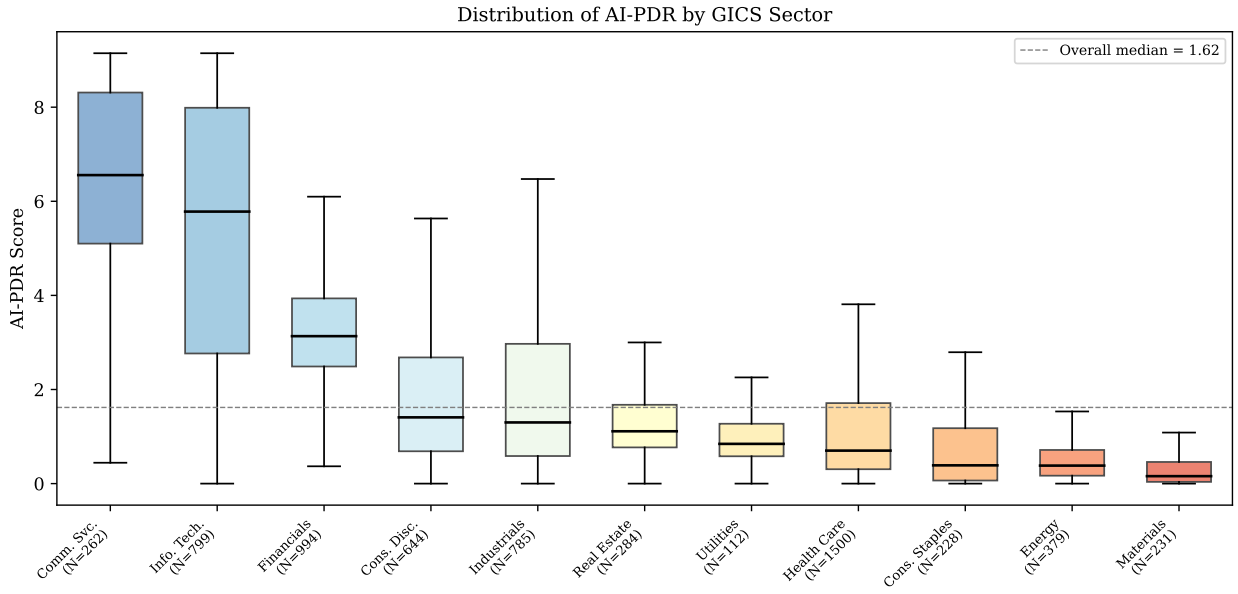


Figure 1: Distribution of AI-PDR by GICS Sector

Note: Box plots show the distribution of firm-level AI-PDR within each GICS sector, ordered by sector median. Boxes span the interquartile range; whiskers extend to 1.5× IQR. The dashed grey line marks the overall sample median. Outliers are suppressed for visual clarity. N denotes the number of firms in each sector.

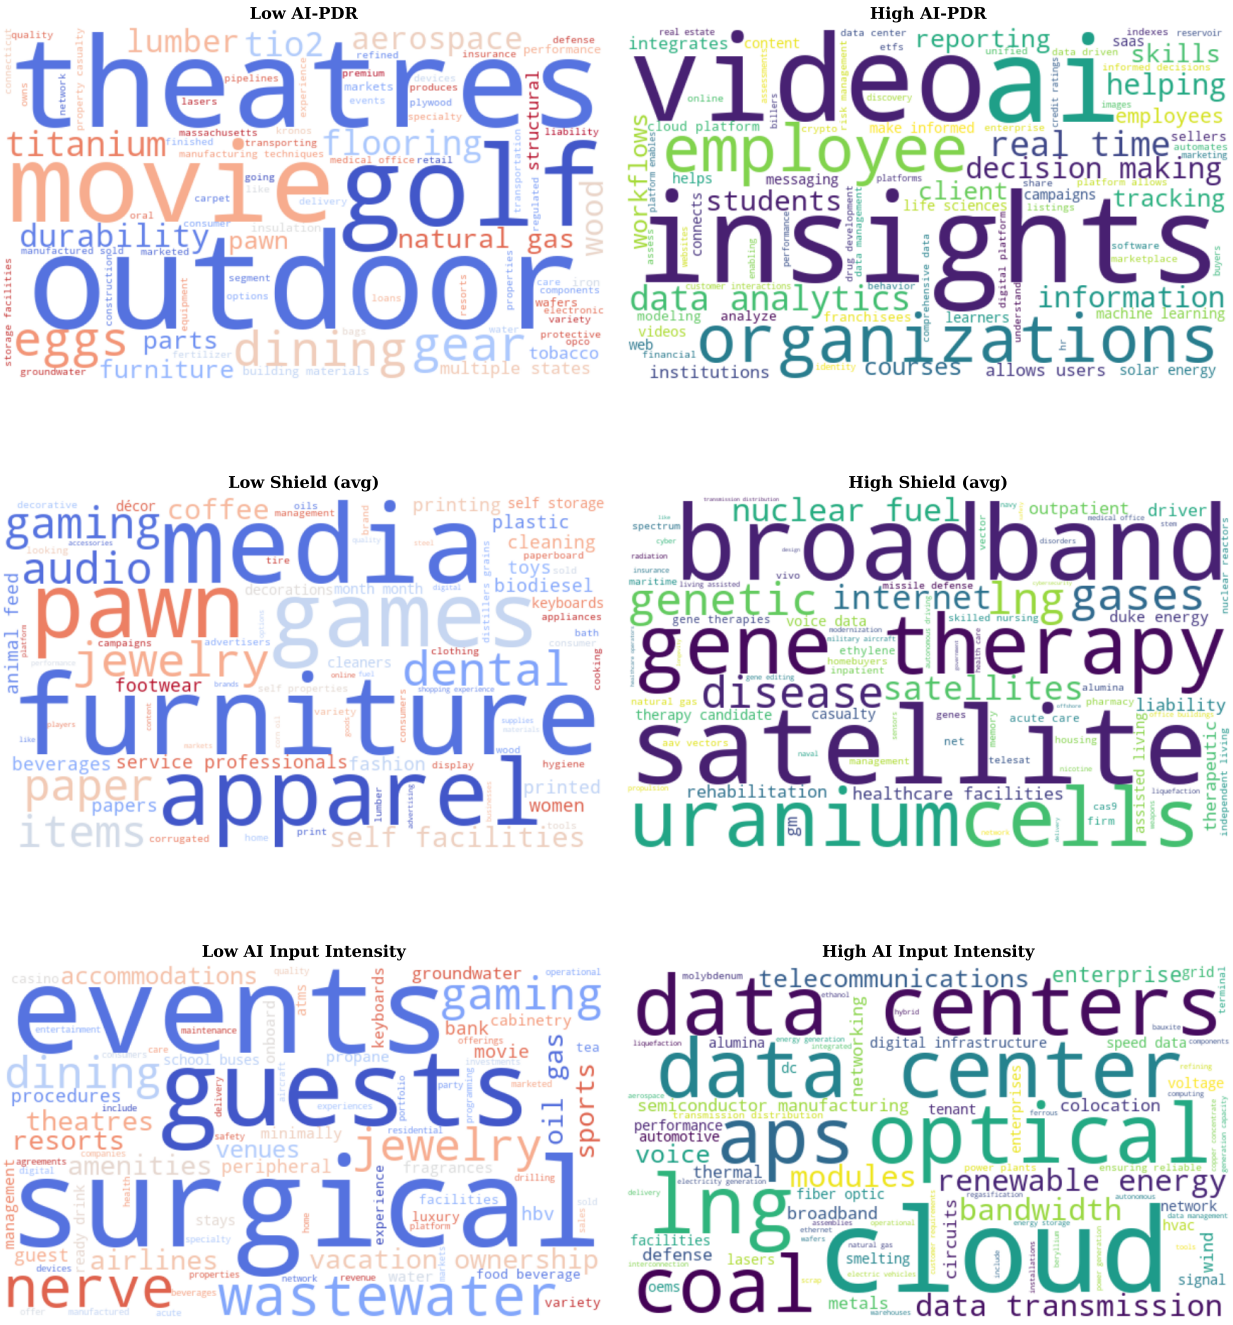


Figure 2: Word Clouds of Product Descriptions by Scoring Dimension

Note: Each panel is a TF-IDF-weighted word cloud of product descriptions from 10-K filings. For each scoring dimension, firms are sorted by their firm-level importance-weighted score, and the ten highest- and ten lowest-scoring firms within each 2-digit GICS sector are pooled (sectors with fewer than twenty firms are dropped). Rows from top to bottom: AI-PDR (Digital Deliverability), the composite Shield (average of the four Shield sub-dimensions), and AI Input Intensity; the left column shows low-score firms, the right column high-score firms. Bigrams are prioritized over subsumed unigrams and a standard 10-K boilerplate stop-word list is applied.

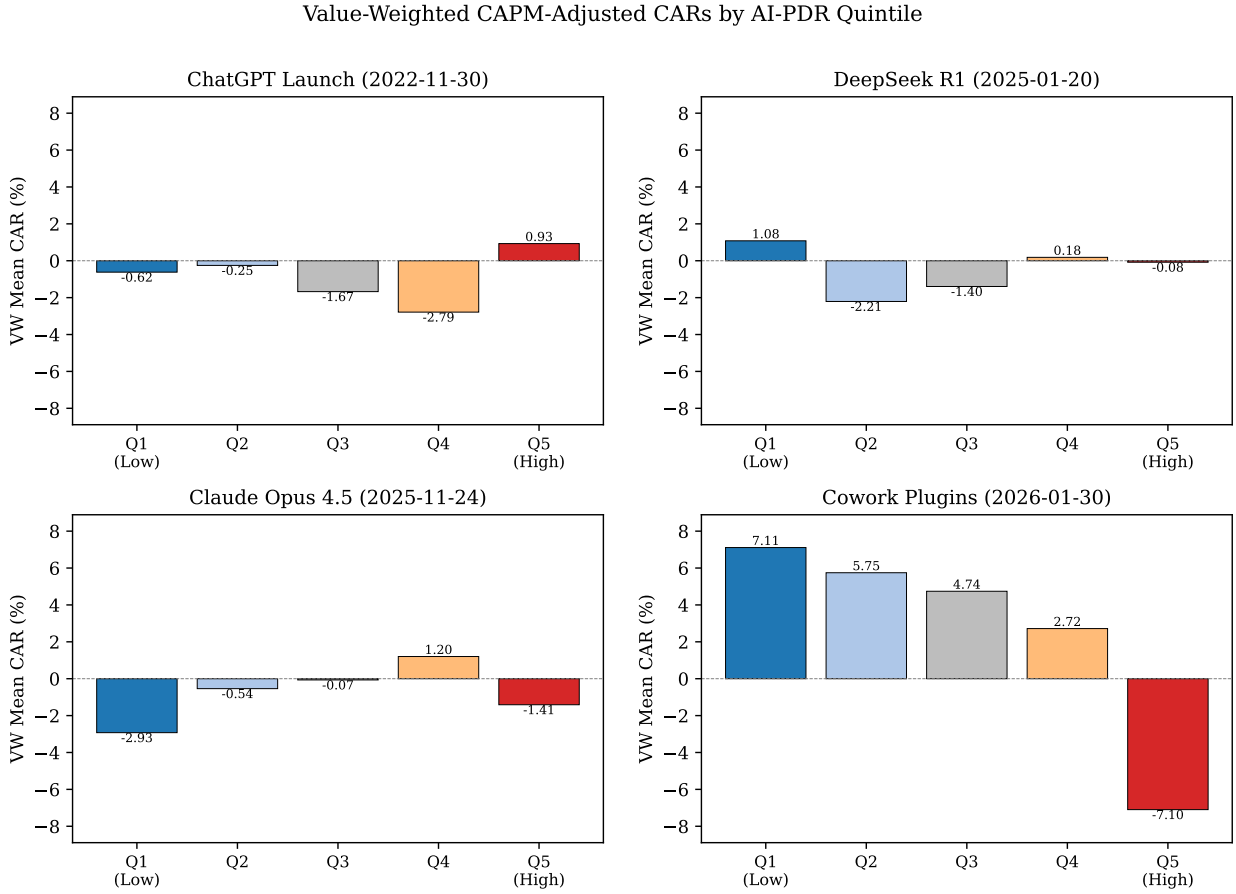


Figure 3: Value-Weighted Mean CARs by AI-PDR Quintile

Note: Firms are sorted into quintiles based on AI-PDR. Bars show value-weighted mean CAPM-adjusted CARs over the $[-1, +10]$ trading-day window around each event. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. Value weights use pre-event market capitalization. CARs are computed as $\text{CAR}_i = \sum_t \text{AR}_{i,t}$ with $\text{AR}_{i,t} = (r_{i,t} - r_{f,t}) - \hat{\alpha}_i - \hat{\beta}_i(r_{m,t} - r_{f,t})$, where $(\hat{\alpha}_i, \hat{\beta}_i)$ are from daily excess-return regressions over $[-252, -11]$.

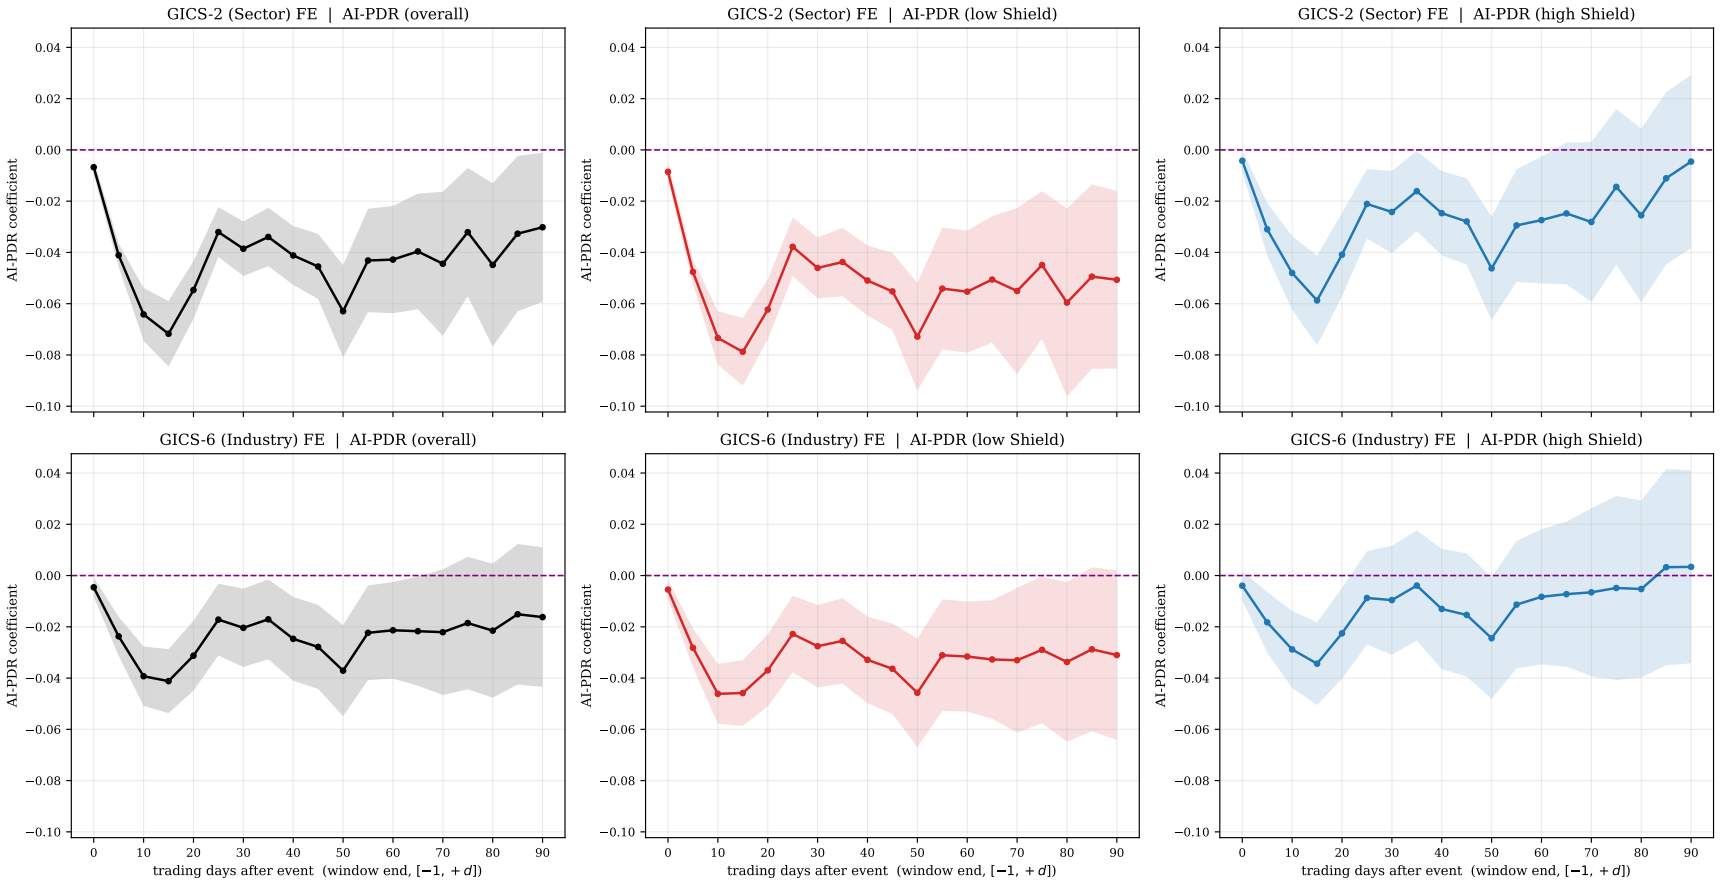


Figure 4: AI-PDR Effect on Cumulative Abnormal Returns across Event Windows (Cowork Plugins)

Note: Each panel plots the coefficient on standardized AI-PDR from a cross-sectional OLS regression of the CAPM-adjusted $[-1, +d]$ CAR around the Cowork Plugins event (January 30, 2026) on AI-PDR and the baseline controls, for each window end $d = 0, 5, \ldots, 90$ trading days; the shaded band is the 95% confidence interval. The top and bottom rows use GICS 2-digit (sector) and 6-digit (industry) fixed effects. The left panels plot the overall (pooled) coefficient; the middle and right panels split firms at the median Shield within each GICS 2-digit sector and plot the slope for low- and high-Shield firms separately. Controls, the estimation universe, and clustering (GICS 6-digit industry, also used for the bands) follow the baseline specification in Equation (1).

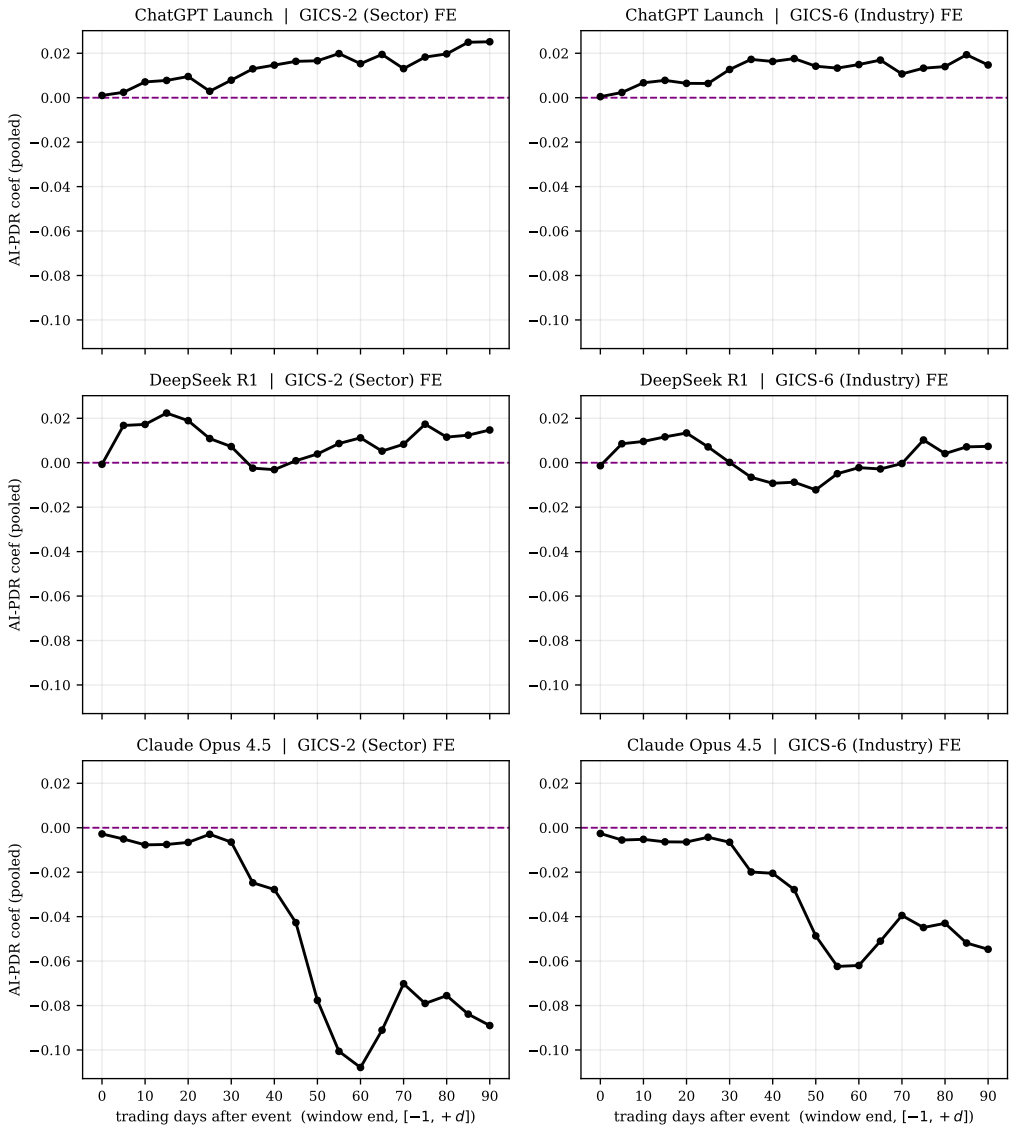


Figure 5: AI-PDR Effect on Cumulative Abnormal Returns across Event Windows, Other Main-Sample Events (ChatGPT, DeepSeek, Claude Opus)

Note: Each panel plots the pooled AI-PDR coefficient from a cross-sectional OLS regression of the CAPM-adjusted $[-1, +d]$ cumulative abnormal return on AI-PDR and the baseline controls of Equation (1), for window ends $d = 0, 5, 10, \ldots, 90$ trading days. Rows are the three other main-sample events, ChatGPT Launch (November 30, 2022), DeepSeek R1 (January 20, 2025), and Claude Opus 4.5 (November 24, 2025). The left and right columns use GICS 2-digit (sector) and GICS 6-digit (industry) fixed effects. Controls, the estimation universe, and clustering are as in Figure 4. The Claude Opus 4.5 window spans the Cowork Plugins event at +45 trading days, so Opus windows beyond roughly +45 absorb the Cowork repricing.

Table 1: Summary Statistics

	N	Mean	SD	P10	P25	Median	P75	P90
Panel A — AI-PDR, ESZT, and Firm Controls								
AI-PDR (Digital Deliverability)	3,995	2.67	2.61	0.00	0.47	2.00	4.00	7.09
Legal/Regulatory	3,995	5.22	2.75	1.76	2.70	4.91	7.60	9.20
Production Moats	3,995	4.95	1.82	2.71	3.43	4.71	6.40	7.67
Customer Trust	3,995	6.69	2.13	3.75	5.09	6.75	8.55	9.45
Customer Embeddedness	3,995	5.14	1.70	2.70	4.08	5.31	6.32	7.25
AI Input Intensity	3,995	0.93	1.19	0.00	0.00	0.44	1.47	2.67
Cowork Overlap	2,985	0.24	0.64	0.00	0.00	0.00	0.00	0.77
ESZT Labor Exp.	2,474	0.36	0.07	0.28	0.32	0.37	0.41	0.44
log(ME)	4,031	7.56	1.78	5.33	6.14	7.39	8.70	10.07
Tobin's Q	4,031	2.15	1.99	0.97	1.05	1.44	2.42	4.15
ROA	4,031	0.03	0.19	-0.21	0.01	0.07	0.12	0.18
Asset Growth	4,031	0.16	0.64	-0.13	-0.03	0.04	0.13	0.39

Panel B — Pairwise Correlations

	Dig.	Legal	Moats	Trust	Embed.	AI In.	Cowork	ESZT
Dig.	1.00							
Legal	-0.20	1.00						
Moats	-0.27	0.33	1.00					
Trust	0.09	0.79	0.34	1.00				
Embed.	0.06	0.47	0.59	0.67	1.00			
AI In.	0.12	-0.31	0.32	-0.11	0.32	1.00		
Cowork	0.55	-0.27	-0.09	-0.10	0.05	0.11	1.00	
ESZT	0.39	0.09	0.03	0.23	0.20	0.13	0.18	1.00

Note: Panel A reports cross-sectional distributions for the firm sample. AI-PDR (Digital Deliverability), the four Shield sub-dimensions (Legal/Regulatory, Production Moats, Customer Trust, Customer Embeddedness), AI Input Intensity, and Cowork Overlap are event-year (FY2025) importance-weighted firm-level averages reported on the 0–10 scale, using the most recent 10-K filing whose fiscal-year end strictly precedes the Cowork Plugins launch (2026-01-30). This matches the scoring used in the regression tables for the Cowork Plugins event. Eisfeldt et al. (2025) labor exposure is on its original 0–1 scale from the published source. In the regression tables, all scoring dimensions and ESZT are standardized to unit variance within each estimation sample. For the descriptive panel above, firm controls are the most recent Compustat fiscal-year observation as of the Cowork Plugins event. In the regression tables, controls are merged event-by-event to the most recent pre-event fiscal year. Firm controls are winsorized at the 1st and 99th percentiles. Panel B reports pairwise Pearson correlations among the six firm-characteristic scores (AI-PDR, the four Shield sub-dimensions, and AI Input Intensity), the Cowork Overlap score, and ESZT, computed on the subsample where all eight variables are non-missing ($N = 1,913$).

Table 2: Event Study: Cross-Sectional Regressions of CARs on AI-PDR

Event	AI-PDR	N	R^2
ChatGPT Launch	0.0067	2,922	0.1028
	(1.33)		
DeepSeek R1	0.0096*	2,894	0.1377
	(1.78)		
Claude Opus 4.5	−0.0052	2,769	0.1108
	(-1.24)		
Cowork Plugins	−0.0395***	2,803	0.3325
	(-6.68)		

Note: Each row reports a cross-sectional OLS regression of event-window CARs on AI-PDR (Digital Deliverability, standardized to unit variance within each sample). CARs are CAPM-adjusted with firm-specific market betas estimated on $[-252, -11]$ trading days and are winsorized at the 1st and 99th percentiles. All regressions use the $[-1, +10]$ event window, include GICS 6-digit industry fixed effects (GICS 6-digit industries with fewer than five firms are dropped), and control for log(ME), Tobin's Q, ROA, and asset growth, all winsorized at the 1st and 99th percentiles. AI-PDR is event-year-specific: for each event we use the firm's importance-weighted Digital Deliverability score from the single 10-K filing year immediately before the event (FY2021 for ChatGPT, FY2024 for DeepSeek R1 and Claude Opus 4.5, FY2025 for Cowork Plugins). The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table 3: Event Study: Cross-Sectional Regressions of CARs on AI-PDR across Industry-FE Granularities

	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
ChatGPT Launch	0.0055	0.0071*	0.0070*	0.0067	0.0062
	(1.08)	(1.82)	(1.72)	(1.33)	(1.06)
	$N = 2,946$	$N = 2,946$	$N = 2,946$	$N = 2,922$	$N = 2,847$
DeepSeek R1	0.0147***	0.0172***	0.0109**	0.0096*	0.0141**
	(4.95)	(4.26)	(2.24)	(1.78)	(2.34)
	$N = 2,921$	$N = 2,921$	$N = 2,921$	$N = 2,894$	$N = 2,827$
Claude Opus 4.5	0.0006	−0.0077**	−0.0077**	−0.0052	−0.0041
	(0.23)	(-2.23)	(-2.36)	(-1.24)	(-0.84)
	$N = 2,795$	$N = 2,795$	$N = 2,795$	$N = 2,769$	$N = 2,690$
Cowork Plugins	−0.0482***	−0.0646***	−0.0498***	−0.0395***	−0.0335***
	(-8.72)	(-12.38)	(-8.10)	(-6.68)	(-5.52)
	$N = 2,827$	$N = 2,827$	$N = 2,827$	$N = 2,803$	$N = 2,722$

Note: Each column reports a cross-sectional regression of CAPM-adjusted CARs over $[-1, +10]$ on AI-PDR (Digital Deliverability), standardized to unit variance within the estimation sample. Controls (not reported): log(ME), Tobin's Q, ROA, asset growth, each winsorized at the 1st and 99th percentiles. CARs are also winsorized at 1/99. Industry fixed effects vary across columns as indicated. GICS 6-digit and 8-digit specifications drop industries with fewer than five firms. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses across all columns. *, **, *** denote significance at 10%, 5%, 1%.

Table 4: Window Robustness: CAR Regressions across Event Windows

AI-PDR	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
ChatGPT Launch	0.0036	0.0067	0.0065	0.0127
	(1.44)	(1.33)	(1.11)	(1.59)
DeepSeek R1	0.0022	0.0096*	0.0134**	0.0002
	(0.93)	(1.78)	(2.04)	(0.02)
Claude Opus 4.5	−0.0058*	−0.0052	−0.0064	−0.0061
	(-1.81)	(-1.24)	(-1.39)	(-1.03)
Cowork Plugins	−0.0174***	−0.0395***	−0.0316***	−0.0204**
	(-4.72)	(-6.68)	(-4.49)	(-2.57)

Note: Each cell reports AI-PDR from a cross-sectional OLS regression of CAPM-adjusted CARs over the indicated window on AI-PDR (Digital Deliverability, standardized to unit variance within the sample). Controls (not reported): log(ME), Tobin's Q, ROA, asset growth, each winsorized at 1/99. CARs are also winsorized at 1/99. GICS 6-digit industry fixed effects throughout (GICS 6-digit industries with fewer than five firms are dropped). The estimation universe is firms with market capitalization above $100M and Compustat total assets above $100M. GICS 6-digit industry-clustered t-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table 5: Event Study: Interaction of AI-PDR with Shield Composite

Panel A — Composite Shield, by Event

Event	AI-PDR	Shield	AI-PDR × Shield	N	R^2
ChatGPT Launch	0.0047	−0.0037	−0.0050*	2,922	0.1038
	(0.94)	(-0.72)	(-1.70)		
DeepSeek R1	0.0090*	−0.0052	0.0001	2,894	0.1384
	(1.84)	(-1.25)	(0.03)		
Claude Opus 4.5	−0.0065	−0.0055	−0.0020	2,769	0.1116
	(-1.43)	(-1.36)	(-0.71)		
Cowork Plugins	−0.0366***	−0.0099	0.0118***	2,803	0.3371
	(-5.37)	(-1.49)	(2.85)		

Panel B — Supply vs. Demand Shield Decomposition (Cowork Plugins)

	(1) Shield = Supply Side	(2) Shield = Demand Side
AI-PDR	−0.0367***	−0.0377***
	(-4.87)	(-6.14)
Shield	−0.0138**	−0.0054
	(-2.04)	(-0.97)
AI-PDR × Shield	0.0136***	0.0082**
	(3.45)	(2.23)
N	2,803	2,803
R^2	0.3387	0.3348

Note: Panel A reports, for each of the four events, a cross-sectional OLS regression of CAPM-adjusted CARs over the $[-1, +10]$ window on AI-PDR (Digital Deliverability), the composite Shield (simple average of Legal/Regulatory, Production Moats, Customer Trust, and Customer Embeddedness), and the AI-PDR×Shield interaction. Panel B decomposes the Cowork Plugins regression by replacing the composite Shield with either the Supply-side shield (average of Legal/Regulatory and Production Moats) or the Demand-side shield (average of Customer Trust and Customer Embeddedness) and the corresponding AI-PDR × Side interaction. All score regressors are standardized to unit variance within the estimation sample and demeaned; interactions are products of the standardized-demeaned regressors. Each regression includes log(ME), Tobin's Q, ROA, asset growth (each winsorized at 1/99), and GICS 6-digit industry fixed effects. CARs are also winsorized at 1/99. The estimation universe is firms with market capitalization above $100M and Compustat total assets above $100M. GICS 6-digit industry-clustered t-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table 6: AI-PDR Effect by Cowork Plugin Overlap Subsample (Cowork Plugins event only)

Sample	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
Overlap $= 0$	-0.0231^{***}	-0.0337^{***}	-0.0281^{***}	-0.0266^{***}	-0.0195^{***}
	(-3.43)	(-8.88)	(-5.76)	(-4.97)	(-3.53)
	$N = 2,203$	$N = 2,203$	$N = 2,203$	$N = 2,167$	$N = 2,058$
Overlap > 0	-0.0546^{***}	-0.0743^{***}	-0.0563^{***}	-0.0358^{***}	-0.0364^{***}
	(-7.66)	(-11.82)	(-5.68)	(-3.46)	(-2.74)
	$N = 624$	$N = 624$	$N = 624$	$N = 585$	$N = 510$

Note: This table re-estimates the Cowork Plugins regression on two subsamples defined by Cowork Overlap, across five industry fixed-effect granularities. "Overlap = 0" restricts to firms with zero Cowork Overlap (about 78% of the sample). "Overlap > 0" restricts to firms with positive Cowork Overlap. All regressions use CAPM-adjusted CARs over the $[-1, +10]$ window with controls (log(ME), Tobin's Q, ROA, asset growth, each winsorized at 1/99). CARs are also winsorized at 1/99. AI-PDR is standardized to unit variance within the estimation sample. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses across all columns. *, **, *** denote significance at 10%, 5%, 1%.

Table 7: Analyst Forecast Revisions (Cowork Plugins, December FY-End Firms)

Panel A — Scaled Revisions across Four Outcomes

	(1)–(2) EPS revision		(3)–(4) Revenue revision		(5)–(6) LTG change		(7)–(8) Target Price revision	
AI-PDR	0.003	0.010	0.000	0.009**	0.859	0.455	−0.060***	−0.026***
	(0.44)	(1.25)	(0.09)	(2.08)	(1.47)	(0.53)	(-6.01)	(-2.67)
GICS 2-digit FE	Yes		Yes		Yes		Yes	
GICS 6-digit FE		Yes		Yes		Yes		Yes
N	1,452	1,452	1,923	1,923	651	651	2,016	2,016
R^2	0.050	0.134	0.035	0.082	0.044	0.193	0.108	0.202

Panel B — Target Price: Shield Decomposition

	(1)–(2) Composite Shield		(3)–(4) Shield = Supply Side		(5)–(6) Shield = Demand Side	
AI-PDR	−0.047***	−0.018*	−0.041***	−0.016	−0.057***	−0.023**
	(-4.52)	(-1.78)	(-4.18)	(-1.64)	(-5.34)	(-2.22)
Shield	0.031***	0.015*	0.028***	0.012	0.024***	0.012
	(3.18)	(1.65)	(2.76)	(1.24)	(2.93)	(1.52)
AI-PDR × Shield	0.019***	0.017***	0.023***	0.019***	0.008	0.011*
	(3.24)	(3.39)	(4.56)	(4.22)	(1.34)	(1.87)
GICS 2-digit FE	Yes		Yes		Yes	
GICS 6-digit FE		Yes		Yes		Yes
N	2,016	2,016	2,016	2,016	2,016	2,016
R^2	0.121	0.207	0.122	0.208	0.116	0.205

Note: The sample is December fiscal year-end firms with market capitalization above \$100M and Compustat total assets above \$100M. Panel A reports eight cross-sectional OLS regressions of scaled analyst-forecast revisions, one per (outcome, industry-FE granularity) pair. Columns (1)–(2) and (3)–(4) report revisions in EPS and revenue forecasts for fiscal year 2026; columns (5)–(6) report the change in the long-term growth (LTG) forecast, which is not tied to a specific fiscal year; columns (7)–(8) report the revision in the 12-month target price. Within each outcome, the first sub-column uses GICS 2-digit (sector) fixed effects and the second uses GICS 6-digit (industry) fixed effects, as indicated by the GICS 2-digit FE and GICS 6-digit FE rows. The revision is $y_i = (\text{post}_i - \text{pre}_i)/\text{pre}_i$ for EPS, Revenue, and Target Price, and the level change $y_i = \text{post}_i - \text{pre}_i$ for LTG (already in percentage points). The pre-event observation is the IBES mean forecast for December 2025; the post-event observation is the IBES mean forecast for March 2026. Each Panel A regression enters AI-PDR (Digital Deliverability). Panel B reports six cross-sectional regressions on the Target Price outcome. Each column pair uses a different Shield definition. Columns (1)–(2) use the composite Shield (average of Legal/Regulatory, Production Moats, Customer Trust, and Customer Embeddedness). Columns (3)–(4) use only the Supply-side shield (average of Legal/Regulatory and Production Moats), and columns (5)–(6) use only the Demand-side shield (average of Customer Trust and Customer Embeddedness). In Panel B, the rows labeled "Shield" and "AI-PDR $\times$ Shield" refer to whichever Shield definition the column pair uses. Within each shield spec, the first sub-column uses GICS 2-digit (sector) fixed effects and the second uses GICS 6-digit (industry) fixed effects. All score regressors are standardized to unit variance within the estimation sample and demeaned; interactions are products of the standardized-demeaned regressors. Firm controls (log(ME), Tobin's Q, ROA, asset growth, winsorized at 1/99) are included in every regression. EPS observations with pre-event EPS below \$0.50 (including all negative-EPS firms) are dropped to avoid the sign ambiguity of scaling by a negative denominator. Revenue observations with pre-event revenue below \$10M are likewise dropped, on the same small-denominator rationale. Scaled revisions are winsorized at 1/99. GICS 6-digit industry-clustered t-statistics in parentheses across all columns. *, **, *** denote significance at 10%, 5%, 1%.

Internet Appendix

Listing A1. Shield scoring prompts (full, four per-dimension calls).

System prompt (shared across all four shield API calls). The model is told that it is a risk scoring engine and that its output must be valid JSON following the schema exactly, with no markdown and no commentary. Each per-dimension user message ends with the JSON array of products to score. Each call writes a short `reasoning` string before emitting a 0–10 integer `score`.

The four shield dimensions below give the per-dimension text appended to that shared system prompt. In each, the *anchors* describe categories (examples in parentheses are illustrative, not exhaustive), and the *scoring rules* are the per-dimension instructions handed to the model alongside the rubric.

Dimension A. Legal and Regulatory Requirements. Score each product on its legal and regulatory requirements (0–10 integer). The question is whether selling the product requires specific legal licenses, regulatory approvals, or professional certifications, focusing on product-specific licensing (general business regulations that apply to every firm, such as tax or employment law, do not count). Higher scores indicate a higher legal barrier to new entrants.

Anchors.

0 No product-specific licensing or regulatory requirement (commodity digital content, template documents, consumer apparel).

2 Minimal regulatory friction (general retail, consumer software).

4 Light industry-specific licensing (private security, cosmetology, general contractor trades).

6 Moderate licensing or compliance regime (real estate brokerage, food-service establishments).

8 Heavy professional licensing required per practitioner or per firm (CPA audit, broker-dealer services, retail pharmacy).

10 Prohibited without specific government approval or licensure (prescription drugs, banking deposit products, insurance underwriting).

Scoring rules.

1. Write reasoning before the score. Think step by step about what legal or regulatory requirements are tied specifically to selling this product.
2. Infer obvious industry requirements even if the description does not spell them out ("audit" implies CPA and PCAOB licensure, "banking" implies a banking charter). Do not invent requirements the description does not support.

3. Do not default to 5. There is no anchor at 5, that integer is a gap, so choose the nearest anchor above or below.
4. Use the full 0–10 range. Odd scores (1, 3, 7, 9) are valid when a product falls between anchors.
5. Differentiate products within the same batch. Giving them all the same score is almost always wrong.
6. Base the score on the product description, not the name.

Dimension B. Production Moats. Score each product on its production moats (0–10 integer). The question is whether a new, well-funded entrant could build or acquire the production-side assets needed to produce the product within a few years. Relevant assets include proprietary physical infrastructure, irreplaceable proprietary datasets used in production, patents and process know-how, specialized manufacturing or capacity assets, and structural network effects embedded in the market. Higher scores indicate a harder-to-replicate production side.

Anchors.

0 No moat (commodity agricultural products, basic content creation, generic consumer services such as cleaning or landscaping).

2 Minimal moat (commodity wholesale distribution, house-label or unbranded consumer goods, generic IT services using off-the-shelf tools).

4 Weak moat (standard commodity manufactured goods such as basic steel or commodity chemicals, generic web hosting, commercial staffing and temporary placement).

6 Moderate moat (specialty chemicals with proprietary processes, specialty industrial equipment with proprietary engineering).

8 Strong moat (specialized proprietary data services such as financial terminals or credit-bureau data, advanced-node semiconductor products, core telecom backbone services).

10 Massive moat (oil and gas pipeline transportation, stock-exchange trading and clearing, commercial aviation aircraft manufacturing).

Scoring rules.

1. Write reasoning before the score. Identify the specific production-side assets, physical infrastructure, patents, unique data, specialized capacity, or network effects, that an entrant would need.

2. Infer obvious production-side assets from the industry even when the description does not spell them out ("pipeline transportation" implies right-of-way and installed physical infrastructure, "advanced-node semiconductor foundry" implies multi-billion-dollar fab capacity, "core telecom backbone" implies fiber and spectrum holdings). Do not invent moats the description does not support.
3. Do not default to 5. There is no anchor at 5, that integer is a gap, so choose the nearest anchor above or below.
4. Use the full 0–10 range. Odd scores (1, 3, 7, 9) are valid when a product falls between anchors.
5. Differentiate products within the same batch. Giving them all the same score is almost always wrong.
6. Base the score on the product description, not the name.

Dimension C. Customer Trust Requirements. Score each product on its customer trust requirements (0–10 integer). The question is whether a customer would accept the product from a new, unbranded provider with no track record, reputation, or institutional backing. The focus is on demand-side trust, the customer's willingness to purchase from an unknown provider. Higher scores indicate a higher trust threshold for new providers.

Anchors.

0 Commodity, customers indifferent to provider identity (generic translations, commodity agricultural products, basic construction materials such as lumber or drywall).

2 Light trust, customers evaluate mostly on functionality and price (commodity bottled water, unbranded mobile games from unknown developers, generic cleaning services).

4 Moderate credibility needed (consumer health apps, online tutoring services).

6 Meaningful trust signals needed (accounting software, accredited medical practices, news journalism).

8 High trust threshold, multi-year contracts, SLAs, or performance warranties matter materially (enterprise cloud databases, industrial equipment with performance warranties, wealth management).

10 Customers require established brand or long institutional track record (custodian banking, national defense prime contracting, long-term annuity products).

Scoring rules.

1. Write reasoning before the score. Think step by step about whether a buyer would switch to a new provider based on functionality and price alone, or whether they need assurance of the provider's reliability and longevity.
2. Do not default to 5. There is no anchor at 5, that integer is a gap, so choose the nearest anchor above or below.
3. Use the full 0–10 range. Odd scores (1, 3, 7, 9) are valid when a product falls between anchors.
4. Differentiate products within the same batch. Giving them all the same score is almost always wrong.
5. Base the score on the product description, not the name.

Dimension D. Customer Embeddedness. Score each product on its customer embeddedness (0–10 integer). The question is how costly, disruptive, or time-consuming it would be for an existing customer to move to a functionally equivalent alternative. Sources of embeddedness include accumulated user data stored with the provider, deep integrations with other systems, trained user habits or workflows, contractual or multi-year commitments, installed infrastructure at the customer's site, and regulatory or compliance processes built around the incumbent. Higher scores indicate that it is harder for an existing customer to leave.

Anchors.

0 Trivial to switch, one-off or commodity purchase with no embedding (a book, a single stock photo, a one-time consulting report).

2 Minor friction, recurring commodity with portable delivery and cancel-anytime terms (office supplies subscription, commodity newsletter subscription, consumer streaming service with user playlists).

4 Moderate friction, days to a week or two, portable data, manageable user retraining (consumer cloud storage with several years of files and shared links, personal primary checking with direct deposit and recurring autopayments).

6 Notable friction, weeks of planned migration, data conversion, or user retraining (small-business accounting software with multi-year transaction history and chart of accounts, primary care provider with accumulated medical records, commercial property management service for a multi-tenant building).

8 High lock-in, switching disrupts operations and takes months (enterprise CRM with customer-specific customizations, enterprise HR or payroll platform with benefits and time-tracking integrations).

10 Near-permanent embeddedness, data or regulatory irrecoverability makes leaving impractical for most customers (hospital EHR tied to decade-plus clinical workflow, regulated utility with no alternative provider).

Scoring rules.

1. Write reasoning before the score. State specifically what locks the customer in (data, integrations, network, contract, habit, or infrastructure), then commit to a number.
2. Do not default to 5. There is no anchor at 5, that integer is a gap.
3. Use the full 0–10 range. Odd scores (1, 3, 7, 9) are valid between anchors.
4. Differentiate products within the same batch. Giving them all the same score is almost always wrong.
5. Base the score on the product description, not the name.

Listing A2. Overlap scoring prompt for the Cowork Plugins.

System message. The model is told that it is a product-market overlap scoring engine and that its output must be valid JSON following the schema exactly, with no markdown and no commentary.

User message. Below are 10 functional bundles. Each bundle is defined by the specific list of capabilities it performs.

1. Productivity. Personal knowledge-worker task management (markdown task lists, status tracking, stale-item triage), workplace memory that learns the user's people, projects, and terminology, action-item extraction from email, calendar, and chat to surface missed todos, and a visual task dashboard.

2. Sales. Prospect and account research, call preparation and post-call summaries, pipeline review and weighted forecasting, outreach drafting (email, LinkedIn), competitive intelligence with battlecards, and custom sales asset creation (decks, one-pagers, landing pages).

3. Customer support. Support ticket triage with priority assessment (P1 through P4) and routing, multi-source customer research, response drafting tailored to channel and urgency, escalation packaging with reproduction steps and business impact, and knowledge base article authoring from resolved issues.

4. Product management. PRD and feature spec writing with user stories and acceptance criteria, roadmap planning (Now/Next/Later, OKR alignment, dependency mapping), stakeholder updates tailored by audience, user research synthesis from interviews, surveys, and support tickets, competitive briefs with feature matrices, and product metrics review.

5. Marketing. Content drafting (blog posts, social, email, landing pages, press releases, case studies), campaign planning with content calendars and success metrics, brand voice review against style guides, SEO audits (keyword research, on-page, technical), email nurture and drip sequences, marketing performance reporting, and competitive positioning analysis.

6. Legal. Contract review against an organizational negotiation playbook with clause-level redline generation, NDA triage (green/yellow/red classification), vendor agreement status checks, compliance workflows (GDPR/CCPA, DPA reviews), legal briefings (daily, topic, or incident), and templated responses for data subject requests and discovery holds.

7. Finance. Month-end close sequencing, journal entry preparation (accruals, prepaids, fixed assets, payroll, revenue) with proper debits and credits, account reconciliation (GL-to-subledger, bank recs, intercompany), financial statement generation with period-over-period comparison, variance and flux analysis with waterfall decomposition, and SOX 404 control testing workpapers.

8. Data. Natural-language data analysis backed by SQL query writing across dialects (Snowflake, BigQuery, Databricks), dataset profiling and quality assessment, publication-quality visualization in Python, interactive HTML dashboards with filters, and pre-share validation (methodology, sanity checks, bias detection).

9. Enterprise search. Unified cross-tool search across email, chat, cloud storage, wikis, project tracking, CRM, and ticketing, with query decomposition per source, deduplication and source attribution across results, and daily or weekly activity digests grouped by topic.

10. Bio research. Biomedical literature search (PubMed, bioRxiv, Wiley journals), single-cell RNA-seq quality control and analysis (`scvi-tools`), bioinformatics pipelines (Nextflow nf-core for RNA-seq, variant calling, ATAC-seq), drug-target prioritization (OpenTargets) and AI-driven drug discovery (histopathology, bioactive compound databases), clinical trial protocol generation, lab data management with instrument-output standardization (Allotrope), scientific illustration generation, and research strategy frameworks (Fischbach-Walsh).

Two types of products can overlap with a bundle, both of which count. (a) Software overlap, the product is a software tool whose functions are also performed by a bundle (a CRM platform overlaps with Sales, a help desk SaaS overlaps with Customer support). (b) Service overlap, the product is a human or professional service whose work output is also produced by a bundle (a bookkeeping firm's core service overlaps with Finance, a contract-review law firm's service overlaps with Legal).

Task. For each product or service below, score the degree of functional overlap with each of the 10 bundles independently. The model returns 10 integer scores per product, one per bundle.

Score (0–10 integer). The score depends jointly on (i) whether the bundle covers the product's core function, and (ii) whether the non-covered functions are themselves core to the product. Anchors are at odd integers, with even scores (2, 4, 6, 8) valid when a product falls between two anchors. The same 0–10 scale applies to every bundle.

0 No overlap. Physical products, physical infrastructure, capital-intensive services, or domains entirely outside the 10 bundles (steel, oil, airlines, semiconductor fab).

1 Negligible. Bundle covers a trivial piece of the product, and the product's core function is something else.

3 Weak. Bundle covers a peripheral feature of the product, and the product's core function is not covered.

5 Moderate. Bundle covers the product's core function, but the non-covered functions are themselves core to the product.

7 Strong. Bundle covers the product's core function, and the non-covered functions are peripheral.

9 Near-complete. Bundle covers the product's core function and most peripheral functions, with remaining differences being minor (branding, legacy integrations).

10 Exact match. The product is itself a bundle's core function with no meaningful differentiator.

Calibration examples (used to anchor judgment, not scored).

"Glean unified workplace search" → Enterprise search = 9. Bundle covers core and most peripherals (cross-tool search across email, chat, storage, wiki, and ticketing, source attribution, daily digests, query decomposition per source). Remaining differences from the bundle are minor.

"BlackLine close management software" → Finance = 7. Bundle covers core (close sequencing, journal entries, reconciliation, intercompany matching, variance analysis). Does not cover integration tooling or audit trail features, but those are peripheral to the close workflow.

"HubSpot Marketing Hub" → Marketing = 7. Bundle covers core (content drafting, email automation, SEO audits, campaign planning). Does not cover landing-page hosting, but that is peripheral to the product.

"Outsourced bookkeeping services" → Finance = 5. Bundle covers part of core (close, JE, reconciliation). Does not cover AP/AR, payroll preparation, or tax work, which are themselves core to bookkeeping engagements.

"Tableau data visualization" → Data = 5. Bundle covers part of core (chart and dashboard creation). Does not cover the governed data layer, Tableau Server, or access control, which are themselves core to the product.

"Mixpanel product analytics" → Product management = 3. Bundle covers product metrics review, a peripheral feature of the product. Does not cover the product core (event instrumentation and behavioural analytics).

"Notion workspace" → Productivity = 3. Bundle covers lightweight task lists and reminders, peripheral features of the product. Does not cover the product core (workspace docs, wiki, and databases).

"DocuSign e-signature platform" → Legal = 1. Bundle trivially covers post-execution contract status tracking. Does not cover the product core (signature workflow and document execution). The bundle is contract review and compliance, not contract execution.

"LinkedIn Premium subscription" → Sales = 1. Bundle trivially covers InMail outreach. Does not cover the product core (professional networking and profile visibility).

"Commercial banking services" (loans, deposits, payments, cash management) → all bundles 0. Banking is not accounting work (close, JE, reconciliation), it is lending, treasury, and cash management. Domain is not function for FINANCE, and no other bundle has functional overlap.

RULES.

1. Write reasoning before the scores. Base scores on what the product actually does, not what the name suggests. Read the product description carefully. Marketing names can be misleading.

2. Score each bundle independently against the same 0–10 rubric. A product can have meaningful overlap with more than one bundle (a sales enablement tool may score on both SALES and MARKETING).

3. Most products will score 0 on most bundles. When in doubt on any single bundle, pick the lower anchor. A product that overlaps with one bundle and is unrelated to the other nine should have one non-zero score and nine zeros. Do not spread non-zero scores across multiple bundles.

4. Sold-to is not overlap. The score depends on what the product itself does, not who buys it. A commercial liability insurance product scores 0 on LEGAL even when its buyers are all law firms. A contract management software product scores 7 on LEGAL because the product itself does what the bundle does.

5. Domain is not function. Being in the same industry as a bundle does not create overlap, only matching functions do. Banking and financial-services products score 0 on

FINANCE because the bundle is accounting work (close, JE, reconciliation), not lending, treasury management, cash management, wealth management, retirement services, loans, or investment advisory. A pharmaceutical drug or therapy in clinical trials scores 0 on BIO RESEARCH because the bundle is research workflow, not drug products.

6. If all products in a batch are physical goods or infrastructure, every bundle for every product should be 0. Never inflate to create variation.

The user message closes with the instruction to return one object per product, echoing back the `idx` and `name` received, plus a brief `reasoning` (one to three sentences) and the `scores` object holding all 10 bundle scores. The firm-level Cowork Overlap is the max across the 10 plugin scores, aggregated to firm level via the importance-weighted average used by AI-PDR.

Listing A3. AI Input Intensity taxonomy classification prompt.

This listing reproduces the LLM prompt used to construct AI Input Intensity. The prompt is applied to the same product set extracted from 10-K filings (Listing 1) and uses GPT-4.1 with a strict JSON schema, temperature 0, and no markdown or commentary in the response.

System message. The model is told it is a product classification engine that assigns each product to one category from the AI-input taxonomy below or to `NONE` if no category fits, and that output must be valid JSON following the schema exactly.

TAXONOMY. The taxonomy contains 23 categories grouped into four tiers, plus a residual `NONE`.

Tier 1 (AI core).

`T1_1` Data center AI accelerators (training and inference chips for hyperscale).

`T1_2` High-bandwidth memory (HBM).

`T1_3` AI-class networking (InfiniBand, NVLink, AI-fabric Ethernet, AI optical interconnects).

`T1_4` AI server systems and rack-scale designs (DGX/HGX class).

`T1_5` GPU-dedicated AI cloud whose primary business is selling AI compute.

`T1_6` GPU-accelerated programming libraries, ML runtimes, MLOps platforms, training and fine-tuning platforms, and inference serving systems.

`T1_7` Foundation models and inference APIs.

Tier 2 (AI-adjacent infrastructure).

`T2_1` High-performance AI storage and storage-management software (parallel file systems, AI-marketed flash arrays).

`T2_2` AI-specific data center cooling (liquid, immersion, direct-to-chip).

`T2_3` Data center power distribution and backup (UPS, switchgear, fuel cell modules).

`T2_4` Data center colocation and interconnection services.

`T2_5` Data center construction and engineering services.

`T2_6` Hyperscale cloud.

`T2_7` Edge AI silicon and modules for consumer devices, automotive, and robotics.

Tier 3 (semiconductor manufacturing).

`T3_1` EDA software and chip design IP licensing.

`T3_2` Fab and packaging materials and consumables (silicon wafers, photomasks, substrates, chemicals, gases).

`T3_3` Semiconductor capital equipment (lithography, etch, deposition, metrology, test, gas/fluid delivery, process tool components).

`T3_4` Contract wafer fabrication and foundry services.

`T3_5` Advanced chip packaging services (CoWoS, FOWLP, 2.5D/3D stacking).

Tier 4 (general inputs).

`T4_1` Electric power generation and grid operations.

`T4_2` General industrial cooling and power equipment.

`T4_3` Commodity memory and storage products (DDR DRAM, NAND flash, NVMe SSDs, hard drives).

`T4_4` Enterprise and telecom networking and optics.

NONE Product is not an AI input.

Rules (applied in order).

1. Pick exactly one category code.
2. Classify by what the product is, not by what it is branded. A generic data center colocation is T2_4, a generic contract foundry is T3_4, a generic fab materials supplier is T3_2, any electric power or grid operations product is T4_1, even if the description does not say "AI".
3. AI inputs, not AI applications. Products that ML engineers and AI infrastructure teams buy to train, fine-tune, deploy, or serve AI models (chips, networking, GPU capacity, runtimes, foundation model APIs) fit the taxonomy. Products that non-technical business customers buy to receive an AI-produced outcome (SaaS, AI-powered analytics, medical AI apps, AI chatbots) go to NONE, even when branded with AI.
4. Highest-tier match wins. If a product genuinely fits categories in multiple tiers, classify to the highest tier (T1 > T2 > T3 > T4).
5. Within the same tier, match the category that best fits the product's primary economic function.
6. Emit NONE when no category matches. Do not force-fit.

The user message asks the model to classify the input batch and return one object per product with `idx`, `name`, `reasoning`, integer `tier` in 0–4 (where 0 means NONE), and `category` from the 23-code enumeration plus NONE. The firm-level taxonomy-based AI Input Intensity is the importance-weighted fraction of the firm's products classified into any Tier 1–4 category.

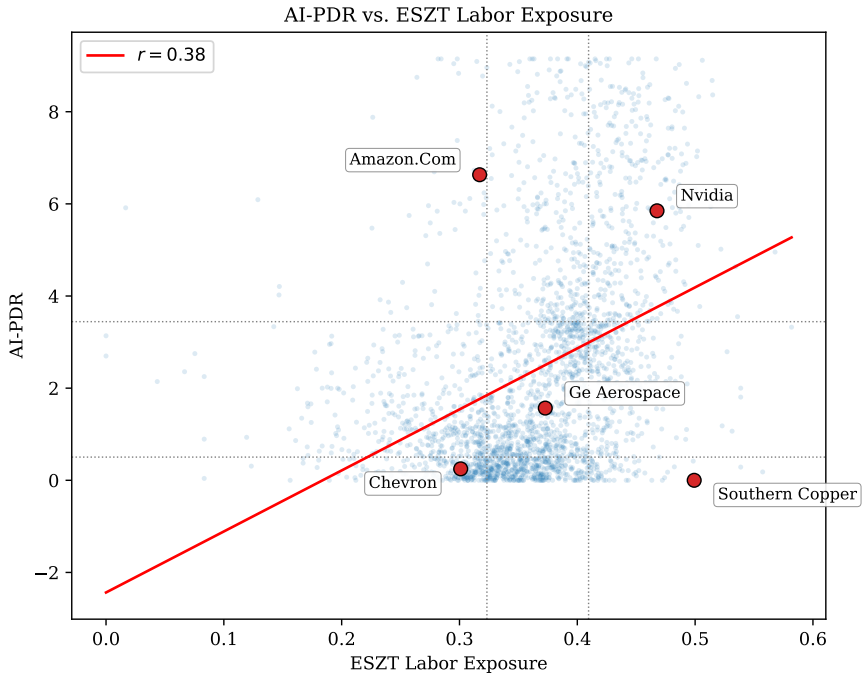


Figure A1: AI-PDR vs. ESZT Labor Exposure

Note: Each dot represents a firm. The horizontal axis is the Eisfeldt et al. (2025) generative AI labor exposure measure; the vertical axis is AI-PDR. The red line is a linear fit; the displayed Pearson correlation is computed on the plotted sample. The moderate positive correlation confirms that the two measures capture related but distinct dimensions of AI exposure. Five named firms are highlighted as representative examples. Cutoffs at the 25th and 75th percentiles of AI-PDR and of ESZT define four corner quadrants. A central quadrant is defined as the firms whose AI-PDR and ESZT both lie within the central 25% of the sample. Within each quadrant, the labeled firm is the largest by market equity at the most recent pre-event observation.

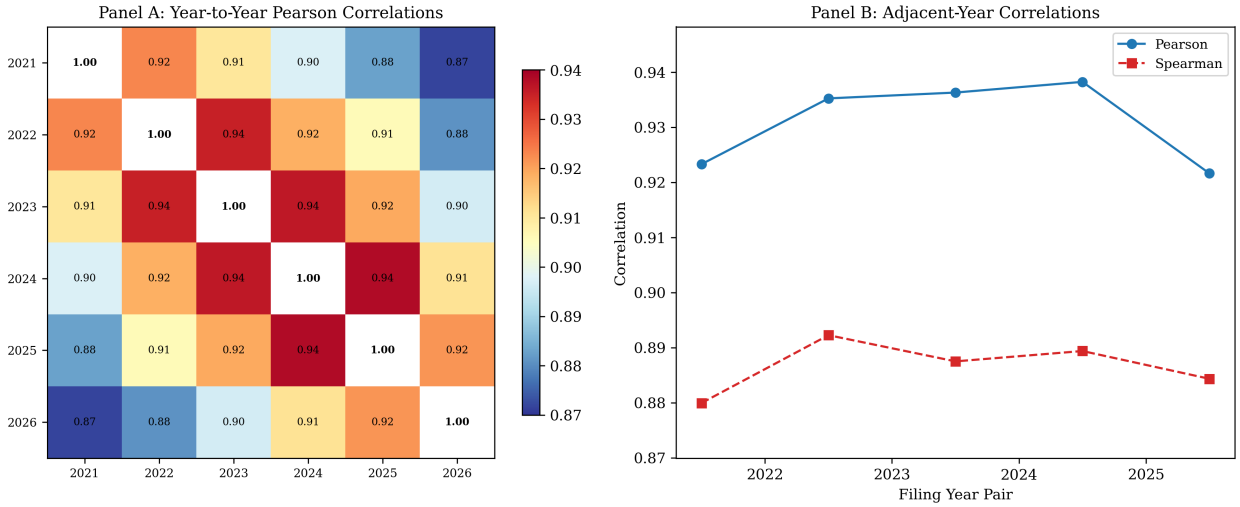


Figure A2: Temporal Stability of AI-PDR Scores

Note: Panel A displays a heatmap of pairwise Pearson correlations between firm-level AI-PDR scores computed from 10-K filings in each fiscal year. Panel B plots adjacent-year Pearson and Spearman correlations over time. AI-PDR for each firm-year is computed as the importance-weighted average of product-level scores from that year's 10-K filing.



Figure A3: Word Clouds of Product Descriptions by Shield Sub-Dimension

Note: Each panel is a TF-IDF-weighted word cloud of product descriptions drawn from 10-K filings. For each Shield sub-dimension, firms are sorted by their firm-level importance-weighted score, and the ten highest- and ten lowest-scoring firms within each 2-digit GICS sector are pooled (sectors with fewer than twenty firms are dropped). Columns from left to right: Legal/Regulatory, Production Moats, Customer Trust, Customer Embeddedness. Top row shows low-score firms; bottom row shows high-score firms. Bigrams that subsume single words are prioritized, and a standard 10-K boilerplate stop-word list is applied.



Figure A4: Word Clouds of Product Descriptions: SaaS vs Non-SaaS Firms

Note: Each panel is a TF-IDF-weighted word cloud of product descriptions from 10-K filings, with the left panel pooling firms classified as SaaS and the right panel pooling non-SaaS firms. The firm-level SaaS classification is produced by GPT-4.1 using the Wikipedia SaaS definition; a firm is flagged Yes if any of its top-importance product lines in its most recent 10-K is classified as Software-as-a-Service. Bigrams that subsume single words are prioritized, and a standard 10-K boilerplate stop-word list is applied.

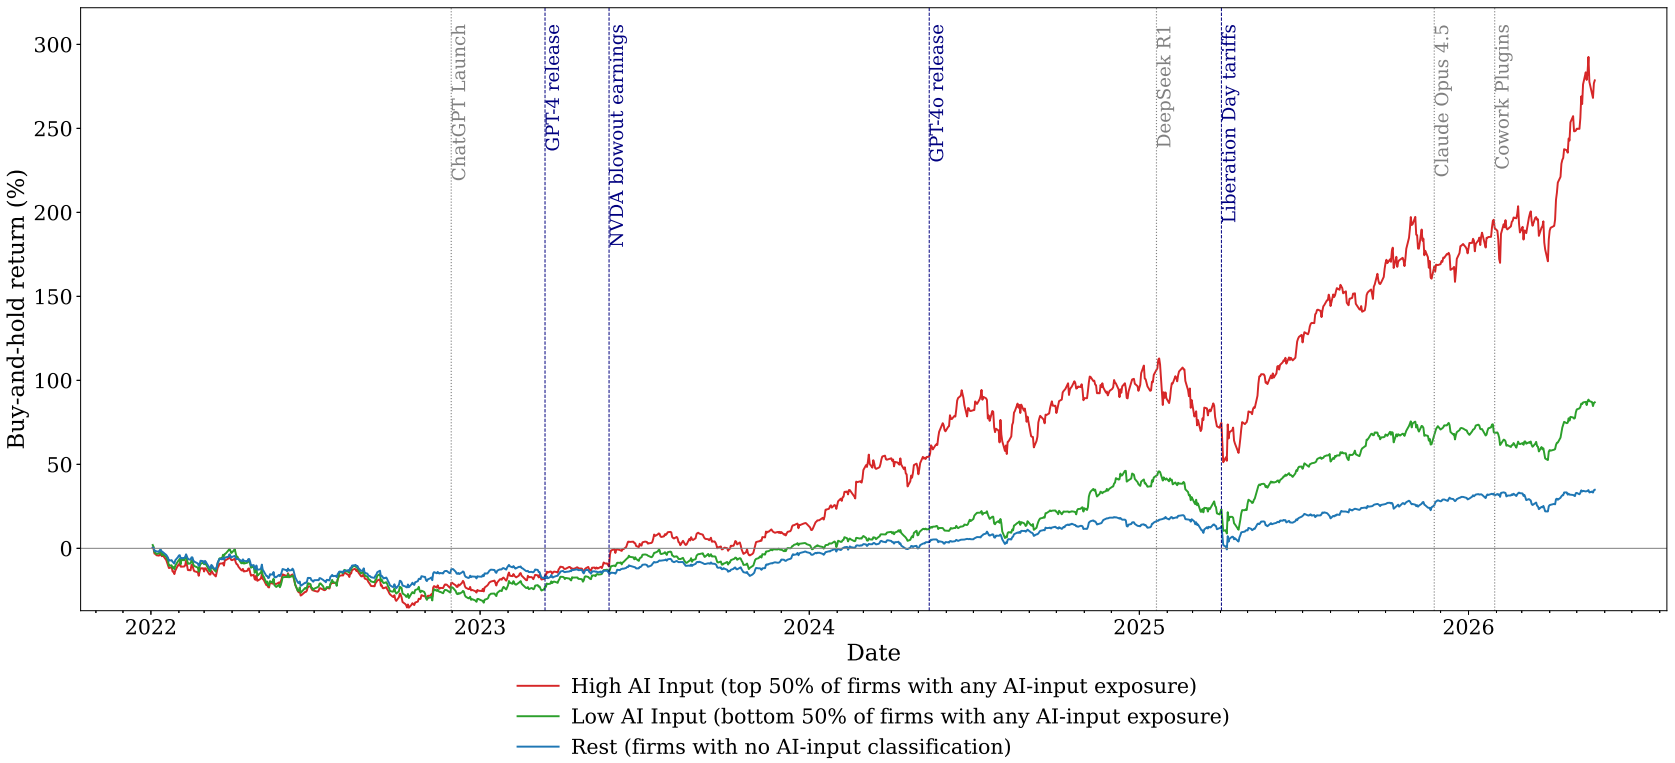


Figure A5: Value-Weighted Buy-and-Hold Cumulative Returns by AI Input Intensity, January 2022 to May 2026

Note: The figure plots value-weighted cumulative buy-and-hold returns for three portfolios of US-listed firms over January 3, 2022 through May 21, 2026. *High AI Input* is the top 50% of firms with any non-zero taxonomy-based AI Input Intensity. *Low AI Input* is the bottom 50% of firms with any non-zero AI Input Intensity. *Rest* is all other firms in CRSP. Portfolios are reconstituted annually. Each year's cutoff is computed from the cross-section of firms with a 10-K filing in the prior year, so a firm enters the High basket only after a pre-period 10-K classifies it as such. Within each year, portfolio returns are value-weighted each trading day by prior-day market capitalization. Dotted grey vertical lines mark the four event-study events in the paper. Dashed navy lines mark four additional AI- and market-relevant milestones, the GPT-4 release on March 14, 2023, NVIDIA's blowout Q1 FY24 earnings on May 24, 2023, the GPT-4o release on May 13, 2024, and the "Liberation Day" reciprocal-tariff announcement on April 2, 2025.

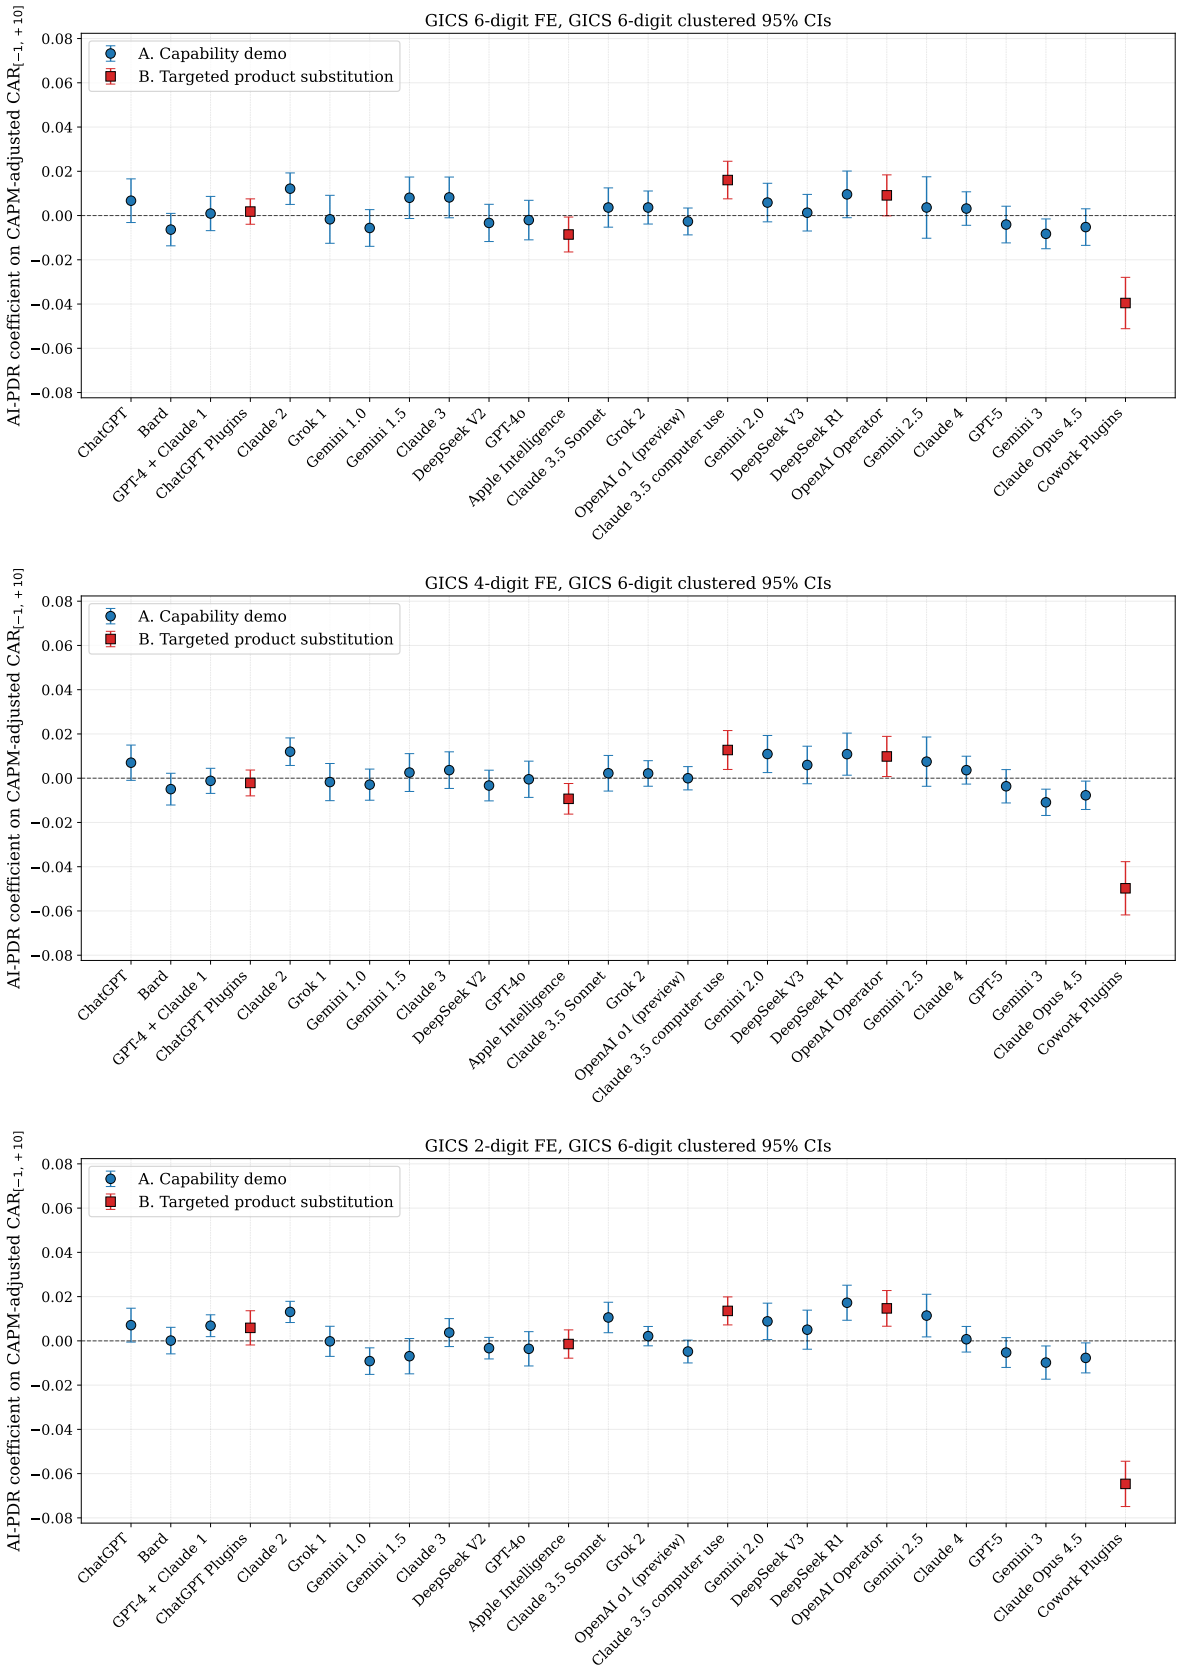


Figure A6: AI-PDR Coefficient on $\text{CAR}_{[-1,+10]}$ across 26 AI Release Events

Notes: Each panel plots the AI-PDR coefficient with 95% confidence intervals at each of 26 AI release events from late 2022 through January 2026. The specification matches Table 2, a cross-sectional OLS regression of CAPM-adjusted $CAR_{[-1,+10]}$ on standardized AI-PDR (Digital Deliverability) with log(ME), Tobin's Q, ROA, and asset growth as controls (all winsorized at 1/99). For each event, AI-PDR is taken from the most recent preceding filing year. The estimation universe is firms with market capitalization and Compustat total assets both above $100M. Industries with fewer than five firms at the FE level are dropped. The three panels differ only in the industry fixed-effect granularity, GICS 6-digit (top), GICS 4-digit (middle), GICS 2-digit (bottom). Standard errors are GICS 6-digit industry-clustered across all three panels. Events are classified as Category A (capability demonstration, blue circles) or Category B (targeted product substitution, red squares) by GPT-4.1 using the launching lab's official blog post and three contemporaneous news articles per event. The classification prompt is reproduced below.

System message. The model is told it is classifying AI release events and that output must be valid JSON following the schema exactly, with no markdown and no commentary.

User message. You are classifying AI release events into Category A (Capability demo) or Category B (Targeted product substitution).

Category A. Capability demo. The announcement describes general model improvements (benchmarks, modalities, reasoning, cost, context length, agentic capability, open weights) without naming specific commercial product categories or third-party services the model is positioned to replace. The substance is "the AI got better, cheaper, or faster," not "the AI replaces X."

Category B. Targeted product substitution. The announcement explicitly names commercial product categories or specific third-party services that the AI is positioned to replace or compete with. This includes language identifying named third-party services or product categories the AI now substitutes for, platform integration that subsumes downstream apps, or productized agents launched as competitors to specific software categories. Competition between AI products from different AI labs does not count as targeted substitution. The substituted product or category must be a non-AI commercial product or service.

Evidence rules. (1) Primary source, the launching company's official blog post. (2) Secondary source, contemporaneous news articles (cross-check). (3) Ignore market reaction. Stock-price moves are the dependent variable in this study and must not enter classification. (4) Ignore speculation about future products that have not launched.

Output schema. Return a JSON object with keys `event_name`, `category` ("A" or "B"), `confidence` ("high", "medium", or "low"), `named_substitutes` (list of exact-wording commercial products or categories named in the source, empty when category is A), `key_evidence_quotes` (one to three short verbatim quotes), and `reasoning` (a 2–3 sentence explanation).

Table A1: Variable Definitions

Variable	Definition
Panel A: Dependent Variables	
CAR (CAPM-adjusted)	Cumulative abnormal return over the event window using firm-specific pre-event CAPM intercept and beta. The daily abnormal return is $\mathrm{AR}_{i,t} = (r_{i,t} - r_{f,t}) - \hat{\alpha}_i - \hat{\beta}_i\,(r_{m,t} - r_{f,t})$, and the CAR is the sum of $\mathrm{AR}_{i,t}$ over the event window. The pre-event intercept and beta $(\hat{\alpha}_i, \hat{\beta}_i)$ are estimated by OLS of $(r_{i,t} - r_{f,t})$ on $(r_{m,t} - r_{f,t})$ over $[-252, -11]$ trading days before each event; firms with fewer than 100 estimation-window observations are dropped. The market excess return $r_{m,t} - r_{f,t}$ and the risk-free rate $r_{f,t}$ are the published Fama–French daily factors (the FF market factor is constructed from the CRSP value-weighted index in excess of the one-month Treasury yield). Baseline window is $[-1, +10]$ trading days. Winsorized at 1st/99th percentiles.
Panel B: Key Independent Variables	
AI-PDR	AI Product Displacement Risk. Firm-level importance-weighted average of product-level Digital Deliverability scores, the degree to which the product can be produced and delivered in purely digital form, with no physical goods, human labor, or financial capital. Baseline uses the firm's most recent 10-K filing before each event date (FY 2021 for ChatGPT, FY 2024 for DeepSeek R1 and Claude Opus 4.5, FY 2025 for Cowork Plugins). The all-years-averaged variant (FY 2018–2026) is the robustness specification in Table A15. Winsorized at 1st/99th percentiles.
Legal/Regulatory	Supply-side shield. Licensing, permitting, and compliance friction that a new entrant would have to clear to produce and sell the product. Higher = greater regulatory barrier.
Production Moats	Supply-side shield. Proprietary production assets—fabrication capacity, specialized data, installed infrastructure—that a new entrant cannot easily replicate. Higher = greater asset moat.

Continued on next page

Table A1 (continued)

Variable	Definition
Customer Trust	Demand-side shield. Extent to which customers require an established track record, reputation, or institutional backing to accept the product from a new, unbranded provider. Higher = greater trust threshold.
Customer Embeddedness	Demand-side shield. Switching costs that an existing customer would incur to move to a functionally equivalent alternative (data lock-in, deep integrations, installed infrastructure, long-term contracts). Higher = greater lock-in.
AI Input Intensity	Degree to which scaling AI training and inference requires more units of the firm's product, such that the firm captures the additional revenue. Higher = stronger AI-input position.
Shield (composite)	Firm-level variable equal to the simple average of Legal/Regulatory, Production Moats, Customer Trust, and Customer Embeddedness on the 0–10 scale.
Supply-side shield	Firm-level variable equal to the simple average of Legal/Regulatory and Production Moats on the 0–10 scale.
Demand-side shield	Firm-level variable equal to the simple average of Customer Trust and Customer Embeddedness on the 0–10 scale.
Cowork Overlap	Firm-level measure of how closely the firm's products compete with the 10 Cowork Plugins launched January 30, 2026. Each extracted product is scored (0–10) by GPT-4.1 on competitive overlap with the specific plugin capabilities (see Listing A2), and scores are aggregated to firm level using importance weights computed from the firm's most recent pre-event 10-K filing (FY 2025). Plugin definitions are from the Anthropic launch announcement (`anthropic.com/research/cowork-plugins`, accessed February 2026). Winsorized at 1st/99th percentiles.

Continued on next page

Table A1 (continued)

Variable	Definition
ESZT AI exposure	Firm-level workforce exposure to generative AI from Eisfeldt et al. (2025). Measures the share of a firm's workforce tasks (weighted by occupational time) that could be executed more efficiently using generative AI, on a 0–1 scale. Higher values indicate greater scope for firm-value upside from labor-AI substitution and complementarity.
Panel C: Firm-Level Controls	
log(ME)	Natural log of market equity (price × shares outstanding), measured at the most recent fiscal year-end before the event. Winsorized at 1st/99th percentiles.
Tobin's Q	Market value of assets (book assets minus book equity plus market equity) divided by book assets. Winsorized at 1st/99th percentiles.
ROA	Return on assets: operating income before depreciation scaled by book assets. Winsorized at 1st/99th percentiles.
Asset growth	One-year percentage change in total assets. Winsorized at 1st/99th percentiles.
Panel D: Heterogeneity Interactions	
SA index	Size-age financial constraints index from Hadlock and Pierce (2010).
WW index	Financial constraints index from Whited and Wu (2006).
KZ index	Financial constraints index from Kaplan and Zingales (1997).
Panel E: Analyst Forecast Variables	
EPS Forecast	Consensus mean earnings per share for the target fiscal year (December 2026 for December FY-end firms). Source: IBES (pre-event, December 2025), IBES (post-event, March 2026).
Revenue Forecast	Consensus mean revenue forecast in millions of dollars for the target fiscal year. Source: IBES (pre-event), IBES (post-event).
LTG	Long-term EPS growth rate, 3–5 year horizon (%). Source: IBES (pre-event), IBES (post-event).

Continued on next page

Table A1 (continued)

Variable	Definition
Target Price	12-month analyst consensus price target ($). Source: IBES (pre-event), IBES (post-event).

Table A2: GPT-4.1 Rescore Validation, Random 1000-Product Sample

Dimension	N	Spearman ρ	Pearson r	% within ± 1
Digital Deliverability	1,000	0.856	0.869	68.4%
Legal/Regulatory	1,000	0.845	0.852	47.3%
Production Moats	1,000	0.750	0.762	41.8%
Customer Trust	1,000	0.719	0.720	47.9%
Customer Embeddedness	1,000	0.630	0.637	49.0%

Note. Rescores a random 1,000-product sample (drawn from the full product universe with a fixed seed) using GPT-4.1, and compares the scores against the GPT-4o-mini production baselines used in the paper for five PDR dimensions (Digital Deliverability, Legal/Regulatory, Production Moats, Customer Trust, Customer Embeddedness). Spearman ρ is the rank correlation between the two scorers at the product level, Pearson r is the linear correlation, and % within ± 1 is the share of products for which the two scorers assign integer scores within one point of each other on the 0–10 scale.

Table A3: Summary Statistics: Analyst Forecast Data

	N	Mean	SD	Min	P5	P25	Median	P75	P95	Max
Panel A: December 2025 (pre-event, IBES)										
EPS	2,139	3.58	5.83	-5.83	-2.26	0.30	1.95	4.95	14.86	30.19
Revenue (millions)	2,136	6560.69	16110.04	0.00	13.33	313.14	1276.17	4776.88	28674.54	111636.39
LTG (%)	678	14.16	17.93	-20.80	-6.52	6.50	10.92	16.40	40.67	110.80
Target Price ($)	2,130	82.52	117.13	2.50	5.00	16.61	39.00	93.84	325.06	673.43
Panel B: March 2026 (post-event, IBES)										
EPS	2,139	3.52	5.83	-5.38	-2.26	0.23	1.91	4.94	14.57	30.36
Revenue (millions)	2,136	6642.95	16350.99	0.00	13.14	313.80	1291.96	4781.56	29836.95	113780.24
LTG (%)	678	14.78	15.76	-14.50	-2.46	7.50	11.50	16.87	42.56	94.00
Target Price ($)	2,130	85.91	122.59	2.19	5.00	16.25	40.15	98.08	334.47	721.16
Panel C: Scaled Revisions Used as Dependent Variables in Table 7 Panel A										
EPS revision	1,536	-0.03	0.16	-0.80	-0.29	-0.05	-0.00	0.03	0.15	0.41
Revenue revision	2,032	0.00	0.10	-0.38	-0.12	-0.02	0.00	0.02	0.12	0.45
LTG change	678	0.69	9.37	-43.60	-8.86	-1.21	0.00	1.80	14.01	41.80
Target Price revision	2,130	0.03	0.17	-0.45	-0.23	-0.04	0.02	0.10	0.33	0.65

Note: Summary statistics for the analyst consensus forecasts used in Table 7. Panels A and B report the pre-event (IBES, December 2025) and post-event (IBES, March 2026) forecast levels for each outcome. Panel C reports the scaled revisions constructed from these levels and used as the dependent variables in the Cowork Plugins analyst-revision cross-sectional regressions. For EPS, revenue, and target price, the revision is $y = (\text{post} - \text{pre})/\text{pre}$; for LTG (already in percentage points), the change is $y = \text{post} - \text{pre}$. Sample restricted to December fiscal year-end firms with market capitalization above $100M, Compustat total assets above $100M, both pre-event and post-event forecasts available, and a non-missing AI-PDR score. For the proportional revisions, Panel C additionally drops firms with pre-event EPS below $0.50 or pre-event revenue below $10M, so the Panel C EPS and revenue counts fall below the Panel A and B levels. Revenue is in millions of dollars. LTG is long-term earnings growth rate (%). Target price is in dollars. All variables are winsorized at the 1st and 99th percentiles.

Table A4: Examples of High- and Low-AI-PDR Firms

Firm	Top Product	AI-PDR	Shield	Sector
High AI-PDR				
Palantir Technologies Inc	Palantir Foundry	9.15	5.58	Info. Tech.
Adobe Inc	Creative Cloud Pro	9.15	4.67	Info. Tech.
Atlassian Corp	Service Management Solutions	9.15	3.77	Info. Tech.
Reddit, Inc	Community Engagement Platform	9.15	3.32	Comm. Svc.
Roblox Corp	Roblox Client	9.15	3.11	Comm. Svc.
AppLovin Corp	AI Advertising Solutions	9.09	3.87	Info. Tech.
Zscaler, Inc	Zero Trust Network Access	9.08	5.51	Info. Tech.
Cloudflare, Inc	Connectivity Cloud Services	9.00	4.49	Info. Tech.
Datadog, Inc	Infrastructure Monitoring	9.00	4.23	Info. Tech.
Snowflake Inc	Data Engineering Services	8.97	4.80	Info. Tech.
Low AI-PDR				
Howmet Aerospace Inc	Engine Products	0.00	6.00	Industrials
Altria Group, Inc	Smokeable Tobacco Products	0.00	5.53	Cons. Staples
Freeport-McMoRan Inc	Copper Products	0.00	5.22	Materials
Tyson Foods, Inc	Beef Products	0.00	3.97	Cons. Staples
Newmont Corp	Gold Production	0.00	3.80	Materials
Martin Marietta Materials Inc	Aggregates	0.00	3.26	Materials
Southern Copper Corp	Copper Products	0.00	2.97	Materials
Amrize Ltd	Cement	0.00	2.86	Materials
Kraft Heinz Co	Easy Ready Meals	0.00	2.53	Cons. Staples
Smurfit Westrock plc	Consumer Packaging	0.00	2.46	Materials

Note: Selection procedure: firms are ranked by AI-PDR; the top and bottom 50 form the high-PDR and low-PDR pools. Within each pool we list the 10 largest firms by market equity. "Top Product" is the highest-importance product line extracted from the firm's most recent 10-K filing.

Table A5: Correlation of AI-PDR with Firm Characteristics

Variable	N	Pearson r	Spearman ρ	$\hat{\beta}$
ESZT Labor Exposure	2,602	0.384	0.432	6.3489^{***} (3.34)
log(ME)	4,183	0.079	0.042	0.0701^{**} (2.25)
R&D/Assets	2,134	-0.016	0.070	0.3833 (0.34)
Intangibles/Assets	4,088	0.279	0.217	2.1456^{***} (5.00)
PPE/Assets	4,006	-0.348	-0.376	-1.8879^{***} (-3.14)
Profit Margin	3,414	0.039	-0.007	0.0002^{**} (2.03)

Note: Columns 2–3 report Pearson and Spearman correlations between firm-level AI-PDR and each characteristic. Column 4 reports the coefficient from a univariate regression of AI-PDR on each characteristic with GICS 6-digit industry fixed effects, with GICS 6-digit industry-clustered t-statistics in parentheses. All variables enter in their raw units; no variable is standardized by sample standard deviation. ESZT Labor Exposure is the Eisfeldt et al. (2025) generative AI labor exposure measure. All Compustat variables are from the most recent fiscal year. *, **, *** denote significance at 10%, 5%, 1%.

Table A6: Timeline of Anthropic Product Announcements Around the Event Window

Date	Announcement	PS?	Justification
Feb 2025	Claude Code preview	No	CLI coding tool for software developers; no SaaS workflow automation.
May 22, 2025	Claude Code GA	No	Broader availability of same coding tool. No domain-specific automation.
Sep 29, 2025	Claude Sonnet 4.5	No	Model upgrade. No product-market entry.
Nov 24, 2025	Claude Opus 4.5	No	Frontier model upgrade. Improves capability but does not package domain-specific workflow tools. Classified as capability demonstration in our study.
Jan 12, 2026	Cowork preview (Max)	No	General-purpose desktop agent without plugins or domain-specific automation. Max only ($100–200/mo).
Jan 16, 2026	Cowork Pro ($20/mo)	No	Widens access to same general-purpose agent. No plugins or domain-specific functionality.
Jan 30, 2026	**Cowork Plugins**	**Yes**	**10 open-source plugins (sales, legal, finance, marketing, customer support, data analysis, etc.) bundling domain knowledge, connectors, and sub-agents to automate workflows sold as standalone SaaS subscriptions. First entry into identifiable incumbent product markets.**
Feb 5, 2026	Claude Opus 4.6	No	Model upgrade (1M context, agent teams). Postdates event window. No new product-market entry.
Feb 24, 2026	Cowork enterprise update	No	Connectors (Google Workspace, DocuSign, FactSet), private plugin marketplaces, 12 new plugins. Distribution expansion, not new product-substitution event.

Note: PS? = Product Substitution, indicating whether the announcement constitutes direct entry into incumbent product markets with packaged, domain-specific workflow automation. Model releases improve capability (inputs) but do not distribute tools competing with specific SaaS products (outputs). The Cowork preview (Jan 12) and Pro expansion (Jan 16) provide a general-purpose platform, not a product-market entrant. Only the January 30 plugin launch targets identifiable SaaS revenue streams. Sources: TechCrunch, Engadget, Axios (Jan 2026); CNN, CNBC, Fast Company (Feb 2026); Anthropic GitHub (`anthropics/knowledge-work-plugins`).

Table A7: High-PDR Firms Outside Technology and Communications

Firm	Top Product	Is SaaS	AI-PDR	Sector
Amazon Com Inc	E-Commerce Retail	No	6.63	Cons. Disc.
Mercadolibre Inc	Digital Payment Solutions	Yes	7.08	Cons. Disc.
Flutter Entertainment plc	Igaming	Yes	6.53	Cons. Disc.
Maplebear Inc	E-Commerce Solutions	Yes	6.55	Cons. Staples
Visa Inc	Payment Processing Services	No	6.98	Financials
Mastercard Inc	Payment Processing Services	No	7.07	Financials
S&P Global Inc	Data and Analytics Solutions	No	8.50	Financials
Iqvia Holdings Inc	Technology & Analytics Solutions	Yes	6.54	Health Care
Veeva Systems Inc	Industry Cloud Solutions	Yes	8.29	Health Care
Doximity, Inc	Digital Medical Professional Platform	Yes	8.36	Health Care
Automatic Data Processing Inc	Human Capital Management Solutions	Yes	6.42	Industrials
Verisk Analytics, Inc	Underwriting Solutions	Yes	6.94	Industrials
Broadridge Financial Solutions, Inc	Investor Communication Solutions	Yes	6.47	Industrials
Costar Group, Inc	Commercial Real Estate Information and Analyt	No	6.86	Real Estate

Note: Selection procedure: after excluding GICS sectors 45 (Information Technology) and 50 (Communication Services), firms are ranked by AI-PDR; the top 100 form the high-PDR pool. From this pool we display, for each sector, up to the 3 firms with the largest market equity, with sectors ordered alphabetically and firms within each sector ordered by market equity descending. "Top Product" shows one of the firm's highest-importance product lines from its most recent 10-K filing. "Is SaaS" is the firm-level classification: a firm is flagged Yes if any of its top-importance product lines in the most recent year is classified as Software-as-a-Service using OpenAI's GPT-4.1, applying the following definition. *SaaS is a cloud computing service model in which the provider's primary product is application software running on the provider's infrastructure; customers access it over the internet through a thin client (web browser, mobile app, or desktop client) rather than running the application themselves on their own hardware. The same instance serves many customers with updates released centrally, and revenue is recurring (subscription, freemium, pay-per-use, or ad-supported) rather than perpetual licensing.*

Table A8: Robustness: Excluding Technology Firms (Cowork Plugins)

Sample	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
Full	-0.0482^{***}	-0.0646^{***}	-0.0498^{***}	-0.0395^{***}	-0.0335^{***}
	(-8.72)	(-12.38)	(-8.10)	(-6.68)	(-5.52)
	$N = 2,827$	$N = 2,827$	$N = 2,827$	$N = 2,803$	$N = 2,722$
Ex-IT (45)	-0.0436^{***}	-0.0573^{***}	-0.0478^{***}	-0.0381^{***}	-0.0307^{***}
	(-6.00)	(-9.23)	(-8.60)	(-6.83)	(-5.29)
	$N = 2,484$	$N = 2,484$	$N = 2,484$	$N = 2,460$	$N = 2,379$
IT Only	-0.0714^{***}	-0.0714^{***}	-0.0281^{**}	-0.0270^{**}	-0.0315^{**}
	(-23.69)	(-23.69)	(-2.09)	(-1.97)	(-2.09)
	$N = 343$	$N = 343$	$N = 343$	$N = 343$	$N = 343$
Ex-IT/Comm	-0.0423^{***}	-0.0522^{***}	-0.0438^{***}	-0.0363^{***}	-0.0301^{***}
	(-5.20)	(-9.00)	(-8.20)	(-6.90)	(-5.60)
	$N = 2,374$	$N = 2,374$	$N = 2,374$	$N = 2,350$	$N = 2,273$
IT/Comm Only	-0.0673^{***}	-0.0683^{***}	-0.0340^{***}	-0.0246^{**}	-0.0251^{**}
	(-11.20)	(-15.95)	(-3.48)	(-2.44)	(-2.10)
	$N = 453$	$N = 453$	$N = 453$	$N = 453$	$N = 449$

Note: This table re-estimates the Cowork Plugins regression on five subsamples defined by technology-sector inclusion, across five industry fixed-effect granularities. "Full" is the baseline sample. "Ex-IT" excludes GICS sector 45 (Information Technology). "IT Only" restricts to GICS sector 45. "Ex-IT/Comm" excludes both sector 45 and sector 50 (Communication Services, including Interactive Media and Entertainment). "IT/Comm Only" restricts to those two sectors. All regressions use CAPM-adjusted CARs over the $[-1, +10]$ window with controls (log(ME), Tobin's Q, ROA, asset growth, each winsorized at 1/99). CARs are also winsorized at 1/99. AI-PDR is standardized to unit variance within the estimation sample. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses across all columns. *, **, *** denote significance at 10%, 5%, 1%.

Table A9: Examples of SaaS and Non-SaaS Firms

Firm	Top Product
Panel A. SaaS firms (random)	
Braze, Inc.	Customer Engagement Platform
PROS Holdings, Inc.	AI-Driven Pricing Solutions
Appian Corp	Process Automation Platform
Paylocity Holding Corp	Human Capital Management Software
Commvault Systems Inc	Operational Recovery
Panel B. Non-SaaS firms (random)	
Global Partners Lp	Gasoline Distribution
Grid Dynamics Holdings, Inc.	Enterprise AI Services
OneMain Holdings, Inc.	Personal Loans
BlueLinx Holdings Inc.	Specialty Building Products
PENTAIR plc	Fluid Treatment and Pump Products
Panel C. Non-SaaS firms with the highest AI-PDR	
Shutterstock, Inc.	Content Licensing
Getty Images Holdings, Inc.	Premium Visual Content Solutions
Exodus Movement, Inc.	Self-Custodial Wallet Services
Circle Internet Group, Inc.	Stablecoin Services
MongoDB, Inc.	Developer Data Platform

Note: A random sample of five firms classified as Software-as-a-Service (Panel A) and five classified as non-SaaS (Panel B), plus the five non-SaaS firms with the highest firm-level AI-PDR (Panel C). A firm is flagged SaaS if any of its top-importance product lines in its most recent year is classified as Software-as-a-Service using OpenAI's GPT-4.1, applying the following definition. *SaaS is a cloud computing service model in which the provider's primary product is application software running on the provider's infrastructure; customers access it over the internet through a thin client (web browser, mobile app, or desktop client) rather than running the application themselves on their own hardware. The same instance serves many customers with updates released centrally, and revenue is recurring (subscription, freemium, pay-per-use, or ad-supported) rather than perpetual licensing.* The displayed product is one of the firm's top-importance product lines.

Table A10: Event Study by SaaS Classification: CARs on AI-PDR (Cowork Plugins)

	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
SaaS	-0.0289^{***}	-0.0442^{***}	-0.0262^{**}	-0.0201	-0.0218
	(-3.34)	(-4.26)	(-2.14)	(-1.38)	(-1.49)
	$N = 393$	$N = 393$	$N = 393$	$N = 352$	$N = 308$
Non-SaaS	-0.0323^{***}	-0.0458^{***}	-0.0358^{***}	-0.0278^{***}	-0.0223^{***}
	(-4.37)	(-7.57)	(-6.32)	(-5.10)	(-4.40)
	$N = 2,434$	$N = 2,434$	$N = 2,434$	$N = 2,404$	$N = 2,310$

Note: Re-estimates the Cowork Plugins regression on two subsamples partitioned by firm-level SaaS classification, across five industry fixed-effect granularities. The first row group restricts to firms classified as SaaS; the second drops those firms. A firm is flagged SaaS if any of its top-importance product lines in its most recent year is classified as Software-as-a-Service using OpenAI's GPT-4.1, applying the following definition. *SaaS is a cloud computing service model in which the provider's primary product is application software running on the provider's infrastructure; customers access it over the internet through a thin client (web browser, mobile app, or desktop client) rather than running the application themselves on their own hardware. The same instance serves many customers with updates released centrally, and revenue is recurring (subscription, freemium, pay-per-use, or ad-supported) rather than perpetual licensing.* CARs are CAPM-adjusted with firm-specific betas on $[-252, -11]$, regressed on standardized AI-PDR (Digital Deliverability), with log(ME), Tobin's Q, ROA, and asset growth as controls (each winsorized at 1/99). CARs are also winsorized at 1/99. Industry fixed effects vary across columns as indicated. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses across all columns. *, **, *** denote significance at 10%, 5%, 1%.

Table A11: Robustness: Excluding Firms Reporting Earnings in the Event Window (Cowork Plugins)

	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
Full Sample	-0.0482^{***}	-0.0646^{***}	-0.0498^{***}	-0.0395^{***}	-0.0335^{***}
	(-8.72)	(-12.38)	(-8.10)	(-6.68)	(-5.52)
	$N = 2,827$	$N = 2,827$	$N = 2,827$	$N = 2,803$	$N = 2,722$
Excluding Earners	-0.0456^{***}	-0.0595^{***}	-0.0479^{***}	-0.0382^{***}	-0.0282^{***}
	(-7.69)	(-10.14)	(-7.49)	(-5.29)	(-3.61)
	$N = 1,942$	$N = 1,942$	$N = 1,942$	$N = 1,906$	$N = 1,820$

Note: This table tests, for the Cowork Plugins event over the $[-1, +10]$ trading-day window, whether the AI-PDR coefficient is confounded by earnings announcements landing inside that window. "Full Sample" is the baseline regression reported as Cowork Plugins in Table 3. "Excluding Earners" drops firms whose Compustat `rdq` (the calendar-date earnings press release ahead of the 10-Q filing) for a fiscal quarter falls within the $[-1, +10]$ trading-day window around January 30, 2026, the Cowork Plugins event date. For most firms this is the Q4 2025 earnings announcement. All regressions use CAPM-adjusted CARs with controls (log(ME), Tobin's Q, ROA, asset growth, each winsorized at 1/99). CARs are also winsorized at 1/99. AI-PDR is standardized to unit variance within the estimation sample. Industry fixed effects vary across columns as indicated. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses across all columns. *, **, *** denote significance at 10%, 5%, 1%.

Table A12: Moderators of the AI-PDR Effect: Size and Financial Constraints (Cowork Plugins)

Interaction	AI-PDR	Interaction	N	R^2
AI-PDR × log(ME)	-0.0393^{***}	-0.0018	2,803	0.3331
	(-6.62)	(-1.48)		
AI-PDR × R&D/Assets	-0.0509^{***}	-0.0373	1,400	0.3200
	(-6.40)	(-0.76)		
AI-PDR × SA	-0.0387^{***}	0.0173	2,802	0.3341
	(-6.47)	(1.21)		
AI-PDR × WW	-0.0453^{***}	0.0189	2,120	0.3602
	(-6.98)	(0.61)		
AI-PDR × KZ	-0.0430^{***}	-0.0000	2,372	0.3360
	(-6.55)	(-0.68)		

Note: Each row adds an interaction between AI-PDR and a demeaned firm characteristic to the baseline specification for the Cowork Plugins event ($[-1, +10]$). AI-PDR is standardized to unit variance and demeaned within the estimation sample; the AI-PDR main effect is evaluated at the mean of the interacted variable. R&D intensity is R&D expenditure (Compustat `xrd`) scaled by total assets; firms with missing R&D are dropped. SA is the Hadlock and Pierce (2010) size-age index. WW is the Whited and Wu (2006) index. KZ is the Kaplan and Zingales (1997) index. Higher values of SA, WW, and KZ indicate greater financial constraints. All specifications include GICS 6-digit industry fixed effects and controls (log(ME), Tobin's Q, ROA, asset growth). The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M, so the size-interaction term reflects within-sample variation rather than absolute smallness. GICS 6-digit industry-clustered t-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A13: Shield Moderation of AI-PDR: Individual Shield Sub-Dimensions (Cowork Plugins)

	(1) Legal	(2) Moats	(3) Trust	(4) Embeddedness
AI-PDR	-0.0400^{***}	-0.0375^{***}	-0.0376^{***}	-0.0384^{***}
	(-5.81)	(-5.91)	(-5.88)	(-6.47)
Shield	-0.0197^{**}	-0.0055	-0.0050	-0.0042
	(-2.56)	(-0.89)	(-0.79)	(-0.97)
AI-PDR × Shield	0.0078	0.0143^{***}	0.0091^{**}	0.0055^{*}
	(1.53)	(3.06)	(2.17)	(1.84)
N	2,803	2,803	2,803	2,803
R^2	0.3376	0.3377	0.3349	0.3339

Note: Each column reports a separate cross-sectional OLS regression of CAPM-adjusted CARs on the Cowork Plugins sample over the $[-1, +10]$ event window. Each column enters one Shield sub-dimension and its interaction with AI-PDR in isolation, with the sub-dimension shown in the column header. The sub-dimensions are Legal (Legal/Regulatory), Moats (Production Moats), Trust (Customer Trust), and Embeddedness (Customer Embeddedness). AI-PDR and each Shield sub-dimension are standardized to unit variance within the estimation sample and demeaned; interactions are products of the standardized-demeaned regressors. Every regression includes log(ME), Tobin's Q, ROA, asset growth (each winsorized at 1/99), and GICS 6-digit industry fixed effects. CARs are also winsorized at 1/99. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A14: Robustness: Controlling for Interest-Rate Sensitivity (Cowork Plugins)

	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
Baseline	-0.0482^{***}	-0.0646^{***}	-0.0498^{***}	-0.0395^{***}	-0.0335^{***}
	(-8.72)	(-12.38)	(-8.10)	(-6.68)	(-5.52)
	$N = 2,827$	$N = 2,827$	$N = 2,827$	$N = 2,803$	$N = 2,722$
+ Rate Control	-0.0483^{***}	-0.0646^{***}	-0.0492^{***}	-0.0392^{***}	-0.0332^{***}
	(-8.72)	(-12.26)	(-7.92)	(-6.62)	(-5.48)
	$N = 2,827$	$N = 2,827$	$N = 2,827$	$N = 2,803$	$N = 2,722$

Note: Dependent variable is the CAPM-adjusted CAR over $[-1, +10]$ around Cowork Plugins. "Baseline" is the standard specification reported as Cowork Plugins in Table 3. "+ Rate Control" additionally controls for interest-rate sensitivity. Rate β is the coefficient on Δy_{10} from a regression of daily excess returns on the market excess return and the daily change in the 10-year Treasury yield over $[-252, -11]$ trading days before the event, capturing duration exposure. AI-PDR and Rate β are each standardized to unit variance and demeaned within the estimation sample. Rate β is winsorized at 1/99 before standardization. Controls (log(ME), Tobin's Q, ROA, asset growth) enter in raw units. CARs are winsorized at 1/99. Industry fixed effects vary across columns as indicated. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses across all columns. *, **, *** denote significance at 10%, 5%, 1%.

Table A15: Robustness: AI-PDR Averaged Across 2021–2026 Filings, across Industry-FE Granularities

	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
ChatGPT Launch	0.0060	0.0086**	0.0071*	0.0076	0.0075
	(1.25)	(2.35)	(1.65)	(1.39)	(1.17)
	$N = 3,236$	$N = 3,236$	$N = 3,236$	$N = 3,216$	$N = 3,142$
DeepSeek R1	0.0147***	0.0177***	0.0110**	0.0095	0.0143**
	(4.90)	(4.42)	(2.04)	(1.52)	(1.98)
	$N = 2,972$	$N = 2,972$	$N = 2,972$	$N = 2,944$	$N = 2,879$
Claude Opus 4.5	0.0002	−0.0085**	−0.0085**	−0.0050	−0.0044
	(0.09)	(−2.53)	(−2.49)	(−1.18)	(−0.91)
	$N = 2,869$	$N = 2,869$	$N = 2,869$	$N = 2,841$	$N = 2,768$
Cowork Plugins	−0.0498***	−0.0687***	−0.0559***	−0.0469***	−0.0414***
	(−8.87)	(−12.92)	(−8.50)	(−6.77)	(−5.87)
	$N = 2,847$	$N = 2,847$	$N = 2,847$	$N = 2,823$	$N = 2,746$

Note: Same specification as Table 3 but replacing the event-year-specific AI-PDR with the within-firm mean of AI-PDR over the firm's 10-K filings from fiscal years 2021–2026 (discarding earlier years to keep the measure close to the firm's recent product mix). The same averaged AI-PDR is used at all four events. AI-PDR is standardized to unit variance within the estimation sample and demeaned. Controls (not reported): log(ME), Tobin's Q, ROA, asset growth (each winsorized at 1/99). CARs are also winsorized at 1/99. Industry fixed effects vary across columns as indicated. The estimation universe is firms with market capitalization above $100M and Compustat total assets above $100M. GICS 6-digit industry-clustered t-statistics in parentheses across all columns. *, **, *** denote significance at 10%, 5%, 1%.

Table A16: Event Study: Cross-Sectional Regressions of CARs on AI-PDR and ESZT Labor Exposure

Event	AI-PDR	ESZT	N	R^2
ChatGPT Launch	0.0050	0.0168***	2,226	0.3502
	(0.86)	(2.99)		
DeepSeek R1	0.0005	−0.0022	2,052	0.4980
	(0.08)	(-0.34)		
Claude Opus 4.5	−0.0069	−0.0089*	1,946	0.4144
	(-1.03)	(-1.69)		
Cowork Plugins	−0.0328***	−0.0146*	1,914	0.5771
	(-2.69)	(-1.83)		

Note: This table re-estimates the joint regression of CARs on AI-PDR and Eisfeldt et al. (2025) labor exposure (ESZT) in the ESZT-style specification. Dependent variable is the market-adjusted CAR over $[-1, +10]$ trading days ($r_{i,t} - r_{m,t}$, cumulated). Estimator is weighted least squares, weighting each firm by its market capitalization on the day before the event (Nov 29, 2022 for ChatGPT, and the corresponding day -1 for each other event), following ESZT's value-weighted convention. The estimation universe is the full Compustat-CRSP cross section with non-missing ESZT, without the \$100M market-cap and \$100M total-assets filter used elsewhere in the paper. Industry fixed effects are at the NAICS 3-digit level, matching ESZT. Firm controls are log(ME), Tobin's Q, ROA, and asset growth (all winsorized at 1/99). CARs are also winsorized at 1/99. ESZT is the published Eisfeldt et al. (2025) firm-level labor exposure. Both regressors are standardized to unit variance and demeaned within the estimation sample, so reported coefficients are in p.p. of CAR per one-SD increase. HC1-robust t-statistics in parentheses (following ESZT's reporting). *, **, *** denote significance at 10%, 5%, 1%.

Table A17: Event Study: FF5-Adjusted CARs on AI-PDR across Industry-FE Granularities

	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
ChatGPT Launch	0.0021	0.0027	0.0022	0.0024	0.0023
	(0.58)	(0.72)	(0.54)	(0.49)	(0.41)
	$N = 2,941$	$N = 2,941$	$N = 2,941$	$N = 2,918$	$N = 2,838$
DeepSeek R1	0.0079***	0.0112***	0.0075	0.0060	0.0119
	(2.97)	(2.64)	(1.25)	(0.86)	(1.51)
	$N = 2,915$	$N = 2,915$	$N = 2,915$	$N = 2,888$	$N = 2,821$
Claude Opus 4.5	−0.0014	−0.0103***	−0.0126***	−0.0089**	−0.0072*
	(-0.61)	(-3.17)	(-3.80)	(-2.17)	(-1.68)
	$N = 2,788$	$N = 2,788$	$N = 2,788$	$N = 2,762$	$N = 2,683$
Cowork Plugins	−0.0379***	−0.0586***	−0.0405***	−0.0330***	−0.0269***
	(-7.34)	(-9.13)	(-6.22)	(-5.71)	(-4.53)
	$N = 2,818$	$N = 2,818$	$N = 2,818$	$N = 2,794$	$N = 2,714$

Note: Same specification as Table 3 but using Fama–French five-factor adjusted CARs. Abnormal returns are computed relative to a five-factor model (MKTRF, SMB, HML, RMW, CMA) estimated over $[-252, -11]$ trading days before each event. Event window $[-1, +10]$. AI-PDR is standardized to unit variance within the estimation sample and demeaned. Controls (log(ME), Tobin's Q, ROA, asset growth) are winsorized at the 1st and 99th percentiles. CARs are also winsorized at 1/99. Industry fixed effects vary across columns as indicated. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses across all columns. *, **, *** denote significance at 10%, 5%, 1%.

Table A18: FF5 Window Robustness: CAR Regressions across Event Windows

	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
ChatGPT Launch	0.0015	0.0024	0.0045	0.0068
	(0.62)	(0.49)	(0.78)	(0.83)
DeepSeek R1	0.0018	0.0060	0.0128*	0.0004
	(0.71)	(0.86)	(1.76)	(0.04)
Claude Opus 4.5	−0.0076***	−0.0089**	−0.0076	−0.0074
	(−2.67)	(−2.17)	(−1.47)	(−1.21)
Cowork Plugins	−0.0103***	−0.0330***	−0.0287***	−0.0174**
	(−3.27)	(−5.71)	(−3.45)	(−2.01)

Note: Same specification as Table 4 but using Fama–French five-factor adjusted CARs. Abnormal returns are computed relative to a five-factor model (MKTRF, SMB, HML, RMW, CMA) estimated over $[-252, -11]$ trading days before each event. AI-PDR is standardized to unit variance within the estimation sample and demeaned. All regressions include GICS 6-digit industry fixed effects and firm controls (log(ME), Tobin's Q, ROA, asset growth) all winsorized at 1/99. CARs are also winsorized at 1/99. The estimation universe is firms with market capitalization above $100M and Compustat total assets above $100M. Significance stars based on GICS 6-digit industry-clustered t-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A19: Robustness to Additional Controls: Cowork Plugins

Specification	AI-PDR	N	R^2
Baseline	-0.0395^{***}	2,803	0.3325
	(-6.68)		
+ Firm characteristics	-0.0372^{***}	2,277	0.3502
	(-6.37)		
+ Leverage, cash	-0.0375^{***}	2,258	0.3531
	(-6.11)		
+ Momentum, volatility	-0.0343^{***}	2,258	0.3931
	(-5.86)		
+ Firm age	-0.0343^{***}	2,258	0.3931
	(-5.83)		
+ R&D intensity	-0.0441^{***}	1,217	0.3866
	(-4.39)		

Note: Each row reports the AI-PDR coefficient from a cross-sectional regression of $[-1, +10]$ CAPM-adjusted CARs around Cowork Plugins (January 30, 2026) on AI-PDR and progressively richer control sets. AI-PDR is standardized to unit variance within the estimation sample and demeaned; the additional controls enter in their raw units. All specifications include GICS 6-digit industry fixed effects and baseline controls (log(ME), Tobin's Q, ROA, asset growth). "Firm characteristics" adds log(Assets), asset tangibility (PPE/Assets), intangibles ratio, and profit margin. Each subsequent row adds controls cumulatively. "Leverage, cash" adds long-term debt ratio and cash/assets. "Momentum, volatility" adds prior 12-month return and daily return volatility. "Firm age" adds log(Age) from first Compustat appearance. "R&D intensity" adds R&D/Assets; firms with missing R&D are dropped. The estimation universe is firms with market capitalization above $100M and Compustat total assets above $100M. GICS 6-digit industry-clustered t-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A20: Event Study with GPT-3.5 AI-PDR (Chronologically Consistent Scorer)

Event	AI-PDR	N	R^2
ChatGPT Launch	0.0081**	2,922	0.1039
	(2.33)		
DeepSeek R1	0.0037	2,894	0.1353
	(1.10)		
Claude Opus 4.5	−0.0023	2,769	0.1103
	(-0.74)		
Cowork Plugins	−0.0321***	2,803	0.3291
	(-5.45)		

Note. GPT-3.5 robustness mirror of Table 2. AI-PDR is re-estimated using GPT-3.5-Turbo-0125 (training cutoff September 2021, predating the ChatGPT launch) in place of the baseline GPT-4o mini, with all other choices held fixed. Each row reports a cross-sectional OLS regression of CAPM-adjusted CARs over $[-1, +10]$ on standardized AI-PDR (Digital Deliverability), with GICS 6-digit industry fixed effects and controls (log(ME), Tobin's Q, ROA, asset growth, all winsorized at the 1st and 99th percentiles). CARs are also winsorized at 1/99. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses. The firm-level correlation between the all-years averaged GPT-3.5 and GPT-4o mini AI-PDR is $\rho = 0.94$. *, **, *** denote significance at 10%, 5%, 1%.

Table A21: Example Products by AI Input Taxonomy Category

Category	Product	Firm
Panel A. Tier 1 (AI core)		
Data center AI accelerators (training and inference chips for hyperscale)	Data Center Solutions AI Computing Solutions Custom ASICs	Nvidia Corp Advanced Micro Devices Inc Intel Corp
High-bandwidth memory (HBM)	High-Bandwidth Memory (HBM) High Bandwidth Memory (HBM) Components	Micron Technology Inc Netlist Inc
AI-class networking, including InfiniBand, NVLink, AI-fabric Ethernet, and AI optical interconnects	Networking Connectivity Solutions Networking Solutions AI Center Networking Solutions	Broadcom Inc. Advanced Micro Devices Inc Arista Networks, Inc.
AI server systems and rack-scale designs (DGX/HGX class)	Accelerated Computing Platforms Data Center Solutions Server Solutions	Nvidia Corp Qualcomm Inc Hewlett Packard Enterprise Co
GPU-dedicated AI cloud whose primary business is selling AI compute	DGX Cloud Services CoreWeave Cloud Platform AI Cloud Services	Nvidia Corp CoreWeave, Inc. IREN Ltd
GPU-accelerated programming libraries, ML runtimes, MLOps platforms, model training and fine-tuning platforms, and inference serving systems	AI Solutions and Software AI Development Tools Artificial Intelligence Platform (AIP)	Nvidia Corp Alphabet Inc. Palantir Technologies Inc.
Foundation models and inference APIs	Generative and Agentic AI Solutions AI Services Adobe Firefly Solutions	Nvidia Corp Alphabet Inc. Adobe Inc.
Panel B. Tier 2 (AI-adjacent infrastructure)		

Category	Product	Firm
High-performance AI storage and storage-management software, including parallel file systems and AI-marketed flash arrays	Unistore Services Lyve Edge-to-Cloud Platform AI-Driven Storage Solutions	Snowflake Inc. Seagate Technology Holdings plc Hewlett Packard Enterprise Co
AI-specific data center cooling, including liquid, immersion, and direct-to-chip	Liquid Cooling Solutions Immersion Cooling Technology Data Center Cooling Solutions	Super Micro Computer, Inc. Bitmine Immersion Technologies, Inc. Modine Manufacturing Co
Data center power distribution and backup, including UPS, switchgear, and fuel cell modules	Shared Backup Power Solutions Power Management Solutions Embedded and Critical Power Solutions	American Tower Corp Vertiv Holdings Co Flex Ltd.
Data center colocation and interconnection services	Co-location Services Data Center Leasing Colocation Services	Cme Group Inc. American Tower Corp Nasdaq, Inc.
Data center construction and engineering services	Integrated Modular Solutions Development and Construction Services Rack Level Services	Vertiv Holdings Co Digital Realty Trust, Inc. Super Micro Computer, Inc.
Hyperscale cloud	Cloud Computing Services Cloud Infrastructure Services Cloud Computing Services	Nvidia Corp Microsoft Corp Amazon Com Inc
Edge AI silicon and modules for consumer devices, automotive, and robotics	Automotive Solutions Embedded Computing Products Integrated Circuit Products	Nvidia Corp Advanced Micro Devices Inc Qualcomm Inc
Panel C. Tier 3 (semiconductor manufacturing)		
EDA software and chip design IP licensing	Custom Silicon Solutions Custom Silicon Solutions Factory Automation Software	Broadcom Inc. Advanced Micro Devices Inc Applied Materials Inc

Category	Product	Firm
Fab and packaging materials and consumables (inputs that are consumed during fabrication), including silicon wafers, photomasks, substrates, chemicals, and gases	Clean Services Wafer Manufacturing Solutions Electronics Materials Solutions	Lam Research Corp Kla Corp 3M Co
Semiconductor capital equipment, including lithography, etch, deposition, metrology, test, gas/fluid delivery subsystems, and process tool components	Wafer Fabrication Equipment Semiconductor Fabrication Equipment Metrology Solutions	Lam Research Corp Applied Materials Inc Kla Tencor Corp
Contract wafer fabrication and foundry services (chip manufacturing for other firms' designs)	Semiconductor Solutions Foundry Services Custom ASICs	Broadcom Inc. Intel Corp Marvell Technology, Inc.
Advanced chip packaging services, including CoWoS, FOWLP, and 2.5D/3D stacking	Advanced Packaging Solutions Advanced Packaging Systems Multi-Chip Modules (MCM)	Lam Research Corp Applied Materials Inc MACOM Technology Solutions Holdings, Inc.
Panel D. Tier 4 (general inputs)		
Electric power generation and grid operations	Energy Generation Systems Hydrogen Production Energy Generation and Distribution	Tesla, Inc. Exxon Mobil Corp Berkshire Hathaway Inc
General industrial cooling and power equipment	Power Conversion Energy Solutions HVAC and Refrigeration Solutions	General Electric Co Caterpillar Inc United Technologies Corp
Commodity memory and storage products, including DDR DRAM, NAND flash, NVMe SSDs, and hard drives	Servers and Storage System Solutions Low-Power DRAM (LPDDR) Storage Controllers	Broadcom Inc. Micron Technology Inc Marvell Technology, Inc.

Enterprise and telecom networking and optics	Networking Solutions	Nvidia Corp
	Fibre Channel Products	Broadcom Inc.
	Networking Solutions	Cisco Systems, Inc.

Note: The table lists example products from each of the 23 AI-input taxonomy categories (Listing A3), grouped into the four tiers (Panels A–D). Each (firm, product) pair is kept once across the full panel of 10-K filings, using the most recent year's classification when the same product appears in multiple years. Within each category, firms are ranked by their most recent monthly CRSP market capitalization and the top three are shown, with one product per firm (the firm's first product classified into the category by the GPT-4.1 taxonomy classifier). Categories with fewer than three firms classified into them show all available firms.

Table A22: AI-PDR and AI Input Intensity across Event Windows (OLS vs WLS weighted by market capitalization, GICS 6-digit and GICS 2-digit fixed effects)

Panel A — ChatGPT Launch, OLS, GICS-6 FEs				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	0.0045*	0.0071	0.0065	0.0129
	(1.72)	(1.34)	(1.05)	(1.53)
AI Input	0.0049	0.0014	−0.0367	−0.0215
	(0.35)	(0.09)	(-1.56)	(-0.87)

Panel B — ChatGPT Launch, WLS, GICS-6 FEs				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	0.0065*	0.0155**	0.0156**	0.0164**
	(1.76)	(2.39)	(2.44)	(2.42)
AI Input	−0.0042	0.0034	−0.0204	0.0312
	(-0.38)	(0.16)	(-1.31)	(0.87)

Panel C — ChatGPT Launch, OLS, GICS-2 FEs				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	0.0036**	0.0073*	0.0083*	0.0070
	(2.25)	(1.95)	(1.86)	(1.03)
AI Input	0.0115	0.0017	−0.0366	−0.0083
	(1.14)	(0.10)	(-1.62)	(-0.34)

Panel D — ChatGPT Launch, WLS, GICS-2 FEs				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	0.0006	0.0083**	0.0114**	0.0096
	(0.23)	(2.00)	(1.98)	(1.32)
AI Input	0.0068	0.0088	−0.0237	0.0402
	(0.57)	(0.35)	(-1.10)	(1.12)

Panel E — DeepSeek R1, OLS, GICS-6 FEs				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	0.0027	0.0094*	0.0143**	0.0009
	(1.05)	(1.74)	(2.00)	(0.11)
AI Input	−0.0177	−0.0875***	−0.0418*	−0.0584*
	(-1.10)	(-4.28)	(-1.66)	(-1.80)

Panel F — DeepSeek R1, WLS, GICS-6 FEs				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	−0.0002	0.0020	0.0135*	0.0068
	(-0.04)	(0.29)	(1.65)	(0.52)
AI Input	0.0254	−0.0206	−0.0362	−0.0683*
	(1.25)	(-0.59)	(-0.98)	(-1.72)

Panel G — DeepSeek R1, OLS, GICS-2 FEs				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	0.0039**	0.0153***	0.0182***	0.0046
	(2.03)	(3.86)	(3.13)	(0.65)
AI Input	−0.0098	−0.0967***	−0.0548**	−0.1067***
	(-0.58)	(-4.96)	(-2.08)	(-3.49)

Panel H — DeepSeek R1, WLS, GICS-2 FEs				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	0.0052	0.0138***	0.0100*	0.0002
	(1.51)	(2.79)	(1.89)	(0.02)
AI Input	0.0062	−0.0702*	−0.0764**	−0.1171***
	(0.34)	(-1.89)	(-2.02)	(-2.96)

Table A22: AI-PDR and AI Input Intensity across Event Windows (OLS vs WLS weighted by market capitalization, GICS 6-digit and GICS 2-digit fixed effects) (continued)

Panel I — Claude Opus 4.5, OLS, GICS-6 FEs					*Panel J — Claude Opus 4.5, WLS, GICS-6 FEs*				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]		[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	−0.0060*	−0.0050	−0.0074	−0.0076	AI-PDR	−0.0011	−0.0036	−0.0101	−0.0135
	(-1.72)	(-1.08)	(-1.52)	(-1.26)		(-0.27)	(-0.38)	(-0.87)	(-1.00)
AI Input	−0.0063	0.0132	−0.0242	−0.0075	AI Input	−0.0113	−0.0327	−0.0223	−0.0194
	(-0.54)	(0.61)	(-0.66)	(-0.16)		(-0.75)	(-1.17)	(-0.86)	(-0.41)
Panel K — Claude Opus 4.5, OLS, GICS-2 FEs					*Panel L — Claude Opus 4.5, WLS, GICS-2 FEs*				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]		[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	−0.0061**	−0.0074**	−0.0081**	−0.0081	AI-PDR	−0.0033	−0.0006	−0.0062	−0.0056
	(-2.45)	(-2.26)	(-2.52)	(-1.55)		(-0.92)	(-0.10)	(-0.95)	(-0.60)
AI Input	−0.0024	0.0103	−0.0557*	−0.0309	AI Input	0.0071	−0.0128	−0.0270	−0.0021
	(-0.25)	(0.52)	(-1.81)	(-0.63)		(0.53)	(-0.66)	(-1.19)	(-0.05)
Panel M — Cowork Plugins, OLS, GICS-6 FEs					*Panel N — Cowork Plugins, WLS, GICS-6 FEs*				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]		[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	−0.0194***	−0.0401***	−0.0319***	−0.0186**	AI-PDR	−0.0230***	−0.0482***	−0.0293**	−0.0240**
	(-5.81)	(-6.95)	(-4.06)	(-2.12)		(-3.95)	(-5.01)	(-2.45)	(-2.17)
AI Input	−0.0635**	0.0302	0.0482	0.1307***	AI Input	−0.0469	0.0400	0.0400	0.0968***
	(-2.11)	(0.97)	(1.15)	(3.49)		(-0.98)	(0.93)	(0.94)	(3.40)
Panel O — Cowork Plugins, OLS, GICS-2 FEs					*Panel P — Cowork Plugins, WLS, GICS-2 FEs*				
	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]		[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
AI-PDR	−0.0340***	−0.0643***	−0.0546***	−0.0353***	AI-PDR	−0.0341***	−0.0621***	−0.0470***	−0.0299***
	(-14.41)	(-11.68)	(-8.74)	(-5.91)		(-7.02)	(-8.38)	(-5.26)	(-3.61)
AI Input	−0.0599**	0.0300	0.0176	0.1095***	AI Input	−0.0328	0.0527	0.0499	0.0973**
	(-2.13)	(1.14)	(0.59)	(3.59)		(-0.78)	(1.17)	(1.09)	(2.49)

Note: Each of the sixteen panels reports, for a single (event, fixed-effects level, estimator) combination, two coefficients (AI-PDR and AI Input) from a cross-sectional regression of CAPM-adjusted CARs on AI-PDR (Digital Deliverability) and AI Input Intensity, over four event windows. AI-PDR is standardized to unit variance and demeaned within each sample. AI Input Intensity enters in raw share units so its coefficient is the change in CAR per unit change in the share of a firm's products classified into any AI-input category (Tiers 1–4) by the GPT-4.1 taxonomy classifier (Listing A3). Both regressors use firm values from the fiscal year aligned with the event date. *OLS* panels use unweighted ordinary least squares. *WLS* panels weight each observation by firm market capitalization. *GICS-6 FEs* denotes GICS 6-digit industry fixed effects with industries having fewer than five firms dropped, *GICS-2 FEs* denotes GICS 2-digit sector fixed effects. Other controls (not reported), log(ME), Tobin's Q, ROA, asset growth, each winsorized at 1/99. CARs winsorized at 1/99. The estimation universe is firms with market capitalization and Compustat total assets both above \$100M. GICS 6-digit industry-clustered t-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.