

Generative AI Product Displacement Risk

Brett Green, Renping Li, and Dimitris Papanikolaou*

ABSTRACT

Generative AI can raise firm value by lowering labor costs but destroy it by displacing firm products. We introduce Generative AI Product Displacement Risk (AI-PDR), constructed as a measure of how digitally deliverable a firm's products are. At the January 2026 launch of Cowork Plugins, high-AI-PDR firms exhibit large negative cumulative abnormal returns. Top-quintile firms underperform the bottom quintile by 12 percentage points. Competitive shields to displacement significantly attenuate the repricing, with supply-side shields delivering the larger moderation. Analysts revise down target prices, near-term earnings forecasts, and long-term growth forecasts for high-AI-PDR firms. In the second quarter of 2026, revenue growth slows for high-AI-PDR firms and gross margins fall. Over longer horizons, SaaS firms recover to their pre-Cowork-Plugins level, while non-SaaS high-AI-PDR firms stay repriced.

*Green is with the Olin Business School, Washington University in St. Louis, Li is with the Freeman School of Business, Tulane University, and Papanikolaou is with the Kellogg School of Management, Northwestern University, and NBER. This is a preliminary draft. Results and interpretations are subject to revision. We thank Cyrus Mevorach, Donghwa Shin, Sheng-Jun Xu, and seminar and conference participants at NFA 2026, the AI in Finance Conference (Montréal), the AI in Finance Conference (Dresden), and the Next-Gen AI and Autonomous Finance (NGAI) Workshop for helpful comments.

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, Kogan, and Panageas, 2012a](#); [Gârleanu, Panageas, and Yu, 2012b](#); [Kogan, Papanikolaou, Seru, and Stoffman, 2017](#)). Generative AI (Gen-AI) may drive this displacement at unprecedented scale because it produces digital output at near-zero marginal cost, creating close substitutes for many digitally deliverable products and services.

A growing literature examines how Gen-AI displaces labor ([Acemoglu, Autor, Hazell, and Restrepo, 2022](#); [Eloundou, Manning, Mishkin, and Rock, 2024](#); [Hampole, Papanikolaou, Schmidt, and Seegmiller, 2025](#); [Webb, 2020](#)) and how Gen-AI adoption raises firm value ([Babina, Fedyk, He, and Hodson, 2024](#); [Eisfeldt, Schubert, Zhang, and Taska, 2025](#)), but this work studies which workers can be replaced or which firms adopt Gen-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 Gen-AI may benefit through lower costs, but a firm whose product is replaceable faces revenue erosion and margin compression regardless of how efficiently it operates internally. Which firms are most vulnerable to this product-level threat, and whether and when the equity market prices this risk, remain open questions.

We introduce *Gen-AI Product Displacement Risk* (AI-PDR) to measure this threat. Using a large language model, we extract individual product and service lines from the business description of each firm’s 10-K filing. We then 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 labor-exposure measure of [Eisfeldt et al. \(2025\)](#), hereafter ESZT. Their index captures which firms’ *workers* are replaceable by Gen-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, Consumer Staples, and Energy 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, Moody’s, Morningstar, and LegalZoom, IT consulting firms such as

Accenture and Cognizant, online education firms such as Duolingo and Chegg, perpetual-license software vendors such as PTC and Progress Software, video-game publishers such as Take-Two Interactive and Electronic Arts, and freelancer marketplaces that intermediate digitally deliverable knowledge work such as Upwork and Fiverr. These firms have high gross margins and growth rates, making their valuations susceptible to margin compression and growth deceleration induced by competition.

We study whether and when AI-PDR is priced by examining equity market reactions to four major AI releases, namely ChatGPT (November 2022), DeepSeek R1 (January 2025), Claude Opus 4.5 (November 2025), and Cowork Plugins (January 2026), hereafter the four main-sample releases. The first three are generative AI model releases. The fourth, Cowork Plugins, launched by Anthropic on January 30, 2026, is not a generative AI model, but an AI product that performs a specified set of work tasks. It offered 11 open-source agentic AI plugins covering productivity, sales, customer support, legal, finance, data analysis, and other enterprise functions, all accessible through a \$20-per-month subscription.

The distinguishing feature of Cowork Plugins is the unmatched breadth and depth of its overlap with the functionality of incumbent Compustat firms' products. We compare Cowork Plugins to all other major AI product releases from 2023 to 2026, such as M365 Copilot GA, Claude Code, Claude for Financial Services, and Claude Managed Agents. We score every Compustat firm's products on functional overlap with each release. Cowork Plugins overlaps with far more incumbent products than any other release, and the overlap is overwhelmingly stronger. At moderate overlap it covers 3.6 times as many products as the largest other release, and at the deepest overlap the lead grows to about eight times. Cowork Plugins also generated a large ex-post market response. Within days, software and professional-services stocks shed approximately \$285 billion in market capitalization globally ([CNN Business, 2026](#)). The plugins saw widespread adoption, making Cowork Plugins the first commercially successful agentic AI release for general knowledge work. By September 2026, Anthropic's revenue was pacing at more than ten times its end-of-2025 level, driven by broad enterprise adoption of its Cowork and Claude Code products ([Berkowitz, 2026](#)).

Our first finding is that AI-PDR strongly predicts the equity market repricing at the Cowork Plugins release. In a cross-sectional regression of $[-1, +10]$ cumulative abnormal returns on AI-PDR, the coefficient is large and negative, -0.047 ($t = -8.34$), implying that top-quintile firms underperform bottom-quintile firms by approximately 12 percentage

points (p.p.). In contrast, we do not observe a comparable repricing at the earlier AI releases. ChatGPT and DeepSeek R1 produce small positive AI-PDR coefficients under the baseline GICS 6-digit industry fixed effects, attributable to a Gen-AI tailwind, and Claude Opus 4.5 produces a small, insignificant negative one. All three are less than a sixth of the size of the Cowork Plugins effect.

The contrast across AI releases is consistent with a mechanism in which model releases and broad entry into incumbent product space resolve different forms of uncertainty. Generative AI model releases leave the path from raw technology to actual product competition uncertain. A release that enters the product space broadly collapses that uncertainty by directly competing with the incumbents. Within-event variation sharpens the interpretation. A Cowork Overlap score measures how closely each firm’s products compete with the launched plugins. Interacting AI-PDR with this score reveals that the AI-PDR effect is significantly stronger among overlapping firms.

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 score with AI-PDR. A one-standard-deviation (one-SD) higher supply-side shield attenuates the AI-PDR effect by about 32 percent, while a one-SD higher demand-side shield attenuates it by about 16 percent. Both classes therefore attenuate the repricing, with supply-side shields delivering the larger moderation.

While the media dubbed the market response to Cowork Plugins a “SaaS apocalypse” (CNBC, 2026a), the effect holds for firms whose products are not classified as SaaS, such as financial and legal data providers, IT consulting firms, and online education firms. It varies little with firm size, R&D intensity, or financial constraints, consistent with broad-based repricing. In addition, the Cowork Plugins effect is robust along several dimensions. It holds when we control for pre-event interest-rate sensitivity and when we exclude firms that reported earnings within the event window. The repricing is robust to alternative industry-fixed-effects granularities, to cumulative addition of firm-level controls, to adding the ESZT labor-exposure measure as a regressor, and to Fama-French five-factor-adjusted abnormal returns. It is also robust to re-scoring with GPT-3.5, whose September 2021 training cutoff predates all of our events.

Extending the window end to 160 trading days after the launch, the AI-PDR effect stays

negative and significant at the 1% level through 120 trading days and attenuates only at 160. The attenuation is not a uniform recovery. We split firms into SaaS firms, with any product line classified as SaaS, and non-SaaS firms, with none. SaaS firms recover to their pre-Cowork-Plugins level. For non-SaaS firms the repricing persists, at about -6 p.p. through 120 trading days and -4.7 p.p. at 160, significant at the 5% level at every horizon.

The divergence reflects one-sided exposure to the benefits of agentic adoption. Agent-driven workflows are reshaping the application-software market, for example, Salesforce’s agentic AI offering reaching roughly \$800 million in annual recurring revenue in fiscal 2026 (Salesforce, 2026b), and ServiceNow’s AI annual contract value crossing \$1 billion in mid-2026 (ServiceNow, 2026). SaaS firms benefit from agentic adoption. They sell the platform layer an agent runs on and bill it as recurring usage, so they benefit both when agents developed by other firms run on their platforms and when they release their own agentic AI products that complement existing applications. What non-SaaS high-AI-PDR firms sell, curated data, ratings, reports, instruction, programs, and brokered expertise, is the output generative AI also produces, not a platform layer that bills agentic usage. They therefore do not benefit and stay repriced.

The stock performance of the firms that launch agentic AI products themselves supports this interpretation. We identify agentic AI launches from the LSEG news archive with a keyword search of headlines followed by LLM screening. Launches accelerate after Cowork Plugins, 310 by 183 firms within eight months against about 8 per month before. As a group, launchers finish our sample period about 8 p.p. above industry peers with the same AI-PDR. We also examine the announcement returns of launches. A firm’s first launch after Cowork Plugins earns about 3.3 p.p. in the days around the announcement.

Turning to other AI product releases, we find that the displacement effects are not limited to Cowork Plugins, but the repricing around other releases is smaller and short-lived. Using the product-level overlap scores, we examine returns around the other major AI product releases from 2023 to 2026. We divide them into two groups, business-oriented (such as M365 Copilot GA, Claude for Financial Services, and Claude Managed Agents) and consumer-oriented (such as Apple Intelligence, Sora, and OpenAI Operator). Displacement effects appear at the business-oriented group too. Firms with a one-SD higher overlap experience a 0.31 p.p. lower CAR over $[-1, +3]$. These effects, however, are short-lived and much smaller compared to Cowork Plugins, which lose significance within twenty trading days of

the release. The Cowork Plugins repricing, by contrast, stays in place through 120 trading days, and for non-SaaS firms longer still.

We further document changes in analyst and investor beliefs, the latter of which provide supporting evidence on the mechanism behind the repricing. After Cowork Plugins, consensus target prices for high-AI-PDR firms significantly fall, with a magnitude in line with the Cowork Plugins CAR effect. Near-term EPS, revenue, and long-term growth forecasts fall as well. The shield effect carries over to the analyst evidence. Supply-side shields moderate the target-price revisions more than demand-side shields. The investor-belief evidence comes from StockTwits messages. From the 7.5 million messages posted around the four main-sample releases, we flag those endorsing the belief that Gen-AI displaces the firm’s products. We find an increase in investor belief about product displacement that is unique to Cowork Plugins. The average firm posts 3.3 times as many endorsing messages per week after the release as before, and the rise loads on AI-PDR.

Realized fundamentals begin to move the same way. In the second quarter of 2026, year-over-year revenue growth slows significantly more for high-AI-PDR firms than for other firms. A one-SD higher AI-PDR is associated with revenue growth slowing by 3.7 p.p. Their gross margins also fall, with a one-SD higher AI-PDR associated with a 0.9 p.p. larger year-over-year decline. Estimating the same specification quarter by quarter, the post-Cowork-Plugins quarter is the first time in which both outcomes decelerate significantly.

As a complementary test, we construct an AI Input Intensity measure that captures whether scaling Gen-AI requires more units of a firm’s product. The event-study window around Cowork Plugins is not the ideal test for this measure. The high demand for compute, memory, and data-center infrastructure had already been priced in well before the event window through multi-year compute contracts signed by hyperscalers. We therefore look at the long horizon. From January 2022 to September 2026, value-weighted buy-and-hold returns are 257 percent for firms in the top half of the non-zero AI Input Intensity distribution, 104 percent for firms in the bottom half, and 42 percent for firms with no AI-input products.

We contribute to three literatures. The first studies Gen-AI’s displacement effects, which has focused 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 the labor-market effects of robots and AI. [Felten, Raj, and Seamans \(2021\)](#), [Felten, Raj, and Seamans \(2023\)](#), [Webb \(2020\)](#), and [Eloundou et al. \(2024\)](#)

construct occupation-level exposure indices, while [Hampole et al. \(2025\)](#) trace occupational exposure to firm-level employment effects. At the macro level, [Acemoglu \(2025\)](#) and [Aghion, Jones, and Jones \(2019\)](#) model the aggregate productivity and growth effects of AI, and [Bick, Blandin, and Deming \(2026\)](#) document the rapid pace of Gen-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 Gen-AI and firm value. The closest antecedent is [Eisfeldt et al. \(2025\)](#), who construct a labor-exposure measure and show that firms with more Gen-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 Gen-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 Gen-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 repricing arrives with broad entry into incumbent product space.

Finally, we connect to the broader technological displacement literature. [Kogan and Papanikolaou \(2014\)](#) show that technology shocks reallocate value across firms, and [Kogan, Papanikolaou, and Stoffman \(2020\)](#) show that creative destruction is priced in the stock market, with innovation by competitors lowering the market value of incumbent firms. [Kakhbod, Kogan, Li, and Papanikolaou \(2024\)](#) 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 and the threats to them builds on [Hoberg and Phillips \(2010\)](#), [Hoberg, Phillips, and Prabhala \(2014\)](#), and [Hoberg and Phillips \(2016\)](#). We extract and score individual product lines with an LLM along dimensions of displacement vulnerability, and AI-PDR retains its effect after controlling for a digital keyword share of the same filings.

The remainder of the paper proceeds as follows. Section [I](#) describes the data and measure construction. Section [II](#) presents the results on repricing. Section [III](#) examines analyst forecast revisions and investor beliefs. Section [IV](#) examines realized firm fundamentals.

Section V provides additional evidence. Section VI concludes.

I. Data and Measure Construction

A. AI-PDR Measure

We construct AI-PDR from business descriptions (Item 1) of 10-K annual filings retrieved from SEC EDGAR for fiscal years 2021–2026. The construction proceeds in two stages (product extraction and product scoring), both implemented with GPT-5.6 Terra 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 product and service lines this company sells, from the 10-K Item 1 (Business Description) below.

Rules:

1. Each entry is a distinct product line the company sells to external customers. Do NOT list internal capabilities, support functions, or infrastructure unless they are sold externally as a product or service.
2. Merge aggressively. Combine variants of the same business line into a single entry. Examples:
 - All staffing categories (IT, clinical, light industrial) -> "Staffing Services"
 - All commercial loan types (term, revolver, SBA) -> "Commercial Lending"
 - All grades of rolled steel -> "Flat-Rolled Steel Products"

- All modules of one software platform -> one product line (e.g., "Enterprise Workflow Platform")
 - All generic oncology drugs -> "Oncology Therapeutics"
3. Aim for 3-15 distinct product lines. If your count exceeds 15, merge further until it fits. For diversified conglomerates with heterogeneous business segments (e.g., aviation + healthcare + power), keep segments separate but merge aggressively within each segment.
 4. Description: 1-2 sentences grounded in the text. Focus on HOW the product is produced and delivered (e.g., requires physical manufacturing, financial capital, in-person human labor, regulated infrastructure, or is digital-only), not just what it is called.
 5. Importance: "high" if clearly a primary revenue driver; "medium" if meaningful but not dominant; "low" if ancillary. If unsure, "medium".
 6. Evidence: 1-3 short quotes (≤ 20 words each) from the text that support inclusion and importance. Must be near-exact substrings.

Item 1 text:
{item1_text}

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 139,114 product entries from roughly 8,600 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 Gen-AI-agnostic by design, with no mentions of AI or displacement, asking only about digital deliverability to limit look-ahead bias from the LLM scorer.

Listing 2. Digital Deliverability scoring prompt

SYSTEM:

You are a risk scoring engine. Output must be valid JSON following the schema exactly. No markdown, no commentary.

USER:

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 -- it is a gap between defined
   levels. 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 your score on the product DESCRIPTION, not just the name. Marketing
   names can be misleading.

Return one object per product in the input. Echo back the same "idx" and
"name" for each product so results can be matched.

Products to score:
{product_batch_json}

```

We also score each product on four competitive-shield dimensions (legal and regulatory requirements, production moats, customer trust, and customer embeddedness) that capture 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. We set the product-level AI-PDR score equal to the Digital Deliverability rating. 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):

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

with ω_p the importance weight of product p . 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 $\text{AI-PDR}_{i,t}$ 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 Cework Plugins). Section II.C 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. Sectors whose outputs are primarily digital (Communication Services and Information Technology) rank highest, while those producing physical goods (Materials, Consumer Staples, and Energy) rank 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 AI-PDR. High-scoring products on digital deliverability feature “workflows,” “search,” “reports,” “SaaS,” and “insights.” Low-scoring products feature “theatre,” “transports,” “tissue,” “dining,” and “drinks.” Table A2 lists the largest firms among the highest- and lowest-scoring. High-AI-PDR firms include design and creative software (Autodesk, Adobe), enterprise and developer software (Atlassian, Salesforce, Datadog, Zoom), and social and gaming platforms (Reddit, Roblox, Electronic Arts, Take-Two). Low-AI-PDR firms cluster in pharmaceuticals and biologics (Vertex Pharmaceuticals, Bristol Myers Squibb, Amgen, AbbVie), energy and metals (Chevron, Southern Copper), rail transportation (Union Pacific), food, beverage, and tobacco (PepsiCo, Philip Morris), and electronic components (Amphenol), all firms whose products require physical extraction, manufacturing, processing, or distribution.

Table A3 reports correlations between AI-PDR and observables. The correlation with the ESZT labor-exposure measure is 0.38, positive but far from collinear, confirming that product displacement and labor exposure capture distinct dimensions. AI-PDR correlates negatively with tangibility (-0.34) and only weakly with R&D intensity (0.05). The raw correlation with firm size is small (0.05). AI-PDR is positively correlated with gross margin. In particular, a one-point higher AI-PDR is associated with a 2.8 p.p. higher gross margin, which implies that the median-AI-PDR firm in the top quintile has an 18 p.p. higher gross margin than one in the low AI-PDR quintile. AI-PDR is also positively associated with revenue growth, at 0.8 p.p. per AI-PDR point. The high margin and fast growth of high-AI-PDR firms make their valuations susceptible to margin compression and growth deceleration induced by competition.

Because AI-PDR derives from product descriptions that change slowly, it should be stable over time. Figure A1 confirms this. Adjacent-year Pearson correlations range from 0.94 to

0.95, and non-adjacent years (e.g., 2021 versus 2026) have correlations of about 0.89. This stability confirms that AI-PDR captures durable product characteristics rather than filing-year idiosyncrasies.

B. Other Data and Summary Statistics

Daily stock returns through 2025 are from the CRSP Stock Files. For 2026, we extend CRSP with daily returns constructed from Yahoo Finance prices, screened for stock-split artifacts and validated against CRSP over 2025. We merge firms with Compustat via the CRSP-Compustat Merged (CCM) link table. We obtain firm-level Gen-AI labor exposure scores from [Eisfeldt et al. \(2025\)](#). Control variables are constructed with no look-ahead. For each event, we use Compustat data from each firm’s most recent fiscal year ending before the event year to construct $\log(\text{ME})$, Tobin’s Q, ROA, and asset growth. Table [I](#) reports summary statistics and Table [A1](#) provides definitions for all variables used in the analysis.

II. Main Results

A. Events

We study equity market reactions to four major 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), the four main-sample releases.¹ At the first three events, the launching lab released a general-purpose generative AI model. Cowork Plugins, by contrast, is a set of packaged agents made available for purchase that directly compete with specific incumbent digitally deliverable products.

As 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 Gen-AI tool on January 12, 2026. On January 30, Anthropic added plugin support and open-sourced 11 plugins on GitHub ([TechCrunch, 2026](#)). Ten are role-specific plugins covering productivity, enterprise search, marketing, sales, finance, data analysis, legal,

¹Section [V.A](#) extends the analysis to the major AI product releases of 2023 to 2026, and Section [V.B](#) to 25 LLM model releases from late 2022 through July 2026.

customer support, product management, and biology research, and the eleventh is a plugin-management utility for building custom plugins. 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.

We compare Cowork Plugins to all other major AI product releases from 2023 to 2026, releases in which an AI provider shipped a product that performs a specified set of work tasks. Model releases are excluded, since they do not ship a delimited set of tasks whose overlap with existing products can be measured. Thirteen releases qualify, ChatGPT Plugins, Code Interpreter, M365 Copilot GA, Apple Intelligence, Claude 3.5 computer use, ChatGPT Search, Sora, Gemini Deep Research, OpenAI Operator, Claude Code, Claude for Financial Services, Claude Managed Agents, and Claude Design, besides Cowork Plugins itself. We summarize each release into functional bundles based on official documentation, and score every Compustat-firm product line in the release’s filing-year vintage (the 10-K filing year preceding the release year) on functional overlap with each bundle on the 0–10 scale. A product’s overlap with the release is the maximum across the release’s bundles. Firm-level overlap is the importance-weighted mean across the firm’s products.

Figure A3 compares the fourteen releases, Cowork Plugins included. At moderate overlap, 663 products held by 450 firms overlap with Cowork Plugins, 3.6 times the products and 3.1 times the firms of the largest other release. At deep overlap, the lead grows to 8.3 times the products and 7.6 times the firms. The releases with the second and third most overlap are ChatGPT Plugins and Code Interpreter, trailing far behind Cowork Plugins. This unmatched breadth and depth of overlap with the functionality of incumbent Compustat firms’ products is the distinguishing feature of Cowork Plugins.

The pricing structure also 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](#),

2026a), 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). 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 Gen-AI agents perform the operational tasks, the number of paid human users for a SaaS platform shrinks. Businesses can then reduce their SaaS license counts, 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, 2026a). 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 A4 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 upgrade. The Opus 4.6 release on February 5 was a further model upgrade. 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 incumbent Compustat-firm 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, 2026a).

B. Cross-Sectional Regressions

B.1. Event Study Design

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

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

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. The pre-event market beta $\hat{\beta}_i$ comes from OLS regressions of $(r_{i,t} - r_{f,t})$ on $(r_{m,t} - r_{f,t})$ over $[-252, -11]$ trading days, and the estimated intercept is excluded from the benchmark. Using the beta-only CAPM benchmark rather than the statistical market model follows the treatment of [Ramelli and Wagner \(2020\)](#). Subtracting a firm's estimated intercept over an event window extrapolates its estimation-period drift, which compounds with the horizon ([Kothari and Warner, 2007](#)). As a robustness test, in Appendix Table [A27](#) we report findings under a characteristics-matched benchmark ([Daniel, Grinblatt, Titman, and Wermers, 1997](#)) and under the alpha-inclusive market-model construction.

We then estimate cross-sectional regressions:

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

where δ_j are GICS 6-digit industry fixed effects (with industries having fewer than five firms dropped) and $\mathbf{X}_i$ includes $\log(\text{ME})$, Tobin's Q, ROA, and asset growth. All right-hand-side variables, and the CARs, 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. Appendix Table [A5](#) reports coefficient sensitivity across alternative FE granularities (no industry FE through GICS 8-digit). Table [A15](#) reports a complementary specification that includes the ESZT labor exposure as a regressor.

B.2. Baseline Results

Table II reports the baseline cross-sectional regression of CAPM-adjusted $[-1, +10]$ CARs on AI-PDR, separately for each of the four main-sample events. AI-PDR is standardized to mean zero and unit variance within each estimation sample. At Cowork Plugins, the AI-PDR coefficient is -0.0466 ($t = -8.34$). The three generative AI model releases produce coefficients less than a sixth of the size under the same specification. ChatGPT Launch yields 0.0068 ($t = 1.82$) and DeepSeek R1 yields 0.0075 ($t = 2.01$), small positive coefficients consistent with a Gen-AI tailwind, while Claude Opus 4.5 yields -0.0054 ($t = -1.58$), small and insignificant. As the standardized AI-PDR spread from quintile 1 to quintile 5 is about 2.5 sample standard deviations, the Cowork coefficient implies a Q5–Q1 CAR spread of roughly 12 p.p. over $[-1, +10]$ within narrowly defined industries.

Appendix Table A5 shows that this contrast survives across five industry fixed-effect choices, from no industry FE through GICS 8-digit sub-industry, with the Cowork Plugins coefficient negative and significant at the 1% level in every column. The Claude Opus 4.5 coefficient is negative and significant at coarse granularities (-0.0116 , $t = -3.76$ at GICS 2-digit). It is possible that by then the market had partially incorporated Gen-AI product displacement risk, since Claude Opus 4.5’s improved performance had made automated agent workflows a more credible threat.

Figure 3 presents the estimated effects across different windows up to 160 trading days after the event, its two panels plotting the AI-PDR coefficient continuously at GICS 2-digit and 6-digit fixed effects. The coefficients are tabulated in Appendix Table A6. For Cowork Plugins the effect remains in place far beyond the baseline window. Under GICS 6-digit fixed effects the coefficient is -0.0466 at $[-1, +10]$ and remains between -0.0352 and -0.0526 through $[-1, +120]$, significant at the 1% level at every tabulated window through $[-1, +120]$. The standard errors naturally widen as the window grows, as each added day contributes noise to the cumulative abnormal return. Appendix Figure A6 traces the same regression backward into the pre-event weeks over $[-40, +40]$ trading days. High-AI-PDR firms begin to slip about twelve trading days before the launch, around the January 12 Cowork desktop preview, while the decline over the preview weeks is small next to the drop at the launch itself. January 30 remains the right event date, as it is the date on which the released product directly overlaps identifiable incumbent products, since the preview carried no domain-specific functionality.

Only at $[-1, +160]$ does the effect attenuate, to -0.0270 ($t = -1.43$) under GICS 6-digit fixed effects, while remaining significant at GICS 2-digit (-0.0369 , $t = -2.48$). The attenuation at the longest horizon is not a uniform recovery. Section II.B.5 shows that it is the SaaS side of the sample recovering to its pre-Cowork-Plugins level while the repricing of non-SaaS firms persists. Figure 4 repeats the plot for the three earlier Gen-AI releases and finds no comparable pattern. The one exception is 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.

The divergence between the repricing at Cowork Plugins and the muted response at the model releases reflects how the two kinds of release resolve different forms of uncertainty. Generative AI model releases resolve uncertainty about what the technology *can* do, but the path from a model 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. Cowork Plugins collapses this chain of uncertainty. Markets observe entry into the product space and direct competition with the incumbents. The repricing is consistent with the removal of this uncertainty discount.

B.3. Effect Heterogeneity

We exploit within-event variation to discipline the interpretation. We construct a Cowork Overlap score measuring how closely each incumbent Compustat firm’s products compete with the Cowork Plugins. The plugin definitions are from the Anthropic launch announcement.² An LLM scores each incumbent-firm product on a 0–10 scale against each of the 10 role-specific plugin bundles (Listing A2).³ The product-level Cowork Overlap is the maximum across the 10 plugin-specific scores, and the firm-level Cowork Overlap is the importance-weighted average across products. The firm-level distribution is right-skewed, as about 68% of firms have zero overlap. Cowork Overlap correlates with AI-PDR at a Pearson $\rho = 0.66$.

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

³The eleventh plugin, a plugin-management utility that automates the creation of custom plugins rather than any incumbent product category, is not scored.

Table III brings the overlap score into the regression in two ways, each at GICS 2-digit and GICS 6-digit fixed effects. Columns (1)–(2) add the standardized firm-level Cowork Overlap next to AI-PDR. At the baseline GICS 6-digit fixed effects the overlap main effect is -0.0179 ($t = -5.30$), so a one-SD higher overlap predicts a 1.8 p.p. lower CAR beyond what AI-PDR explains, while the AI-PDR coefficient remains at -0.0359 ($t = -5.41$). Columns (3)–(4) interact AI-PDR with an indicator that the firm has at least one overlapping product, a maximum product-level overlap score of 3 or above, which holds for about 22% of the firms. The interaction is -0.0175 ($t = -2.86$), implying that the AI-PDR slope is about 50% steeper for firms with an overlapping product. That firms whose products overlap the launched plugins bear a larger repricing, both directly and through a steeper AI-PDR slope, supports that the cross-sectional return effect operates through the product-displacement channel.

Table A13 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). The interactions are all insignificant. The displacement pattern is broad-based rather than concentrated among small, R&D-intensive, or financially constrained firms.

B.4. Is the Effect Limited to SaaS?

A natural challenge is that the Cowork Plugins repricing is just a software story. We restrict the sample by excluding GICS sectors 45 (Information Technology) and 50 (Communication Services), and re-estimate the regressions (Table A8). The AI-PDR coefficient is -0.0426 ($t = -8.09$) excluding Information Technology and -0.0367 ($t = -7.35$) excluding Communication Services as well.

Even outside the two technology sectors, many high-AI-PDR firms still sell software-as-a-service (SaaS) or broadly software, such as Veeva Systems in Health Care or Verisk Analytics in Industrials. We therefore ask an LLM to classify each incumbent-firm product line as software and, nested within software, as SaaS under a cloud-delivery, multi-tenant, recurring-revenue definition (Listing A5). A firm is software if any of its product lines is software and non-software only if none is, and likewise for SaaS. Table A9 shows examples. Table A10 re-estimates the regression separately for software versus non-software firms, and for SaaS versus non-SaaS firms. The non-software subsample carries an AI-PDR coefficient

of -0.0209 ($t = -4.54$) at the GICS 6-digit FE, and the non-SaaS subsample -0.0276 ($t = -6.44$). The effect extends beyond software and SaaS.

The two tests above identify software firms by classification. As an alternative that does not rely on any classification, we identify software-like firms from returns themselves, by their comovement with the iShares Expanded Tech-Software Sector ETF (IGV). We estimate each firm’s pre-event loading on the IGV excess return from a two-factor regression of daily excess returns on the market and IGV over $[-252, -11]$ trading days. The software beta enters at -0.0288 ($t = -6.39$) alone and stays significant jointly, yet AI-PDR survives at -0.0415 ($t = -7.50$, Table A11).

The reason the effect is not limited to SaaS or software is that what makes a product displaceable is digital deliverability. High-AI-PDR firms span financial and legal data providers such as S&P Global, Moody’s, Morningstar, and LegalZoom, IT consulting firms such as Accenture and Cognizant, online education firms such as Duolingo and Chegg, perpetual-license software vendors such as PTC and Progress Software, video-game publishers such as Take-Two Interactive and Electronic Arts, and freelancer marketplaces that intermediate digitally deliverable knowledge work such as Upwork and Fiverr. Their products, curated data, ratings, reports, instruction, programs, and brokered expertise, are exactly the outputs a Gen-AI agent can deliver. The repricing follows the digital deliverability of the product and is not limited to software.

B.5. Beneficiaries of Agentic AI Products

The previous subsection shows that the repricing at Cowork Plugins holds for SaaS and non-SaaS firms alike over the baseline window (Table A10). Over longer horizons the two groups part ways, which is shown in Figure 5 and Table IV. The repricing unwinds for SaaS firms and persists for non-SaaS firms. SaaS firms recover to their pre-Cowork-Plugins level, the SaaS AI-PDR slope returns to zero by 160 trading days (-0.001 , $t = -0.03$), and a characteristics-matched benchmark agrees (0.005 , $t = 0.16$). For non-SaaS firms the slope is -0.064 ($t = -3.74$) at 120 trading days and -0.047 ($t = -2.31$) at 160, significant at the 5% level at every tabulated window. The pattern holds under a majority rule that classifies a firm as SaaS only when at least half of its importance-weighted product lines are SaaS (Appendix Table A30).

The rebound on the SaaS side lines up with benefits from agentic AI products. SaaS firms

sell the platform layer an agent runs on. An agent does its work inside their applications, records, and workflows, so agent adoption comes back to them as usage revenue. SaaS firms thus benefit both when agents developed by other firms run on their platforms and when they release their own agentic AI products that complement existing applications. Bloomberg Intelligence’s 2026 outlook describes AI agents reshaping the application-software market as enterprises adopt agent-driven workflows, with pricing shifting from seats toward usage (Bloomberg Intelligence, 2026). Salesforce reported that Agentforce, its agentic AI offering, reached roughly \$800 million in annual recurring revenue in fiscal 2026, up 169 percent year over year (Salesforce, 2026b; Fortune, 2026), and ServiceNow’s AI annual contract value crossed \$1 billion in mid-2026, with agentic deployments growing ninefold over nine months (ServiceNow, 2026). Both firms fell at the release and recovered, their CARs moving from -15.4% and -14.8% over $[-1, +10]$ to $+2.2\%$ and $+7.2\%$ over $[-1, +160]$.

Non-SaaS products offer no such complementarity. What their sellers sell, curated data, ratings, reports, instruction, programs, and brokered expertise, is the output generative AI also produces, not a platform layer that bills agentic usage. The group spans financial and legal data providers such as S&P Global, Moody’s, Morningstar, and LegalZoom, IT consulting firms such as Accenture and Cognizant, online education firms such as Duolingo and Chegg, perpetual-license software vendors such as PTC and Progress Software, video-game publishers such as Take-Two Interactive and Electronic Arts, and freelancer marketplaces that intermediate digitally deliverable knowledge work such as Upwork and Fiverr. A business model that sells curated answers or unmetered programs rather than a metered platform does not directly benefit from an agentic workflow.

Moreover, such curated answers fall squarely on the strengths of generative AI and are among the products most susceptible to displacement, as witnessed in the legal and medical industries by the success of new generative-AI-native entrants such as Harvey and OpenEvidence.⁴ The new entrants, along with incumbent providers such as Westlaw and UpToDate, are at the same time themselves at risk as general-purpose generative AI models advance. A study in Nature Medicine asked physicians to blind-rate answers to real questions raised in

⁴Harvey, a legal AI startup founded in 2022, grew from \$100 million in annual recurring revenue in August 2025 to roughly \$300 million by May 2026 and was valued at \$11 billion in March 2026 (CNBC, 2026b). OpenEvidence, a medical AI startup founded in 2022, reached about \$150 million in annualized revenue in 2025, is used by more than 40 percent of US physicians, and doubled its valuation to \$12 billion in January 2026 (CNBC, 2026c).

clinical work, and the frontier general-purpose models outscored both UpToDate’s Expert AI and OpenEvidence (Vishwanath et al., 2026).

We collect agentic AI product release announcements from the LSEG news archive over January 2023 through September 2026 in three steps, a keyword search over the news headlines, an LLM screen of the headlines, and a full-story LLM verification of the announcer, the timing, and the offering. We find 554 launches by 229 firms. Examples include Intuit’s Anthropic-powered AI agent platform for financial workflows, Oracle’s agentic AI platform for corporate banking, Paychex’s Wise AI platform for an agentic digital workforce, and Broadridge’s agentic AI deployment across its capital-markets and wealth-management workflows. Launches accelerate sharply after Cowork Plugins, 310 by 183 firms in the eight months after the release against 244 by 127 firms in the three years before.

We show that agentic AI product launchers outperform by examining both their performance as a group following Cowork Plugins and their own product announcement returns. First, we regress CARs on AI-PDR and an indicator for the firm launching agentic AI products after Cowork Plugins (Table V). Launchers sit modestly below their industry peers through the first months (-0.028 , $t = -1.64$, at $[-1, +20]$), converge to peers by 120 trading days, and finish about 7.7 p.p. above peers with the same AI-PDR by $[-1, +160]$ ($t = 2.55$). Under a characteristics-matched benchmark (Daniel et al., 1997), it is 0.096 ($t = 3.07$). The outperformance accrues to the launchers as a group late in the window.

Second, the dates of agentic AI product launch announcements themselves carry news. Appendix Table A28 reports mean CARs around each firm’s first agentic AI launch, with the firm’s own launch date as the event date. Before Cowork Plugins, a firm’s first launch shows no announcement return at the 5% level at any tabulated window. After Cowork Plugins, a firm’s first launch earns 1.4 p.p. over $[-1, +1]$ ($t = 2.01$) and 3.3 p.p. over $[-1, +10]$ ($t = 2.07$), with similar magnitudes under the characteristics-matched benchmark. Both pieces of evidence are consistent with launchers benefiting from agentic AI adoption. Moreover, the market systematically rewards agentic AI product launchers after the release of Cowork Plugins but not before, consistent with Cowork Plugins collapsing the uncertainty over the path from raw technology to actual product competition and making the commercial success of agentic AI products credible.

B.6. Moderation by Competitive Shields

Table VI extends the specification by adding a *Shield* composite, the simple average of four firm-level shield scores (legal and regulatory requirements, production moats, customer trust, and customer embeddedness), each the importance-weighted average of its product-level scores, together with its interaction with AI-PDR,

$$\text{CAR}_i = \delta_j + \beta_1 \cdot \text{AI-PDR}_i + \beta_2 \cdot \text{Shield}_i + \beta_3 \cdot \text{AI-PDR}_i \times \text{Shield}_i + \gamma' \mathbf{X}_i + \varepsilon_i, \quad (2)$$

with AI-PDR and Shield standardized to unit variance and demeaned within the sample, so the interaction is the product of the standardized-demeaned terms, and δ_j and $\mathbf{X}_i$ as in Equation (1). For Cowork Plugins at the baseline $[-1, +10]$ window, $\hat{\beta}_{\text{AI-PDR} \times \text{Shield}} = 0.0121$ ($t = 3.29$). This implies that while a one-SD higher AI-PDR is associated with about 4.5 p.p. lower CAR, a one-SD higher composite Shield attenuates the effect by about 27 percent.

Columns (3) through (6) of Table VI decompose the composite into supply-side shields and demand-side shields. The former comprises legal and regulatory requirements and production moats, which capture barriers a new entrant would face producing and selling the same output. 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 0.0135 ($t = 2.92$) at the $[-1, +10]$ window. The demand-side interaction $\hat{\beta}_{\text{AI-PDR} \times \text{Demand}}$ is 0.0077 ($t = 2.14$) at the same window. Both supply-side and demand-side shields therefore moderate the displacement effect, with the supply-side coefficient larger in magnitude than the demand-side coefficient.

The pattern is consistent with a mechanism in which entrant-side production assets and customer-side behavior respond differently as Gen-AI advances. Production moats such as proprietary data, fabrication assets, and network effects govern whether a Gen-AI entrant can produce the substitute at all, so a high-AI-PDR firm sitting behind such a barrier is partially protected even as Gen-AI advances. Legal and regulatory barriers also work this way, as product-specific licenses, regulatory approvals, and professional certifications govern whether a Gen-AI entrant can sell the substitute at all. Customer trust and customer embeddedness rest on customer behavior that is itself responsive to the substitute. Trust thresholds fall as Gen-AI-produced output accumulates a track record, and switching costs erode as Gen-AI-native alternatives mature their interfaces and integrations. Markets price

both classes of shield, with greater weight on the production-side barriers that are more likely to persist over time.

C. 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.

C.1. Identification

Placebo dates. Appendix Figure A5 compares the Cowork Plugins estimate with 169 placebo dates across 2023 to 2026, none of whose windows overlap an event window. The Cowork Plugins coefficient lies about five placebo standard deviations below the placebo mean, and no placebo date produces a slope below it.

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 (-0.0466 , $t = -8.28$ versus the -0.0466 , $t = -8.34$ baseline).

Coincident earnings announcements. The Cowork Plugins event window overlaps with the late-January earnings season. Table A12 excludes firms whose Compustat earnings-announcement date falls within the $[-1, +10]$ window, reducing N from 2,845 to 1,951 at the GICS 6-digit baseline. The AI-PDR coefficient remains negative and significant at the 1% level, at -0.0435 ($t = -6.96$) versus the -0.0466 ($t = -8.34$) baseline.

C.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). Appendix Table A29 replaces 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 across the five industry-FE granularities. For

Cowork Plugins at GICS 6-digit FE the AI-PDR coefficient is -0.0530 ($t = -9.58$), slightly larger in magnitude than the -0.0466 baseline, and the three model releases keep their small baseline pattern under the averaged measure. AI-PDR captures a stable firm-level product characteristic rather than a one-year aberration.

Chronologically consistent scorer. A separate concern is that the scoring LLM’s training data may encode which products Gen-AI will displace, so that digital deliverability ratings reflect learned associations rather than inherent product characteristics. [He, Lv, Manela, and Wu \(2025a,b\)](#) 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 main-sample event. [Table A18](#) reports the results. The all-years GPT-3.5 AI-PDR correlates with the baseline at a Pearson $\rho = 0.91$, indicating broad agreement on which products are digitally deliverable despite a multi-year gap in training data. For Cowork Plugins at GICS 6-digit FE, the AI-PDR coefficient is -0.0356 ($t = -6.46$), replicating the main finding.

Value of LLM-based scoring. A simpler alternative to LLM scoring is a dictionary word count in the tradition of text-based measures. To test whether the LLM adds information beyond counting words, we compute for each firm the share of Item 1 words matching a 30-term digital dictionary (software, cloud, platform, and the like, listed in full in the note to [Table A11](#)). The share is computed on the same filing vintage as the event AI-PDR. The keyword share correlates with AI-PDR at 0.73 and predicts the Cowork Plugins repricing on its own (-0.0280 , $t = -3.40$ at the GICS 6-digit FE). In the joint regression, however, it collapses to -0.0099 ($t = -1.29$) while AI-PDR is little changed at -0.0412 ($t = -7.08$). A word-count measure is therefore not sufficient, and its explanatory power is fully subsumed by the LLM-based score.

C.3. Specification and Controls

Adding the ESZT labor-exposure control. [Table A15](#) adds the ESZT labor-exposure measure as a control in a regression weighted by market capitalization, 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 (0.0191 , $t = 3.32$), consistent with their

original finding that labor-exposed firms outperformed on news of Gen-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 -0.0451 ($t = -2.72$) and the ESZT coefficient is -0.0125 ($t = -1.75$). 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 A16 replaces CAPM-adjusted CARs with Fama–French five-factor adjusted CARs. For Cowork Plugins at GICS 6-digit FE, the AI-PDR coefficient is -0.0443 ($t = -5.92$). 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 remains large and highly significant.

Additional controls. Table A17 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 -0.0381 to -0.0466 and remains highly significant throughout. Adding R&D intensity last reduces the sample to 1,347 firms, where the point estimate grows in magnitude to -0.0477 ($t = -4.19$). The stability across specifications indicates that AI-PDR does not proxy for standard firm characteristics.

III. Analyst Forecast Revisions and Investor Beliefs

A. Analyst Forecast Revisions

Section II 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. The repricing should appear most strongly in long-horizon valuations such as target prices, which embed expectations of cash flows over the indefinite future. Near-term earnings and revenue forecasts are anchored on existing contracts and order pipelines, so any near-term response should be smaller.

We test this with a cross-sectional regression of scaled forecast revisions on standardized AI-PDR.

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

where δ_j are industry fixed effects, $\mathbf{X}_i$ are the same firm controls as in Equation (1), 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 from pre-event to post-event since it is already in p.p. For each firm, the pre-event and post-event EPS and revenue forecasts are all forecasts for the same fiscal year, that is, the first fiscal year ending on or after December 2026. The pre-event forecast is the firm’s last IBES mean consensus before January 30, 2026. The post-event forecast is its last consensus after the event through June 30, 2026. We drop EPS observations with pre-event EPS below \$0.50 (which includes all negative-EPS firms) and revenue observations with pre-event revenue below \$10 million to avoid sign ambiguity from scaling by a near-zero or negative denominator. Each firm corresponds to one observation in the regression sample, that is, the post-event forecast minus the pre-event forecast. Table A19 reports summary statistics for the analyst forecast sample.

Table VII reports results for four forecast outcomes, namely consensus mean EPS, revenue, long-term growth (LTG), and target price, under GICS 2-digit sector fixed effects. Target prices show the strongest revision, with $\hat{\beta} = -0.105$ ($t = -5.32$). A one-SD higher AI-PDR is associated with a roughly 10.5 p.p. more negative proportional revision in the target price within sectors. Within narrow GICS 6-digit industries the revision is -0.055 ($t = -3.58$, untabulated), broadly in line with the event-study CAR of about -4.7 p.p. per one-SD AI-PDR at Cowork Plugins (Table II).

The near-term forecasts move in the same direction. The EPS revision is -0.029 ($t = -2.66$), and the LTG change is -1.5 p.p. ($t = -2.06$).⁵ The revenue revision is marginal (-0.009 , $t = -1.85$). Analysts therefore mark down the forecast profile of high-AI-PDR firms, with the largest revision in the long-horizon target price and smaller effects on revenue and EPS. The pattern is consistent with displacement risk operating primarily on cash flows over a long horizon while already denting near-term earnings expectations. Appendix Table A20 repeats the EPS and revenue regressions for forecast years further into the future, that is, fiscal years ending on or after December 2026, 2027, and 2028. There are negative revisions across the board. The EPS revision extends to the December 2027 horizon with little change, and the revenue revision strengthens beyond the first fiscal year.

⁵LTG coverage is sparse and the forecasts are revised infrequently. Fewer than 40% of the firms with a target-price revision also have a long-term growth revision, and analysts produce LTG less consistently than EPS or revenue.

Appendix Table A21 tests whether Shields moderate the target-price revision using the same interaction structure as Table VI. The composite-Shield specification yields $\hat{\beta}_{\text{AI-PDR}} = -0.101$ ($t = -5.20$) and $\hat{\beta}_{\text{AI-PDR} \times \text{Shield}} = 0.017$ ($t = 1.69$). When the composite is split into supply-side and demand-side components, the supply-side interaction is statistically significant while the demand-side is not, with $\hat{\beta}_{\text{AI-PDR} \times \text{Supply}} = 0.046$ ($t = 4.69$) versus $\hat{\beta}_{\text{AI-PDR} \times \text{Demand}} = -0.008$ ($t = -0.87$). The supply-side coefficient is far larger than the demand-side coefficient, mirroring the event-study finding (Table VI), so markets and sell-side analysts agree on which class of shield provides more durable protection against Gen-AI substitution. The target-price revision also remains negative in AI-PDR after controlling for the firm’s realized return over the same period (Table A22). The target-price revision therefore does not simply reflect analysts anchoring their targets to the realized price move.

B. Text-Based Investor Beliefs around Releases

We argue that the three generative AI model releases leave the path from raw technology to actual product competition uncertain. Cwork Plugins collapses that uncertainty as it enters the product space and directly competes with the incumbents. If this reading is correct, investors should have experienced an increase in their beliefs about product displacement at Cwork Plugins and not at the other releases. The repricing and the analyst revisions are both consistent with such an update, but neither observes beliefs directly. This subsection provides direct evidence using StockTwits messages.

StockTwits is a social platform on which investors post short messages tagged to individual stocks. Because the messages are timestamped, firm-specific, and written by investors themselves, the platform has become a standard source for measuring investor beliefs in finance research (Antweiler and Frank, 2004; Cookson and Niessner, 2020; Cookson, Engelberg, and Mullins, 2023). We collect every message posted on any firm over the four calendar weeks on each side of each of the four main-sample releases, 7.5 million messages in total, with a typical message of around 100 characters. We provide summary statistics and example messages in Table A23. In our analysis, we focus on messages with at least 20 characters.

We identify messages that endorse the belief that generative AI displaces the firm’s products in three steps. First, a keyword prefilter removes messages with no generative-AI,

displacement, software, or obsolescence language. Second, we let a GPT-4.1 classifier decide whether each message endorses the belief that generative AI displaces the firm’s products, under the prompt in Listing A3. The rubric admits three forms of displacement, an AI provider’s own product substituting for the firm’s product, AI lowering the cost of building software so that rivals multiply, and customers building substitutes in-house. We admit the latter two forms, in which the displacer is not the Gen-AI provider itself. Third, because StockTwits messages are brief, conversational, and unstructured, machine classification is more error-prone here than when applied to curated 10-K text. A verification pass therefore re-judges every message the classifier labels as endorsing under the same rubric with a stronger model, GPT-5.5.⁶ We count a message as endorsing when it passes the keyword prefilter and both classification rounds. We also manually review every message that passes the second step, replacing the third step with the manual review, and the findings are robust (Table A24).

We find 14 endorsing messages in the ChatGPT window, 70 at DeepSeek R1, 155 at Claude Opus 4.5, and 656 at Cowork Plugins (Table A23). Direct substitution by an AI provider’s own product is the dominant form throughout, while the form in which the displacer is the customer concentrates disproportionately in the Cowork Plugins window. Panel B shows an example message of each form. Figure 6 plots the average firm-level number of endorsing messages by week around each release, among firms with at least 5 qualifying messages of at least 20 characters in the week. The raw series is flat around the three earlier releases. At Cowork Plugins there is an increase around the release, and the series stays elevated through the end of the window. The post-period average is 3.3 times the pre-period average. The increase in investor belief about product displacement is visible in the raw data before any regression, and it is unique to Cowork Plugins.

We then estimate cross-sectional regressions analogous to the CAR evidence,

$$\Delta\text{Endorse}_i = \delta_j + \beta_1 \cdot \text{AI-PDR}_i + \gamma' \mathbf{X}_i + \varepsilon_i, \quad (3)$$

with one observation per firm, where $\Delta\text{Endorse}_i$ is the pre-to-post change in firm i ’s endorsement of displacement, δ_j are industry fixed effects, and $\mathbf{X}_i$ are the event-study controls,

⁶Rejected messages most often lack a displacement claim about the tagged firm in the message’s own words, discuss an industry outside the rubric’s scope, restate price moves, or merely relay news headlines.

augmented in the log-count and calibrated specifications with the change in the log of the firm’s total message count. AI-PDR is the event-year score of the event study, and standard errors are clustered by GICS 6-digit industry as before. The sample consists of firms with at least 10 messages in both periods. We use three groups of outcome variables. First, the Level columns use the change in the endorsement fraction itself (in p.p.). Second, the Log count columns use the change in the log of one plus the firm’s number of endorsing messages, controlling for the change in the log of its total number of messages. Third, the Calibrated log columns use $m(N) = \log N$ for $N \geq 1$ and $m(0) = -1$, with the same message-growth control. This addresses the concerns of [Cohn, Liu, and Wardlaw \(2022\)](#) and [Chen and Roth \(2024\)](#) about log-like transformations of counts with zeros, making the weight on the extensive margin an explicit choice. Under this calibration, going from zero to one endorsing message contributes one log point, the same weight as a 2.7-fold increase in the count at the intensive margin.

We present our findings in Table [VIII](#). At Cowork Plugins the endorsement of high-AI-PDR firms rises significantly more than that of their sector peers across the three outcome variables under GICS 2-digit fixed effects. The level estimate implies that a one-SD higher AI-PDR raises the endorsement fraction by 0.085 p.p. ($t = 3.64$), roughly four times the pre-event mean fraction of 0.023 p.p. The log-count estimate implies that the same shift raises the firm’s endorsing-message count by 0.091 log points ($t = 4.69$), and the calibrated estimate by 0.126 log points ($t = 4.87$). One log point is the weight of the move from zero to one endorsing message. Under GICS 6-digit fixed effects, the log-count and calibrated estimates remain positive and significant at the 1% level, while the level estimate stays positive and marginally significant, pointing to the role of cross-industry variation in identifying the effect on the level of the endorsement fraction. ChatGPT Launch, DeepSeek R1, and Claude Opus 4.5 are null in every cell.

Taken together, the StockTwits messages provide direct evidence for the belief channel. Investor belief in product displacement increased at Cowork Plugins and the increase concentrated in high-AI-PDR firms. We argue that only Cowork Plugins moved this belief because it entered the product space and directly competed with the incumbents. The earlier releases leave the path to actual product competition uncertain, so investors had little concrete to write about. A \$20-per-month agent that reviews contracts, reconciles accounts, and runs sales pipelines competes with identifiable products, and the messages show investors writing

about exactly those products.

IV. Realized Firm Fundamentals

The evidence so far concerns prices, forecasts, and beliefs. This section asks whether realized fundamentals have started to move the same way. We merge the AI-PDR measure to quarterly Compustat and study two outcomes, the change in the revenue growth rate and the year-over-year change in gross margin. For the growth rate outcome, we ask whether growth is slowing for high-AI-PDR firms. We take each firm’s fiscal quarter ending inside calendar Q2 2026, which comprises fully post-Cowork-Plugins-release fiscal quarters. The outcome is the year-over-year growth rate in that quarter minus the year-over-year growth rate one year earlier. Each firm contributes one observation, and we regress each outcome on standardized AI-PDR with the same controls as in Equation (1) and industry fixed effects.

Table IX reports the results at GICS 2-digit fixed effects. Growth is slowing for high-AI-PDR firms. A one-SD higher AI-PDR is associated with revenue growth slowing by 3.7 p.p. ($t = -3.81$) and a 0.9 p.p. ($t = -4.14$) larger decline in gross margin. Columns (3) to (6) split the AI-PDR slope by SaaS and non-SaaS, with each pair of columns estimated in one pooled regression as in Table IV. The deterioration appears on both sides of the divide. Revenue growth slows by 4.4 p.p. for SaaS firms ($t = -3.73$) and 3.7 p.p. for non-SaaS firms ($t = -2.36$), and the gross-margin declines are 1.2 p.p. ($t = -3.13$) and 0.9 p.p. ($t = -2.60$). The SaaS and non-SaaS slopes are statistically indistinguishable for both outcomes. In contrast to the repricing, which diverges along the SaaS divide over longer horizons, the early realized deterioration is shared across the divide.

Figure 7 traces the timing. We estimate the same specification cohort by cohort, where the cohort for a calendar quarter contains every firm with a fiscal quarter ending inside that quarter, from 2024 Q1 through 2026 Q2. The coefficients sit near zero through 2024 and the first half of 2025 for both outcomes. Revenue growth begins to slow for high-AI-PDR firms in the 2025 Q3 cohort. The coefficient under revenue growth rate is significantly negative in the 2026 Q1 cohort, the first cohort containing post-release fiscal quarters. 2026 Q2 is the first cohort in which both outcomes deteriorate significantly for high-AI-PDR firms. Realized fundamentals thus move in the direction of the repricing within two quarters of the release.

V. Additional Tests

A. Other Major AI Product Releases

The headline contrast sets Cowork Plugins against generative AI model releases. This subsection asks what happens at the thirteen other major AI product releases of 2023 to 2026, using the firm-level overlap scores of Section II.A. We divide the thirteen releases into business-oriented and consumer-oriented sets and estimate stacked event studies of CARs on standardized firm-level overlap, with AI-PDR and the controls interacted with event dummies and industry $\times$ event fixed effects (Table A25). For the business-oriented releases we find a robust but short-lived negative announcement return for firms whose products overlap with the functionality of the AI product. In the narrow business group (M365 Copilot GA, Claude for Financial Services, and Claude Managed Agents), a one-SD higher overlap is associated with a 0.31 p.p. lower CAR over $[-1, +1]$ ($t = -2.67$) and 0.31 p.p. over $[-1, +3]$ ($t = -2.65$) at GICS 6-digit $\times$ event fixed effects, and the effect loses significance by $[-1, +20]$. The broad business group, which adds Claude 3.5 computer use and Claude Code, shows the same short-lived pattern (-0.0018 , $t = -3.29$ at $[-1, +1]$).⁷ No effects are detected for the consumer-oriented releases through $[-1, +10]$, and only a small negative estimate appears at $[-1, +20]$.

The product-release evidence generalizes the displacement channel beyond Cowork Plugins while sharpening its uniqueness. Business-oriented product releases move overlapping firms briefly. Cowork Plugins, ex post the first commercially successful agentic AI release and ex ante the release with the widest and strongest overlap with incumbents, moves them several times more strongly and durably. In the narrow business group, a one-SD higher overlap is associated with a 0.31 p.p. lower CAR over $[-1, +3]$, an effect gone within twenty trading days. At Cowork Plugins, the standardized-overlap effect is -0.0179 per one-SD at $[-1, +10]$ (Table III), nearly six times as large.

⁷The narrow business-oriented group contains the releases whose products are business-oriented throughout. The broad group adds releases whose products, at the time of release, could not be accessed through a consumer-oriented web interface but only through an API or developer tooling. Claude 3.5 computer use therefore falls in the broad business-oriented group, available at issuance only through the API, while OpenAI Operator, released three months later with a consumer-oriented web interface, falls in the consumer-oriented group.

B. LLM Model Releases

The remaining AI releases are LLM model releases, upgrades to the underlying models rather than products that perform a specified set of work tasks. To test the AI-PDR result at these events, we re-estimate the Table II cross-sectional regression at 25 LLM model releases from late 2022 through July 2026. The set covers every flagship model release from OpenAI, Google, Anthropic, xAI, and DeepSeek over this period, including the two most talked-about releases of summer 2026 (OpenAI’s GPT-5.6 and Moonshot’s Kimi K3, both July 2026).

Figure A4 plots the AI-PDR coefficient with 95% confidence intervals at every model release, in the baseline specification with GICS 6-digit fixed effects and GICS 6-digit clustered standard errors. No model release produces a Cowork-style response. The largest coefficient in absolute value, -0.0228 at GPT-5.6, is about half of the Cowork coefficient of -0.0466 , and of the four releases that clear the 5 percent bar, one is positive (DeepSeek R1) and three are negative (OpenAI o1 (preview, reasoning model) at -0.0096 , Gemini 3 at -0.0072 , and GPT-5.6 at -0.0228). That includes the model releases that follow Cowork Plugins in 2026, GPT-5.4, Claude Mythos, GPT-5.6, and Kimi K3, upgrades that could strengthen agent abilities, consistent with the displacement repricing having already occurred at Cowork Plugins. Together with the previous subsection, the scan separates the two kinds of AI news. Product releases that ship work tasks move the firms whose products they overlap, and model releases move the cross-section little.

C. AI Input Intensity

AI-PDR captures the threat that Gen-AI substitutes for a firm’s output. By contrast, scaling Gen-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 A4) that assigns each extracted product to one of 23 AI-input categories grouped into four tiers (Gen-AI core, Gen-AI-adjacent infrastructure, semiconductor manufacturing, and general inputs), with NONE for products that fit no category. Firm-level AI Input Intensity is the importance-weighted fraction of the firm’s products classified into any category.

The classification has face validity. The top firms by market capitalization within each

category include both the major Gen-AI suppliers (e.g., NVIDIA, Broadcom, Alphabet, Microsoft) and specialized names in each specific category (e.g., CoreWeave for GPU-dedicated Gen-AI cloud, Vertiv for Gen-AI-specific data center cooling, Arista Networks for Gen-AI-class networking). Application-layer firms that brand themselves as “Gen-AI-powered” but do not supply an upstream input to building or running Gen-AI are classified as **NONE**, consistent with the construct capturing upstream inputs to Gen-AI rather than downstream Gen-AI consumption.

Table A26 adds AI Input Intensity to the event-study specification for all four events over four event windows, at GICS 2-digit and GICS 6-digit fixed effects. At ChatGPT Launch and Claude Opus 4.5, the AI Input Intensity coefficient is statistically indistinguishable from zero in nearly every cell.

DeepSeek R1 is materially different. At $[-1, +10]$, AI Input Intensity is significantly negative at both FE levels, and moving from 0 to 100 percent share of products classified into AI-input categories is associated with roughly a 5 p.p. lower CAR at GICS 6-digit FE. The coefficient stays significantly negative at $[-1, +30]$ at the sector level. The negative coefficient is consistent with a Gen-AI-cost-deflation narrative. More efficient reasoning models reduce the units of compute, memory, networking, and physical infrastructure required to deliver a unit of Gen-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 loss for any U.S.-listed stock on record.

Cowork Plugins shows a different pattern. The AI Input Intensity coefficient is negative at $[-1, +3]$ and turns positive and significant at both FE levels by $[-1, +10]$. Moving from 0 to 100 percent share of products classified into AI-input categories is associated with roughly a 4 p.p. lower CAR over $[-1, +3]$ and a 14 p.p. higher CAR over $[-1, +30]$ at GICS 6-digit FE. One explanation is that, before Cowork Plugins, the market was pricing in a scenario where high-AI-PDR firms would each train their own Gen-AI models in-house, sustaining strong Gen-AI-infrastructure demand. The release prompted a revision toward a winner-take-all situation in which only a few Gen-AI leaders train models, implying lower per-firm AI-input demand. By $[-1, +30]$ 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 Gen-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 agentic 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 Gen-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 A2 reports value-weighted buy-and-hold cumulative returns for three portfolios sorted on AI Input Intensity, over January 3, 2022 through September 21, 2026. Firms in the top half of the non-zero AI Input Intensity distribution returned 256.7% over the period, firms in the bottom half returned 103.8%, and firms with no AI-input classification returned 41.9%, against 60.7% for the value-weighted market benchmark. 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.

VI. Conclusion

Prior work has documented that Gen-AI can raise firm value by lowering labor costs. This paper shows that Gen-AI can destroy firm value by displacing products, and that equity markets distinguish between the two channels. We introduce Gen-AI Product Displacement Risk (AI-PDR), the digital deliverability of a firm’s product mix, to measure this threat. We focus on the release of Cowork Plugins, which overlaps with several times more incumbent Compustat-firm products than any other release. The repricing shows up at Cowork Plugins, and not at the generative AI model releases that precede it. Other business-oriented AI product releases move overlapping firm stock prices only briefly.

The repricing does not stand alone. Analysts revise down target prices, earnings forecasts, and long-term growth forecasts for high-AI-PDR firms. Investor messages on StockTwits endorsing displacement shoot up for those firms. In the second quarter of 2026, the revenue growth and gross margins of high-AI-PDR firms deteriorate relative to industry peers. Prices, forecasts, beliefs, and realized fundamentals move in the same direction when an agentic AI product enters incumbent product space at an unprecedented breadth and depth.

Over the medium to long horizon, the repricing diverges along a SaaS divide. SaaS firms sell the platform layer an agent runs on and can bill it as recurring usage, benefit from agentic adoption, and recover to their pre-Cowork-Plugins level within 160 trading days. Firms that release agentic AI products of their own finish the sample period significantly above industry peers. Non-SaaS firms sell curated data, ratings, reports, instruction, programs, and brokered expertise, the output generative AI also produces. They do not directly benefit from agentic AI adoption and their repricing persists.

The results speak to the strategic question for affected firms, which face two broad paths. They can race, investing in Gen-AI-augmented versions of existing products and pivoting to Gen-AI-native delivery, or retreat, harvesting remaining revenue and returning cash to shareholders. The first eight months after the release already contain an answer for many firms. Launches of agentic AI products accelerate to about 40 per month against 8 before, and launchers finish above peers with the same AI-PDR. For non-SaaS firms the choice is harder, since what they sell is not the layer agents run on, and often is the very output generative AI produces. 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. Whether high-AI-PDR non-SaaS firms can race at all, or must retreat, is a natural direction for future research.

A separate question is the structure of the displacer industry. Displacement could arrive from a small number of large Gen-AI providers substituting directly for incumbent products, or from a wave of Gen-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 shift profits from incumbent firms to them, 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 is a net welfare gain or a redistribution of market power.

A complete accounting of Gen-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. Today Gen-AI threatens mainly digital products and services. As it extends into physical goods and regulated services, the frontier of displacement risk will expand. The distinction between labor displacement and product displacement applies wherever a new technology creates close substitutes for existing goods, making the framework portable

beyond the Gen-AI setting studied here.

REFERENCES

- Acemoglu, Daron, 2025, The simple macroeconomics of AI, *Economic Policy* 40, 13–58.
- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, 2022, Artificial intelligence and jobs: Evidence from online vacancies, *Journal of Labor Economics* 40, S293–S340.
- Acemoglu, Daron, and Pascual Restrepo, 2019, Automation and new tasks: How technology displaces and reinstates labor, *Journal of Economic Perspectives* 33, 3–30.
- Acemoglu, Daron, and Pascual Restrepo, 2020, Robots and jobs: Evidence from US labor markets, *Journal of Political Economy* 128, 2188–2244.
- Adobe, 2026, Creative cloud all apps for teams pricing, adobe.com/creativecloud/business/teams, accessed May 2026.
- Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones, 2019, Artificial intelligence and economic growth, in *The Economics of Artificial Intelligence: An Agenda*, 237–290 (University of Chicago Press).
- Andrews, Isaiah, and Maryam Farboodi, 2025, Do markets believe in transformative AI?, Working paper 34243, National Bureau of Economic Research.
- Antweiler, Werner, and Murray Z. Frank, 2004, Is all that talk just noise? The information content of internet stock message boards, *Journal of Finance* 59, 1259–1294.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024, Artificial intelligence, firm growth, and product innovation, *Journal of Financial Economics* 151, 103745.
- Berkowitz, Ben, 2026, Anthropic tops \$100 billion revenue pace, report says, Axios, September 18, 2026, axios.com/2026/09/18/anthropic-100-billion-revenue.
- Bick, Alexander, Adam Blandin, and David Deming, 2026, The rapid adoption of generative AI, *Management Science* .
- Bloomberg Intelligence, 2026, Agentic AI 2026 outlook, Bloomberg Professional Services.
- Chen, Jiafeng, and Jonathan Roth, 2024, Logs with zeros? some problems and solutions, *Quarterly Journal of Economics* 139, 891–936.
- CNBC, 2026a, AI fears pummel software stocks: Is it “illogical” panic or a SaaS apocalypse?, February 6, 2026, cnbc.com/2026/02/06/ai-anthropic-tools-saas-software-stocks-selloff.html.
- CNBC, 2026b, Legal AI startup Harvey valued at \$11 billion, March 25, 2026, cnbc.com/2026/03/25/legal-ai-startup-harvey-raises-200-million-at-11-billion-valuation.html.
- CNBC, 2026c, OpenEvidence, the “ChatGPT for doctors,” doubles valuation to \$12 billion,

- January 21, 2026, [cnn.com/2026/01/21/open-evidence-chatgpt-for-doctors-doubles-valuation-to-12-billion.html](https://www.cnn.com/2026/01/21/open-evidence-chatgpt-for-doctors-doubles-valuation-to-12-billion.html).
- CNN Business, 2026, Anthropic’s new AI tool sends shudders through software stocks, February 4, 2026, [cnn.com/2026/02/04/investing/us-stocks-anthropic-software](https://www.cnn.com/2026/02/04/investing/us-stocks-anthropic-software).
- Cohn, Jonathan B., Zack Liu, and Malcolm I. Wardlaw, 2022, Count (and count-like) data in finance, *Journal of Financial Economics* 146, 529–551.
- Cookson, J. Anthony, Joseph E. Engelberg, and William Mullins, 2023, Echo chambers, *Review of Financial Studies* 36, 450–500.
- Cookson, J. Anthony, and Marina Niessner, 2020, Why don’t we agree? Evidence from a social network of investors, *Journal of Finance* 75, 173–228.
- Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* 52, 1035–1058.
- Eisfeldt, Andrea L., and Gregor Schubert, 2025, Generative AI and finance, *Annual Review of Financial Economics* 17, 363–393.
- Eisfeldt, Andrea L., Gregor Schubert, Miao Ben Zhang, and Bledi Taska, 2025, Generative AI and firm values, *Journal of Finance* forthcoming.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, 2024, GPTs are GPTs: Labor market impact potential of LLMs, *Science* 384, 1306–1308.
- Fast Company, 2026, Why one Anthropic update wiped billions off software stocks, [fastcompany.com/91487960/why-one-anthropic-update-wiped-billions-off-software-stocks](https://www.fastcompany.com/91487960/why-one-anthropic-update-wiped-billions-off-software-stocks).
- Felten, Edward, Manav Raj, and Robert Seamans, 2021, Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses, *Strategic Management Journal* 42, 2195–2217.
- Felten, Edward, Manav Raj, and Robert Seamans, 2023, How will language modelers like ChatGPT affect occupations and industries?, Working paper.
- Fortune, 2026, AI’s next act, how Salesforce is turning efficiency gains into revenue, April 18, 2026, fortune.com/2026/04/18/salesforce-agentforce-ai-efficiency-revenue-growth/.
- Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas, 2012a, Displacement risk and asset returns, *Journal of Financial Economics* 105, 491–510.
- Gârleanu, Nicolae, Stavros Panageas, and Jianfeng Yu, 2012b, Technological growth and asset pricing, *Journal of Finance* 67, 1265–1292.
- Habib-Valdhorn, Shiri, 2026, Surviving the software storm, *Globes*, May 26, 2026, [en.globe](https://www.globe.co.il/en)

s.co.il/en/article-surviving-the-software-storm-1001543960.

- Hadlock, Charles J., and Joshua R. Pierce, 2010, New evidence on measuring financial constraints: Moving beyond the KZ index, *Review of Financial Studies* 23, 1909–1940.
- Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller, 2025, Artificial intelligence and the labor market, Working paper 33509, National Bureau of Economic Research.
- He, Songrun, Linying Lv, Asaf Manela, and Jimmy Wu, 2025a, Chronologically consistent large language models, Working paper, Washington University in St. Louis.
- He, Songrun, Linying Lv, Asaf Manela, and Jimmy Wu, 2025b, Instruction tuning chronologically consistent language models, Working paper, Washington University in St. Louis.
- Hoberg, Gerard, and Gordon Phillips, 2010, Product market synergies and competition in mergers and acquisitions: A text-based analysis, *Review of Financial Studies* 23, 3773–3811.
- Hoberg, Gerard, and Gordon Phillips, 2016, Text-based network industries and endogenous product differentiation, *Journal of Political Economy* 124, 1423–1465.
- Hoberg, Gerard, Gordon Phillips, and Nagpurnanand Prabhala, 2014, Product market threats, payouts, and financial flexibility, *Journal of Finance* 69, 293–324.
- Kakhbod, Ali, Leonid Kogan, Peiyao Li, and Dimitris Papanikolaou, 2024, Measuring creative destruction, Working paper 7234-24, MIT Sloan School of Management.
- Kaplan, Steven N., and Luigi Zingales, 1997, Do investment-cash flow sensitivities provide useful measures of financing constraints?, *Quarterly Journal of Economics* 112, 169–215.
- Kogan, Leonid, and Dimitris Papanikolaou, 2014, Growth opportunities, technology shocks, and asset prices, *Journal of Finance* 69, 675–718.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *Quarterly Journal of Economics* 132, 665–712.
- Kogan, Leonid, Dimitris Papanikolaou, and Noah Stoffman, 2020, Left behind: Creative destruction, inequality, and the stock market, *Journal of Political Economy* 128, 855–906.
- Kothari, S. P., and Jerold B. Warner, 2007, Econometrics of event studies, in B. Espen Eckbo, ed., *Handbook of Empirical Corporate Finance*, volume 1, chapter 1, 3–36 (Elsevier, Amsterdam).
- Pástor, Luboš, and Pietro Veronesi, 2009, Technological revolutions and stock prices, *American Economic Review* 99, 1451–1483.
- Ramelli, Stefano, and Alexander F. Wagner, 2020, Feverish stock price reactions to COVID-

- 19, *Review of Corporate Finance Studies* 9, 622–655.
- Salesforce, 2026a, Sales cloud pricing, salesforce.com/sales/pricing, accessed May 2026.
- Salesforce, 2026b, Salesforce delivers record fourth quarter fiscal 2026 results, Press release, February 25, 2026, salesforce.com/news/press-releases/2026/02/25/fy26-q4-earnings/.
- Schumpeter, Joseph A., 1942, *Capitalism, Socialism and Democracy* (Harper and Brothers, New York).
- ServiceNow, 2026, ServiceNow reports second quarter 2026 financial results, Press release, July 22, 2026, newsroom.servicenow.com/press-releases/details/2026/ServiceNow-Reports-Second-Quarter-2026-Financial-Results/default.aspx.
- TechCrunch, 2026, Anthropic brings agentic plug-ins to Cowork, January 30, 2026, techcrunch.com/2026/01/30/anthropic-brings-agentic-plugins-to-cowork/.
- Vishwanath, Krithik, Anton Alyakin, Mrigayu Ghosh, Ali Hage, Sean N. Neifert, Cordelia Orillac, Nataniel J. Mandelberg, Hammad A. Khan, Jin Vivian Lee, Jie J. Yao, William Robert Small, Aakaash Varma, D. Brock Hewitt, Yindalon Aphinyanaphongs, Daniel Alexander Alber, and Eric Karl Oermann, 2026, General-purpose large language models outperform specialized clinical AI tools on medical benchmarks, *Nature Medicine* 32, 2405–2409.
- Webb, Michael, 2020, The impact of artificial intelligence on the labor market, Working paper, Stanford University.
- Whited, Toni M., and Guojun Wu, 2006, Financial constraints risk, *Review of Financial Studies* 19, 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, and whiskers extend to $1.5 \times$ 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). The top row is AI-PDR (Digital Deliverability) and the bottom row the composite Shield (average of the four Shield sub-dimensions). 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. AI-PDR Effect on CARs 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 $d = 0, 5, \dots, 160$ trading days, and the shaded band is the 95% confidence interval. The top panel uses GICS 2-digit (sector) and the bottom panel GICS 6-digit (industry) fixed effects. Controls, the estimation universe, and clustering follow the baseline specification in Equation (1).




Figure 4. AI-PDR Effect on CARs across Event Windows, Other Main-Sample Events

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 $d = 0, 5, 10, \dots, 160$ trading days. Columns 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 top and bottom rows use GICS 2-digit (sector) and GICS 6-digit (industry) fixed effects. Controls and the estimation universe are as in Figure 3. Shaded bands are 95% confidence intervals from GICS 6-digit clustered standard errors. 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.




Figure 5. AI-PDR Effect across Event Windows, SaaS and Non-SaaS Firms (Cowork Plugins)

Note. Each panel plots the SaaS or non-SaaS AI-PDR slope from the pooled regressions of Table IV at GICS 6-digit fixed effects, one regression per window end d from 0 to 160 trading days in steps of 10. The top row uses the any-line SaaS rule and the bottom row the majority rule, as defined in Table IV and Appendix Table A30, whose controls, estimation universe, and scaling of AI-PDR also apply here. Each panel title reports the group's firm count at the longest window. Shaded bands are 95% confidence intervals from GICS 6-digit clustered standard errors.




Figure 6. Text-Based Investor Endorsement of Generative AI Displacement by Week

Note. Each panel plots the average across firms of the firm-level number of endorsing messages in that week, keeping firms with at least 5 qualifying messages of at least 20 characters in the week. Endorsing messages pass the keyword prefilter and two rounds of LLM classification. Shaded bands are 95% confidence intervals for the cross-firm mean. Week 0 is the seven days beginning on the event date, and the dashed line marks the event.




Figure 7. Quarterly Path of the AI-PDR Coefficient on Firm Fundamentals

Note. Each point is the AI-PDR coefficient from the full-sample GICS 2-digit specification in columns (1) and (2) of Table IX, estimated cohort by cohort from 2024 Q1 through 2026 Q2 and reported in percentage points. The cohort for a calendar quarter contains every firm with a fiscal quarter ending inside that calendar quarter, with outcomes defined as in Table IX. Shaded bands are 95% confidence intervals from GICS 6-digit clustered standard errors. The dashed line marks the Cowork Plugins release (January 30, 2026), which falls between the 2025 Q4 and 2026 Q1 cohorts.

Table I
Summary Statistics

Panel A — AI-PDR and Firm Controls								
	<i>N</i>	Mean	SD	P10	P25	P50	P75	P90
AI-PDR (Digital Deliverability)	3,984	2.08	2.54	0.00	0.00	1.20	2.75	6.62
Legal/Regulatory	3,984	5.64	3.06	1.85	2.67	5.64	8.62	10.00
Production Moats	3,984	4.85	1.80	2.50	3.50	4.80	6.14	7.15
Customer Trust	3,984	6.36	2.08	3.14	4.95	6.88	8.00	8.71
Customer Embeddedness	3,984	4.84	2.30	1.40	3.11	5.02	6.67	7.75
AI Input Intensity	3,984	0.05	0.16	0.00	0.00	0.00	0.00	0.10
Cowork Overlap	2,950	0.34	0.79	0.00	0.00	0.00	0.21	1.17
log(ME)	4,059	7.59	1.77	5.39	6.19	7.44	8.73	10.11
Tobin's Q	4,059	2.17	1.97	0.96	1.05	1.47	2.47	4.24
ROA	4,059	0.03	0.19	-0.20	0.01	0.07	0.12	0.18
Asset Growth	4,059	0.22	0.97	-0.13	-0.02	0.05	0.15	0.43

Panel B — Pairwise Correlations					
	Dig.	Legal	Prod.	Trust	Embed.
Dig.	1.00				
Legal	-0.37	1.00			
Prod.	-0.16	0.28	1.00		
Trust	0.05	0.31	0.49	1.00	
Embed.	0.29	-0.05	0.38	0.56	1.00

Note. Panel A reports cross-sectional distributions for the firm sample. AI-PDR, the four Shield sub-dimensions, and Cowork Overlap are filing-year 2025 importance-weighted firm-level averages on the 0–10 scale, from the most recent 10-K filing preceding the Cowork Plugins launch. AI Input Intensity is the filing-year 2025 share of the firm's products in an AI-input taxonomy category. Controls are the most recent pre-event Compustat fiscal year, winsorized at the 1st and 99th percentiles. Panel B reports pairwise Pearson correlations among AI-PDR and the Shield sub-dimensions ($N = 3,984$).

Table II
Event-Study Regressions of CARs on AI-PDR

Event	AI-PDR	N	R^2
ChatGPT Launch	0.0068* (1.82)	2,832	0.0911
DeepSeek R1	0.0075** (2.01)	2,887	0.1446
Claude Opus 4.5	-0.0054 (-1.58)	2,763	0.1400
Cowork Plugins	-0.0466*** (-8.34)	2,845	0.3633

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 $\log(\text{ME})$, 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 AI-PDR from the single 10-K filing year immediately before the event (filing year 2021 for ChatGPT, 2024 for DeepSeek R1 and Claude Opus 4.5, 2025 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 III
AI-PDR Effect by Cowork Plugin Overlap (Cowork Plugins event only)

	(1)	(2)	(3)	(4)
AI-PDR	−0.0535*** (−10.07)	−0.0359*** (−5.41)	−0.0520*** (−9.25)	−0.0351*** (−5.30)
Cowork Overlap (std)	−0.0240*** (−6.23)	−0.0179*** (−5.30)		
1(Max Overlap ≥ 3)			−0.0206 (−1.47)	−0.0121 (−1.43)
AI-PDR × 1(Max Overlap ≥ 3)			−0.0235*** (−2.62)	−0.0175*** (−2.86)
Controls	Yes	Yes	Yes	Yes
FE	GICS-2	GICS-6	GICS-2	GICS-6
N	2,785	2,766	2,783	2,764
R^2	0.292	0.379	0.286	0.375

Note. All columns use the Cowork Plugins regression sample with CAPM-adjusted CARs over the $[-1, +10]$ window and the baseline controls (log(ME), Tobin’s Q, ROA, asset growth, each winsorized at 1/99). CARs are also winsorized at 1/99. Columns (1) and (2) add the firm-level Cowork Overlap (the importance-weighted mean of product-level overlap scores, standardized to unit variance and demeaned within the sample) next to AI-PDR, so its coefficient is the overlap main effect beyond AI-PDR. Columns (3) and (4) interact AI-PDR with an indicator that the firm’s highest product-level overlap score is 3 or above. AI-PDR is standardized to unit variance and demeaned within the estimation sample, and the interaction is the product of standardized AI-PDR and the indicator. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M. Within each specification the first column uses GICS 2-digit and the second GICS 6-digit industry fixed effects, dropping fixed-effect cells with fewer than 5 firms. GICS 6-digit industry-clustered t -statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table IV
AI-PDR across Event Windows by SaaS Classification (Cowork Plugins)

Window	CAR				Char.-matched			
	SaaS × AI-PDR	Non-SaaS × AI-PDR	SaaS × AI-PDR	Non-SaaS × AI-PDR	SaaS × AI-PDR	Non-SaaS × AI-PDR	SaaS × AI-PDR	Non-SaaS × AI-PDR
[−1, +10]	−0.0767*** (−10.54)	−0.0598*** (−8.19)	−0.0460*** (−5.86)	−0.0413*** (−7.00)	−0.0679*** (−8.90)	−0.0563*** (−8.56)	−0.0400*** (−4.90)	−0.0401*** (−6.52)
[−1, +20]	−0.0800*** (−10.65)	−0.0581*** (−7.01)	−0.0470*** (−4.97)	−0.0420*** (−5.15)	−0.0725*** (−9.61)	−0.0557*** (−7.34)	−0.0429*** (−4.80)	−0.0424*** (−5.38)
[−1, +30]	−0.0666*** (−8.31)	−0.0468*** (−4.90)	−0.0389*** (−3.63)	−0.0365*** (−3.22)	−0.0567*** (−6.88)	−0.0437*** (−5.00)	−0.0335*** (−3.45)	−0.0375*** (−3.55)
[−1, +40]	−0.0796*** (−9.15)	−0.0531*** (−5.56)	−0.0485*** (−4.78)	−0.0438*** (−3.89)	−0.0665*** (−7.16)	−0.0494*** (−5.38)	−0.0423*** (−4.42)	−0.0467*** (−4.52)
[−1, +80]	−0.1129*** (−4.26)	−0.0664*** (−3.69)	−0.0489* (−1.88)	−0.0457*** (−2.70)	−0.1040*** (−3.70)	−0.0623*** (−3.22)	−0.0401 (−1.42)	−0.0409** (−2.20)
[−1, +120]	−0.0816*** (−4.02)	−0.0747*** (−4.44)	−0.0330 (−1.32)	−0.0643*** (−3.74)	−0.0727*** (−3.30)	−0.0713*** (−4.11)	−0.0235 (−0.86)	−0.0608*** (−3.28)
[−1, +160]	−0.0389* (−1.70)	−0.0443** (−2.39)	−0.0009 (−0.03)	−0.0474** (−2.31)	−0.0375 (−1.50)	−0.0430** (−2.22)	0.0047 (0.16)	−0.0431* (−1.94)
FE	GICS-2	GICS-2	GICS-6	GICS-6	GICS-2	GICS-2	GICS-6	GICS-6
<i>N</i>	526	2,327	525	2,308	526	2,327	525	2,308

Note. Each pair of columns reports the SaaS and non-SaaS AI-PDR slopes from one pooled cross-sectional OLS regression of the CAPM-adjusted CAR over the indicated window around Cowork Plugins on a SaaS indicator, SaaS × AI-PDR, non-SaaS × AI-PDR, and the firm controls. Under the any-line rule a firm is SaaS when any product line in its most recent 10-K filing year is classified as SaaS and non-SaaS only when none is. All columns use the any-line rule, and Appendix Table A30 repeats the table under the majority rule. Within each benchmark the first column pair uses GICS 2-digit and the second GICS 6-digit industry fixed effects, dropping fixed-effect cells with fewer than five firms. AI-PDR is scaled by its full-sample standard deviation everywhere. Controls, winsorization, and the estimation universe follow Table V, and the *N* row reports each column's firm count at the longest window. Columns (5) to (8) repeat columns (1) to (4) with the characteristics-matched abnormal return of Appendix Table A27 as the dependent variable. GICS 6-digit industry-clustered *t*-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table V
Agentic AI Launches and CARs (Cowork Plugins)

Window	CAR				Char.-matched			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
[−1, +10]	−0.0273* (−1.71)	−0.0176 (−1.02)	−0.0279* (−1.79)	−0.0185 (−1.11)	−0.0225 (−1.28)	−0.0139 (−0.76)	−0.0234 (−1.35)	−0.0149 (−0.83)
[−1, +20]	−0.0433*** (−3.05)	−0.0281 (−1.64)	−0.0426*** (−3.15)	−0.0288* (−1.79)	−0.0371*** (−2.85)	−0.0242 (−1.64)	−0.0368*** (−2.97)	−0.0249* (−1.81)
[−1, +30]	−0.0304 (−1.32)	−0.0188 (−0.74)	−0.0298 (−1.30)	−0.0199 (−0.80)	−0.0257 (−1.33)	−0.0160 (−0.80)	−0.0250 (−1.30)	−0.0171 (−0.87)
[−1, +40]	−0.0312 (−1.03)	−0.0212 (−0.68)	−0.0319 (−1.07)	−0.0232 (−0.76)	−0.0288 (−1.21)	−0.0215 (−0.96)	−0.0290 (−1.24)	−0.0233 (−1.08)
[−1, +80]	−0.0730** (−2.19)	−0.0301 (−0.77)	−0.0704** (−2.11)	−0.0306 (−0.79)	−0.0499 (−1.52)	−0.0097 (−0.25)	−0.0484 (−1.49)	−0.0106 (−0.28)
[−1, +120]	−0.0245 (−0.82)	0.0090 (0.25)	−0.0255 (−0.86)	0.0069 (0.19)	−0.0027 (−0.09)	0.0287 (0.84)	−0.0046 (−0.16)	0.0264 (0.78)
[−1, +160]	0.0552* (1.94)	0.0772** (2.55)	0.0523* (1.89)	0.0746** (2.51)	0.0734** (2.54)	0.0955*** (3.07)	0.0696** (2.50)	0.0924*** (3.05)
Controls	AI-PDR	AI-PDR	AI-PDR × Shield	AI-PDR × Shield	AI-PDR	AI-PDR	AI-PDR × Shield	AI-PDR × Shield
FE	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6
<i>N</i>	2,865	2,845	2,865	2,845	2,865	2,845	2,865	2,845

Note. Each cell reports the coefficient on an indicator equal to one for firms with at least one agentic AI product launch dated on or after January 30, 2026, from the LSEG launch list described in the text, in a cross-sectional OLS regression of the CAPM-adjusted CAR over the indicated window around Cowork Plugins. All columns include AI-PDR (standardized to unit variance and demeaned within the sample) and the firm controls (log(ME), Tobin’s Q, ROA, asset growth, winsorized at 1/99). Columns (3)–(4) add the composite Shield, standardized and demeaned, and its AI-PDR interaction, the product of the standardized, demeaned terms. Within each pair the first column uses GICS 2-digit and the second GICS 6-digit industry fixed effects, dropping fixed-effect cells with fewer than five firms. CARs use each firm’s market beta estimated over [−252, −11] and are winsorized at 1/99. The estimation universe is firms with market capitalization above \$100M and Compustat total assets above \$100M, with the *N* row reporting each column’s firm count. Columns (5) to (8) repeat columns (1) to (4) with the characteristics-matched abnormal return of Appendix Table A27 as the dependent variable. GICS 6-digit industry-clustered *t*-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table VI
Interaction of AI-PDR with Competitive Shields (Cowork Plugins)

	Composite		Supply Side		Demand Side	
	(1)	(2)	(3)	(4)	(5)	(6)
AI-PDR	−0.0679*** (−13.15)	−0.0446*** (−8.09)	−0.0613*** (−10.26)	−0.0419*** (−6.17)	−0.0731*** (−13.94)	−0.0491*** (−9.32)
Shield	0.0056 (1.02)	0.0014 (0.31)	0.0032 (0.40)	−0.0067 (−1.04)	0.0059 (1.16)	0.0055 (1.44)
AI-PDR × Shield	0.0126*** (3.00)	0.0121*** (3.29)	0.0147*** (3.63)	0.0135*** (2.92)	0.0065 (1.48)	0.0077** (2.14)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
FE	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6
N	2,865	2,845	2,865	2,845	2,865	2,845
R^2	0.2745	0.3667	0.2753	0.3678	0.2718	0.3652

Note. Each column is a cross-sectional OLS regression of CAPM-adjusted CARs over the $[-1, +10]$ window around the Cowork Plugins event (January 30, 2026) on AI-PDR (Digital Deliverability), a Shield score, and their interaction. Columns (1) and (2) use the composite Shield, the simple average of Legal/Regulatory, Production Moats, Customer Trust, and Customer Embeddedness. Columns (3) and (4) use the supply-side shield, the average of Legal/Regulatory and Production Moats, and columns (5) and (6) the demand-side shield, the average of Customer Trust and Customer Embeddedness. The rows labeled Shield and AI-PDR × Shield refer to whichever Shield definition the column uses. Within each pair the first column uses GICS 2-digit and the second GICS 6-digit industry fixed effects, dropping fixed-effect cells with fewer than five firms. All score regressors are standardized to unit variance within the estimation sample and demeaned, and interactions are products of the standardized-demeaned regressors. Each regression includes $\log(\text{ME})$, Tobin's Q, ROA, and asset growth (each 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. GICS 6-digit industry-clustered t -statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table VII
Analyst Forecast Revisions (Cowork Plugins)

	(1) EPS revision	(2) Revenue revision	(3) LTG change	(4) Target Price revision
AI-PDR	−0.029*** (-2.66)	−0.009* (-1.85)	−1.541** (-2.06)	−0.105*** (-5.32)
Controls	Yes	Yes	Yes	Yes
FE	GICS-2	GICS-2	GICS-2	GICS-2
N	1,960	2,533	1,028	2,655
R^2	0.099	0.098	0.045	0.149

Note. The sample is all firms with market capitalization above \$100M and Compustat total assets above \$100M, a non-missing AI-PDR score, GICS classification, and firm controls. Each column is a cross-sectional OLS regression of a scaled analyst-forecast revision around the Cowork Plugins event (January 30, 2026) on AI-PDR, the firm controls (log(ME), Tobin’s Q, ROA, asset growth, winsorized at 1/99), and GICS 2-digit sector fixed effects. Columns (1) and (2) are revisions in EPS and revenue forecasts for the first fiscal year ending on or after December 2026. Column (3) is the change in the long-term growth (LTG) forecast, not tied to a specific fiscal year, and column (4) the revision in the 12-month target price. For each firm the pre-event observation is its last IBES summary consensus before January 30, 2026, and the post-event observation its last consensus after the event through June 30, 2026. The revision is $y_i = (\text{post}_i - \text{pre}_i)/\text{pre}_i$ for EPS, revenue, and target price, and the level change $y_i = \text{post}_i - \text{pre}_i$ (capped at ± 50 percentage points) for LTG. AI-PDR is standardized to unit variance and demeaned within the estimation sample. EPS observations with pre-event EPS below \$0.50 (including all negative-EPS firms) and revenue observations with pre-event revenue below \$10M are dropped, and scaled revisions are winsorized at 1/99. GICS 6-digit industry-clustered t -statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table VIII
Text-Based Investor Endorsement of Generative AI Displacement

	Level (p.p.)		Log count		Calibrated log	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A — ChatGPT Launch (pre-event mean 0.000 p.p., $N = 2,434$)						
AI-PDR	0.014	0.017	0.004	0.003	0.005	0.005
	(1.04)	(0.99)	(1.52)	(1.07)	(1.50)	(1.07)
Panel B — DeepSeek R1 (pre-event mean 0.005 p.p., $N = 1,384$)						
AI-PDR	0.002	0.001	0.005	0.002	0.008	0.004
	(0.70)	(0.33)	(1.42)	(0.60)	(1.54)	(0.72)
Panel C — Claude Opus 4.5 (pre-event mean 0.029 p.p., $N = 1,484$)						
AI-PDR	-0.033	-0.032	-0.015	-0.007	-0.020	-0.009
	(-1.40)	(-1.21)	(-1.02)	(-0.68)	(-1.04)	(-0.66)
Panel D — Cowork Plugins (pre-event mean 0.023 p.p., $N = 1,547$)						
AI-PDR	0.085***	0.036*	0.091***	0.059***	0.126***	0.082***
	(3.64)	(1.87)	(4.69)	(2.76)	(4.87)	(2.84)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
FE	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6

Note. Each panel reports cross-sectional OLS regressions of the pre-to-post change in the firm-level endorsement of product displacement on AI-PDR, one observation per firm. The sample contains firms with at least 10 qualifying messages (i.e., at least 20 characters) in both the four weeks from the event day and the four weeks before. Endorsing messages pass the keyword prefilter and two rounds of LLM classification. The Level columns use the pre-to-post change in the endorsement fraction, $100 \times (\text{endorsing messages} / \text{qualifying messages})$, in p.p. The Log count columns use the change in the log of one plus the firm's endorsing messages and add the change in the log of its qualifying messages as a control. The Calibrated log columns use the change in $m(N)$, with $m(N) = \log N$ for a firm with $N \geq 1$ endorsing messages and $m(0) = -1$, also controlling for the change in the log of its qualifying messages (Chen and Roth, 2024). AI-PDR is the event-year score of Table II, standardized to unit variance within each event sample, and firm controls and their winsorization also follow Table II. Within each specification the first column includes GICS 2-digit fixed effects and the second GICS 6-digit fixed effects, and t -statistics in parentheses are clustered by GICS 6-digit industry. *, **, *** denote significance at 10%, 5%, 1%.

Table IX
Post-Release Change in Firm Fundamentals and AI-PDR

	Full sample		Δ Revenue growth		Δ Gross margin	
	Δ Revenue growth	Δ Gross margin	SaaS	Non-SaaS	SaaS	Non-SaaS
	(1)	(2)	(3)	(4)	(5)	(6)
AI-PDR	−0.0372*** (−3.81)	−0.0089*** (−4.14)	−0.0436*** (−3.73)	−0.0367** (−2.36)	−0.0122*** (−3.13)	−0.0090*** (−2.60)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
FE	GICS-2	GICS-2	GICS-2	GICS-2	GICS-2	GICS-2
N	2,217	2,190	2,211		2,184	
R^2	0.039	0.043	0.040		0.043	

Note. The sample is the paper's estimation universe (firms with market capitalization above \$100M and Compustat total assets above \$100M) merged to quarterly Compustat, one observation per firm, taken in the firm's fiscal quarter ending inside the calendar quarter April 1 to June 30, 2026. Revenue is quarterly revenue (revtq). Gross margin is $1 - \text{cogsq}/\text{revtq}$. Growth rates compare a fiscal quarter with the same fiscal quarter one year earlier, revenue growth is $\ln(\text{revtq}_q/\text{revtq}_{q-4})$, requiring revenue to exceed \$1M in the two quarters, and the margin columns require revenue above \$5M. For the revenue growth columns the outcome is the change in the year-over-year growth rate, the growth rate in the outcome quarter minus the growth rate in the same fiscal quarter one year earlier, and for the margin columns it is the year-over-year margin change in the outcome quarter. Quarterly outcomes are winsorized at 1/99 across the firm-quarter panel and again within the regression sample. Coefficients are in decimal units. Columns (1) and (2) report cross-sectional OLS regressions of the indicated outcome on AI-PDR (the event-year score from filing year 2025, standardized to unit variance within the sample), controls ($\log(\text{ME})$, Tobin's Q, ROA, asset growth, each winsorized at 1/99), and GICS 2-digit sector fixed effects, dropping GICS 6-digit industry cells with fewer than five firms. Columns (3) to (6) report the SaaS and non-SaaS AI-PDR slopes for the indicated outcome, where each pair of columns comes from one pooled regression of the same specification that replaces AI-PDR with $\text{SaaS} \times \text{AI-PDR}$ and non-SaaS $\times$ AI-PDR and adds a SaaS indicator, on the subsample of firms covered by the SaaS classification. Under the any-line rule a firm is SaaS when any product line in its most recent 10-K filing year is classified as SaaS and non-SaaS only when none is. The N and R^2 entries spanning a pair of columns refer to that pooled regression. GICS 6-digit industry-clustered t -statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Internet Appendix

Listing A1. Shield scoring prompts.

The four prompts are re-typeset from the deployed text with the anchor and rule content preserved. In the deployed prompts every anchor list is preceded by the instruction “Each level below describes a category, examples in parentheses are illustrative, not exhaustive. Score by matching the category, not the specific examples,” each prompt closes with “Return one object per product in the input, echoing the same index and name,” and the rule text varies slightly by dimension.

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.

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.

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, 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 bank-

ing, 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.
6. Score what the customer risks by choosing an unknown provider. A borrower receiving funds risks little, so lending products score low. A customer who hands over money, data, or safety and depends on the provider's future performance (deposits, custody, insurance, annuities) risks a great deal, so such products score high.

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 numbers [0, the odd integers 1 through 9, and 10], 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. Bundle and product share no meaningful functions.
- 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. Do not feel obliged to 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.

Rules 1 and 2 of the deployed prompt are merged above, so the numbering differs from the code's 1 through 7, and the deployed prompt closes with the same products-to-score input block as Listing [A5](#).

Listing A3. StockTwits displacement classification prompt.

Messages first pass a keyword prefilter requiring at least one match from a broad vocabulary in four families. Keywords are matched as word stems, so that, for example, the stem *displac* captures *displace*, *displaced*, and *displacement*, and common misspellings of model and product names are also included.

Generative-AI terms. 4o, agentic, AI, Anthropic, Apple Intelligence, artificial intelligence, Bard, bot, chatbot, ChatGPT, Claude, Cohere, Copilot, Cowork, Cursor, DALL-E, deep learning, DeepSeek, Fable, Gemini, GenAI, generative, GPT, GPT-5, Grok, Haiku, Hugging Face, image, video, content, or text generation, intelligence tools, Lingonaut, Llama, LLM, machine learning, Midjourney, Mistral, ML, Mythos, Nano Banana, neural net, o3, o4, OpenAI, Opus, Perplexity, prompt, Qwen, R1, Replit, robot, Seedance, Skynet, Sonnet, Sora, Stable Diffusion, the machines, Thumio, and xAI.

Displacement and competition language. chicken little, commoditize, compete, competition, cooked, disappear, displace, disrupt, doomsday, dying, easily, erode, genie, gold standard, in-house, migrate, moat, obsolete, outcompete, overblown, replace, replicate, sticky, substitute, threat, toast, translate, vibe, vibe-code, wrapper, wreck, and phrases such as anyone can, build an app, build your own, can't touch, don't need to pay, go to zero, going away, is dead, it's over for, just fine, make your own, no going back, no need for, not one company can, rip out, sky is falling, so easy, that easy, who needs, why would anyone pay, will die, and within a day (also an hour, a minute, a weekend, or a month).

Software and subscription terms. code, coding, SaaS, seat, software, and subscription.

Obsolescence analogies. Blockbuster, buggy whip, Chegg, dinosaur, dodo, fax machine, horse and buggy, Kodak, MySpace, Nokia, and typewriter.

Non-matching messages are labeled not relevant without a classifier call. The system prompt applied to the passing messages is as follows.

You classify StockTwits messages. Each message is tagged to one company, whose name, ticker, and a short product profile are given. Use the company name and the product profile to judge what business the company is in. Classify each message independently.

The question: does the message claim that this company's own products or services are being or will be displaced (replaced, made obsolete, or out-competed) because of generative AI?

Generative AI here means large language models (such as ChatGPT, Claude, or Gemini) and the chatbots, agents, assistants, and coding tools built on them.

Generative AI can cause this in three forms, and any of them counts:

- (1) an AI provider’s product directly does what the company’s product does;
- (2) generative AI’s strong coding ability makes software so cheap to build that many more rivals can compete with the company;
- (3) the company’s customers use AI tools to build their own replacement for the company’s product instead of paying for it.

In forms (2) and (3) the direct displacer is a new competitor or the customer, not the AI provider itself, and it still counts as long as generative AI is what made it possible.

LABELS:

- “**endorses_displacement**”: the messenger states, predicts, or agrees that generative AI will (or already does) displace THIS company’s products, via one or more forms. The message must contain an actual statement that THIS company’s own product is being or will be replaced, made obsolete, or out-competed because of generative AI. A bare AI mention, generic bearishness, or “AI is bad for this sector” is not enough. If you cannot point to the phrase where the messenger says this company’s product is displaced, label **not_relevant**.
- “**not_relevant**”: anything else.

RULES. Apply every rule below. When a rule applies, it decides the label, no matter how the message superficially sounds.

1. Judge the messenger’s expressed stance, not whether displacement is truly likely.

The company must be one generative AI can displace:

2. If the company is NOT a software/digital/knowledge-work business (use the name and profile), label **not_relevant**. Chips and other hardware, data centers, pharma/biotech, cars, autonomous driving/robotaxis, robots, airlines, retail, banks, REITs, crypto/miners, etc. are NOT displaced by generative AI here.
3. If the company BENEFITS from or SELLS AI (GPUs/chips, data centers, cloud, or other AI infrastructure and tooling), label **not_relevant**. This covers two message types.
 - The message says that cheaper or more efficient AI (for example an open-source or low-cost model) REDUCES DEMAND for the company’s chips, other hardware, compute, or infrastructure, or that AI providers or models compete with each other. Falling demand for an AI input the company sells is NOT generative AI displacing the company’s product.

- The message is simply bullish that AI demand HELPS the company because it sells the inputs AI needs.

The displaced product must not be the company's own generative AI:

4. If the displaced product is the company's OWN generative AI model or the chatbots, agents, assistants, and coding tools built on them, beaten or made worthless by a rival's AI model, label **not_relevant**. A message that the company's spending on or investments in generative AI was wasted because of a better or cheaper generative AI model built by a rival is also **not_relevant**. A claim that generative AI displaces the company's non-generative-AI products (e.g., software, search, content) still counts as endorsing displacement even when the company also builds AI models.

The message must make a real generative-AI displacement claim:

5. The message itself must attribute the displacement to generative AI in the message's own words. A prediction of decline, obsolescence, or lost business with no generative-AI reference in the message text is **not_relevant**, even when an AI cause is plausible from context.
6. The displacement claim must be about this company's own product or service, the tagged company, not the sector in general. A message may cashtag several tickers, and if the displacement is about a different cashtagged ticker, it is **not_relevant** for this company.
7. If the message is about AI being overvalued, a bubble, or AI stocks crashing, that is AI valuation, not generative AI displacing this company's product, label **not_relevant**.
8. If the message is generic bearishness/bullishness, price/chart/options talk, short-seller fights, regulatory rumors, or a bare AI mention with no displacement claim, label **not_relevant**.

The claim must be the messenger's own stance:

9. If the message is sarcastic, joking, or mocking the displacement fear, often by absurd over-generalization like "AI can do everything, no software company needed", the messenger is dismissing the fear, not endorsing it, label **not_relevant**. Judge the messenger's actual intended stance, not the literal words.
10. If the messenger merely reports that others believe in displacement (the market, analysts, news media) or only asks a question about it, that is not their own endorsement, label **not_relevant**, unless the messenger clearly signals agreement with the fear. This includes formatted news digests and recap-bot posts that summarize guidance, analyst notes, headlines, or market moves without the messenger's own displacement claim.

And the cashtag must really mean the company:

11. StockTwits auto-cashtags any word that is also a ticker symbol, so a cashtag does not always

refer to the company. When the ticker is used as the ordinary word it spells rather than as the company (for example “\$AI” meaning artificial intelligence in general, as in “the \$AI race”, not the company C3.ai), the message is not about the company at all, label `not_relevant`.

Examples, messages with the correct outputs, for calibration.

1. **Company:** Adobe (\$ADBE) | **Message:** *\$ADBE What no one want to talk about: AI and Vibe coding is killing BIG Software*
→ label `endorses_displacement`, forms [1, 2]
2. **Company:** SoundHound AI (\$SOUN) | **Message:** *\$SPY \$QQQ \$STUDY I think \$SOUN is another good example. Whats the moat? How do you keep competitors out of your space??? Literally anyone can now create an AI voice assistant for drive throughs etc. The barrier to entry for AIaaS is rapidly falling. I'd encourage people to expirment and play with these tools.*
→ label `endorses_displacement`, forms [2]
3. **Company:** Duolingo (\$DUOL) | **Message:** *\$DUOL I vibe coded a Duolingo replacement overnight*
→ label `endorses_displacement`, forms [3]
4. **Company:** Super Micro Computer (\$SMCI) | **Message:** *\$NVDA \$smci \$mu \$DELL \$AVGO UC Berkeley replicated open source DeepSeek for \$30. This is a huge blow at OpenAI who is raising hundreds of billions of dollars. AI no need huge \$\$ for its hardware or software dev. and uses. Bearish in short and intermediate term. <https://tribune.net.ph/2025/01/31/uc-berkeley-researchers-claim-to-replicate-deepseek-ai-for-30>*
→ label `not_relevant`, forms [] (SMCI sells AI hardware, cheaper AI cutting demand for its gear is rule 3, not displacement of its product)
5. **Company:** C3.ai (\$AI) | **Message:** *\$NVDA when does the \$AI bubble pop?*
→ label `not_relevant`, forms [] (\$AI here is the word, not the company, rule 11, and bubble talk is valuation, rule 7)
6. **Company:** Intuit (\$INTU) | **Message:** *\$INTU 11:12 AM EST, 01/28/2026 (MT Newswires) – Intuit (INTU) shares increasingly appear dislocated, and the pullback in recent months on artificial intelligence fears could present a buying opportunity, RBC Capital Markets said in a Tuesday research report. Intuit can develop and embed AI capabilities into TurboTax to drive retention, but there remains a risk that more clients use AI instead of consulting TurboTax or experts. In the long run, the company could build its own tax agents using its data, domain expertise, and client trust, analysts wrote.*
→ label `not_relevant`, forms [] (a news digest reporting an analyst note, rule 10, not the

messenger's own displacement claim)

7. **Company:** Meta Platforms (\$META) | **Message:** *\$META China's deep seek is deemed better than metas llama for a fraction of the price, and the stock goes up because they're going to spend 10 billion more than consensus... Welcome to the upside down*
→ label **not_relevant**, forms [] (a rival AI model beating the company's own AI model is rule 4, not displacement of a non-generative-AI product)
8. **Company:** Microsoft (\$MSFT) | **Message:** *\$MSFT Not low enough. Claude AI replaces most MSFT apps. For FREE.*
→ label **endorses_displacement**, forms [1] (the displaced products are the company's apps, ordinary software, so rule 4 does not apply even though the company also builds AI models)

If “**endorses_displacement**”, list which forms apply, and the list must not be empty, because a genuine displacement claim always fits at least one of the three forms. If you cannot name a form that fits, the message is not a real displacement claim about this company, so label **not_relevant** instead. For “**not_relevant**”, forms = []. Return one result per input message, each carrying the same “index” it was given. Output JSON only.

Listing A4. AI Input Intensity taxonomy classification prompt.

The taxonomy is re-typeset from the deployed prompt. Each category in the deployed prompt is followed by three to six named example products, omitted here for space.

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.

Listing A5. Software and SaaS classification prompt.

System message. You are a careful classifier. Return valid JSON only.

User message. For each product below, decide first whether the product is software, then, only if it is software, whether it is SaaS (Software-as-a-Service), answering Yes or No to each. A product that is not software is never SaaS.

Software. A product is software when what the customer pays for is or includes the use of a computer program, either an application the customer uses or a system program such as an operating system, a database, or a security or developer tool. A program the customer operates as part of the product is enough, even when the rest of the product is a good, a service, or an account. Judge by what the customer operates, not by how the product is labeled.

A product is not software when the customer pays for something else and the program only delivers it, for example

- a physical device or piece of equipment with at most internal firmware,
- computing, storage, hosting, or connectivity capacity,
- work performed by the provider's staff for that customer,
- data, ratings, research, media, or instructional content, whether the provider's own or licensed from others.

SaaS. A software product is SaaS when all of the following hold. The application runs on the provider's infrastructure and customers reach it over the internet through a browser, mobile app, or thin client rather than running it on their own hardware. One centrally maintained instance serves many customers, with updates released to all at once. Revenue is recurring (subscription, free tiers with paid upgrades, pay-per-use, or ad-supported) rather than perpetual licensing.

Return JSON of the form `{"classifications": [{"idx": int, "product": "echo of input product name", "is_software": "Yes" or "No", "is_saas": "Yes" or "No", "reasoning": "one-sentence justification"}, ...]}` with one entry per input, `is_saas` "No" whenever `is_software` is "No". Echo the exact `product` string back so each result can be matched to its input.

Products: *[the JSON array of product names and descriptions]*




Figure A1. 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 A2. Value-Weighted Buy-and-Hold Cumulative Returns by AI Input Intensity, January 2022 to September 2026

Note. The figure plots value-weighted cumulative buy-and-hold returns for three portfolios of US-listed firms over January 3, 2022 through September 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. The dashed black line is the value-weighted return of all sample firms under the same weighting, the market benchmark. Dotted grey vertical lines mark the four event-study events in the paper. Dashed navy lines mark two additional Gen-AI- and market-relevant milestones, NVIDIA's blowout Q1 FY24 earnings on May 24, 2023, and the "Liberation Day" reciprocal-tariff announcement on April 2, 2025.




Figure A3. Moderate and Deep Overlap across the Fourteen Releases

Note. Each release is represented by functional bundles, and every product line in the release’s filing-year vintage is scored 0–10 on overlap with each bundle by the LLM, the product’s score being the maximum across bundles in a single scoring pass, Cowork Plugins’ pass being its production scoring. Panel A counts products with at least moderate overlap, a score of 5 or higher, stacked by severity band (5 to 7 and 7 to 10). Panel B counts firms with at least one such product, stacked by the band of the firm’s highest-scoring product. On the scoring scale, 5 means the bundle covers the product’s core function but the non-covered functions are themselves core to the product, 7 means the bundle covers the core function and the non-covered functions are peripheral, and 10 means the product is a bundle’s core function with no meaningful differentiator. Cowork Plugins’ axis label is set in bold.




Figure A4. AI-PDR Coefficient on $CAR_{[-1, +10]}$ across 25 LLM Model Releases

Note. The figure plots the AI-PDR coefficient with 95% confidence intervals at each of 25 LLM model releases from late 2022 through July 2026, the release scan of Section V.B excluding the product releases covered by Section V.A. The specification matches Table II, a cross-sectional OLS regression of CAPM-adjusted $CAR_{[-1, +10]}$ on standardized AI-PDR (Digital Deliverability) with $\log(\text{ME})$, 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, and industries with fewer than five firms at the FE level are dropped. The regression uses GICS 6-digit fixed effects, and standard errors are GICS 6-digit industry-clustered.




Figure A5. Placebo-Date Distribution of the AI-PDR Coefficient

Note. Each observation is the AI-PDR coefficient from the baseline specification of Table II, the cross-sectional regression of the CAPM-adjusted CAR over $[-1, +10]$ on AI-PDR and the firm controls with GICS 6-digit fixed effects and GICS 6-digit clustering, estimated at each of 169 placebo event dates, every fifth trading day from January 16, 2023 through August 31, 2026, excluding dates within 11 trading days of any of the four events so that no placebo window overlaps an event window. Pre-event betas are estimated per placebo date over $[-252, -11]$ trading days. At placebo dates AI-PDR is the firm-level score averaged across filing years, standardized within each sample, and the four event lines use the event-year scores of Table II.




Figure A6. AI-PDR Coefficient from 40 Trading Days before to 40 after Cowork Plugins

Note. The panels repeat the specification of Figure 3 over window ends d from -40 to $+40$ trading days. For $d \geq 0$ each point is the coefficient on standardized AI-PDR from the regression of the CAPM-adjusted $[-1, +d]$ CAR around Cowork Plugins on AI-PDR and the baseline controls. For negative d the plotted value is minus the coefficient from the same regression on the CAR cumulated over $[d, -1]$, so the curve reads as the path of high-AI-PDR firms relative to the day before the event. The top panel uses GICS 2-digit and the bottom panel GICS 6-digit fixed effects, the shaded bands are 95% confidence intervals from GICS 6-digit clustered standard errors, and the dotted vertical line marks the event.

Table A1
Variable Definitions

Variable	Definition
<i>Panel A: Dependent Variables</i>	
CAR (CAPM-adjusted)	Cumulative abnormal return over the event window using the firm-specific pre-event CAPM beta. The daily abnormal return is $AR_{i,t} = (r_{i,t} - r_{f,t}) - \hat{\beta}_i (r_{m,t} - r_{f,t})$, with the estimated intercept not subtracted, and the CAR is the sum of $AR_{i,t}$ over the event window. The pre-event beta $\hat{\beta}_i$ is 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. Baseline window is $[-1, +10]$ trading days. Winsorized at 1st/99th percentiles.
<i>Panel B: Key Independent Variables</i>	
AI-PDR	Gen-AI Product Displacement Risk. Firm-level importance-weighted average of product-level Digital Deliverability scores, the degree to which the core value of the product could be delivered as a purely digital output, without physical goods, in-person service, or transfer of financial capital. It 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). 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	Importance-weighted share of the firm’s products classified into any category of the Gen-AI-input taxonomy (Listing A4), on a 0–1 scale. A product is in scope when scaling Gen-AI training and inference requires more units of it. 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 role-specific Cowork Plugins launched January 30, 2026. Each extracted product is scored (0–10) by an LLM on competitive overlap with each plugin bundle separately (see Listing A2), the product-level overlap is the maximum across the 10 bundles, and product scores are aggregated to firm level using importance weights, computed from the firm’s most recent pre-event 10-K filing (FY 2025). Winsorized at 1st/99th percentiles.

Continued on next page

Table A1 (continued)

Variable	Definition
ESZT Gen-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-Gen-AI substitution and complementarity.
<i>Panel C: Firm-Level Controls</i>	
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.
<i>Panel D: Heterogeneity Interactions</i>	
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) .
<i>Panel E: Analyst Forecast Variables</i>	
EPS Forecast	Consensus mean earnings per share for the first fiscal year ending on or after December 2026, matched on the fiscal-period end date. Source is the IBES summary consensus, taken pre-event (the firm’s last consensus before January 30, 2026) and post-event (its last consensus after the event through June 30, 2026).

Continued on next page

Table A1 (continued)

Variable	Definition
Revenue Forecast	Consensus mean revenue forecast in millions of dollars for the first fiscal year ending on or after December 2026, matched on the fiscal-period end date. Source is the IBES summary consensus, pre-event and post-event as defined for the EPS forecast.
LTG	Long-term EPS growth rate, 3–5 year horizon (%), not tied to a specific fiscal year. Source is the IBES summary consensus, pre-event and post-event as defined for the EPS forecast.
Target Price	12-month analyst consensus price target (\$), not tied to a specific fiscal year. Source is the IBES summary consensus, pre-event and post-event as defined for the EPS forecast.
<i>Panel F: Firm Fundamentals Variables (Table IX)</i>	
Δ Revenue growth	Change in the firm’s year-over-year quarterly revenue growth rate, the firm’s fiscal quarter ending inside April 1 to June 30, 2026 against the same fiscal quarter one year earlier (<code>revtq</code>). Reported in decimals in Table IX and in p.p. in Figure 7.
Δ Gross margin	Change in the firm’s quarterly gross margin, revenue minus cost of goods sold over revenue (<code>revtq</code> , <code>cogsq</code>), the firm’s fiscal quarter ending inside April 1 to June 30, 2026 against the same fiscal quarter one year earlier. Reported in decimals in Table IX and in p.p. in Figure 7.

Table A2
Examples of High- and Low-AI-PDR Firms

Firm	Top Product	AI-PDR	Sector
<i>High AI-PDR</i>			
Autodesk, Inc	Software Subscriptions and Enterprise Access Agreements	8.80	Info. Tech.
Adobe Inc	Firefly Generative AI Solutions	8.80	Info. Tech.
Roblox Corp	Robux Virtual Currency and Roblox Premium Subscriptions	8.80	Comm. Svc.
Atlassian Corp	IT and Enterprise Service Management Software	8.80	Info. Tech.
Electronic Arts Inc	Live Game Services and Downloadable Content	8.80	Comm. Svc.
Take Two Interactive Software Inc	Interactive Entertainment Software Titles	8.80	Comm. Svc.
Zoom Communications, Inc	Zoom AI Companion	8.75	Info. Tech.
Datadog, Inc	Observability and Performance Monitoring	8.71	Info. Tech.
Reddit, Inc	Digital Advertising Platform	8.67	Comm. Svc.
Salesforce, Inc	Agentforce Commerce	8.63	Info. Tech.
<i>Low AI-PDR</i>			
Vertex Pharmaceuticals Inc	Cystic Fibrosis Medicines	0.00	Health Care
Union Pacific Corp	Premium Intermodal and Automotive Rail Transportation	0.00	Industrials
Bristol Myers Squibb Co	Immuno-Oncology Checkpoint Inhibitors	0.00	Health Care
Amgen Inc	Nplate	0.00	Health Care
AbbVie Inc	Immunology Therapies	0.00	Health Care
Chevron Corp	Refined Petroleum Fuels and Other Petroleum Products	0.00	Energy
Philip Morris International Inc	Cigarettes	0.00	Cons. Staples
Amphenol Corp	Value-Add Interconnect Assemblies and Power Distribution Systems	0.00	Info. Tech.
Southern Copper Corp	Refined Copper and Copper Rod	0.00	Materials
PepsiCo Inc	Convenient Foods	0.00	Cons. Staples

Note. Firms are ranked by AI-PDR, the top and bottom 50 form the high-PDR and low-PDR pools, and 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 A3
Correlation of AI-PDR with Firm Characteristics

Variable	N	Pearson r	Spearman ρ	$\hat{\beta}$
ESZT Labor Exposure	2,587	0.377	0.434	0.0076*** (4.55)
log(ME)	4,218	0.054	0.030	0.0268 (1.02)
R&D/Assets	2,270	0.050	0.128	0.0027 (0.90)
Intangibles/Assets	4,169	0.315	0.247	0.0276*** (6.95)
PPE/Assets	4,058	-0.339	-0.388	-0.0194*** (-4.72)
Gross Margin	4,067	0.244	0.251	0.0277*** (3.54)
Net Margin	4,067	0.000	0.019	0.0015 (0.31)
Revenue Growth	4,106	0.059	0.087	0.0075** (2.37)

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 each characteristic on AI-PDR with GICS 2-digit sector fixed effects, with GICS 6-digit industry-clustered t -statistics in parentheses. ESZT Labor Exposure is the [Eisfeldt et al. \(2025\)](#) generative AI labor exposure measure. All Compustat variables are from the most recent fiscal year. Gross Margin is sales minus cost of goods sold over sales, and Net Margin is net income over sales. Revenue Growth is the one-year growth rate of sales. Gross Margin, Net Margin, and Revenue Growth are bounded between -1 and 1 . All characteristics are winsorized at the 1st and 99th percentiles. *, **, *** denote significance at 10%, 5%, 1%.

Table A4
Timeline of Anthropic Product Announcements Around the Event Window

Date	Announcement	PS?	Justification
Feb 2025	Claude Code preview	No	CLI coding tool for developers, no workflow automation.
May 22, 2025	Claude Code GA	No	Broader availability of the same coding tool.
Sep 29, 2025	Claude Sonnet 4.5	No	Model upgrade.
Nov 24, 2025	Claude Opus 4.5	No	Frontier model upgrade, no packaged workflow tools.
Jan 12, 2026	Cowork preview (Max)	No	General-purpose desktop agent, no plugins. Max plans only.
Jan 16, 2026	Cowork Pro (\$20/mo)	No	Widens access to the same agent, still no plugins.
Jan 30, 2026	Cowork Plugins	Yes	11 open-source plugins, 10 of them role-specific (sales, legal, finance, marketing, customer support, data analysis, and others), bundling domain knowledge, connectors, and sub-agents to automate workflows sold as standalone SaaS subscriptions. First entry into identifiable incumbent product markets, and the first commercially successful agentic AI release for general knowledge work.
Feb 5, 2026	Claude Opus 4.6	No	Model upgrade.
Feb 24, 2026	Cowork enterprise update	No	New connectors, private plugin marketplaces, 12 new plugins, extending the January 30 program.

Note. PS? = Product Substitution, direct entry into incumbent product markets with packaged, domain-specific workflow automation. Model releases improve the underlying models but do not distribute tools competing with specific SaaS products. Sources: TechCrunch, Engadget, Axios (Jan 2026), CNN, CNBC, Fast Company (Feb 2026), Anthropic GitHub ([anthropics/knowledge-work-plugins](https://github.com/anthropics/knowledge-work-plugins)).

Table A5
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	0.0040 (1.07) $N = 2,850$	0.0057* (1.92) $N = 2,850$	0.0062** (2.04) $N = 2,850$	0.0068* (1.82) $N = 2,832$	0.0072* (1.88) $N = 2,754$
DeepSeek R1	0.0141*** (4.86) $N = 2,907$	0.0164*** (4.94) $N = 2,907$	0.0098** (2.44) $N = 2,907$	0.0075** (2.01) $N = 2,887$	0.0096** (2.37) $N = 2,817$
Claude Opus 4.5	-0.0013 (-0.53) $N = 2,781$	-0.0116*** (-3.76) $N = 2,781$	-0.0074*** (-2.60) $N = 2,781$	-0.0054 (-1.58) $N = 2,763$	-0.0042 (-1.10) $N = 2,677$
Cowork Plugins	-0.0557*** (-11.28) $N = 2,865$	-0.0705*** (-13.68) $N = 2,865$	-0.0556*** (-9.33) $N = 2,865$	-0.0466*** (-8.34) $N = 2,845$	-0.0421*** (-7.20) $N = 2,771$

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): $\log(\text{ME})$, 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 A6
CAR Regressions across Event Windows

Window	ChatGPT Launch		DeepSeek R1		Claude Opus 4.5		Cowork Plugins	
[-1, +10]	0.0057*	0.0068*	0.0164***	0.0075**	-0.0116***	-0.0054	-0.0705***	-0.0466***
	(1.92)	(1.82)	(4.94)	(2.01)	(-3.76)	(-1.58)	(-13.68)	(-8.34)
[-1, +20]	0.0079**	0.0063	0.0183***	0.0116*	-0.0112***	-0.0081	-0.0670***	-0.0435***
	(2.19)	(1.25)	(3.73)	(1.88)	(-3.56)	(-1.37)	(-10.16)	(-5.74)
[-1, +30]	0.0023	0.0110*	0.0092*	0.0004	-0.0150**	-0.0099	-0.0530***	-0.0352***
	(0.45)	(1.89)	(1.67)	(0.06)	(-2.54)	(-1.63)	(-7.58)	(-3.74)
[-1, +40]	0.0095	0.0171**	0.0017	-0.0035	-0.0376***	-0.0222***	-0.0602***	-0.0414***
	(1.42)	(2.55)	(0.33)	(-0.54)	(-4.00)	(-2.68)	(-8.10)	(-4.67)
[-1, +80]	0.0068	0.0083	0.0069	-0.0032	-0.0949***	-0.0571***	-0.0872***	-0.0478***
	(0.74)	(0.82)	(0.89)	(-0.41)	(-6.13)	(-3.82)	(-4.19)	(-2.62)
[-1, +120]	0.0082	0.0058	-0.0176*	-0.0158	-0.1175***	-0.0777***	-0.0767***	-0.0526***
	(0.67)	(0.41)	(-1.93)	(-1.52)	(-4.55)	(-3.86)	(-5.26)	(-3.15)
[-1, +160]	0.0082	0.0008	-0.0333***	-0.0246*	-0.1027***	-0.0691***	-0.0369**	-0.0270
	(0.58)	(0.05)	(-2.96)	(-1.65)	(-4.74)	(-3.28)	(-2.48)	(-1.43)
FE	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6

Note. Each cell reports the coefficient on AI-PDR (Digital Deliverability, standardized to unit variance within the sample) from a cross-sectional OLS regression of the CAPM-adjusted cumulative abnormal return over the indicated $[-1, +d]$ window on AI-PDR. For each event the first column uses GICS 2-digit (sector) and the second GICS 6-digit (industry) fixed effects. Controls (not reported): $\log(\text{ME})$, Tobin's Q, ROA, asset growth, each winsorized at 1/99. CARs are also winsorized at 1/99. GICS 6-digit industries with fewer than five firms are dropped from the GICS 6-digit columns. 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 A7
High-PDR Firms Outside Technology and Communications

Firm	Top Product	Software	AI-PDR	Sector
Booking Holdings Inc.	Online Accommodation Reservation Services	Yes	6.92	Cons. Disc.
eBay Inc	Online Marketplace Services	Yes	6.47	Cons. Disc.
Duolingo, Inc.	Duolingo for Schools	Yes	8.80	Cons. Disc.
S&P Global Inc.	Financial Market Data, Analytics & Research	Yes	7.76	Financials
Moody's Corp	Banking Risk Workflow Solutions	Yes	7.68	Financials
Nasdaq, Inc.	Financial Crime Management Technology	Yes	6.17	Financials
IQVIA Holdings Inc.	Life Sciences Technology Platforms	Yes	5.75	Health Care
Veeva Systems Inc	Veeva Development Cloud	Yes	7.90	Health Care
Waystar Holding Corp.	Financial Clearance Solutions	Yes	7.86	Health Care
Automatic Data Processing Inc	Cloud HCM Platforms	Yes	5.74	Industrials
Verisk Analytics, Inc.	P&C Underwriting Data and Analytics Solutions	Yes	7.92	Industrials
Broadridge Financial Solutions, Inc.	Capital Markets Trade Lifecycle Platform	Yes	6.78	Industrials
CoStar Group, Inc.	Commercial Real Estate Information and Analytics Platform	Yes	8.22	Real Estate
RE/MAX Holdings, Inc.	Mortgage Loan Processing Services and Loan Broker System	Yes	5.95	Real Estate

Note. After excluding GICS sectors 45 (Information Technology) and 50 (Communication Services), the 100 highest-AI-PDR firms form the pool, and we display up to the 3 largest firms by market equity in each sector. “Top Product” is one of the firm’s product lines from its most recent 10-K filing, its software-classified line of highest importance when the firm has one. “Software” is a firm-level flag from an LLM classification of every product line in that filing year, Yes if any line is classified software, No only if none is.

Table A8
Robustness to Excluding Technology Firms (Cowork Plugins)

Sample	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
Full	−0.0557*** (−11.28) $N = 2,865$	−0.0705*** (−13.68) $N = 2,865$	−0.0556*** (−9.33) $N = 2,865$	−0.0466*** (−8.34) $N = 2,845$	−0.0421*** (−7.20) $N = 2,771$
Ex-IT (45)	−0.0498*** (−7.87) $N = 2,516$	−0.0611*** (−9.78) $N = 2,516$	−0.0515*** (−9.49) $N = 2,516$	−0.0426*** (−8.09) $N = 2,496$	−0.0362*** (−6.14) $N = 2,422$
IT Only	−0.0856*** (−20.30) $N = 349$	−0.0856*** (−20.30) $N = 349$	−0.0398*** (−3.04) $N = 349$	−0.0421*** (−2.92) $N = 349$	−0.0506*** (−3.82) $N = 349$
Ex-IT/Comm	−0.0473*** (−8.26) $N = 2,408$	−0.0515*** (−9.05) $N = 2,408$	−0.0431*** (−8.68) $N = 2,408$	−0.0367*** (−7.35) $N = 2,388$	−0.0309*** (−5.58) $N = 2,318$
IT/Comm Only	−0.0840*** (−10.81) $N = 457$	−0.0859*** (−14.05) $N = 457$	−0.0546*** (−4.30) $N = 457$	−0.0456*** (−4.26) $N = 457$	−0.0508*** (−4.97) $N = 453$

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 Classification

Firm	Top Product
<i>Panel A. SaaS firms (random)</i>	
American Express Co	Travel, Dining and Lifestyle Membership Services
Aware Inc	AwareSDK Biometric Authentication and Identity Verification Platform
OptimizeRx Corp	Financial Messaging
Ww International, Inc.	Digital Behavioral Weight Management Subscriptions
Knightscope, Inc.	Security Operations and Emergency Device Management Software
<i>Panel B. Non-SaaS firms with the highest AI-PDR</i>	
SolarWinds Corp	IT Infrastructure and Network Management Software
SciPlay Corp	Social Casino Games
CuriosityStream Inc.	Subscription Video-on-Demand Streaming
Playtika Holding Corp.	Casual Mobile Games
Couchbase, Inc.	Couchbase Capella

Note. A random sample of five firms classified as SaaS (Panel A), plus the five non-SaaS firms with the highest firm-level AI-PDR (Panel B). Every product line in a firm’s most recent 10-K filing year is classified with an LLM, and a firm is SaaS if any of its product lines is classified SaaS and non-SaaS only if none is. A product is SaaS when the application software the customer uses runs on the provider’s infrastructure, hosted centrally and accessed over the internet, serves many customers from the same instance, and is sold on recurring terms, and not SaaS when the primary value the customer pays for is the underlying data, ratings, or measurements rather than the software through which they are delivered. The displayed product is one of the firm’s product lines in its most recent 10-K filing, its SaaS-classified line of highest importance when the firm has one.

Table A10
CARs on AI-PDR by Software and SaaS Classification (Cowork Plugins)

	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
<i>Panel A: Software versus non-software</i>					
Software	−0.0707*** (−7.84) $N = 928$	−0.0878*** (−11.55) $N = 928$	−0.0617*** (−6.68) $N = 928$	−0.0496*** (−4.85) $N = 886$	−0.0538*** (−4.55) $N = 804$
Non-software	−0.0272*** (−4.66) $N = 1,925$	−0.0343*** (−5.42) $N = 1,925$	−0.0261*** (−4.77) $N = 1,925$	−0.0209*** (−4.54) $N = 1,887$	−0.0169*** (−3.88) $N = 1,787$
<i>Panel B: SaaS versus non-SaaS</i>					
SaaS	−0.0686*** (−6.98) $N = 526$	−0.0867*** (−10.49) $N = 526$	−0.0644*** (−6.15) $N = 526$	−0.0433*** (−3.55) $N = 481$	−0.0515*** (−3.80) $N = 414$
Non-SaaS	−0.0331*** (−6.17) $N = 2,327$	−0.0423*** (−7.47) $N = 2,327$	−0.0328*** (−6.27) $N = 2,327$	−0.0276*** (−6.44) $N = 2,306$	−0.0233*** (−5.20) $N = 2,217$

Note. Re-estimates the Cowork Plugins regression on subsamples partitioned by firm-level classification, across five industry fixed-effect granularities. Panel A splits firms into software and non-software, Panel B into SaaS and non-SaaS. Every product line in a firm’s most recent 10-K filing year is classified with an LLM, and a firm is software (SaaS) if any of its product lines is classified software (SaaS) and non-software (non-SaaS) only if none is. A product is software when what the customer pays for is or includes the use of a computer program, an application or a system program, a program the customer operates as part of the product being enough, and not software when the customer pays for something else and the program only delivers it (a device with at most internal firmware, computing capacity, staff work, or data and content). SaaS is software hosted on the provider’s infrastructure and reached over the internet, one centrally updated instance serving many customers, with recurring revenue rather than perpetual licensing. CARs are CAPM-adjusted with firm-specific betas on $[-252, -11]$, regressed on standardized AI-PDR (Digital Deliverability), with $\log(\text{ME})$, 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

Horse Race of AI-PDR, Digital Keyword Share, and Software Beta (Cowork Plugins event only)

	Software Beta		Beta + AI-PDR		Keyword Share		Keyword + AI-PDR	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI-PDR			−0.0608*** (−10.38)	−0.0415*** (−7.50)			−0.0610*** (−8.29)	−0.0412*** (−7.08)
Software Beta (IGV)	−0.0468*** (−8.16)	−0.0288*** (−6.39)	−0.0287*** (−5.97)	−0.0242*** (−5.41)				
Digital Keyword Share					−0.0572*** (−8.01)	−0.0280*** (−3.40)	−0.0161** (−1.96)	−0.0099 (−1.29)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FE	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6
N	2,865	2,845	2,865	2,845	2,865	2,845	2,865	2,845
R^2	0.193	0.354	0.296	0.380	0.205	0.345	0.274	0.365

Note. CAPM-adjusted CARs over $[-1, +10]$ around Cowork Plugins on a constant sample of firms with all three measures. *Digital Keyword Share* is the number of dictionary matches in Item 1 divided by its word count, on the same filing vintage as the event AI-PDR, with the dictionary: digital, software, online, internet, cloud, platform, data, analytics, app, application, website, web, mobile, e-commerce, subscription, streaming, virtual, cyber, algorithm, artificial intelligence, machine learning, saas, api, content, media, database, electronic, download, marketplace, portal (lowercase word-boundary matching, simple plurals, and hyphenated, spaced, and joined variants of multiword terms). *Software Beta (IGV)* is the loading on the IGV excess return from a regression of daily excess returns on the market and IGV excess returns over $[-252, -11]$ trading days before the event (minimum 100 observations). Both benchmarks are winsorized at 1/99, and all three measures are standardized to unit variance and demeaned, so coefficients are one-SD effects. All columns include the baseline controls, CARs are winsorized at 1/99, the universe is firms with market capitalization and Compustat total assets above \$100M, and GICS 6-digit cells with fewer than 5 firms are dropped in the GICS 6-digit FE columns. GICS 6-digit industry-clustered t -statistics in parentheses in all columns. *, **, *** denote significance at 10%, 5%, 1%.

Table A12

Robustness to 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	-0.0557*** (-11.28) $N = 2,865$	-0.0705*** (-13.68) $N = 2,865$	-0.0556*** (-9.33) $N = 2,865$	-0.0466*** (-8.34) $N = 2,845$	-0.0421*** (-7.20) $N = 2,771$
Excluding Earners	-0.0532*** (-11.94) $N = 1,979$	-0.0646*** (-11.42) $N = 1,979$	-0.0522*** (-8.63) $N = 1,979$	-0.0435*** (-6.96) $N = 1,951$	-0.0359*** (-4.90) $N = 1,853$

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 A5. “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 ($\log(\text{ME})$, 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 A13
Size and Financial-Constraint Moderators of the AI-PDR Effect (Cowork Plugins)

Moderator	AI-PDR	AI-PDR $\times$ Moderator	N	R^2
log(ME)	-0.0466*** (-8.40)	-0.0030 (-1.43)	2,845	0.3637
R&D/Assets	-0.0575*** (-5.71)	0.0003 (0.03)	1,481	0.3659
SA	-0.0457*** (-8.01)	0.0040 (1.34)	2,844	0.3702
WW	-0.0460*** (-8.03)	0.0001 (0.04)	2,667	0.3805
KZ	-0.0529*** (-7.36)	-0.0026 (-1.47)	2,142	0.3604

Note. Each row adds a demeaned firm characteristic and its interaction with AI-PDR to the baseline specification for the Cowork Plugins event $([-1, +10])$, so every regression carries the moderator's main effect alongside the interaction (for firm size the standardized demeaned log(ME) control is the main effect). Every moderator is standardized to unit variance and demeaned within its estimation sample, so the interaction coefficients share a per-one-SD scale. AI-PDR is standardized to unit variance and demeaned within the estimation sample, and 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 with ratio terms scaled by lagged net PP&E, with financial firms (SIC 6000-6999) and firms with lagged net PP&E under \$1M dropped. Its winsorized in-sample 10th and 90th percentiles are -13.1 and 16.9. 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 A14
Robustness to Controlling for Interest-Rate Sensitivity (Cowork Plugins)

	No FE	GICS 2-digit	GICS 4-digit	GICS 6-digit	GICS 8-digit
Baseline	−0.0557*** (−11.28) $N = 2,865$	−0.0705*** (−13.68) $N = 2,865$	−0.0556*** (−9.33) $N = 2,865$	−0.0466*** (−8.34) $N = 2,845$	−0.0421*** (−7.20) $N = 2,771$
+ Rate Control	−0.0557*** (−11.31) $N = 2,865$	−0.0705*** (−13.63) $N = 2,865$	−0.0554*** (−9.36) $N = 2,865$	−0.0466*** (−8.28) $N = 2,845$	−0.0422*** (−7.12) $N = 2,771$

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 A5. “+ 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
Cross-Sectional Regressions of CARs on AI-PDR and ESZT Labor Exposure

Event	AI-PDR	ESZT	N	R^2
ChatGPT Launch	-0.0068 (-0.99)	0.0191*** (3.32)	2,218	0.3515
DeepSeek R1	0.0017 (0.26)	-0.0015 (-0.26)	2,049	0.4959
Claude Opus 4.5	-0.0007 (-0.11)	-0.0105* (-1.85)	1,942	0.4113
Cowork Plugins	-0.0451*** (-2.72)	-0.0125* (-1.75)	1,924	0.6179

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 last trading day before the event window opens. The estimation universe is the full Compustat-CRSP cross section with non-missing ESZT. Industry fixed effects are at the NAICS 3-digit level. Firm controls are $\log(\text{ME})$, 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 the change in CAR (in decimal units) per one-standard-deviation increase. HC1-robust t -statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A16
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	0.0016 (0.53) $N = 2,850$	0.0023 (0.76) $N = 2,850$	0.0026 (0.83) $N = 2,850$	0.0037 (0.96) $N = 2,832$	0.0045 (1.14) $N = 2,754$
DeepSeek R1	0.0079*** (2.99) $N = 2,907$	0.0113*** (3.67) $N = 2,907$	0.0083* (1.80) $N = 2,907$	0.0065 (1.27) $N = 2,887$	0.0098* (1.72) $N = 2,817$
Claude Opus 4.5	-0.0042* (-1.69) $N = 2,781$	-0.0147*** (-4.77) $N = 2,781$	-0.0119*** (-3.65) $N = 2,781$	-0.0086** (-2.39) $N = 2,763$	-0.0065* (-1.65) $N = 2,677$
Cowork Plugins	-0.0449*** (-7.94) $N = 2,865$	-0.0648*** (-10.04) $N = 2,865$	-0.0482*** (-7.57) $N = 2,865$	-0.0443*** (-5.92) $N = 2,845$	-0.0406*** (-4.72) $N = 2,771$

Note. Same specification as Table A5 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 ($\log(\text{ME})$, 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 A17
Robustness to Additional Controls (Cowork Plugins)

Specification	AI-PDR	N	R^2
Baseline	−0.0466*** (−8.34)	2,845	0.3633
+ Firm characteristics	−0.0423*** (−7.33)	2,659	0.3782
+ Leverage, cash	−0.0415*** (−7.30)	2,650	0.3804
+ Momentum, volatility	−0.0381*** (−6.39)	2,650	0.4039
+ Firm age	−0.0381*** (−6.38)	2,650	0.4049
+ R&D intensity	−0.0477*** (−4.19)	1,347	0.4118

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 ($\log(\text{ME})$, Tobin’s Q , ROA, asset growth). “Firm characteristics” adds $\log(\text{Assets})$, 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 $\log(\text{Age})$ from first Compustat appearance. “R&D intensity” adds $\text{R\&D}/\text{Assets}$; 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 A18
Event Study with GPT-3.5 AI-PDR (Chronologically Consistent Scorer)

Event	AI-PDR	N	R^2
ChatGPT Launch	0.0029 (0.79)	2,832	0.0902
DeepSeek R1	0.0074** (2.08)	2,887	0.1451
Claude Opus 4.5	0.0021 (0.53)	2,763	0.1393
Cowork Plugins	-0.0356*** (-6.46)	2,845	0.3551

Note. GPT-3.5 robustness mirror of Table II. AI-PDR is reconstructed using GPT-3.5-Turbo-0125 (training cutoff September 2021, predating the ChatGPT launch) in place of the baseline LLM, 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 ($\log(\text{ME})$, 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 baseline-LLM AI-PDR is $\rho = 0.91$. *, **, *** denote significance at 10%, 5%, 1%.

Table A19
Summary Statistics for the Analyst Forecast Data

	<i>N</i>	Mean	SD	Min	P5	P25	Median	P75	P95	Max
<i>Panel A: Pre-event consensus level (last IBES consensus before Jan 30, 2026)</i>										
EPS (\$)	1,960	5.72	6.80	0.55	0.74	1.70	3.49	6.93	19.04	43.81
Revenue (millions)	2,533	7935.30	19970.50	22.41	75.89	469.85	1580.21	5616.78	35572.23	140664.91
LTG (%)	1,028	14.22	18.71	-25.50	-6.49	6.10	11.10	16.81	42.86	112.01
Target Price (\$)	2,655	90.51	128.65	2.50	5.69	18.18	41.93	103.48	347.40	768.16
<i>Panel B: Post-event consensus level (last IBES consensus after the event through Jun 30, 2026)</i>										
EPS (\$)	1,960	5.84	7.12	-0.19	0.58	1.67	3.52	7.05	19.47	44.26
Revenue (millions)	2,533	8362.82	22032.49	15.32	73.38	460.32	1619.17	5728.94	37334.58	162239.00
LTG (%)	1,028	15.61	17.69	-24.38	-3.55	7.49	12.00	18.66	48.02	102.40
Target Price (\$)	2,655	96.41	141.84	2.43	5.33	18.48	43.55	107.56	374.99	865.02
<i>Panel C: Scaled Revisions Used as Dependent Variables in Table VII</i>										
EPS revision	1,960	-0.00	0.26	-1.12	-0.37	-0.06	0.01	0.06	0.29	1.20
Revenue revision	2,533	0.01	0.13	-0.64	-0.16	-0.02	0.01	0.03	0.18	0.56
LTG change	1,028	1.27	13.07	-50.00	-13.06	-1.44	0.00	3.30	23.11	50.00
Target Price revision	2,655	0.05	0.26	-0.49	-0.31	-0.07	0.02	0.14	0.53	1.14

Note. Summary statistics for the IBES summary-consensus forecasts used in Table VII. Panels A and B report the pre-event and post-event forecast levels for each outcome, where the pre-event level is the firm's last IBES consensus before January 30, 2026, and the post-event level is its last IBES consensus after the event through June 30, 2026. Panel C reports the scaled revisions used as the dependent variables in the Cowork Plugins analyst-revision 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 $y = \text{post} - \text{pre}$, capped at ± 50 percentage points. EPS and revenue are for the first fiscal year ending on or after December 2026, matched on the fiscal-period end date. The sample is firms with market capitalization above \$100M, Compustat total assets above \$100M, both a pre-event and a post-event forecast, and a non-missing AI-PDR score. EPS observations with pre-event EPS below \$0.50 and revenue observations with pre-event revenue below \$10M and target-price observations with a non-positive pre-event target are dropped. Revenue is in millions of dollars, LTG in percent, and target price in dollars. All variables are winsorized at the 1st and 99th percentiles.

Table A20
Analyst EPS and Revenue Revisions by Forecast Horizon (Cowork Plugins)

	(1) EPS revision	(2) Revenue revision
December 2026	−0.029*** (-2.66)	−0.009* (-1.85)
N	1,960	2,533
R^2	0.099	0.098
December 2027	−0.030*** (-3.04)	−0.022*** (-5.22)
N	1,943	2,416
R^2	0.134	0.105
December 2028	−0.039*** (-2.92)	−0.020*** (-3.54)
N	1,258	1,637
R^2	0.097	0.106
Controls	Yes	Yes
FE	GICS-2	GICS-2

Note. Each cell reports a cross-sectional OLS regression of a scaled analyst-forecast revision around the Cowork Plugins event (January 30, 2026) on AI-PDR. Each row matches EPS and revenue forecasts on the fiscal-period end date to the firm’s first fiscal year ending on or after the indicated December, so the pre-event and post-event consensus refer to the same fiscal period and differ only in when the forecast was made. The pre-event and post-event timing, the revision definition, the EPS and revenue denominator filters, the standardization of AI-PDR, the firm controls, the winsorization, and the sample are as in Table VII. Both columns use GICS 2-digit (sector) fixed effects. GICS 6-digit industry-clustered t -statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A21
Target-Price Revisions and Competitive Shields (Cowork Plugins)

	(1) Composite Shield	(2) Supply Side	(3) Demand Side
AI-PDR	−0.101*** (−5.20)	−0.074*** (−4.92)	−0.106*** (−6.03)
Shield	0.029*** (2.64)	0.023* (1.82)	0.024** (2.43)
AI-PDR × Shield	0.017* (1.69)	0.046*** (4.69)	−0.008 (−0.87)
Controls	Yes	Yes	Yes
FE	GICS-2	GICS-2	GICS-2
N	2,655	2,655	2,655
R^2	0.156	0.163	0.156

Note. Each column is a cross-sectional OLS regression of the scaled target-price revision around the Cowork Plugins event on AI-PDR, a Shield score, and their interaction, with the firm controls and GICS 2-digit sector fixed effects. The sample, the consensus timing, the revision definition, and the winsorization follow Table VII. Column (1) uses the composite Shield (average of Legal/Regulatory, Production Moats, Customer Trust, and Customer Embeddedness), column (2) the supply-side shield (average of Legal/Regulatory and Production Moats), and column (3) the demand-side shield (average of Customer Trust and Customer Embeddedness). The rows labeled Shield and AI-PDR × Shield refer to whichever Shield definition the column uses. AI-PDR and the Shield modifier are standardized to unit variance and demeaned within the estimation sample, and the interaction is the product of the standardized-demeaned regressors. GICS 6-digit industry-clustered t -statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A22
Target-Price Revisions and Realized Returns (Cowork Plugins)

	(1) CAR control	(2) Raw-return control
AI-PDR	-0.076*** (-5.98)	-0.074*** (-6.03)
CAR [-1, +110]	0.405*** (15.86)	
Raw return [-1, +110]		0.391*** (18.22)
FE	GICS-2	GICS-2
N	2,648	2,648
R^2	0.382	0.409

Note. Each column is a cross-sectional OLS regression of the scaled target-price revision around the Cowork Plugins event (January 30, 2026) on AI-PDR and a control for the firm's own realized return over a window spanning the revision period. The revision follows Table VII, the last IBES summary consensus before the event against the last after it through June 30, 2026, (post-pre)/pre, winsorized at 1/99. The return window [-1, +110] in trading days spans the revision period. CAR is the sum of daily CAPM-adjusted abnormal returns with each firm's market beta estimated over [-252, -11], and the raw return is the compounded simple total return, both winsorized at 1/99 within each estimation sample. AI-PDR is standardized to unit variance and demeaned within each estimation sample. Firm controls (log(ME), Tobin's Q, ROA, asset growth, winsorized at 1/99) are included in every column, with GICS 2-digit sector fixed effects. The sample is firms with market capitalization above \$100M, Compustat total assets above \$100M, a non-missing AI-PDR score, GICS classification, firm controls, both consensususes, and non-missing return controls. GICS 6-digit industry-clustered t -statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A23
StockTwits Messages, Summary Statistics and Examples

	ChatGPT	DeepSeek R1	Claude Opus 4.5	Cowork Plugins
<i>Panel A. Message coverage and belief content</i>				
Messages	1,810,978	1,878,601	1,980,163	1,823,543
pass keyword prefilter	39,575 (2.2%)	81,952 (4.4%)	97,673 (4.9%)	94,148 (5.2%)
Firms	4,528	4,498	4,335	4,347
Messages per firm-day, mean	16.4	23.2	23.3	21.5
Message length (chars), mean	108.1	99.8	109.2	111.8
Endorse displacement	14 (0.001%)	70 (0.004%)	155 (0.008%)	656 (0.036%)
Form 1, provider substitutes directly	14	61	139	539
Form 2, AI-built rivals	0	10	16	79
Form 3, customers self-build	1	7	27	142

Panel B. Example messages, Cowork Plugins window

	Ticker	Message
Form 1	\$ADBE	\$HUBS \$NOW \$ADBE \$CRM \$INTU software with a weak start to 2026, continuing the 2025 theme that AI Agents will steal massive market share or lead to major drops in seat licensing / consumption
Form 2	\$U	would love to load this, but say what you want but ai can create code, video, software for substantially less expense and faster speed these guys can t compete.
Form 3	\$ACN	This is the reason. As a former programmer and slide presenter at ACN why wouldn t a company spend the money on creating its own AI model that it can use over and over instead of paying high priced Bob s.
Non-endorsing messages	\$NVDA	picked up 100 shares for ER next week

Note. Panel A covers all StockTwits messages over the four calendar weeks each side of each event. A message counts as endorsing displacement when it passes the keyword prefilter and two rounds of LLM classification (Listing A3). Percentages are of the column's message count. Forms 1 to 3 are the displacement mechanisms a message endorses (a message may endorse more than one). Panel B shows one example message per category from the Cowork Plugins window. Messages are reproduced as stored, with apostrophes removed by the text pipeline.

Table A24
Text-Based Investor Endorsement of Generative AI Displacement, Manually Reviewed

	Level (p.p.)		Log count		Calibrated log	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A — ChatGPT Launch (pre-event mean 0.000 p.p., $N = 2,434$)						
AI-PDR	0.019 (1.03)	0.023 (0.99)	0.004 (1.47)	0.003 (1.06)	0.006 (1.45)	0.005 (1.06)
Panel B — DeepSeek R1 (pre-event mean 0.007 p.p., $N = 1,384$)						
AI-PDR	-0.002 (-0.81)	-0.001 (-0.48)	0.005 (1.46)	0.004 (1.02)	0.007 (1.52)	0.006 (1.09)
Panel C — Claude Opus 4.5 (pre-event mean 0.033 p.p., $N = 1,484$)						
AI-PDR	-0.025 (-1.12)	-0.027 (-0.98)	-0.012 (-0.73)	-0.009 (-0.76)	-0.017 (-0.74)	-0.013 (-0.77)
Panel D — Cowork Plugins (pre-event mean 0.025 p.p., $N = 1,547$)						
AI-PDR	0.118*** (4.70)	0.057** (2.53)	0.104*** (4.91)	0.068*** (3.05)	0.143*** (5.06)	0.093*** (3.16)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
FE	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6

Note. Follows Table VIII in every respect, with the manual review replacing the LLM verification as the second classification round. t -statistics in parentheses are clustered by GICS 6-digit industry. *, **, *** denote significance at 10%, 5%, 1%.

Table A25
Stacked Event-Study Regressions of CARs on Product Overlap by Release Group

	(1)–(2)		(3)–(4)		(5)–(6)	
	Business Oriented		Business Oriented		Consumer	
	(Narrow)		(Broad)		Oriented	
$[-1, +1]$	−0.0026*	−0.0031***	−0.0012	−0.0018***	−0.0002	0.0000
	(-1.80)	(-2.67)	(-1.59)	(-3.29)	(-0.57)	(0.12)
$[-1, +3]$	−0.0025	−0.0031***	−0.0016*	−0.0023***	−0.0000	−0.0002
	(-1.58)	(-2.65)	(-1.76)	(-3.31)	(-0.11)	(-0.46)
$[-1, +5]$	−0.0016	−0.0029**	−0.0009	−0.0026***	−0.0006	−0.0011
	(-1.06)	(-2.40)	(-0.85)	(-3.04)	(-0.72)	(-1.35)
$[-1, +10]$	−0.0014	−0.0035*	0.0006	−0.0014	−0.0011	−0.0015
	(-0.61)	(-1.81)	(0.32)	(-0.99)	(-0.98)	(-1.30)
$[-1, +20]$	−0.0043	−0.0053	−0.0020	−0.0032	−0.0027**	−0.0024*
	(-1.35)	(-1.58)	(-0.89)	(-1.63)	(-2.10)	(-1.70)
Controls $\times$ event	Yes	Yes	Yes	Yes	Yes	Yes
FE $\times$ event	GICS-2	GICS-6	GICS-2	GICS-6	GICS-2	GICS-6

Note. Each cell reports the coefficient on standardized overlap from a stacked regression of event-window CARs on that variable across the thirteen major AI product releases of Section V.A, Cowork Plugins excluded throughout. CARs are CAPM-adjusted with firm-specific market betas estimated on $[-252, -11]$ trading days. All regressions include AI-PDR and the baseline controls (log(ME), Tobin’s Q, ROA, and asset growth) interacted with event dummies, and absorb industry $\times$ event fixed effects, GICS 2-digit in the first column of each pair and GICS 6-digit in the second. Standardized overlap is the firm’s importance-weighted mean product overlap, standardized within event, a product’s overlap being scored by the LLM in a single pass. Business Oriented (Narrow) stacks M365 Copilot GA, Claude for Financial Services, and Claude Managed Agents, Business Oriented (Broad) adds Claude 3.5 computer use and Claude Code, and Consumer Oriented stacks the remaining eight releases. GICS 6-digit industry-clustered t -statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A26

AI-PDR and AI Input Intensity across Event Windows

<i>Panel A — ChatGPT Launch, GICS-2 FEs</i>					<i>Panel B — ChatGPT Launch, GICS-6 FEs</i>				
	$[-1, +3]$	$[-1, +10]$	$[-1, +20]$	$[-1, +30]$		$[-1, +3]$	$[-1, +10]$	$[-1, +20]$	$[-1, +30]$
AI-PDR	0.0017 (1.15)	0.0067** (2.14)	0.0080** (2.24)	0.0025 (0.46)	AI-PDR	0.0024 (0.93)	0.0086** (2.16)	0.0083 (1.55)	0.0124** (2.06)
AI Input	0.0106 (1.08)	0.0181 (1.07)	-0.0089 (-0.52)	0.0115 (0.50)	AI Input	0.0068 (0.56)	0.0200 (1.00)	-0.0027 (-0.13)	0.0016 (0.08)
<i>Panel C — DeepSeek R1, GICS-2 FEs</i>					<i>Panel D — DeepSeek R1, GICS-6 FEs</i>				
	$[-1, +3]$	$[-1, +10]$	$[-1, +20]$	$[-1, +30]$		$[-1, +3]$	$[-1, +10]$	$[-1, +20]$	$[-1, +30]$
AI-PDR	0.0032 (1.51)	0.0148*** (4.78)	0.0177*** (3.40)	0.0071 (1.31)	AI-PDR	0.0018 (0.75)	0.0079** (2.17)	0.0128* (1.92)	0.0010 (0.13)
AI Input	-0.0019 (-0.11)	-0.0677*** (-5.25)	-0.0293** (-2.01)	-0.0601*** (-2.79)	AI Input	-0.0055 (-0.31)	-0.0532*** (-3.03)	-0.0145 (-0.71)	-0.0181 (-0.55)
<i>Panel E — Claude Opus 4.5, GICS-2 FEs</i>					<i>Panel F — Claude Opus 4.5, GICS-6 FEs</i>				
	$[-1, +3]$	$[-1, +10]$	$[-1, +20]$	$[-1, +30]$		$[-1, +3]$	$[-1, +10]$	$[-1, +20]$	$[-1, +30]$
AI-PDR	-0.0058** (-2.39)	-0.0104*** (-3.73)	-0.0113*** (-3.40)	-0.0142*** (-2.77)	AI-PDR	-0.0034 (-1.11)	-0.0048 (-1.36)	-0.0087 (-1.41)	-0.0102 (-1.60)
AI Input	0.0110 (1.25)	0.0321** (2.21)	-0.0008 (-0.04)	0.0459 (1.40)	AI Input	0.0050 (0.53)	0.0215 (1.35)	0.0028 (0.13)	0.0393 (1.14)
<i>Panel G — Cowork Plugins, GICS-2 FEs</i>					<i>Panel H — Cowork Plugins, GICS-6 FEs</i>				
	$[-1, +3]$	$[-1, +10]$	$[-1, +20]$	$[-1, +30]$		$[-1, +3]$	$[-1, +10]$	$[-1, +20]$	$[-1, +30]$
AI-PDR	-0.0370*** (-16.79)	-0.0686*** (-12.02)	-0.0646*** (-9.13)	-0.0486*** (-6.51)	AI-PDR	-0.0226*** (-6.23)	-0.0453*** (-8.61)	-0.0423*** (-5.59)	-0.0322*** (-3.31)
AI Input	-0.0388** (-2.22)	0.0525*** (3.59)	0.0601*** (2.93)	0.1135*** (4.51)	AI Input	-0.0400** (-2.46)	0.0604*** (3.86)	0.0900** (2.57)	0.1405*** (4.47)

Note. Each of the eight panels reports, for a single event and fixed-effects level, two coefficients (AI-PDR and AI Input) from a cross-sectional OLS 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 LLM taxonomy classifier (Listing [A4](#)). Both regressors use firm values from the fiscal year aligned with the event date. *GICS-2 FEs* denotes GICS 2-digit sector fixed effects (left panels), *GICS-6 FEs* denotes GICS 6-digit industry fixed effects with industries having fewer than five firms dropped (right panels). 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%.

Table A27
AI-PDR across Event Windows under the Beta-Only, Alpha-Inclusive, and Characteristics-Matched Benchmarks
(Cowork Plugins)

Window	AI-PDR (pooled)			SaaS $\times$ AI-PDR			Non-SaaS $\times$ AI-PDR		
	No int.	Char.-matched	<i>With alpha</i>	No int.	Char.-matched	<i>With alpha</i>	No int.	Char.-matched	<i>With alpha</i>
[−1, +10]	−0.0465*** (−8.34)	−0.0441*** (−7.55)	−0.0424*** (−7.12)	−0.0460*** (−5.86)	−0.0400*** (−4.90)	−0.0387*** (−4.73)	−0.0413*** (−7.00)	−0.0401*** (−6.52)	−0.0396*** (−6.41)
[−1, +20]	−0.0434*** (−5.74)	−0.0430*** (−5.84)	−0.0375*** (−5.37)	−0.0470*** (−4.97)	−0.0429*** (−4.80)	−0.0354*** (−4.28)	−0.0420*** (−5.15)	−0.0424*** (−5.38)	−0.0401*** (−5.10)
[−1, +30]	−0.0351*** (−3.74)	−0.0347*** (−3.98)	−0.0267*** (−3.08)	−0.0389*** (−3.63)	−0.0335*** (−3.45)	−0.0221** (−2.35)	−0.0365*** (−3.22)	−0.0375*** (−3.55)	−0.0341*** (−3.21)
[−1, +40]	−0.0414*** (−4.67)	−0.0421*** (−5.00)	−0.0293*** (−3.24)	−0.0485*** (−4.78)	−0.0423*** (−4.42)	−0.0250** (−2.50)	−0.0438*** (−3.89)	−0.0467*** (−4.52)	−0.0407*** (−3.71)
[−1, +80]	−0.0477*** (−2.62)	−0.0414** (−2.09)	−0.0280* (−1.96)	−0.0489* (−1.88)	−0.0401 (−1.42)	−0.0098 (−0.47)	−0.0457*** (−2.70)	−0.0409** (−2.20)	−0.0404*** (−2.62)
[−1, +120]	−0.0526*** (−3.15)	−0.0469*** (−2.61)	−0.0204 (−1.23)	−0.0330 (−1.32)	−0.0235 (−0.86)	0.0287 (1.28)	−0.0643*** (−3.74)	−0.0608*** (−3.28)	−0.0541*** (−2.99)
[−1, +160]	−0.0270 (−1.43)	−0.0218 (−1.06)	0.0113 (0.46)	−0.0009 (−0.03)	0.0047 (0.16)	0.0761** (2.16)	−0.0474** (−2.31)	−0.0431* (−1.94)	−0.0387* (−1.65)
FE	GICS-6	GICS-6	GICS-6	GICS-6	GICS-6	GICS-6	GICS-6	GICS-6	GICS-6
N	2,845	2,845	2,845	2,833	2,833	2,833	2,833	2,833	2,833

Note. Each row estimates, for the indicated window around Cowork Plugins (January 30, 2026), the pooled AI-PDR slope of the Table II specification and the any-line SaaS and non-SaaS slopes of the Table IV specification, each under three abnormal-return constructions. No intercept is the paper’s definition, subtracting only the firm’s pre-event market beta times the market excess return. Characteristics-matched replaces the CAPM benchmark with a matched-portfolio benchmark in the tradition of Daniel et al. (1997). Firms are sorted into 25 value-weighted portfolios, five size quintiles crossed with five momentum quintiles, formed at trading day −11 from the most recent market capitalization and the buy-and-hold return over trading days [−252, −21], requiring at least 90 daily returns over the formation window. Abnormal returns are raw returns minus the matched portfolio’s return, cumulated with the same gap treatment as the baseline, and available for every sample firm. *With alpha* additionally subtracts the firm’s market-model intercept, estimated over [−252, −11], on every spanned trading day, extrapolating the firm’s estimation-period drift over the window (Ramelli and Wagner, 2020), and is reported only to document that contamination. All columns use GICS 6-digit industry fixed effects and the firm controls, with AI-PDR scaled as in each source table, and the *N* row reports the pooled and classified samples. GICS 6-digit industry-clustered *t*-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A28
Announcement Returns around First Agentic AI Launches

Window	First launch before Cowork Plugins		First launch after Cowork Plugins	
	(1) CAPM	(2) Char.-matched	(3) CAPM	(4) Char.-matched
$[-1, +1]$	-0.0008 (-0.17)	0.0050 (1.08)	0.0136** (2.01)	0.0140** (2.05)
$[-1, +10]$	-0.0110 (-0.93)	0.0118 (1.01)	0.0331** (2.07)	0.0333** (2.17)
$[-1, +30]$	-0.0266* (-1.84)	0.0260* (1.75)	0.0224 (1.04)	0.0436** (2.05)
N	124	124	102	100

Note. The sample is each firm's first agentic AI product launch from the LSEG launch list described in the text, split into launches dated before January 30, 2026 and launches dated on or after it. The event date is the firm's own launch date, and each cell reports the mean cumulative abnormal return over the indicated window across launches, with a launch contributing through its last available trading day. In the CAPM columns the daily abnormal return is the firm's excess return minus the launch-specific market beta times the market excess return, the beta estimated over $[-252, -11]$ trading days relative to the launch with at least 100 observations. In the characteristics-matched columns the daily abnormal return is the firm's return minus the value-weighted return of its size and momentum portfolio, one of 25 portfolios formed from quintiles of market capitalization and prior-year return eleven trading days before Cowork Plugins. Window returns are winsorized at 1/99 within each column, and the N row reports each column's launch count. t -statistics clustered by launch date in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A29
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	0.0045 (1.07) $N = 3,220$	0.0077** (2.24) $N = 3,220$	0.0073* (1.66) $N = 3,220$	0.0089 (1.54) $N = 3,201$	0.0095 (1.46) $N = 3,129$
DeepSeek R1	0.0151*** (5.07) $N = 2,963$	0.0188*** (5.73) $N = 2,963$	0.0136*** (3.12) $N = 2,963$	0.0118** (2.34) $N = 2,942$	0.0149*** (2.71) $N = 2,874$
Claude Opus 4.5	−0.0024 (−0.96) $N = 2,856$	−0.0135*** (−4.36) $N = 2,856$	−0.0096*** (−2.99) $N = 2,856$	−0.0076** (−2.00) $N = 2,837$	−0.0069* (−1.65) $N = 2,761$
Cowork Plugins	−0.0570*** (−11.45) $N = 2,901$	−0.0735*** (−15.92) $N = 2,901$	−0.0607*** (−10.59) $N = 2,901$	−0.0530*** (−9.58) $N = 2,881$	−0.0488*** (−9.01) $N = 2,807$

Note. Same specification as Table A5 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. The same averaged AI-PDR is used at all four events. Filings before fiscal 2021 are excluded because the earlier product mix is less indicative of the firm’s current revenue base. 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, with GICS 6-digit and 8-digit specifications dropping 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 A30
AI-PDR across Event Windows by SaaS Classification, Majority Rule (Cowork Plugins)

Window	CAR				Char.-matched			
	× SaaS AI-PDR	Non-SaaS × AI-PDR	× SaaS AI-PDR	Non-SaaS × AI-PDR	× SaaS AI-PDR	Non-SaaS × AI-PDR	× SaaS AI-PDR	Non-SaaS × AI-PDR
[−1, +10]	−0.0615*** (−5.03)	−0.0635*** (−11.10)	−0.0443*** (−3.94)	−0.0431*** (−7.49)	−0.0550*** (−4.29)	−0.0593*** (−10.73)	−0.0394*** (−3.27)	−0.0414*** (−6.78)
[−1, +20]	−0.0658*** (−4.84)	−0.0615*** (−9.35)	−0.0377*** (−2.73)	−0.0423*** (−5.49)	−0.0619*** (−4.49)	−0.0586*** (−9.26)	−0.0363*** (−2.64)	−0.0424*** (−5.60)
[−1, +30]	−0.0593*** (−3.97)	−0.0488*** (−6.61)	−0.0336** (−1.98)	−0.0347*** (−3.55)	−0.0531*** (−3.53)	−0.0450*** (−6.27)	−0.0307* (−1.90)	−0.0351*** (−3.80)
[−1, +40]	−0.0726*** (−4.44)	−0.0535*** (−6.76)	−0.0430** (−2.52)	−0.0401*** (−4.27)	−0.0642*** (−3.97)	−0.0486*** (−5.98)	−0.0406** (−2.39)	−0.0418*** (−4.55)
[−1, +80]	−0.1159*** (−3.92)	−0.0665*** (−3.93)	−0.0429* (−1.70)	−0.0406** (−2.45)	−0.1093*** (−3.36)	−0.0613*** (−3.33)	−0.0354 (−1.26)	−0.0352* (−1.93)
[−1, +120]	−0.0687** (−2.48)	−0.0710*** (−5.06)	−0.0203 (−0.55)	−0.0549*** (−3.46)	−0.0646** (−2.13)	−0.0653*** (−4.32)	−0.0148 (−0.38)	−0.0496*** (−2.88)
[−1, +160]	−0.0123 (−0.46)	−0.0382** (−2.42)	0.0290 (0.86)	−0.0348* (−1.85)	−0.0153 (−0.50)	−0.0351** (−2.00)	0.0309 (0.82)	−0.0293 (−1.42)
FE	GICS-2	GICS-2	GICS-6	GICS-6	GICS-2	GICS-2	GICS-6	GICS-6
<i>N</i>	174	2,679	174	2,659	174	2,679	174	2,659

Note. The table repeats Table IV under the majority rule, under which a firm is SaaS when at least half of its importance-weighted product lines are classified as SaaS and non-SaaS otherwise. Columns, construction, controls, fixed effects, clustering, and the estimation universe follow Table IV, and the *N* row reports each column's firm count at the longest window. GICS 6-digit industry-clustered *t*-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.