AI Product Displacement Risk

March 26, 2026

Abstract

Generative AI can raise firm value by replacing workers but destroy it by replacing products. We construct AI Product Displacement Risk (AI-PDR), a firm-level measure of vulnerability to product substitution by generative AI. AI-PDR captures the demand-side threat to revenue—the opposite channel from the supply-side labor cost savings measured by existing AI-exposure indices. Event studies around four AI releases show that the displacement signal strengthens as AI advances from demonstration to targeted product substitution: early capability releases produce no within-industry return spread, a later release reveals a significant spread once correlated beta exposure is absorbed, and the 2026 launch of AI agents designed to replace specific incumbent SaaS products generates the largest repricing—top-quintile firms underperform bottom-quintile firms by 4.8 percentage points. Digital deliverability drives the entire effect.

1 Introduction

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

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

We introduce *AI Product Displacement Risk* (AI-PDR) to measure this product-level threat. Using a large language model, we extract individual product and service lines from the business description (Item 1) of each firm's 10-K filing, score each product along three dimensions—digital deliverability, trust and regulatory barriers, and proprietary asset barriers—and aggregate to a firm-level measure weighted by revenue importance. AI-PDR is conceptually distinct from the Eisfeldt et al. (forthcoming) (hereafter ESZT) labor-exposure measure: whereas their index captures which firms' *workers* are replaceable by AI, AI-PDR captures which firms' *products* are replaceable. The two correlate at only 0.27. AI-PDR ranks Communication Services and Information Technology highest, and Health Care, Energy, and Utilities lowest.

We study equity market reactions to four major AI releases: ChatGPT (November 2022), DeepSeek R1 (January 2025), Claude Opus 4.5 (November 2025), and Cowork Plugins (January 2026). The first three released general-purpose models that demonstrated new capabilities without targeting specific product markets. The fourth—Cowork Plugins, launched by Anthropic on January 30, 2026—was different: it offered 11 open-source AI agent plugins covering productivity, sales, customer support, legal, finance, data analysis, and other enterprise functions, all accessible through a $20/month subscription that undercut incumbent

per-seat SaaS fees by an order of magnitude. Within days, approximately $285 billion in market capitalization was erased from software and professional services stocks globally (CNN Business, 2026), with firms whose core products overlapped with these agents experiencing the sharpest declines.

Our central finding is that equity markets price product displacement risk with increasing intensity as AI advances from capability demonstration to targeted product substitution. The three capability releases produce no significant within-industry return spread between high- and low-PDR firms in the baseline specification. However, controlling for market beta—which tends to be higher for high-PDR growth firms—reveals an emerging displacement pricing effect for Claude Opus 4.5 that grows over a 20-day window. Cowork Plugins produces the largest and most robust effect, significant without conditioning on market beta: top-quintile firms underperform bottom-quintile firms by 4.8 percentage points over the five-day event window within industries. The divergence across events is consistent with a simple mechanism: capability demonstrations leave uncertain the path from raw technology to actual product competition, while a targeted substitution event collapses that uncertainty by entering identifiable incumbent revenue streams directly.

Within-event variation sharpens the interpretation. Splitting the sample by a Cowork Overlap score—measuring how closely each firm's products compete with the specific 10 plugins launched—reveals that the AI-PDR effect is concentrated among high-overlap firms. The effect holds excluding all Information Technology and Communication Services firms and does not vary with firm size, R&D intensity, or financial constraints, consistent with broad-based repricing. A decomposition of AI-PDR into its three sub-components shows that digital deliverability drives the effect: entering it alone yields nearly identical explanatory power to the composite measure; trust and proprietary barriers add nothing.

The result is robust along several dimensions. Controlling for pre-event interest-rate sensitivity rules out duration exposure. Excluding firms that reported earnings within the event window preserves significance, ruling out correlated earnings news. Three alternative LLM scorers—Qwen2.5-72B, Llama-3.1-70B, and GPT-3.5-Turbo—each reproduce the main finding. The GPT-3.5 test is particularly informative: its September 2021 training cutoff predates all events in our study, addressing the concern that the scorer's knowledge of realized AI capabilities contaminates the ratings. The effect persists across event windows of varying length, consistent with permanent repricing rather than a transient liquidity effect. Results survive alternative aggregation rules, year-specific scores, and cumulative addition of controls.

We contribute to three literatures. The first studies AI's displacement effects, which has focused almost entirely on labor. Acemoglu and Restrepo (2019) provide the theoretical

framework distinguishing task displacement from task creation. Acemoglu and Restrepo (2020) and Acemoglu et al. (2022) document employment effects of robots and AI. Felten et al. (2021), Felten et al. (2023), Webb (2020), and Eloundou et al. (2024) construct occupation-level exposure indices, while Hampole et al. (2025) study firm-level employment effects. At the macro level, Acemoglu (2025) and Aghion et al. (2019) model equilibrium growth effects of automation, and Bick et al. (2026) document the rapid pace of generative AI adoption. We add the product dimension to this body of work, showing that the demand-side threat to revenue can dominate the supply-side benefit to costs in equity valuations.

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

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

The remainder of the paper proceeds as follows. Section 2 describes the data and measure construction. Section 3 presents the event-study results. Section 4 reports robustness tests. Section 5 concludes.

2 Data and Measure Construction

2.1 AI-PDR Measure

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

Step 1: Product Extraction

For each filing, we provide the Item 1 text to the LLM with the following prompt:

Listing 1: Product extraction prompt

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

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

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

```
   substrings.
```

Each extraction returns a structured list of products with names, descriptions, importance ratings, and supporting evidence quotes. Requiring textual evidence disciplines the extraction by grounding each product in specific filing language. The extraction step yields approximately 257,000 product entries from roughly 9,600 unique firms across all filing years.

Step 2: Product Scoring

We score each extracted product on three dimensions of displacement vulnerability, each on a 0–10 integer scale where higher values indicate greater vulnerability to AI displacement:

Listing 2: Product scoring prompt (abridged; full prompt in Listing A1)

```
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 how inherently vulnerable
it is to displacement by generative AI.

Dimensions (0-10 integers):

A. DIGITAL DELIVERABILITY
   Could a GenAI system produce and deliver this in purely
   digital form, with no physical goods, no human labor, and
   no financial capital required?
   0  = physical product (steel, crude oil, drugs)
   4  = requires financial capital or human workers (bank
        loans, staffing, exchange trading)
   8  = digital output, some human oversight (analytics,
        curated data feeds)
   10 = entirely digital, LLM can produce equivalent output
        (template copy, stock images, basic code)

B. TRUST / REGULATORY BARRIERS
   Would a customer accept this from an AI with no human or
   institutional brand behind it?
```

```
   0  = legally cannot (FDA-approved drugs, licensed devices)
   4  = meaningful trust required (wealth management,
        enterprise software with SLAs)
   10 = no trust barrier (commodity content, templates)

C. PROPRIETARY ASSET BARRIERS
   Does producing this require assets an AI entrant cannot
   replicate?
   0  = massive moat (physical infrastructure, dominant
        network effects, decades of proprietary data)
   4  = moderate moat (specialized datasets, brand reputation)
   10 = no moat (basic content creation, simple translation)

IMPORTANT:
1. Write reasoning BEFORE each score.
2. For digital deliverability, ask: "What is the binding
   input the customer is paying for?" If it is capital,
   physical goods, or regulated infrastructure, score LOW.
3. Do NOT default to 5. Use the FULL 0-10 range.
```

The full prompt (Listing A1) includes calibration anchors at every two points on each scale with concrete product examples. The JSON schema orders the reasoning field before the score field, requiring a chain-of-thought justification before each numerical rating. Products from different firms are mixed within each batch (15 products per batch) to prevent same-firm anchoring.

Aggregation

The product-level AI-PDR score is the simple average of the three dimension scores:

$$\text{AI-PDR}_p = \frac{\text{Digital}_p + \text{Trust}_p + \text{Proprietary}_p}{3}.$$

The firm-level measure is an importance-weighted average across products, where weights reflect the LLM's assessment of each product's revenue importance to the firm (high = 3, medium = 2, low = 1):

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

where ω_p is the importance weight of product p. The simple average treats the three dimensions as substitutes; as a robustness check, we also aggregate using a geometric mean, which penalizes products that score low on any single dimension (Table A13). The two specifications correlate at $\rho = 0.98$ and produce virtually identical results. We winsorize the firm-level measure at the 1st and 99th percentiles. For the baseline specification, we average AI-PDR$_{i,t}$ across all available filing years for each firm (FY 2018–2026), producing a single firm-level score. Because this average includes filings from years after some events, Section 4 verifies robustness using only the most recent filing year before each event (Table A14): FY 2021 for ChatGPT, FY 2024 for DeepSeek R1 and Claude Opus 4.5, and FY 2025 for Cowork Plugins.

Sub-component properties

Table A2 reports summary statistics and pairwise correlations for the three scoring dimensions. Digital deliverability has the highest dispersion (SD = 0.230), reflecting a sharp divide between informational and physical products. Trust and proprietary barriers have lower dispersion (0.192 and 0.158). Pairwise correlations range from 0.825 (digital–proprietary) to 0.961 (trust–proprietary): regulated industries tend to possess both trust requirements and proprietary assets.

Measure validation

We evaluate AI-PDR along four dimensions.

Cross-industry distribution. Figure 1 plots firm-level AI-PDR by GICS sector. Communication Services (median 0.64) and Information Technology (0.54) rank highest; Health Care (0.10), Energy (0.15), and Utilities (0.17) rank lowest. Sectors whose outputs are primarily digital score highest, while those producing physical or regulated goods score lowest. Within every sector there is substantial variation—even within Health Care, firms range from near zero (drug manufacturers) to above 0.4 (health IT platforms)—confirming that AI-PDR captures within-industry heterogeneity, not merely sector membership.

Face validity. Figure 2 displays word clouds of product descriptions by scoring dimension. High-scoring products on digital deliverability feature "SaaS," "software," "platform," and "cloud"—the categories commentators identified as most vulnerable.[1] Low-scoring products

[1]See, e.g., "SaaS in, SaaS out: Here's what's driving the SaaSpocalypse," *TechCrunch*, March 1, 2026; "The $1 Trillion Software Carnage: How AI Agents Broke the SaaS Model," *MarketMinute*, February 24, 2026.

feature "natural gas," "facilities," "proprietary," and "clinical." Table A3 lists the ten highest- and lowest-scoring firms. High-PDR firms span digital advertising, SaaS platforms, content, and productivity software. Low-PDR firms include electric utilities, deepwater drilling, pharmaceuticals, clinical-stage biotechs, and a hospice services provider—firms whose products require physical infrastructure, regulated substances, or in-person delivery.

Correlation with firm characteristics. Table A4 reports correlations between AI-PDR and observables. The correlation with ESZT labor exposure is 0.27—positive but far from collinear (Figure A1)—confirming that product displacement and labor exposure capture distinct dimensions of AI risk. AI-PDR correlates positively with intangibles/assets (0.29) and negatively with tangibility (-0.23) and R&D intensity (-0.15). The negative R&D correlation reflects the concentration of low-PDR firms in pharmaceuticals and biotechnology, where R&D spending is high but products are physical and regulated. Correlations with firm size (0.09) and profit margins (0.05) are near zero, so AI-PDR does not proxy for scale or profitability.

Temporal stability. Because AI-PDR derives from product descriptions that change slowly, it should be stable over time. Figure A2 confirms this: adjacent-year correlations range from 0.89 to 0.92. Even non-adjacent years (e.g., 2018 vs. 2025) have correlations above 0.85. This stability motivates our baseline approach of averaging across all available filing years and confirms that the measure captures durable product characteristics rather than filing-year idiosyncrasies.

2.2 Other Data and Summary Statistics

Daily stock returns are from the CRSP Stock Files (Version 2) for 2018–2025. For 2026, we supplement CRSP with daily returns from Yahoo Finance, validated against CRSP on the December 2025 overlap month. We merge firms with Compustat via the CRSP-Compustat Merged (CCM) link table, restricting to primary links (LINKTYPE $\in$ {LC, LU}, LINKPRIM $\in$ {P, C}). We obtain firm-level generative AI labor exposure from Eisfeldt et al. (forthcoming). Control variables follow the standard Fama–French timing convention: for each firm, we use Compustat data from fiscal year $t-1$ to construct log(ME), log(BM), operating profitability (OP), and investment (INV), matched to returns from July of year t through June of year $t+1$. Table 1 reports summary statistics and Table A1 provides definitions for all variables used in the analysis.

3 Results

3.1 Events

We study equity market reactions to four major generative AI releases— ChatGPT (November 30, 2022), DeepSeek R1 (January 20, 2025), Claude Opus 4.5 (November 24, 2025), and Cowork Plugins (January 30, 2026). We classify events based on whether the announcement introduced AI agents designed to replace specific categories of existing commercial software. The first three are capability demonstrations that each released a general-purpose model. Cowork Plugins, by contrast, launched purchasable AI agents targeting specific incumbent SaaS products. This classification is determined by the content of each announcement, not by the subsequent market reaction.

Because our identification relies heavily on the Cowork Plugins event, we document its institutional features in detail. Anthropic released Claude Cowork on January 12, 2026 as a desktop-based agentic AI tool initially restricted to Max subscribers ($100–200/month), then expanded access to Pro subscribers ($20/month) on January 16 (Engadget, 2026). On January 30, Anthropic added plugin support and open-sourced 11 role-specific plugins on GitHub, covering productivity, enterprise search, marketing, sales, finance, data analysis, legal, customer support, product management, biology research, and a meta-plugin for building custom plugins (TechCrunch, 2026). Each plugin bundles domain knowledge, tool connectors, slash commands, and sub-agents into a single installable package, enabling the AI to function as a department-specific workflow automation engine rather than a general-purpose chatbot.

The pricing structure created a stark asymmetry with incumbent SaaS products. A Claude Pro subscription providing access to all plugins costs $20/month per user, while a Team seat costs $25–125/month—compared with enterprise contracts that can run tens of thousands of dollars annually for comparable functionality. The legal plugin alone automates contract review, NDA triage, compliance workflows, and legal briefings—core revenue functions of Thomson Reuters (Westlaw/CoCounsel) and RELX (LexisNexis). 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. In each case, the plugin targets workflows previously sold as standalone SaaS subscriptions at per-seat prices substantially above the Claude subscription fee.

The market reaction was swift. By February 3–4, approximately $285 billion in market capitalization had been erased from software, financial services, and asset management stocks

globally (CNN Business, 2026). Thomson Reuters suffered its largest single-day decline on record, falling roughly 16%; LegalZoom dropped approximately 20%; RELX fell around 14%; and Salesforce, ServiceNow, and Adobe each declined roughly 7% (Fast Company, 2026). The iShares Expanded Tech-Software Sector ETF (IGV) fell 5.69% on February 4, its worst session since April 2025. The selloff extended internationally: Indian IT services firms Infosys, TCS, and Wipro declined sharply as investors questioned the viability of headcount-based billing models (CNBC, 2026). The core concern was what analysts termed "seat compression": if an AI agent can perform the work of multiple human users of SaaS products, companies need fewer paid seats, directly threatening the per-seat pricing model that underpins most enterprise software valuations.

Table A5 documents the timeline of Anthropic product announcements surrounding our event window to establish that January 30, 2026 is the appropriate event date. Several nearby announcements could confound event attribution. The Cowork desktop agent previewed on January 12 was a general-purpose tool without domain-specific functionality; the Opus 4.5 release on November 24 was a model capability upgrade; and the Opus 4.6 release on February 5 was another model upgrade that postdates the event window. The January 30 plugin launch is qualitatively different: it is the first date on which Anthropic released packaged, role-specific workflow automation tools—spanning legal, finance, sales, marketing, customer support, data analysis, and other functions—that directly compete with identifiable categories of incumbent SaaS products. Consistent with this dating, commentators specifically attributed the February 3–4 selloff to the plugin launch rather than to the earlier Cowork preview or subsequent model releases (Fast Company, 2026; CNBC, 2026).

3.2 Event Study Design and Graphical Evidence

For each event, we compute market-adjusted cumulative abnormal returns (CARs) over a $[-1, +3]$ trading-day window:

$$\text{CAR}_i = \sum_{t=-1}^{3} \left(r_{i,t} - r_{m,t} \right),$$

where $r_{i,t}$ is firm i's return on day t relative to the event and $r_{m,t}$ is the CRSP value-weighted market return. We then estimate cross-sectional regressions:

$$\text{CAR}_i = \delta_j + \beta_1 \cdot \text{AI-PDR}_i + \beta_2 \cdot \text{ESZT}_i + \gamma' \mathbf{X}_i + \varepsilon_i, \tag{1}$$

where δ_j are GICS sub-industry fixed effects (163 categories, excluding industries with fewer than 5 firms), ESZT_i is the Eisfeldt et al. (forthcoming) standardized labor exposure, and $\mathbf{X}_i$ includes log(ME), log(BM), operating profitability, and investment. All right-hand-side variables are winsorized at the 1st and 99th percentiles. We report heteroskedasticity-robust t-statistics in parentheses and GICS sub-industry clustered t-statistics in brackets.

Figure 3 plots value-weighted mean CARs for the top and bottom AI-PDR quintiles over a $[-5, +30]$ window around each event. The five pre-event days serve as a visual pre-trend check. For the three capability demonstrations, high- and low-PDR quintiles show little divergence. For Cowork Plugins, high-PDR firms begin declining at the event date and the spread reaches approximately 11 percentage points by day +10; the divergence partially narrows but persists through day +30 with no full reversal. Figure A3 confirms these patterns using equal-weighted means. Figure 4 reports quintile-sorted CARs over the $[-1, +3]$ window: for Cowork Plugins, CARs decline from Q1 to Q5; other events show no pattern. Figures A4–A5 repeat the analysis using deciles with qualitatively similar patterns.

3.3 Cross-Sectional Regressions

Panel A of Table 2 reports estimates of equation (1). The AI-PDR coefficient is insignificant for all three capability demonstrations: ChatGPT ($t = 1.41$), DeepSeek R1 ($t = 0.76$), Claude Opus 4.5 ($\hat{\beta}_{\text{PDR}} = -0.020$, $t = -1.02$). Cowork Plugins is qualitatively different: $\hat{\beta}_{\text{PDR}} = -0.093$ ($t = -3.44$ robust, $[-3.51]$ clustered), implying a Q5–Q1 spread of approximately 4.8 percentage points (0.093×0.52, the difference in mean PDR between top and bottom quintiles) over the five-day window within industries. The ESZT labor-exposure coefficient moves in the opposite direction. For ChatGPT, $\hat{\beta}_{\text{ESZT}}$ is positive and significant (0.067, $t = 2.62$), consistent with Eisfeldt et al. (forthcoming)'s finding that labor-exposed firms benefited from expected cost savings.[2] For Cowork Plugins, ESZT is insignificant (-0.041, $t = -1.12$), indicating that labor exposure does not predict the product-substitution reaction.

The joint behavior of the two coefficients across events provides the sharpest evidence that product displacement and labor exposure capture distinct economic channels. At ChatGPT, the labor-exposure coefficient is positive and significant while the product-displacement coefficient is zero—consistent with markets pricing expected cost savings from worker replacement. At Cowork Plugins, the pattern reverses: the product-displacement coefficient is negative and significant while the labor-exposure coefficient is zero—consistent with markets

[2]Interestingly, the ESZT labor-exposure result replicates robustly for DeepSeek R1 ($\hat{\beta}_{\text{ESZT}} = 0.082$, $t = 2.47$ robust, [2.36] clustered), suggesting that markets also interpreted the cost-deflation announcement as good news for firms whose workers are replaceable by AI.

pricing revenue erosion from product substitution. The same controls, the same industry fixed effects, but opposite signs on conceptually distinct measures at events with different economic content. Eisfeldt et al. (forthcoming) explicitly demonstrate that their labor-channel results are not driven by product-market effects and note the importance of distinguishing the two channels. Our findings confirm this distinction and show that the product channel operates with the opposite sign when AI shifts from capability demonstration to targeted product entry.

Panel A of Table 3 reports the AI-PDR coefficient across five event windows. For Cowork Plugins, the coefficient is negative and significant at all horizons, growing from -0.040 ($t = -2.10$) at $[-1, +1]$ to -0.230 ($t = -4.57$ clustered) at $[-1, +10]$ before partially reverting to -0.190 ($t = -3.43$) at $[-1, +20]$ and stabilizing at -0.196 ($t = -2.51$) at $[-1, +30]$. The persistence over six weeks without reversal is consistent with permanent repricing. No other event produces a consistently significant pattern.

The apparent insignificance of AI-PDR at capability demonstrations, however, masks a confound between displacement risk and market beta. High-PDR firms tend to be high-beta growth stocks. When markets rally on AI news, their higher betas generate an offsetting positive return that can obscure a negative displacement signal. Panel B of Table 2 adds each firm's pre-event market beta, estimated from daily returns over $[-252, -11]$. For Cowork Plugins, the AI-PDR coefficient is virtually unchanged (-0.085, $t = -3.11$): the displacement signal dominates unconditionally. For Claude Opus 4.5, however, the coefficient shifts from insignificant ($\hat{\beta} = -0.020$, $t = -1.02$) to significant ($\hat{\beta} = -0.049$, $t = -2.72$ robust, $[-2.53]$ clustered) once beta absorbs the correlated market exposure (R^2 rises from 0.24 to 0.35). This suggests markets were beginning to price product displacement risk at capability demonstrations, but the signal was attenuated by high-PDR firms' correlated beta exposure.

Why does Claude Opus 4.5 show displacement pricing while ChatGPT and DeepSeek do not? By November 2025, three years of rapid progress had narrowed the gap between demonstrated capability and actual product competition. Claude Opus 4.5 was widely recognized as approaching human-level performance on professional tasks—coding, legal analysis, financial reasoning—making the threat to digitally deliverable products more concrete than it was at the time of ChatGPT's debut or DeepSeek's cost-reduction announcement. The displacement signal at Opus 4.5 is consistent with markets beginning to price the product channel once AI capabilities became sufficiently close to commercial substitution.

Panel B of Table 3 extends this analysis across event windows. With beta controls, the Claude Opus 4.5 coefficient grows from -0.049 at $[-1, +3]$ to -0.092 at $[-1, +10]$ ($t = -2.99$, clustered $[-3.00]$) and -0.126 at $[-1, +20]$ ($t = -2.89$, $[-3.16]$), a pattern of persistent repricing similar to Cowork Plugins. ChatGPT and DeepSeek remain insignificant across all

windows.

The picture that emerges is not a sharp discontinuity but a gradual buildup. Markets began pricing product displacement risk as early as Claude Opus 4.5, but the signal was too small relative to correlated beta exposure to detect without conditioning. Cowork Plugins was the first event where the displacement signal dominated unconditionally. A Wald test from a stacked regression with event-specific PDR slopes (Table 2) confirms that the Cowork Plugins coefficient is significantly larger than the average capability-demo coefficient ($F = 57.2$, $p < 0.001$).

The divergence reflects how the two event types resolve different forms of uncertainty. General-purpose capability releases resolve uncertainty about what AI *can* do, but the path from demonstrated capability to actual product competition involves uncertain intermediate steps: packaging, distribution, overcoming switching costs, satisfying enterprise procurement requirements. Each step attenuates the present-value threat to incumbents, who retain the option to adapt. A targeted substitution event collapses this chain of uncertainty: markets observe direct entry into identifiable incumbent revenue streams. The repricing reflects the removal of this uncertainty discount, not necessarily a change in beliefs about AI's raw capability. This interpretation is consistent with Andrews and Farboodi (2025), who find that aggregate market beliefs about AI update in discrete jumps around releases; our evidence shows that the *cross-sectional* displacement signal strengthens when a release targets specific incumbent products.

3.4 Effect Heterogeneity

The main result relies on a single product-substitution event, so we exploit within-event variation to discipline the interpretation. We construct a Cowork Overlap score measuring how closely each firm's products compete with the specific 10 plugins launched on January 30, 2026 (Listing A2). The plugin definitions—productivity, sales, customer support, product management, marketing, legal, finance, data analysis, enterprise search, and bio research—are from the Anthropic launch announcement.[3] Each extracted product is scored (0–10) by GPT-4o mini on competitive overlap, aggregated to firm level using importance weights, and rescaled to $[0, 1]$. The measure is right-skewed: 57% of firms have zero overlap, with mean 0.037 (Table 1). Cowork Overlap correlates with AI-PDR at $\rho = 0.67$—they capture related but distinct variation.

Table 4 splits the sample at the median Cowork Overlap (0.011 on the 0–1 scale) and

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

re-estimates the baseline regression in each subsample. For Cowork Plugins, the AI-PDR coefficient is -0.174 ($t = -4.39$ clustered) in the high-overlap half but 0.005 ($t = 0.15$) in the low-overlap half. The displacement signal is concentrated entirely among firms whose products compete with the launched plugins, validating the product displacement channel. The Cowork Plugins launch represents direct competitive entry—AI agents priced to undercut incumbent per-seat fees—rather than gradual insourcing, where clients build AI capabilities internally. This distinction matters because competitive displacement generates a sharper, more immediate market reaction.

Consistent with the overlap findings—the 10 plugins span legal, finance, marketing, and other functions well beyond the technology sector—the effect extends beyond tech firms. Table A6 lists the highest-PDR firms after excluding both Information Technology and Communication Services (GICS sectors 45 and 50). High-PDR firms in the remaining sectors span consumer discretionary (Etsy, eBay, Chegg), financials (Dave, Aether Holdings), and industrials (Planet Labs)—all firms whose core products are digitally delivered and face low regulatory barriers to AI substitution. GICS classifies firms by primary revenue source, not product delivery method, so firms like Chegg (Consumer Discretionary) or Etsy (Consumer Discretionary) appear outside the technology sector despite selling digitally-delivered products. This reinforces the value of a product-level displacement measure that cuts across sector boundaries. Table A7 re-estimates equation (1) in these restricted samples. The AI-PDR coefficient is -0.074 ($t = -2.61$) excluding IT and -0.091 ($t = -3.08$) excluding both, confirming that the displacement signal originates broadly across sectors rather than from technology firms alone.

Table A8 interacts AI-PDR with five firm characteristics: size, R&D intensity, and three financial constraint indices (Hadlock and Pierce, 2010; Whited and Wu, 2006; Kaplan and Zingales, 1997). All interaction terms are insignificant. The displacement effect is broad-based rather than concentrated among small, R&D-intensive, or financially constrained firms.

3.5 Sub-Component Decomposition

We decompose AI-PDR into its three scoring dimensions. Table 5 enters all three dimensions jointly. Digital deliverability is the only significant predictor ($\hat{\beta}_{\text{Digital}} = -0.130$, $t = -4.70$ robust, $[-4.28]$ clustered), while trust ($t = 0.46$) and proprietary ($t = 1.08$) barriers are insignificant. The two barrier dimensions are highly correlated ($\rho = 0.96$; Table A2), so interpreting their individual coefficients requires caution, but the joint insignificance relative to digital deliverability is informative: it indicates that barriers add little beyond what digital deliverability already captures.

Table A9 sharpens this comparison. Entering digital deliverability alone yields a Cowork Plugins coefficient of $\hat{\beta}_{\text{Digital}} = -0.085$ ($t = -4.28$ robust, $[-4.47]$ clustered) with $R^2 = 0.438$, compared with $\hat{\beta}_{\text{PDR}} = -0.085$ ($t = -3.11$, $[-3.10]$) and $R^2 = 0.435$ for the composite measure. The digital-only specification achieves an equal or higher t-statistic and identical fit. The conclusion is that digital deliverability is the active ingredient in AI-PDR: conditional on whether a product can be produced and delivered digitally, trust and proprietary barriers provide no additional information about displacement vulnerability. This finding aligns with independent assessments by credit rating agencies, which identify IT services, customer experience outsourcing, and enterprise software as facing the highest AI disruption risk, while physical service providers—manned security, facilities maintenance—face minimal threat.[4]

Table A10 extends the decomposition across event windows for Cowork Plugins. Digital deliverability is the dominant predictor at all horizons, growing from $\hat{\beta}_{\text{Digital}} = -0.069$ at $[-1, +1]$ to -0.269 at $[-1, +10]$ ($t = -4.39$ clustered) before partially reverting to -0.210 at $[-1, +30]$. Trust barriers remain insignificant throughout. The proprietary barrier coefficient turns marginally positive at $[-1, +10]$ ($\hat{\beta}_{\text{Prop.}} = 0.188$, $t = 2.06$ clustered), suggesting that firms with stronger proprietary asset moats—proprietary data, established distribution networks, regulated licenses—partially recover after the initial repricing. However, this effect is imprecisely estimated at longer horizons ($t = 0.78$ at $[-1, +30]$), so the evidence for moat-based recovery is suggestive rather than definitive.

4 Robustness

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

4.1 Identification

Pre-event trends. Table A11 regresses market-adjusted CARs on AI-PDR using two pre-event windows ending at day -2. The coefficient is insignificant in both windows under industry clustering ($t = -0.73$, -0.68 for $[-10, -2]$, $[-5, -2]$), compared with $t = -3.10$ in the event window. The absence of pre-event drift supports a causal interpretation.

[4]See "AI Impact On Business And Technology Services: Disruption Without Dislocation ... Yet," S&P Global Ratings, February 26, 2026.

Interest-rate sensitivity. Table A12 adds each firm's interest-rate sensitivity, estimated as the coefficient on Δy_{10} (daily change in the 10-year Treasury yield) from a two-factor regression over $[-252, -11]$ trading days. For Cowork Plugins, rate sensitivity is insignificant ($t = -0.96$) and the AI-PDR coefficient is unchanged (-0.086, $t = -3.14$ clustered), ruling out duration exposure.

Coincident earnings announcements. The Cowork Plugins window coincides with earnings season: 341 firms in our regression sample reported earnings within $[-1, +3]$. Table A7 excludes these firms, reducing N from 1,832 to 1,491. The AI-PDR coefficient remains significant ($\hat{\beta} = -0.070$, $t = -2.25$ clustered), ruling out correlated earnings news.

4.2 Measure Construction

A second set of tests verifies that the results are driven by underlying product characteristics, not the specific LLM, aggregation rule, or filing year used to construct AI-PDR.

Alternative AI-PDR specifications. Replacing the equal-weight average with a geometric mean ($\rho = 0.98$; Table A13) or using event-year-specific AI-PDR from the single filing year before each event (Table A14) yields virtually identical results, ruling out stale product descriptions or look-ahead bias.

Alternative LLM scorer. Table A15 re-estimates the baseline regressions using AI-PDR constructed from Qwen2.5-72B-Instruct, an independent open-source LLM that receives the same prompts but differs in architecture and training data. The resulting firm-level measures correlate at $\rho = 0.94$. For Cowork Plugins, the Qwen-based coefficient is $\hat{\beta}_{\text{PDR}} = -0.072$ ($t = -2.24$ robust, $[-1.94]$ clustered), compared with -0.085 ($t = -3.11$) for the baseline. Table A15 provides a third replication using Llama-3.1-70B-Instruct ($\rho = 0.93$), yielding $\hat{\beta}_{\text{PDR}} = -0.068$ ($t = -2.35$ robust, $[-2.31]$ clustered) for Cowork Plugins.

Chronologically consistent scorer. A separate concern is that the scoring LLM's training data may encode information about which products AI can produce, so that digital deliverability ratings reflect learned associations rather than inherent product characteristics. He et al. (2025a,b) formalize this concern and develop chronologically consistent LLMs with strict knowledge cutoffs. Rather than using their ChronoGPT models—which at 1.5 billion parameters lack the instruction-following capacity needed for our scoring rubric—we 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), the first event in our study.

Table A15 reports the results. The GPT-3.5-based AI-PDR correlates at $\rho = 0.95$ with the baseline, indicating broad agreement on which products are digitally deliverable despite a three-year gap in training data. For Cowork Plugins, $\hat{\beta}_{\text{PDR}} = -0.103$ ($t = -2.57$ robust, $[-2.60]$ clustered), replicating the main finding. The consistency across training cutoffs confirms that AI-PDR captures durable product characteristics—whether output is physical or digital, whether production requires human labor or financial capital—rather than ex post knowledge of AI displacement events.

4.3 Specification and Controls

Dropping the ESZT control. Table A16 re-estimates the baseline regression without the ESZT labor-exposure control, nearly doubling the sample from 1,832 to 3,572 firms (ESZT coverage is limited to firms with occupational composition data). The Cowork Plugins coefficient strengthens to -0.113 ($t = -5.59$ robust, $[-3.93]$ clustered), confirming that the result does not depend on the ESZT subsample.

Fama–French five-factor adjustment. Table A17 replaces market-adjusted CARs with Fama–French five-factor adjusted CARs.[5] For Cowork Plugins, $\hat{\beta}_{\text{PDR}} = -0.057$ ($t = -1.94$ robust, $[-1.77]$ clustered). The coefficient is attenuated relative to the market-adjusted baseline because the five-factor model absorbs some cross-sectional return variation correlated with AI-PDR, but the effect strengthens at longer horizons: Table A18 shows the FF5-adjusted coefficient grows from -0.057 at $[-1, +3]$ to -0.178 at $[-1, +10]$ ($t = -3.79$, clustered $[-3.22]$), confirming persistent repricing.

Additional controls. Table A19 cumulatively adds firm characteristics: 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 -0.076 to -0.085 across all

[5]Published daily Fama–French five-factor returns from Ken French's data library end on January 30, 2026—the Cowork Plugins event date itself. Post-event returns for Cowork Plugins therefore require factors beyond the published series. We construct daily FF5 factors from CRSP and Compustat data following the Fama and French (2015) methodology: 2×3 portfolios sorted on size (NYSE median) crossed with book-to-market, operating profitability, or investment (NYSE 30th/70th percentiles), rebalanced each June. On the 2,031-day overlap with published factors, our constructed market factor correlates at 0.999 and SMB at 0.971. HML (0.917), RMW (0.899), and CMA (0.785) show larger discrepancies, likely reflecting minor differences in universe filters and the exclusion of pre-2018 data from portfolio histories. Factor loadings for all events are estimated over the $[-252, -11]$ pre-event window, which falls entirely within the published-factor period for all four events. The constructed factors enter only through the event-window abnormal returns for Cowork Plugins (days +1 through +3 at the baseline window, and further post-event days for wider windows). ChatGPT, DeepSeek R1, and Claude Opus 4.5 use published factors throughout.

broad-sample specifications. Adding R&D intensity last halves the sample to 736 firms; the point estimate (-0.062) remains economically meaningful, though the smaller sample reduces statistical power ($t = -1.55$). The stability across specifications indicates that AI-PDR does not proxy for standard firm characteristics.

Alternative industry classifications. Our baseline uses GICS sub-industry fixed effects (163 categories). Table A20 re-estimates the Cowork Plugins regression under two alternative classifications: SIC 3-digit (102 industries) and Hoberg and Phillips (2016) FIC-50 text-based industries (50 industries derived from 10-K product description similarity). With SIC 3-digit fixed effects, the coefficient strengthens to $\hat{\beta}_{\text{PDR}} = -0.132$ ($t = -5.56$ robust, $[-3.76]$ clustered); with FIC-50, to -0.134 ($t = -6.34$, $[-4.64]$). Both coarser classifications absorb less within-industry variation than GICS sub-industry, leaving more for AI-PDR to explain. The result confirms that the displacement signal is not an artifact of the particular industry classification used.

5 Conclusion

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

The results raise a strategic question for affected firms. Two paths are available: race—invest in AI-augmented versions of existing products and pivot to AI-native delivery—or retreat—harvest remaining revenue and return cash to shareholders. The optimal response depends on whether a firm holds complementary assets that survive the product transition: brand, distribution networks, customer relationships, and proprietary data. Firms with strong complementary assets may integrate AI into their product delivery and preserve their market position; pure-play information goods producers without such assets face existential risk. This parallels historical episodes of creative destruction where incumbents with durable competitive moats adapted while commodity producers did not (Kogan and Papanikolaou, 2014). Whether high-PDR firms race or retreat in the quarters following AI release events is a natural direction for future research.

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

Several limitations point toward future work. AI-PDR is a static measure based on current 10-K filings and does not capture firms' ability to adapt or pivot; it measures exposure, not realized outcomes. The LLM scorer may embed biases from training data about which products are "digital" or "replaceable," though the multi-model robustness tests—including a scorer with a pre-ChatGPT training cutoff—mitigate this concern. The event study identifies market reactions, not realized displacement, which will take years to materialize in revenue and earnings data; examining real effects—changes in investment, hiring, and R&D among high-PDR firms—would test whether the repricing translates into resource reallocation. Our strongest evidence comes from a single product-substitution event, and the generalizability of the escalation pattern depends on whether future AI product launches follow a similar structure. Finally, AI-PDR captures only the demand side and does not account for second-order effects such as firms that benefit from the displacement of their competitors or the potential for displaced firms to become acquisition targets.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Engadget, "Anthropic Opens Up Its Claude Cowork Feature to Anyone with a $20 Subscription," Engadget January 16 2026.

Fama, Eugene F. and Kenneth R. French, "A Five-Factor Asset Pricing Model," *Journal of Financial Economics*, 2015, *116* (1), 1–22.

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

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

_ **,** _ **, and** _ , "How Will Language Modelers Like ChatGPT Affect Occupations and Industries?,"

Working Paper, 2023.

Gârleanu, Nicolae, Stavros Panageas, and Jianfeng Yu, "Technological Growth and Asset Pricing," *Journal of Finance*, 2012, *67* (4), 1265–1292.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

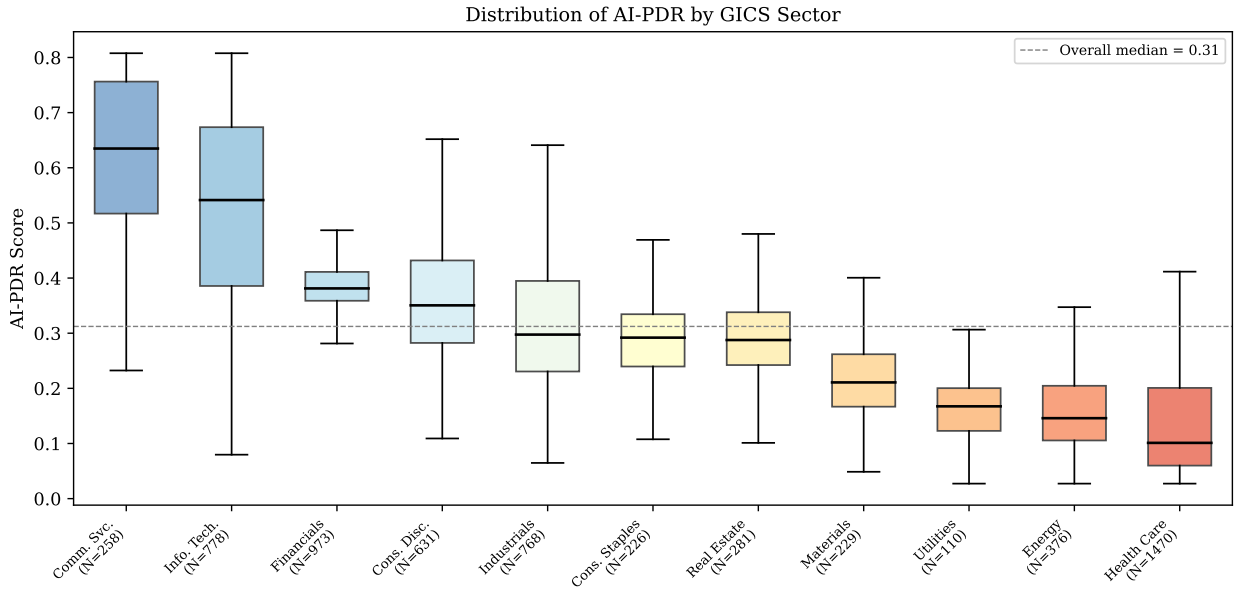


Figure 1: Distribution of AI-PDR by GICS Sector

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

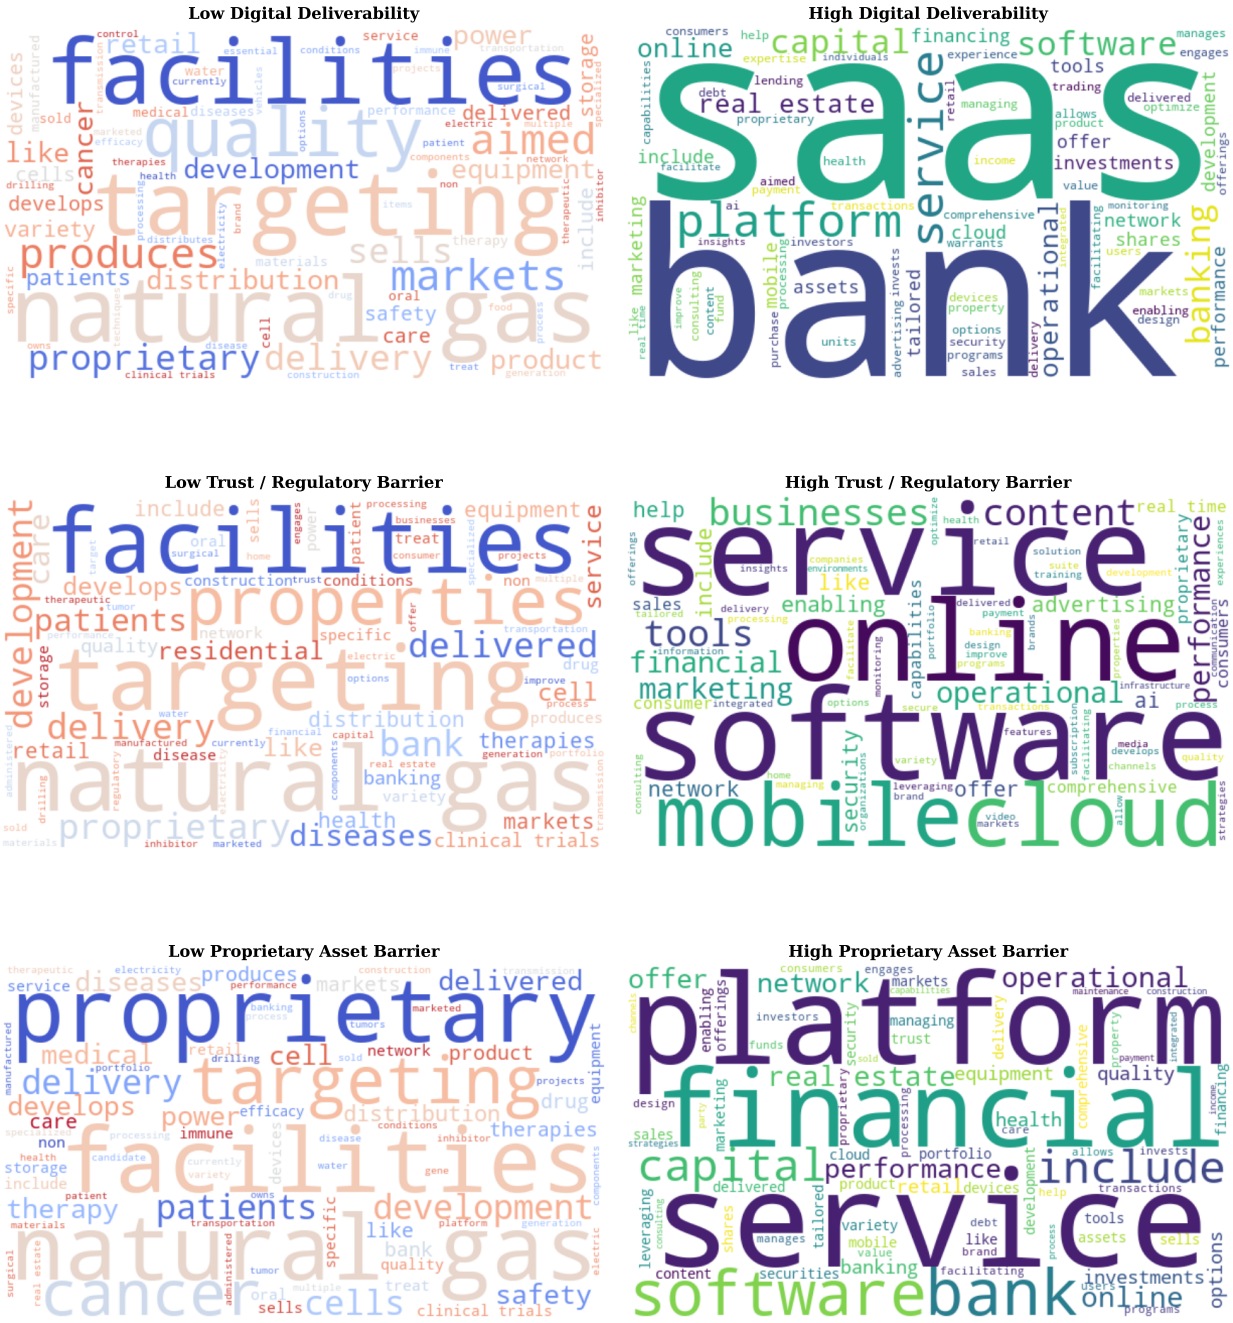


Figure 2: Word Clouds of Product Descriptions by Scoring Dimension

Note: Each panel displays a word cloud of product descriptions from 10-K filings for products scoring in the top quartile ("High") or bottom quartile ("Low") on each scoring dimension. Product descriptions are from the LLM extraction step.

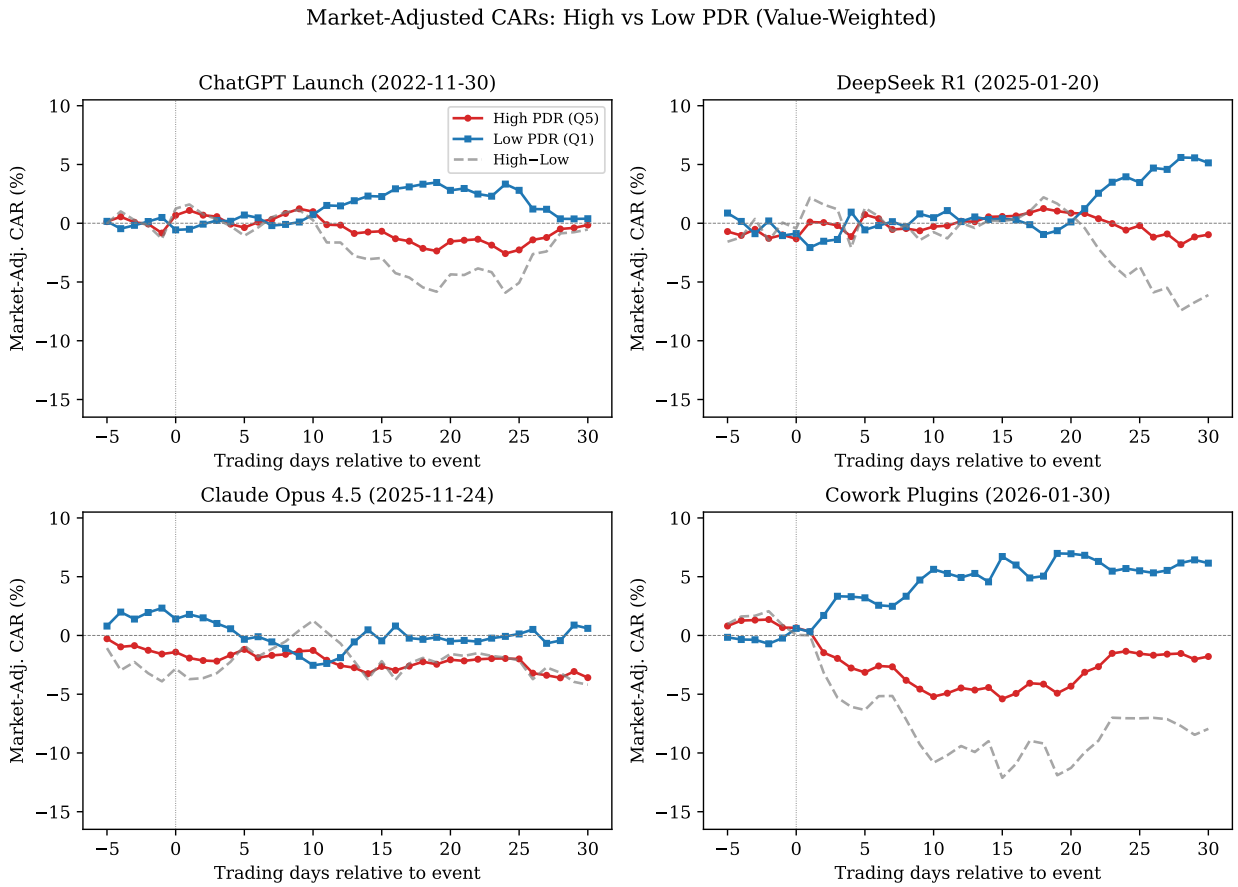


Figure 3: Market-Adjusted Cumulative Abnormal Returns: High vs. Low AI-PDR (Value-Weighted)

Note: Each panel plots value-weighted mean market-adjusted CARs for the top quintile (High PDR, Q5) and bottom quintile (Low PDR, Q1) of firms sorted by AI-PDR. Value weights use pre-event market capitalization. The grey dashed line shows the High−Low spread. The vertical dotted line marks the event date (day 0). All four panels share the same y-axis scale. Sample includes all CRSP firms with non-missing AI-PDR.

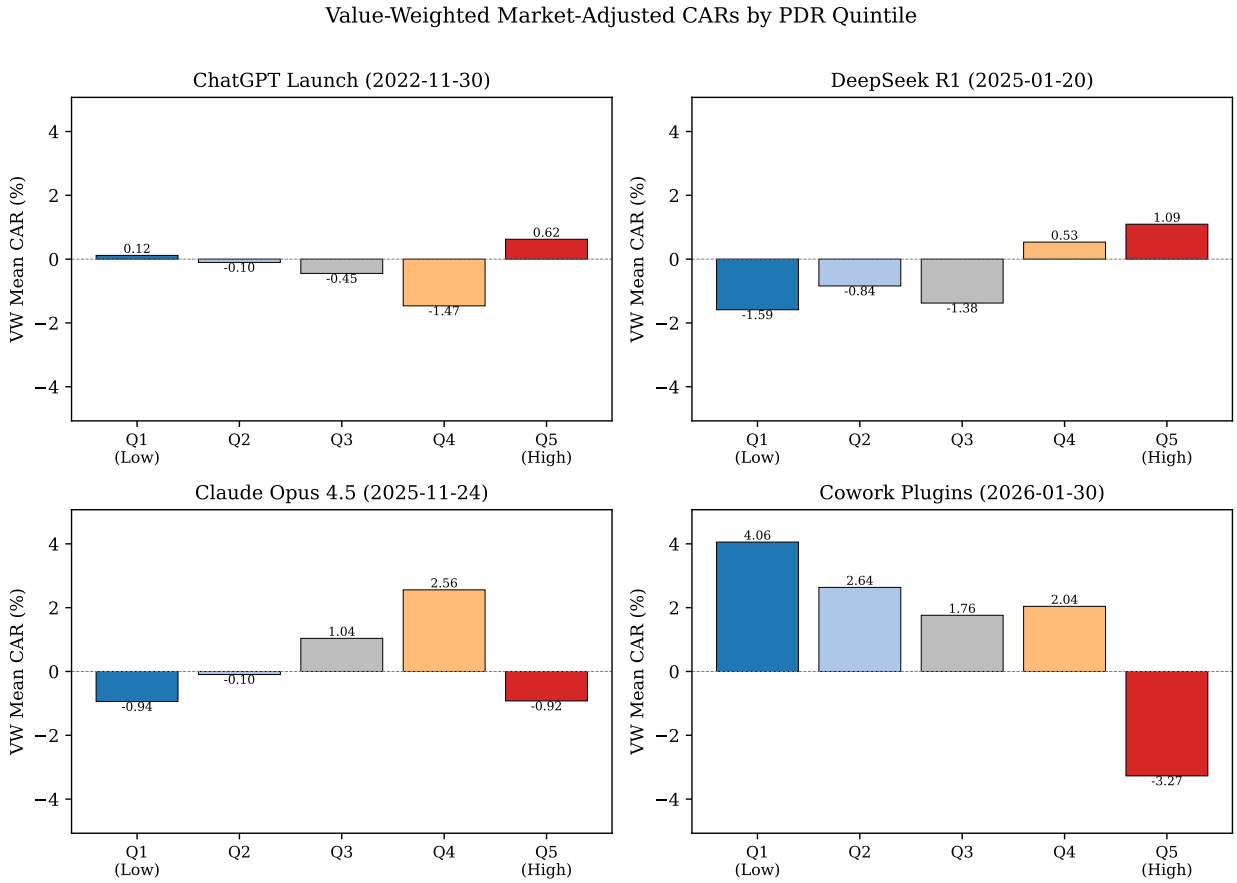


Figure 4: Value-Weighted Mean CARs by AI-PDR Quintile

Note: Firms are sorted into quintiles based on AI-PDR. Bars show value-weighted mean market-adjusted CARs over the [−1, +3] trading-day window around each event, using CRSP market capitalization on the event date as weights. Returns are winsorized at the 1st and 99th percentiles.

Table 1: Summary Statistics

	N	Mean	SD	P10	P25	Median	P75	P90
Panel A — AI-PDR and Firm Characteristics								
AI-PDR	3,950	0.320	0.181	0.083	0.185	0.314	0.409	0.581
Cowork Overlap	3,950	0.037	0.064	0.000	0.000	0.011	0.045	0.102
ESZT AI Exp.	2,079	0.363	0.067	0.281	0.324	0.369	0.408	0.441
log(ME)	3,668	6.764	2.527	3.257	4.998	6.833	8.540	10.010
log(BM)	3,668	-0.927	1.135	-2.452	-1.578	-0.734	-0.106	0.321
Market Beta	3,935	1.030	0.550	0.362	0.658	0.971	1.353	1.752
Panel B — Event-Window CARs (%)								
ChatGPT Launch	4,577	-0.018	7.165	-7.275	-3.234	-0.480	2.734	7.861
DeepSeek R1	4,177	-0.483	8.358	-7.884	-3.996	-1.253	2.113	7.708
Claude Opus 4.5	4,013	2.771	7.852	-4.459	-1.726	1.478	6.186	12.422
Cowork Plugins	3,950	-0.356	10.475	-14.472	-5.781	1.523	6.730	10.666

Note: Panel A reports cross-sectional distributions of AI-PDR and firm characteristics for the Cowork Plugins regression sample. AI-PDR is the equal-weighted average of three sub-dimension scores (digital deliverability, trust sensitivity, proprietary moat). Cowork Overlap is the fraction of a firm's product lines that overlap with Cowork plugin categories. ESZT AI Exp. is the Eisfeldt et al. (forthcoming) standardized labor exposure to AI. Market beta is estimated over $[-252, -11]$ trading days before the event. log(ME) and log(BM) are from the fiscal year ending before each event. Panel B reports distributions of market-adjusted cumulative abnormal returns (CARs, in percent) over the $[-1, +3]$ event window, winsorized at the 1st and 99th percentiles.

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

Event	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{ESZT}}$	N	R^2
Panel A — Baseline				
ChatGPT Launch	0.0260	0.0666***	2,131	0.1487
	(1.41)	(2.62)		
	[1.31]	[2.55]		
DeepSeek R1	0.0159	0.0820**	1,964	0.1626
	(0.76)	(2.47)		
	[0.72]	[2.36]		
Claude Opus 4.5	−0.0204	−0.0261	1,863	0.2427
	(−1.02)	(−0.77)		
	[−0.89]	[−0.81]		
Cowork Plugins	−0.0925***	−0.0405	1,832	0.4295
	(−3.44)	(−1.12)		
	[−3.51]	[−1.04]		
Panel B — Controlling for Market Beta				
ChatGPT Launch	0.0216	0.0654**	2,127	0.1519
	(1.17)	(2.56)		
	[1.05]	[2.51]		
DeepSeek R1	0.0149	0.0811**	1,963	0.1680
	(0.72)	(2.42)		
	[0.70]	[2.33]		
Claude Opus 4.5	−0.0487***	−0.0318	1,862	0.3544
	(−2.72)	(−1.04)		
	[−2.53]	[−1.25]		
Cowork Plugins	−0.0846***	−0.0382	1,832	0.4345
	(−3.11)	(−1.04)		
	[−3.10]	[−0.97]		
Wald test (Panel A): H_0: $\hat{\beta}_{Cowork} = \overline{\beta}_{Demos}$				
F-statistic	57.16***		$p = 0.000$	
	(−7.56)			

Note to Table 2: Dependent variable is the market-adjusted cumulative abnormal return (CAR) over $[-1, +3]$ trading days around each event, winsorized at the 1st and 99th percentiles. $\hat{\beta}_{\text{PDR}}$ is the coefficient on AI-PDR; $\hat{\beta}_{\text{ESZT}}$ is the coefficient on the Eisfeldt et al. (forthcoming) (ESZT) standardized labor exposure. GICS sub-industry fixed effects (163 categories from Compustat). Industries with fewer than 5 firms are excluded. Controls in Panel A are log(ME), log(BM), OP, INV from the fiscal year ending before each event. Panel B adds pre-event market beta (estimated over $[-252, -11]$ trading days, winsorized at 1/99) as an additional control. The Wald test stacks the four AI events in Panel A with event-specific β_{PDR} slopes and tests whether the Cowork Plugins coefficient differs from the average of the three capability-demo coefficients. Signed $t = \text{sgn}(\hat{\delta})\sqrt{F}$. Significance stars based on heteroskedasticity-robust t-statistics (parentheses); industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table 3: Window Robustness: CAR Regressions across Event Windows

	[−1, +1]	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
Panel A — Baseline					
ChatGPT Launch	0.0279**	0.0260	0.0237	−0.0302	−0.0200
	(2.04)	(1.41)	(0.74)	(−0.72)	(−0.41)
	[2.19]	[1.31]	[0.54]	[−0.65]	[−0.39]
DeepSeek R1	0.0081	0.0159	0.0517*	0.0646*	0.0244
	(0.52)	(0.76)	(1.87)	(1.67)	(0.46)
	[0.43]	[0.72]	[1.44]	[1.45]	[0.44]
Claude Opus 4.5	−0.0036	−0.0204	−0.0514	−0.0923**	−0.0853*
	(−0.21)	(−1.02)	(−1.55)	(−2.04)	(−1.65)
	[−0.21]	[−0.89]	[−1.36]	[−2.05]	[−1.53]
Cowork Plugins	−0.0400**	−0.0925***	−0.2297***	−0.1894***	−0.1960***
	(−2.10)	(−3.44)	(−5.11)	(−3.40)	(−2.94)
	[−1.88]	[−3.51]	[−4.57]	[−3.43]	[−2.51]
Panel B — With Market Beta Control					
	[−1, +1]	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
ChatGPT Launch	0.0185	0.0216	0.0242	−0.0162	−0.0312
	(1.43)	(1.17)	(0.76)	(−0.39)	(−0.65)
	[1.42]	[1.05]	[0.56]	[−0.35]	[−0.61]
DeepSeek R1	0.0071	0.0149	0.0527*	0.0651*	0.0302
	(0.47)	(0.72)	(1.91)	(1.68)	(0.59)
	[0.40]	[0.70]	[1.46]	[1.45]	[0.57]
Claude Opus 4.5	−0.0246	−0.0487***	−0.0918***	−0.1256***	−0.1243**
	(−1.53)	(−2.72)	(−2.99)	(−2.89)	(−2.51)
	[−1.67]	[−2.53]	[−3.00]	[−3.16]	[−2.54]
Cowork Plugins	−0.0368*	−0.0846***	−0.2198***	−0.1831***	−0.1818***
	(−1.93)	(−3.11)	(−4.86)	(−3.24)	(−2.70)
	[−1.67]	[−3.10]	[−4.22]	[−3.14]	[−2.24]

Note: Each cell reports $\hat{\beta}_{\text{PDR}}$ from a regression of market-adjusted CARs over the indicated window on AI-PDR, Eisfeldt et al. (forthcoming) labor exposure (suppressed), and controls (log(ME), log(BM), OP, INV, GICS sub-industry FE). Panel A uses the baseline specification. Panel B adds pre-event market beta estimated over $[-252, -21]$ trading days. Industries with fewer than 5 firms are excluded. Heteroskedasticity-robust t-statistics in parentheses; industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table 4: AI-PDR Effect by Cowork Plugin Overlap Subsample

	Baseline	High Overlap $>$ median	Low Overlap $\leq$ median
ChatGPT Launch	0.0216	−0.0245	0.0612*
	(1.17)	(−0.90)	(1.95)
	[1.05]	[−0.97]	[1.42]
N	2,127	1,003	961
DeepSeek R1	0.0149	0.0224	0.0039
	(0.72)	(0.71)	(0.09)
	[0.70]	[0.57]	[0.08]
N	1,963	906	888
Claude Opus 4.5	−0.0487***	−0.0320	−0.0271
	(−2.72)	(−1.44)	(−0.78)
	[−2.53]	[−1.37]	[−0.80]
N	1,862	858	841
Cowork Plugins	−0.0846***	−0.1743***	0.0053
	(−3.11)	(−4.11)	(0.12)
	[−3.10]	[−4.39]	[0.15]
N	1,832	834	828

Note: Each cell reports $\hat{\beta}_{\text{PDR}}$ from a market-adjusted CAR regression over $[-1, +3]$ with GICS sub-industry fixed effects, Eisfeldt et al. (forthcoming) labor exposure, and controls (log(ME), log(BM), OP, INV, pre-event market beta). "Baseline" uses the full sample; overlap subsamples split at the median Cowork Overlap score (see Section 3.4). Industries with fewer than 5 firms are excluded. Heteroskedasticity-robust t-statistics in parentheses; industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table 5: Sub-Component Decomposition: Joint Regression of CARs on PDR Dimensions

Event	$\hat{\beta}_{\text{Digital}}$	$\hat{\beta}_{\text{Trust}}$	$\hat{\beta}_{\text{Prop.}}$	$\hat{\beta}_{\text{ESZT}}$	N	R^2
ChatGPT Launch	0.0033	0.0014	0.0199	0.0662***	2,127	0.1519
	(0.14)	(0.03)	(0.49)	(2.59)		
	[0.14]	[0.03]	[0.49]	[2.49]		
DeepSeek R1	0.0157	−0.0652	0.0737	0.0827**	1,963	0.1694
	(0.57)	(−1.47)	(1.59)	(2.44)		
	[0.52]	[−1.42]	[1.55]	[2.32]		
Claude Opus 4.5	−0.0318	0.0264	−0.0441	−0.0313	1,862	0.3548
	(−1.39)	(0.61)	(−0.97)	(−1.02)		
	[−1.32]	[0.65]	[−1.01]	[−1.22]		
Cowork Plugins	−0.1296***	0.0265	0.0600	−0.0260	1,832	0.4398
	(−4.70)	(0.46)	(1.08)	(−0.71)		
	[−4.28]	[0.49]	[1.12]	[−0.67]		

Note: Each row reports a single regression of market-adjusted CARs over $[-1, +3]$ on three PDR sub-component scores (Digital Deliverability, Trust/Regulatory Barrier, Proprietary Asset Barrier) and Eisfeldt et al. (forthcoming) labor exposure, entered jointly. Controls: pre-event market beta, log(ME), log(BM), OP, INV, and GICS sub-industry fixed effects. Industries with fewer than 5 firms are excluded. Heteroskedasticity-robust t-statistics (parentheses); industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Internet Appendix

Listing A1: Product scoring prompt (full)

```
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 how inherently vulnerable
it is to displacement by generative AI.
Higher scores = easier for generative AI to displace.

Dimensions (0-10 integers):

A. DIGITAL DELIVERABILITY
   Could a GenAI system produce and deliver this product/
   service to the end customer in purely digital form, with
   no physical goods, no human labor, and no financial
   capital required?
   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 with minor digital component
        (e.g., equipment rental with online booking)
   4  = requires significant non-digital input such as
        financial capital, human workers, or regulated
        infrastructure (e.g., bank loans, staffing,
        exchange trading)
   6  = mostly digital but requires real-time human
        expertise (e.g., custom IT consulting, brand
        strategy)
   8  = digital output, some human oversight (e.g.,
        analytics reports, curated data feeds)
   10 = entirely digital, an LLM can produce equivalent
        output (e.g., template marketing copy, stock
        images, basic code)

B. TRUST / REGULATORY BARRIERS
```

```
   Would a customer accept this product from an AI with no
   human or institutional brand behind it? Score by matching
   the category below, not the specific examples.
   0  = legally cannot (e.g., FDA-approved drugs, licensed
        medical devices, regulated insurance products)
   2  = heavy professional gatekeeping (e.g., audited
        financials, licensed engineering)
   4  = meaningful trust required (e.g., wealth management,
        enterprise software with SLAs)
   6  = moderate trust needed (e.g., news journalism,
        mid-tier consulting)
   8  = light trust (e.g., consumer software tools, online
        courses, routine data products)
   10 = no trust barrier (e.g., commodity digital content,
        stock photos, template documents)
   Higher = EASIER to displace

C. PROPRIETARY ASSET BARRIERS
   Does producing this require assets that an AI entrant
   cannot replicate? Score by matching the category below,
   not the specific examples.
   0  = massive moat (e.g., proprietary physical
        infrastructure like pipeline networks; dominant
        network effects like major stock exchanges;
        decades of irreplaceable proprietary data)
   2  = strong moat (e.g., large proprietary dataset,
        established distribution network, regulated
        license that limits entry)
   4  = moderate moat (e.g., specialized but replicable
        datasets, brand reputation, established client
        base)
   6  = weak moat (e.g., general industry expertise,
        customer relationships that could shift)
   8  = minimal moat (e.g., off-the-shelf technology,
        generic service delivery, thin switching costs)
   10 = no moat (e.g., basic content creation, simple
        translation, template generation)
   Higher = EASIER to displace
```

```
IMPORTANT SCORING RULES:
1. Write reasoning BEFORE each score. Think step-by-step
   about what this specific product actually requires to
   produce and deliver, then commit to a number.
2. For digital deliverability, ask: "What is the binding
   input the customer is paying for?" If it is capital,
   physical goods, human workers, or regulated
   infrastructure, the score is LOW regardless of whether
   the transaction happens online.
3. 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.
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 your score on the product DESCRIPTION, not just
   the name. Marketing names can be misleading.
```

Listing A2: Cowork Plugin overlap scoring prompt (full)

```
SYSTEM:
You are a product-market overlap scoring engine. Output must
be valid JSON following the schema exactly. No markdown, no
commentary.

USER:
On January 30, 2026, Anthropic launched 10 "Cowork Plugins"
-- AI agents that automate specific knowledge-work functions.
Each plugin can replace both specialized software products
and human professional services in its domain.

The 10 plugins and their specific capabilities:

1. PRODUCTIVITY -- Task & calendar management, workplace
   memory, daily briefings, meeting summaries, task
   dashboards.
2. SALES -- Prospect research, CRM automation, pipeline
   review, call summaries, sales forecasting, outreach
   drafting, competitive intelligence, custom sales assets.
3. CUSTOMER SUPPORT -- Ticket triage & routing, customer
   research, response drafting, escalation packaging,
   knowledge base article authoring.
4. PRODUCT MANAGEMENT -- PRD/spec writing, roadmap
   development, stakeholder communication, user research
   synthesis, competitive briefs, product metrics analysis.
5. MARKETING -- Content creation, campaign planning, brand
   voice management, SEO audits, email nurture sequences,
   performance reporting, competitive analysis.
6. LEGAL -- Contract review, NDA triage, vendor agreement
   tracking, compliance workflows, legal briefings,
   templated responses for DSARs and discovery holds.
7. FINANCE & ACCOUNTING -- Month-end close sequencing,
   journal entry preparation, account reconciliation,
   financial statement generation, variance/flux analysis,
   SOX 404 control testing.
8. DATA ANALYSIS -- Natural-language SQL queries, data
   exploration & profiling, publication-quality
```

```
   visualization, interactive HTML dashboards, analysis
   validation and QA.
9. ENTERPRISE SEARCH -- Unified cross-platform search
   across chat, email, cloud storage, wikis, project
   management, CRM, and ticketing; activity digests.
10. BIO RESEARCH -- Literature review, single-cell RNA-seq
    analysis, bioinformatics pipelines, clinical trial
    protocol generation, lab data standardization.

TASK: For each product/service below, score how directly it
competes with or could be replaced by these 10 Cowork
Plugins.

Score (0-10 integer):
  0 = no overlap -- physical products, physical
      infrastructure, capital-intensive services
  1 = negligible
  3 = weak -- touches a function a plugin covers, but the
      core value proposition is fundamentally different
  5 = moderate -- meaningfully overlaps with a subset of
      one plugin's capabilities
  7 = strong -- the product's primary function is
      substantially covered by one plugin
  9 = near-complete -- the product is almost exactly what
      one plugin does
  10 = exact match

RULES:
1. Write reasoning BEFORE the score.
2. Identify the single most relevant plugin.
3. Score based on what the PRODUCT does, not the buyer's
   industry.
4. Physical products score 0 regardless of industry.
5. Most products have zero or negligible overlap. Expect
   the majority of scores to be 0-2.
```

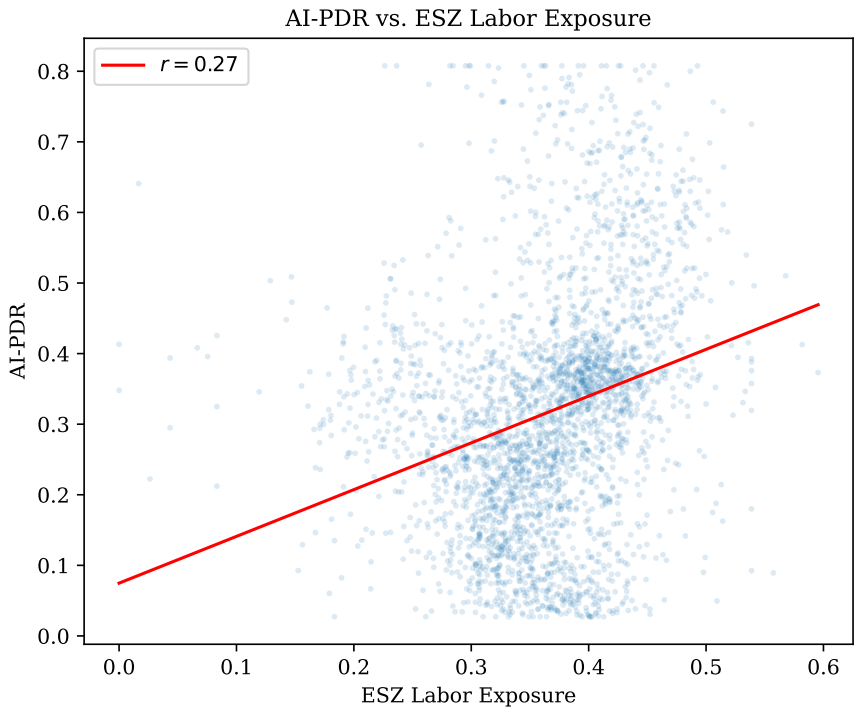


Figure A1: AI-PDR vs. ESZT Labor Exposure

Note: Each dot represents a firm. The horizontal axis is the Eisfeldt et al. (forthcoming) generative AI labor exposure measure; the vertical axis is AI-PDR. The red line is a linear fit with Pearson $r = 0.27$. The moderate positive correlation confirms that the two measures capture related but distinct dimensions of AI exposure.

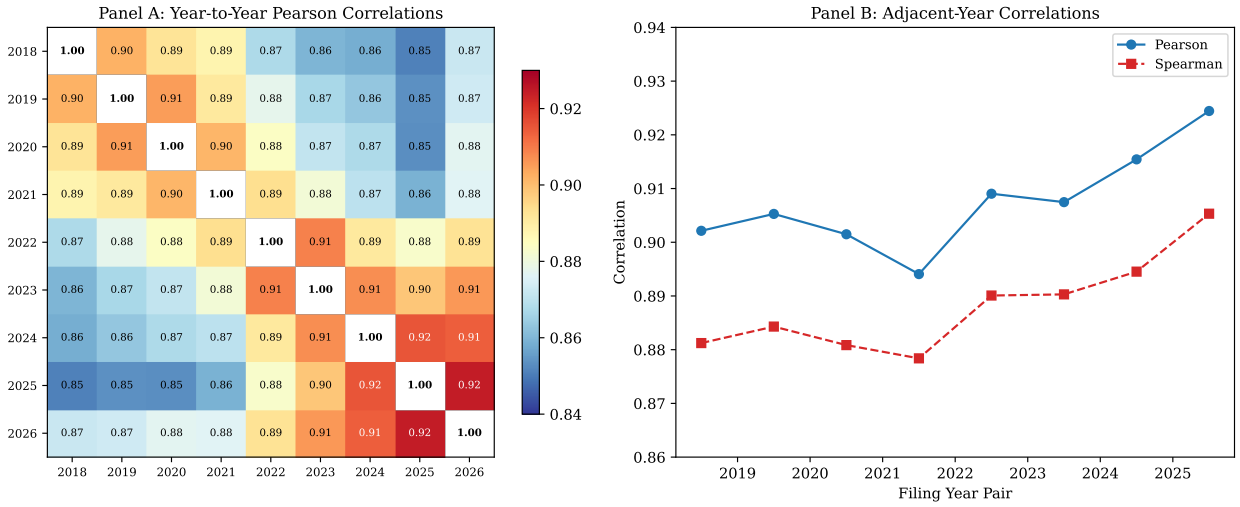


Figure A2: Temporal Stability of AI-PDR Scores

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

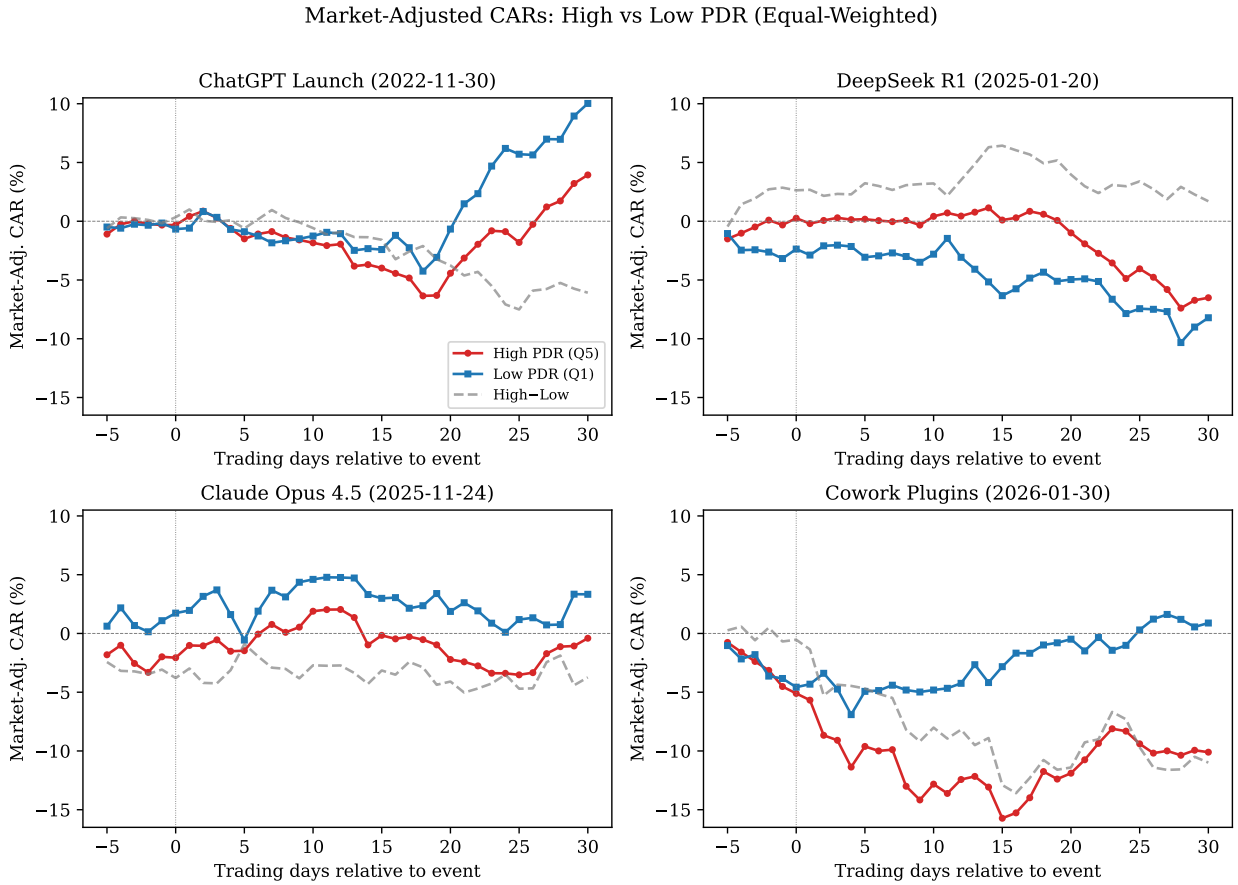


Figure A3: Market-Adjusted Cumulative Abnormal Returns: High vs. Low AI-PDR (Equal-Weighted)

Note: Same as Figure 3 but using equal-weighted (winsorized at the 2nd and 98th percentiles) rather than value-weighted means.

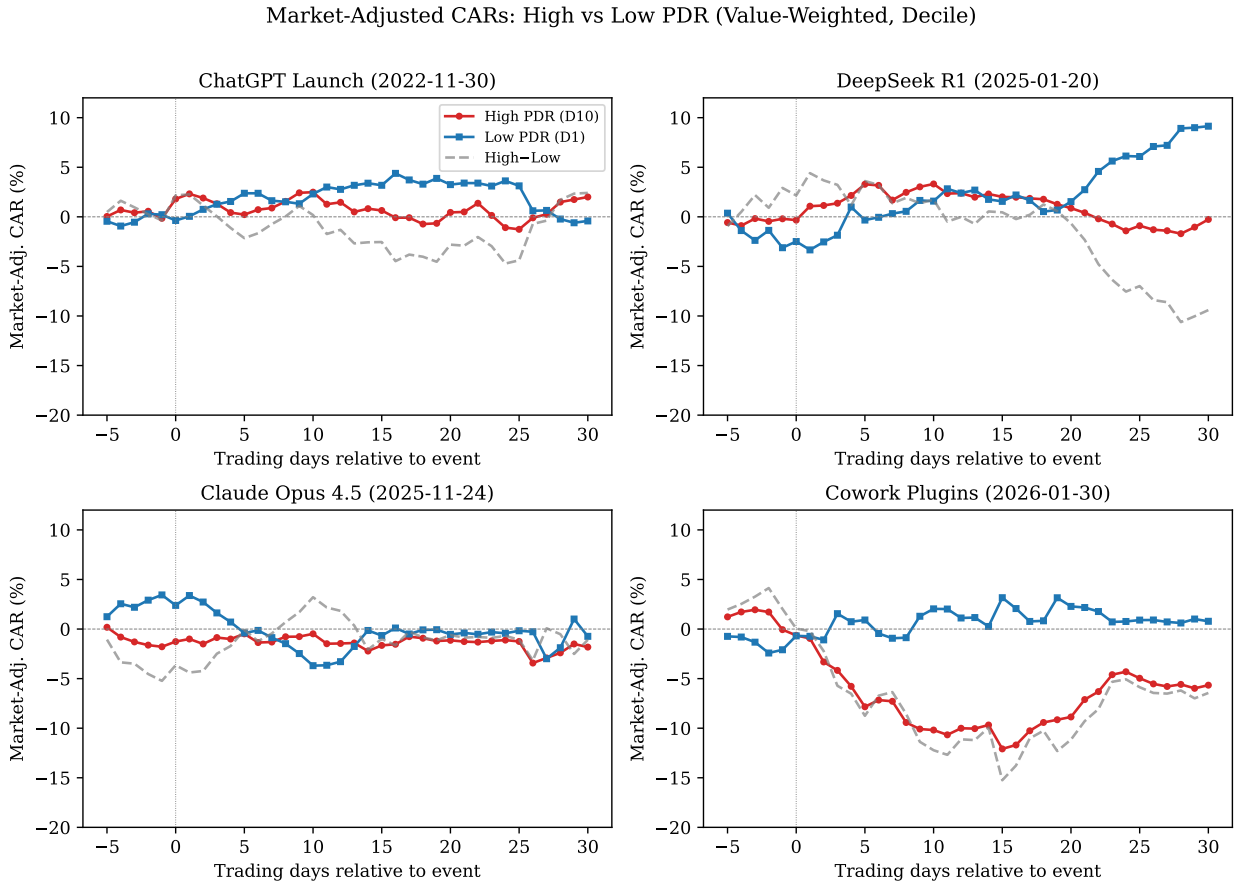


Figure A4: Market-Adjusted Cumulative Abnormal Returns: High vs. Low AI-PDR (Value-Weighted, Decile)

Note: Same as Figure 3 but using the top decile (D10) and bottom decile (D1) of firms sorted by AI-PDR.

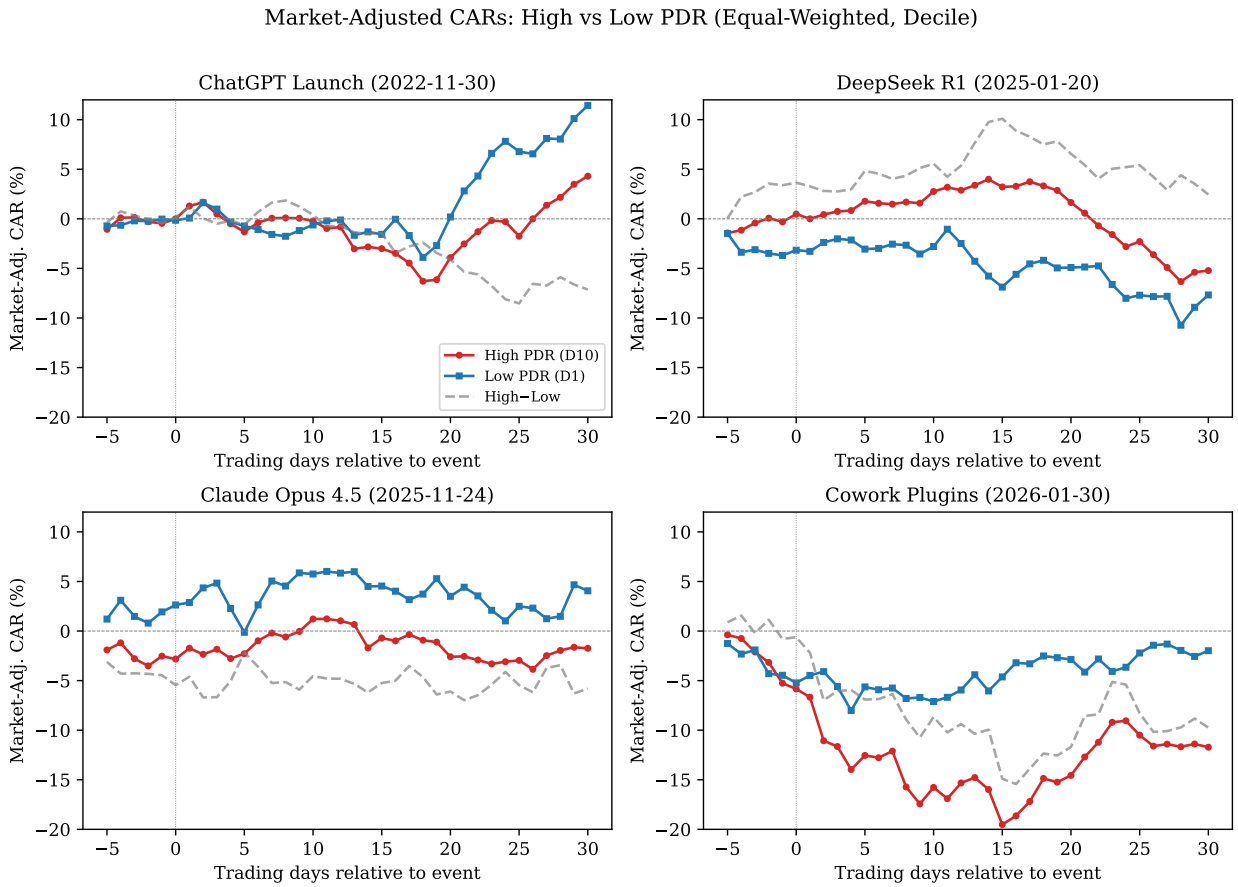


Figure A5: Market-Adjusted Cumulative Abnormal Returns: High vs. Low AI-PDR (Equal-Weighted, Decile)

Note: Same as Figure A3 but using the top decile (D10) and bottom decile (D1) of firms sorted by AI-PDR.

Table A1: Variable Definitions

Variable	Definition
Panel A: Dependent Variables	
CAR (market-adjusted)	Cumulative abnormal return over the event window, computed as $\sum_t (r_{i,t} - r_{m,t})$ where $r_{m,t}$ is the CRSP value-weighted market return. Baseline window is $[-1, +3]$ trading days. Winsorized at 1st/99th percentiles.
CAR (FF5-adjusted)	Cumulative abnormal return using Fama–French five-factor model residuals. Factor loadings estimated over $[-252, -11]$ trading days before each event.
Panel B: Key Independent Variables	
AI-PDR	AI Product Displacement Risk. Importance-weighted average of product-level scores (each the arithmetic mean of digital deliverability, trust barrier, and proprietary barrier) extracted from 10-K Item 1 filings by GPT-4o mini, averaged across all filing years (2018–2026), rescaled to $[0, 1]$. Winsorized at 1st/99th percentiles.
Digital deliverability	Product-level score (0–10): degree to which a generative AI system could produce and deliver the product in purely digital form with no physical goods, human labor, or financial capital.
Trust barrier	Product-level score (0–10): degree to which a customer would accept the product from an AI system with no human or institutional brand behind it. Higher scores indicate lower trust barriers (more vulnerable).
Proprietary barrier	Product-level score (0–10): degree to which producing the product requires assets an AI entrant cannot replicate (proprietary data, regulatory licenses, physical capital). Higher scores indicate lower barriers (more vulnerable).

Continued on next page

Table A1 (continued)

Variable	Definition
Cowork Overlap	Firm-level measure (0–1 scale) of how closely the firm's products compete with the 10 Cowork Plugins launched January 30, 2026. Each extracted product is scored (0–10) by GPT-4o mini on competitive overlap with the specific plugin capabilities (see Listing A2); scores are rescaled to $[0, 1]$, aggregated to firm level using importance weights, and averaged across filing years. Plugin definitions are from the Anthropic launch announcement (`anthropic.com/research/cowork-plugins`, accessed February 2026). Winsorized at 1st/99th percentiles. Correlation with AI-PDR: $\rho = 0.67$.
ESZT AI exposure	Standardized firm-level generative AI labor exposure from Eisfeldt et al. (forthcoming). Mean zero, unit standard deviation.
Panel C: Firm-Level Controls	
log(ME)	Natural log of market equity (price × shares outstanding) from CRSP/Compustat, following Fama–French timing conventions. Winsorized at 1st/99th percentiles.
log(BM)	Natural log of book-to-market ratio. Fama–French timing: fiscal year $t-1$ book equity matched to returns July of year t through June of year $t+1$. Winsorized at 1st/99th percentiles.
OP	Operating profitability, defined following Fama and French (2015). Winsorized at 1st/99th percentiles.
INV	Investment (total asset growth), defined following Fama and French (2015). Winsorized at 1st/99th percentiles.
log(Assets)	Natural log of total assets from Compustat.
PPE/Assets	Property, plant, and equipment scaled by total assets.
Intangibles/Assets	Intangible assets scaled by total assets.
Profit margin	Operating income before depreciation scaled by revenue.
Leverage	Long-term debt scaled by total assets.
Cash/Assets	Cash and short-term investments scaled by total assets.

Continued on next page

Table A1 (continued)

Variable	Definition
Momentum	Cumulative stock return over the prior 12 months.
Volatility	Standard deviation of daily stock returns over the pre-event window.
log(Age)	Natural log of firm age, measured as years since first appearance in Compustat.
R&D/Assets	Research and development expenditure (Compustat `xrd`) scaled by total assets.
Panel D: Systematic Risk Controls	
Market beta	Firm-level CAPM beta estimated from daily excess returns regressed on market excess return and the daily change in the 10-year Treasury yield, over $[-252, -11]$ trading days before each event. Winsorized at 1st/99th percentiles.
Rate sensitivity	Coefficient on the daily change in the 10-year Treasury yield from the same regression as market beta. Captures interest-rate (duration) exposure. Winsorized at 1st/99th percentiles.
Panel E: Heterogeneity Interactions	
SA index	Size-age financial constraints index from Hadlock and Pierce (2010).
WW index	Financial constraints index from Whited and Wu (2006).
KZ index	Financial constraints index from Kaplan and Zingales (1997).
Panel F: Fixed Effects	
GICS sub-industry FE	Fixed effects for GICS sub-industry classifications (163 categories). Industries with fewer than 5 firms in the regression sample are excluded.

Table A2: AI-PDR Sub-Component Properties

Panel A — Summary Statistics

	Mean	Std	P5	P25	P75	P95
AI-PDR (composite)	0.319	0.185	0.052	0.175	0.413	0.685
Digital Deliverability	0.285	0.230	0.023	0.091	0.413	0.771
Trust/Regulatory Barrier	0.330	0.192	0.019	0.191	0.444	0.671
Proprietary Asset Barrier	0.341	0.158	0.092	0.220	0.434	0.621

Panel B — Pairwise Correlations

	Digital	Trust	Proprietary
Digital	1.000		
Trust	0.845	1.000	
Proprietary	0.825	0.961	1.000

Note: Panel A reports summary statistics for the composite AI-PDR score and its three sub-components, each on a 0–1 scale. All scores are importance-weighted averages across product lines within each firm, then averaged across available filing years (2018–2026) and winsorized at the 1st and 99th percentiles. Panel B reports pairwise Pearson correlations among the three sub-components. $N = 6,193$ firms.

Table A3: Examples of High- and Low-AI-PDR Firms

Firm	Top Product	Dig.	Trust	Prop.	PDR	Sector
High AI-PDR						
Klaviyo, Inc.	Omnichannel Messaging Solutions	8	8	8	0.81	Info. Tech.
PINTEREST, INC.	Advertising Services	8	8	8	0.81	Comm. Svc.
ZIPRECRUITER, INC.	Applicant Tracking System (ATS)	8	8	6	0.81	Comm. Svc.
ADOBE INC.	Acrobat Solutions	10	8	8	0.81	Info. Tech.
SEMrush Holdings, Inc.	Search Engine Optimization (SEO) Service	8	8	8	0.81	Info. Tech.
Snap Inc	Visual Messaging Services	10	10	10	0.81	Comm. Svc.
TWILIO INC	Voice Services	10	8	8	0.81	Info. Tech.
DROPBOX, INC.	Security Solutions	6	6	6	0.81	Info. Tech.
Taboola.com Ltd.	Performance Advertising Services	10	8	8	0.81	Comm. Svc.
Roblox Corp	Roblox Cloud	6	6	4	0.81	Comm. Svc.
Low AI-PDR						
GREAT PLAINS ENERGY INC	Electric Utility Services	0	0	0	0.03	Utilities
Enhabit, Inc.	Hospice Services	0	0	0	0.03	Health Care
PACIFIC DRILLING S.A.	Deepwater Drilling Services	0	0	0	0.03	Energy
CALA CORP	Undersea Resort and Casino	0	0	0	0.03	nan
ACADIA PHARMACEUTICALS INC	NUPLAZID	0	0	2	0.03	Health Care
Connect Biopharma Holdings Ltd	COPD Treatment	0	0	0	0.03	Health Care
UNIVERSAL HEALTH REALTY INCOME TRUS	Medical/Office Buildings (MOBs)	0	2	2	0.04	Real Estate
Westlake Chemical Partners LP	Ethylene Production	0	0	0	0.05	Materials
BWX Technologies, Inc.	Environmental Management Services	0	0	2	0.06	Industrials
Carolina Trust BancShares, Inc.	Banking Services	0	0	2	0.07	Financials

Note: Dig. = Digital Deliverability, Trust = Trust/Regulatory Barrier, Prop. = Proprietary Asset Barrier (each 0–10, higher = more vulnerable). PDR is the firm-level AI Product Displacement Risk score (0–1 scale). "Top Product" is the highest-importance product line extracted from the firm's most recent 10-K filing.

Table A4: Correlation of AI-PDR with Firm Characteristics

Variable	N	Pearson r	Spearman ρ	$\hat{\beta}$ (Ind. FE)
ESZ Labor Exposure	2,966	0.271	0.346	0.198^{***}
log(Assets)	6,097	0.095	0.134	-0.003^{***}
R&D/Assets	3,362	-0.152	-0.300	0.006^{**}
Intangibles/Assets	5,984	0.290	0.316	0.028^{***}
PPE/Assets	5,903	-0.229	-0.190	-0.092^{***}
Profit Margin	4,932	0.051	0.081	0.000

Note: Columns 2–3 report Pearson and Spearman correlations between firm-level AI-PDR and each characteristic. Column 4 reports the coefficient from a univariate regression of AI-PDR on each characteristic with GICS sub-industry fixed effects (industries with < 5 firms excluded), using heteroskedasticity-robust standard errors. ESZ Labor Exposure is the Eisfeldt et al. (forthcoming) generative AI labor exposure measure. All Compustat variables are from the most recent fiscal year. *, **, *** denote significance at 10%, 5%, 1%.

Table A5: Timeline of Anthropic Product Announcements Around the Event Window

Date	Announcement	Nature	PS?	Justification
Feb 2025	Claude Code preview	Developer tool	No	CLI coding tool for software developers; no SaaS workflow automation.
May 22, 2025	Claude Code GA	Developer tool GA	No	Broader availability of same coding tool. No domain-specific automation.
Sep 29, 2025	Claude Sonnet 4.5	Model release	No	Model upgrade. No product-market entry.
Oct 15, 2025	Claude Haiku 4.5	Model release	No	Fast/cheap model tier. No product-market entry.
Nov 24, 2025	Claude Opus 4.5	Model release	No	Frontier model upgrade. Improves capability but does not package domain-specific workflow tools. Classified as capability demonstration in our study.
Jan 12, 2026	Cowork preview (Max)	Desktop agent	No	General-purpose desktop agent without plugins or domain-specific automation. Max only ($100–200/mo).
Jan 16, 2026	Cowork Pro ($20/mo)	Pricing/access	No	Widens access to same general-purpose agent. No plugins or domain-specific functionality.
Jan 30, 2026	**Cowork Plugins**	**Product-mkt entry**	**Yes**	**11 open-source plugins (sales, legal, finance, marketing, customer support, data analysis, etc.) bundling domain knowledge, connectors, and sub-agents to automate workflows sold as standalone SaaS subscriptions. First entry into identifiable incumbent product markets ($20+/mo).**
Feb 3–4, 2026	Market reaction	Market event	—	$285B software selloff. Thomson Reuters −16%, LegalZoom −20%, RELX −14%. Market attributed reaction to Jan 30 plugin launch.
Feb 5, 2026	Claude Opus 4.6	Model release	No	Model upgrade (1M context, agent teams). Postdates event window. No new product-market entry.
Feb 17, 2026	Claude Sonnet 4.6	Model release	No	Sonnet-tier model upgrade. No product-market entry.
Feb 24, 2026	Cowork enterprise update	Infrastructure	No	Connectors (Google Workspace, DocuSign, FactSet), private plugin marketplaces, 12 new plugins. Distribution expansion, not new product-substitution event. Stocks partially recovered.

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

Table A6: High-PDR Firms Outside Technology and Communications

Firm	Top Product	Dig.	Trust	Prop.	PDR	Sector
SharpLink Gaming, Inc.	Direct-to-Player (D2P) Websites	10	8	6	0.81	Cons. Disc.
ETSY INC	Etsy Marketplace	10	8	8	0.78	Cons. Disc.
Asset Entities Inc.	Discord Investment Education and Enterta	8	6	8	0.76	Health Care
Scripps Networks Interactive, Inc.	Advertising Sales	10	8	6	0.75	Cons. Disc.
CAMBIUM LEARNING GROUP, INC.	Raz-Kids	10	8	8	0.75	Cons. Disc.
EBAY INC	Marketplace Platforms	10	10	10	0.75	Cons. Disc.
CHEGG, INC	Subscription Services	10	8	6	0.74	Cons. Disc.
Aether Holdings, Inc.	Smart Stock Scanner	10	8	8	0.74	Financials
Planet Labs PBC	Satellite Analytics	6	6	4	0.73	Industrials
Dave Inc./DE	Dave Checking	10	6	4	0.73	Financials

Note: This table lists the 10 highest-PDR firms after excluding GICS sectors 45 (Information Technology) and 50 (Communication Services). Dig. = Digital Deliverability, Trust = Trust/Regulatory Barrier, Prop. = Proprietary Asset Barrier (each 0–10). PDR is the firm-level AI Product Displacement Risk score (0–1 scale). "Top Product" is the highest-importance product line from the firm's most recent 10-K. GICS classifies firms by primary revenue source, not product delivery method; some high-PDR firms in non-technology sectors (e.g., Chegg, Etsy) derive revenue from digitally-delivered products that GICS does not flag as technology.

Table A7: Robustness: Excluding Technology Firms and Earnings Announcers

Panel A — Excluding Technology and Tech-Adjacent Firms

	Full		Ex-IT (45)		Ex-IT/Comm (45,50)	
Event	$\hat{\beta}_{\text{PDR}}$	N	$\hat{\beta}_{\text{PDR}}$	N	$\hat{\beta}_{\text{PDR}}$	N
ChatGPT Launch	0.0216	2,127	0.0130	1,872	0.0145	1,802
	(1.17)		(0.66)		(0.71)	
DeepSeek R1	0.0149	1,963	0.0219	1,727	0.0225	1,650
	(0.72)		(0.98)		(1.01)	
Claude Opus 4.5	-0.0487^{***}	1,862	-0.0406^{**}	1,637	-0.0343^{*}	1,565
	(-2.72)		(-2.27)		(-1.82)	
Cowork Plugins	-0.0846^{***}	1,832	-0.0738^{***}	1,608	-0.0905^{***}	1,540
	(-3.11)		(-2.61)		(-3.08)	

Panel B — Excluding Firms Reporting Earnings in Event Window

	Full Sample		Excluding Earners	
Event	$\hat{\beta}_{\text{PDR}}$	N	$\hat{\beta}_{\text{PDR}}$	N
ChatGPT Launch	0.0216	2,127	0.0153	2,070
	(1.17)		(0.86)	
	[1.05]		[0.76]	
DeepSeek R1	0.0149	1,963	0.0158	1,799
	(0.72)		(0.74)	
	[0.70]		[0.77]	
Claude Opus 4.5	-0.0487^{***}	1,862	-0.0484^{***}	1,822
	(-2.72)		(-2.81)	
	$[-2.53]$		$[-2.41]$	
Cowork Plugins	-0.0846^{***}	1,832	-0.0702^{***}	1,491
	(-3.11)		(-2.72)	
	$[-3.10]$		$[-2.25]$	

Note to Table A7. Panel A tests whether the PDR effect is driven by technology firms. "Full" is the baseline sample. "Ex-IT" excludes GICS sector 45 (Information Technology). "Ex-IT/Comm" additionally excludes sector 50 (Communication Services, including Interactive Media and Entertainment). Panel A reports only HC1 robust t-statistics. Panel B drops firms whose Compustat quarterly earnings announcement date (RDQ) falls within the event window $[-1, +3]$; it reports both HC1 robust t-statistics (parentheses) and industry-clustered t-statistics [brackets]. All regressions use market-adjusted CARs with GICS sub-industry fixed effects, Eisfeldt et al. (forthcoming) labor exposure, pre-event market beta, and controls (log(ME), log(BM), OP, INV). Industries with fewer than 5 firms are excluded. *, **, *** denote significance at 10%, 5%, 1%.

Table A8: Cross-Sectional Heterogeneity: Cowork Plugins

Interaction	$\hat{\beta}_{\text{PDR}}$		$\hat{\beta}_{\text{interaction}}$		N	R^2
PDR × log(ME)	-0.0827^{***}	(-3.03)	-0.0064	(-1.23)	1,832	0.4353
		$[-3.10]$		$[-1.09]$		
PDR × R&D/Assets	-0.0870^{**}	(-2.12)	0.0592	(0.34)	803	0.3932
		$[-2.58]$		[0.34]		
PDR × SA	-0.0818^{***}	(-3.05)	0.0012	(0.03)	1,832	0.4452
		$[-3.01]$		[0.02]		
PDR × WW	-0.0762^{**}	(-2.56)	-0.0204	(-0.17)	1,370	0.4363
		$[-2.46]$		$[-0.16]$		
PDR × KZ	-0.0834^{***}	(-2.82)	-0.0000	(-0.16)	1,510	0.4227
		$[-2.78]$		$[-0.20]$		

Note: Each row adds an interaction between AI-PDR and a demeaned firm characteristic to the baseline specification for the Cowork Plugins event ($[-1, +3]$). $\hat{\beta}_{\text{PDR}}$ is evaluated at the mean of the interacted variable. R&D intensity is R&D expenditure (Compustat `xrd`) scaled by total assets; firms with missing R&D are dropped. SA is the Hadlock and Pierce (2010) size-age index. WW is the Whited and Wu (2006) index. KZ is the Kaplan and Zingales (1997) index. Higher values of SA, WW, and KZ indicate greater financial constraints. All specifications include GICS sub-industry fixed effects, Eisfeldt et al. (forthcoming) labor exposure, pre-event market beta, and controls (log(ME), log(BM), OP, INV). Heteroskedasticity-robust t-statistics in parentheses; industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table A9: Sub-Component Decomposition: Digital Deliverability vs. Composite AI-PDR

	Digital Only		Composite		
Event	$\hat{\beta}_{\text{Digital}}$	R^2	$\hat{\beta}_{\text{PDR}}$	R^2	N
ChatGPT Launch	0.0137	0.1517	0.0216	0.1519	2,127
	(0.92)		(1.17)		
	[0.83]		[1.05]		
DeepSeek R1	0.0117	0.1681	0.0149	0.1680	1,963
	(0.69)		(0.72)		
	[0.65]		[0.70]		
Claude Opus 4.5	-0.0369^{***}	0.3544	-0.0487^{***}	0.3544	1,862
	(-2.60)		(-2.72)		
	$[-2.35]$		$[-2.53]$		
Cowork Plugins	-0.0851^{***}	0.4382	-0.0846^{***}	0.4345	1,832
	(-4.28)		(-3.11)		
	$[-4.47]$		$[-3.10]$		

Note: "Digital Only" enters digital deliverability as the sole PDR dimension; "Composite" enters the overall AI-PDR score (equal-weighted average of digital deliverability, trust/regulatory barrier, and proprietary asset barrier). Both specifications include Eisfeldt et al. (forthcoming) labor exposure, pre-event market beta, log(ME), log(BM), OP, INV, and GICS sub-industry fixed effects. Industries with fewer than 5 firms are excluded. Heteroskedasticity-robust t-statistics (parentheses); industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table A10: Sub-Component Decomposition across Event Windows (Cowork Plugins)

Window	$\hat{\beta}_{\text{Digital}}$	$\hat{\beta}_{\text{Trust}}$	$\hat{\beta}_{\text{Prop.}}$	N	R^2
$[-1,+1]$	-0.0693^{***}	0.0443	0.0062	1,832	0.3264
	(-3.49)	(1.04)	(0.15)		
	$[-2.89]$	[1.09]	[0.16]		
$[-1,+3]$	-0.1296^{***}	0.0265	0.0600	1,832	0.4398
	(-4.70)	(0.46)	(1.08)		
	$[-4.28]$	[0.49]	[1.12]		
$[-1,+10]$	-0.2689^{***}	-0.0427	0.1880^{*}	1,828	0.3972
	(-5.84)	(-0.45)	(1.92)		
	$[-4.39]$	$[-0.50]$	[2.06]		
$[-1,+20]$	-0.2368^{***}	-0.0318	0.1720	1,825	0.2932
	(-3.85)	(-0.26)	(1.35)		
	$[-3.04]$	$[-0.26]$	[1.38]		
$[-1,+30]$	-0.2095^{***}	-0.0134	0.1101	1,800	0.2696
	(-2.85)	(-0.10)	(0.71)		
	$[-2.16]$	$[-0.10]$	[0.78]		

Note: Each row reports a single regression of market-adjusted CARs for the Cowork Plugins event (January 30, 2026) on three PDR sub-component scores (Digital Deliverability, Trust/Regulatory Barrier, Proprietary Asset Barrier), entered jointly with Eisfeldt et al. (forthcoming) labor exposure (suppressed) and controls (pre-event market beta, log(ME), log(BM), OP, INV, GICS sub-industry FE). Industries with fewer than 5 firms are excluded. Heteroskedasticity-robust t-statistics (parentheses); industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table A11: Pre-Trend Test: Cowork Plugins

Window	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{ESZT}}$	N	R^2
Pre-Event Windows				
$[-10, -2]$	-0.0250	0.0286	1,836	0.2354
	(-0.90)	(0.72)		
	$[-0.73]$	$[0.65]$		
$[-5, -2]$	-0.0132	0.0268	1,836	0.2510
	(-0.77)	(1.01)		
	$[-0.68]$	$[0.97]$		
Event Window				
$[-1, +3]$	-0.0846^{***}	-0.0382	1,832	0.4345
	(-3.11)	(-1.04)		
	$[-3.10]$	$[-0.97]$		

Note: Dependent variable is the market-adjusted CAR over each indicated window around the Cowork Plugins launch (January 30, 2026). Pre-event windows end at day -2 to avoid overlap with the event window $[-1, +3]$. All specifications include GICS sub-industry fixed effects, Eisfeldt et al. (forthcoming) labor exposure, pre-event market beta, and controls (log(ME), log(BM), OP, INV). Heteroskedasticity-robust t-statistics (parentheses); industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table A12: Robustness: Controlling for Interest-Rate Sensitivity

	Baseline	+ Rate Control			
Event	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{Rate}}$	N	R^2
ChatGPT Launch	0.0216	0.0181	-0.1946^{***}	2,127	0.1617
	(1.17)	(1.00)	(-3.22)		
	[1.05]	[0.87]			
DeepSeek R1	0.0149	0.0167	0.0924^{*}	1,963	0.1720
	(0.72)	(0.82)	(1.77)		
	[0.70]	[0.80]			
Claude Opus 4.5	-0.0487^{***}	-0.0506^{***}	-0.0796^{*}	1,862	0.3573
	(-2.72)	(-2.82)	(-1.68)		
	[-2.53]	[-2.64]			
Cowork Plugins	-0.0846^{***}	-0.0855^{***}	-0.0492	1,832	0.4351
	(-3.11)	(-3.17)	(-0.96)		
	[-3.10]	[-3.14]			

Note: "Baseline" is the standard specification including pre-event market beta. "+ Rate Control" additionally controls for interest-rate sensitivity. $\hat{\beta}_{\text{Rate}}$ 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 each event, capturing duration exposure. Betas are winsorized at 1/99. All specifications include GICS sub-industry fixed effects, Eisfeldt et al. (forthcoming) labor exposure, pre-event market beta, and controls (log(ME), log(BM), OP, INV). Heteroskedasticity-robust t-statistics in parentheses; industry-clustered t-statistics in brackets (AI-PDR only). *, **, *** denote significance at 10%, 5%, 1%.

Table A13: Robustness: Alternative PDR Aggregation (Geometric Mean)

Event	Aggregation	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{ESZT}}$	N	R^2
ChatGPT Launch	Arithmetic Mean	0.0216	0.0654**	2,127	0.1519
		(1.17)	(2.56)		
	Geometric Mean	0.0152	0.0661**	2,127	0.1516
		(0.92)	(2.57)		
DeepSeek R1	Arithmetic Mean	0.0149	0.0811**	1,963	0.1680
		(0.72)	(2.42)		
	Geometric Mean	0.0091	0.0821**	1,963	0.1679
		(0.48)	(2.43)		
Claude Opus 4.5	Arithmetic Mean	−0.0487***	−0.0318	1,862	0.3544
		(−2.72)	(−1.04)		
	Geometric Mean	−0.0457***	−0.0298	1,862	0.3547
		(−2.78)	(−0.98)		
Cowork Plugins	Arithmetic Mean	−0.0846***	−0.0382	1,832	0.4345
		(−3.11)	(−1.04)		
	Geometric Mean	−0.0838***	−0.0333	1,832	0.4357
		(−3.55)	(−0.91)		

Note: "Arithmetic Mean" is the baseline PDR: equal-weight average of the three dimension scores. "Geometric Mean" uses $\text{PDR}_p = (D_p \times T_p \times P_p)^{1/3}$, which requires all three dimensions to be high for a product to score high. Both measures are importance-weighted, averaged across years, and normalized to $[0, 1]$. Market-adjusted CARs over $[-1, +3]$ with GICS sub-industry fixed effects, Eisfeldt et al. (forthcoming) (ESZT) labor exposure, pre-event market beta, and controls (log(ME), log(BM), OP, INV). Heteroskedasticity-robust t-statistics in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

Table A14: Robustness: Event-Year-Specific AI-PDR

Event	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{ESZT}}$	N	R^2
ChatGPT Launch	0.0029	0.0648***	2,088	0.1497
	(0.19)	(2.59)		
	[0.19]	[2.50]		
DeepSeek R1	0.0214	0.0793**	1,951	0.1666
	(1.22)	(2.37)		
	[1.27]	[2.30]		
Claude Opus 4.5	−0.0363**	−0.0336	1,850	0.3540
	(−2.47)	(−1.11)		
	[−2.34]	[−1.34]		
Cowork Plugins	−0.0666***	−0.0427	1,829	0.4341
	(−3.09)	(−1.17)		
	[−2.99]	[−1.08]		

Note: Same specification as Table 2 but using AI-PDR computed from the single 10-K filing year immediately before each event (FY 2021 for ChatGPT, FY 2024 for DeepSeek R1 and Claude Opus 4.5, FY 2025 for Cowork Plugins) instead of the all-year average. If a firm lacks a filing for the target year, the most recent earlier filing is used. Heteroskedasticity-robust t-statistics in parentheses; industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table A15: Alternative LLM Scorers: CAR Regressions with AI-PDR from Independent Models

	Qwen ($\rho = 0.94$)			Llama ($\rho = 0.93$)			GPT-3.5 ($\rho = 0.95$)		
Event	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{ESZT}}$	N	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{ESZT}}$	N	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{ESZT}}$	N
ChatGPT Launch	0.0267	0.0667***	2,127	0.0147	0.0685***	2,126	0.0197	0.0677***	2,127
	(1.18)	(2.63)		(0.75)	(2.70)		(0.70)	(2.67)	
	[1.05]	[2.62]		[0.70]	[2.70]		[0.71]	[2.72]	
DeepSeek R1	0.0009	0.0848**	1,963	0.0011	0.0841**	1,962	−0.0106	0.0865**	1,963
	(0.04)	(2.55)		(0.05)	(2.49)		(−0.34)	(2.57)	
	[0.04]	[2.50]		[0.06]	[2.43]		[−0.35]	[2.47]	
Claude Opus 4.5	−0.0588***	−0.0344	1,862	−0.0693***	−0.0321	1,861	−0.0732***	−0.0325	1,862
	(−2.58)	(−1.13)		(−3.44)	(−1.06)		(−2.74)	(−1.07)	
	[−2.22]	[−1.33]		[−2.91]	[−1.23]		[−2.60]	[−1.29]	
Cowork Plugins	−0.0722**	−0.0476	1,832	−0.0682**	−0.0483	1,831	−0.1030**	−0.0430	1,832
	(−2.24)	(−1.31)		(−2.35)	(−1.33)		(−2.57)	(−1.19)	
	[−1.94]	[−1.18]		[−2.31]	[−1.19]		[−2.60]	[−1.08]	

Note. Same specification as Table 2 but using AI-PDR scores constructed from three independent LLMs instead of GPT-4o mini. Qwen2.5-72B-Instruct, Llama-3.1-70B-Instruct, and GPT-3.5-Turbo-0125 (training cutoff September 2021). Correlations with the baseline GPT-4o mini PDR are shown in the column headers. R^2 not reported. Dependent variable is the market-adjusted CAR over $[-1, +3]$ trading days. GICS sub-industry fixed effects. Controls as in Table 2. Significance stars based on heteroskedasticity-robust t-statistics (parentheses); industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table A16: Event Study: Cross-Sectional Regressions of CARs on AI-PDR (Without ESZT Control)

Event	$\hat{\beta}_{\text{PDR}}$	N	R^2
ChatGPT Launch	0.0371^{**}	3,837	0.0893
	(2.27)		
	[2.23]		
DeepSeek R1	0.0174	3,772	0.0738
	(0.85)		
	[0.95]		
Claude Opus 4.5	-0.0314^{*}	3,604	0.1930
	(-1.81)		
	$[-1.56]$		
Cowork Plugins	-0.1126^{***}	3,572	0.3780
	(-5.59)		
	$[-3.93]$		

Note: Same specification as Table 2 but without the Eisfeldt et al. (forthcoming) labor exposure control. Dependent variable is the market-adjusted CAR over $[-1, +3]$ trading days. GICS sub-industry fixed effects. Controls: log(ME), log(BM), OP, INV, and pre-event market beta. Heteroskedasticity-robust t-statistics in parentheses; industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table A17: Event Study: FF5-Adjusted CARs on AI-PDR (GICS Sub-Industry FE)

Event	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{ESZT}}$	N	R^2
ChatGPT Launch	0.0196	0.0661**	2,127	0.1136
	(1.06)	(2.49)		
	[0.99]	[2.33]		
DeepSeek R1	0.0220	0.0732**	1,963	0.1292
	(1.06)	(2.10)		
	[1.11]	[1.93]		
Claude Opus 4.5	−0.0486***	−0.0475	1,862	0.1710
	(−2.69)	(−1.53)		
	[−2.85]	[−1.75]		
Cowork Plugins	−0.0565*	0.0126	1,832	0.2732
	(−1.94)	(0.32)		
	[−1.77]	[0.27]		

Note: Same specification as Table 2 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. GICS sub-industry fixed effects. Industries with fewer than 5 firms are excluded. Controls: log(ME), log(BM), OP, INV. Heteroskedasticity-robust t-statistics in parentheses; industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

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

	[−1, +1]	[−1, +3]	[−1, +10]	[−1, +20]	[−1, +30]
ChatGPT Launch	0.0136	0.0193	0.0159	0.0161	0.0236
	(1.03)	(1.04)	(0.49)	(0.36)	(0.43)
	[0.98]	[0.96]	[0.37]	[0.34]	[0.40]
DeepSeek R1	0.0115	0.0215	0.0551*	0.0991**	0.0761
	(0.75)	(1.04)	(1.91)	(2.27)	(1.35)
	[0.72]	[1.10]	[1.65]	[2.29]	[1.42]
Claude Opus 4.5	−0.0274*	−0.0487***	−0.0802**	−0.1039**	−0.0904*
	(−1.65)	(−2.71)	(−2.47)	(−2.23)	(−1.71)
	[−1.94]	[−2.88]	[−2.66]	[−2.58]	[−1.74]
Cowork Plugins	−0.0208	−0.0565*	−0.1777***	−0.1253**	−0.0879
	(−1.10)	(−1.94)	(−3.79)	(−2.14)	(−1.30)
	[−0.91]	[−1.77]	[−3.22]	[−2.08]	[−1.16]

Note: Same specification as Table 3 but using Fama–French five-factor adjusted CARs. Abnormal returns are computed relative to a five-factor model (MKTRF, SMB, HML, RMW, CMA) estimated over $[-252, -11]$ trading days before each event. GICS sub-industry fixed effects, Eisfeldt et al. (forthcoming) labor exposure, log(ME), log(BM), OP, and INV. Industries with fewer than 5 firms are excluded. Significance stars based on heteroskedasticity-robust t-statistics (parentheses); industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table A19: Robustness to Additional Controls: Cowork Plugins

Specification	$\hat{\beta}_{\text{PDR}}$	N	R^2
Baseline	-0.0846^{***}	1,832	0.4345
	(-3.11)		
	$[-3.10]$		
+ Firm characteristics	-0.0762^{**}	1,455	0.4439
	(-2.48)		
	$[-2.54]$		
+ Leverage, cash	-0.0781^{**}	1,444	0.4450
	(-2.54)		
	$[-2.60]$		
+ Momentum, volatility	-0.0789^{***}	1,444	0.4749
	(-2.71)		
	$[-2.49]$		
+ Firm age	-0.0789^{***}	1,444	0.4749
	(-2.71)		
	$[-2.49]$		
+ R&D intensity	-0.0622	736	0.4539
	(-1.55)		
	$[-1.37]$		

Note: Each row reports the AI-PDR coefficient from a cross-sectional regression of $[-1, +3]$ market-adjusted CARs around Cowork Plugins (January 30, 2026) on AI-PDR and progressively richer control sets. All specifications include GICS sub-industry fixed effects, Eisfeldt et al. (forthcoming) labor exposure, pre-event market beta, and baseline controls (log(ME), log(BM), OP, INV). "Firm characteristics" adds log(Assets), asset tangibility (PPE/Assets), intangibles ratio, and profit margin. Each subsequent row adds controls cumulatively. "Leverage, cash" adds long-term debt ratio and cash/assets. "Momentum, volatility" adds prior 12-month return and daily return volatility. "Firm age" adds log(Age) from first Compustat appearance. "R&D intensity" adds R&D/Assets; firms with missing R&D are dropped, reducing the sample to roughly 735 firms. Heteroskedasticity-robust t-statistics in parentheses; industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.

Table A20: Alternative Industry Fixed Effects: Cowork Plugins

Industry FE	$\hat{\beta}_{\text{PDR}}$	$\hat{\beta}_{\text{ESZT}}$	Industries	N	R^2
GICS Sub-Ind.	-0.0846^{***}	-0.0382	107	1,832	0.4345
	(-3.11)	(-1.04)			
	$[-3.10]$	$[-0.97]$			
SIC 3-digit	-0.1323^{***}	-0.0344	102	1,669	0.4159
	(-5.56)	(-0.92)			
	$[-3.76]$	$[-0.84]$			
FIC-50	-0.1341^{***}	-0.0228	40	1,925	0.3248
	(-6.34)	(-0.67)			
	$[-4.64]$	$[-0.47]$			

Note: Each row re-estimates the baseline Cowork Plugins regression (Table 2) using a different industry fixed-effects classification. Row 1 reproduces the baseline with GICS sub-industry FE (163 categories). Row 2 uses SIC 3-digit codes from CRSP. Row 3 uses Hoberg and Phillips (2016) FIC-50 text-based industry classifications (50 industries derived from 10-K product description similarity). Dependent variable is the market-adjusted CAR over $[-1, +3]$ trading days. Controls: log(ME), log(BM), OP, INV, pre-event market beta. Heteroskedasticity-robust t-statistics in parentheses; industry-clustered t-statistics in brackets. *, **, *** denote significance at 10%, 5%, 1%.