

Generative AI Product Displacement Risk

Brett Green¹ Renping Li² Dimitris Papanikolaou³

¹Washington University in St. Louis ²Tulane University ³Northwestern University & NBER

Brown Bag Seminar
Fall 2026

Contents


1. Motivation
2. The AI-PDR measure
3. Events and design
4. Repricing and effect heterogeneities
5. The SaaS divide
6. Beliefs and fundamentals
7. Generality and robustness

Motivation

Gen-AI can cut a firm's costs, or displace its products

- New technologies reallocate value, incumbents whose products become obsolete lose
- Gen-AI produces digitally deliverable products at near-zero marginal cost
- A large literature measures which **workers** Gen-AI can replace
- None of it measures whether the firm's **product** faces direct substitution

Products at risk mean profit loss



Research Question

- Which firms are most vulnerable to Gen-AI substituting for their *products*?
- Whether and when the equity market prices this risk?

Answer. Displacement is priced at product entry

1. High-AI-PDR firms reprice at Cowork Plugins, -4.7 p.p. per SD
2. Competitive shields attenuate the repricing, supply side most
3. Analyst forecasts, investor beliefs, and fundamentals move the same way
4. Over longer horizons, SaaS and non-SaaS high-AI-PDR firms diverge

The AI-PDR measure

Scoring procedure

- SEC EDGAR 10-K Item 1 business descriptions, filing years 2021 to 2026
- Stage 1, extract distinct product lines from each filing
- Stage 2, score each product 0 to 10 on digital deliverability
- 139,114 products, about 8,600 firms

Walt Disney Co, filing year 2025

Product line	Digital
Theme Parks and Resorts	0
Cruises and Vacation Travel	0
Consumer Products Licensing and Retail	4
ESPN Sports Media and Streaming	8
Disney+ and Hulu Streaming Services	10

One AI-agnostic question scores digital deliverability

“Could the core value of this product/service be delivered as a purely digital output (information, software, or digital content), without physical goods, in-person service, or transfer of financial capital?”

- The prompt never mentions AI, limiting look-ahead from the scorer
- Firm-year level **AI-PDR** is the importance-weighted average across firm products

How AI-PDR ranks firms

Vertex Pharmaceuticals Inc

Cystic Fibrosis Medicines

Union Pacific Corp

Premium Intermodal and Automotive Rail Transportation

Chevron Corp

Refined Petroleum Fuels and Other Petroleum Products

Autodesk, Inc

Software Subscriptions and Enterprise Access Agreements

Adobe Inc

Firefly Generative AI Solutions

Roblox Corp


Robux Virtual Currency and Roblox Premium Subscriptions

Low AI-PDR, 0.0

High AI-PDR, 8.8



Digital sectors score high, physical sectors low



AI-PDR is not labor exposure

Variable	Pearson r	$\hat{\beta}$
ESZT Labor Exposure	0.377	0.0076*** (4.55)
log(ME)	0.054	0.0268 (1.02)
R&D/Assets	0.050	0.0027 (0.90)
Intangibles/Assets	0.315	0.0276*** (6.95)
Gross Margin	0.244	0.0277*** (3.54)
Revenue Growth	0.059	0.0075** (2.37)

Events and design

Main-sample events

Generative AI model releases

- ChatGPT (Nov 2022)
- DeepSeek R1 (Jan 2025)
- Claude Opus 4.5 (Nov 2025)
- Ship raw models, the path to product competition stays uncertain

Cowork Plugins (Jan 30, 2026)

- 11 agentic plugins spanning enterprise functions, \$20 per month
- Competing with specific incumbent products

Cowork Plugins, 11 agentic plugins at \$20 per month


- **Jan 12**, Claude Cowork desktop agent, general purpose, not domain specific
- **Jan 30**, plugin support added, 11 plugins open-sourced on GitHub
- Each bundles domain knowledge and tool connectors
→ a department-specific workflow automation tool
- \$20 per user-month on Claude Pro, against \$165 Salesforce Enterprise or \$80 Adobe Creative Cloud seats

Ten role-specific plugins

Productivity	Sales
Customer support	Product management
Marketing	Legal
Finance	Data
Enterprise search	Bio research

Breadth and depth of product overlap is unmatched

- To what extent Cowork Plugins functionally overlap with Compustat firm products?



Strong stock market response

- By Feb 3 to 4,
 - **SaaS firms**, Salesforce, ServiceNow, Adobe each down roughly 7%
 - **Financial and legal data providers**, Thomson Reuters down 16%, RELX 14%, LegalZoom 20%
 - **IT consulting firms**, Accenture down 7.5%, Infosys, TCS, Wipro comparable
 - **Online education firms**, Duolingo down 16.2%, Chegg 8.9%
 - **Video-game publishers**, Take-Two down 16.1%, Electronic Arts 3.5%
 - **Freelancer marketplaces**, Upwork down 18.9%, Fiverr 7.2%
- Media dubbed it a “SaaS apocalypse”

We regress event CARs on AI-PDR

$$\text{CAR}_i^{[-1,+10]} = \alpha_{g(i)} + \beta \text{AI-PDR}_i + \gamma' X_i + \varepsilon_i$$

- CARs CAPM-adjusted, beta over $[-252, -11]$, no estimated intercept
- X_i , log ME, Tobin's Q, ROA, asset growth
- $\alpha_{g(i)}$, GICS 6-digit FE, t clustered the same level
- Sample restrictions, market cap and Compustat assets above \$100M
- AI-PDR vintage, the last filing year before each event
- All variables winsorized 1/99

► score stability

Repricing and effect heterogeneities

Sharp repricing at Cowork Plugins


Event	AI-PDR	<i>N</i>
ChatGPT Launch	0.0068* (1.82)	2,832
DeepSeek R1	0.0075** (2.01)	2,887
Claude Opus 4.5	-0.0054 (-1.58)	2,763
Cowork Plugins	-0.0466*** (-8.34)	2,845
FE	GICS-6	

- CAPM-adjusted $CAR_{[-1,+10]}$ on standardized AI-PDR
- GICS 6-digit FE, GICS 6-digit clustered t

Findings.

- Model-release coefficients under a sixth of the Cowork effect
- Top-quintile AI-PDR firms underperform the bottom quintile by about 12 p.p.


The repricing persists



- Regress $[-1, +t]$ CARs on AI-PDR
- 1% significant at every one of them, attenuates only at +160

► pre-trend

No placebo date in four years comes close




- Baseline specification re-estimated at 169 placebo dates, every fifth trading day, 2023 to 2026

Findings.

- Cowork sits about five placebo SDs below the placebo mean

No comparable repricing at the three model releases



Functionally overlapping firms bear the largest losses

	(1)	(2)
AI-PDR	−0.0359*** (−5.41)	−0.0351*** (−5.30)
Cowork Overlap (std)	−0.0179*** (−5.30)	
1(Max Overlap ≥ 3)		−0.0121 (−1.43)
AI-PDR $\times$ 1(Max Overlap ≥ 3)		−0.0175*** (−2.86)
FE	GICS-6	GICS-6
N	2,766	2,764

Findings.

- One SD of plugin overlap costs a further 1.8 p.p.
- The AI-PDR slope is about half again as steep for overlapping firms

Four shields measure who can resist displacement

Supply side

- **Legal/Regulatory**, licenses, approvals, certifications required to sell
- **Production Moats**, could a well-funded entrant build the production assets within years
- Same 0-10 scoring, composite Shield is the simple average of the four

Demand side

- **Customer Trust**, would customers buy from an unknown, unbranded provider
- **Embeddedness**, cost and disruption of moving to an equivalent alternative

Shields attenuate the repricing, supply side most


	Composite	Supply Side	Demand Side
AI-PDR	−0.0446*** (−8.09)	−0.0419*** (−6.17)	−0.0491*** (−9.32)
Shield	0.0014 (0.31)	−0.0067 (−1.04)	0.0055 (1.44)
AI-PDR × Shield	0.0121*** (3.29)	0.0135*** (2.92)	0.0077** (2.14)
FE	GICS-6	GICS-6	GICS-6
N	2,845	2,845	2,845

Findings.

- A one-SD composite Shield cuts the AI-PDR effect by about 27 percent
- Supply side 32, demand side 16

The SaaS divide

SaaS recovers, non-SaaS stays repriced



- At $[-1, +160]$, SaaS -0.001 ($t = -0.03$), non-SaaS -0.047 ($t = -2.31$)
- Salesforce -15.4% at $+10$ to $+2.2\%$ at $+160$, ServiceNow -14.8% to $+7.2\%$

[Details](#)

Why the divide, SaaS benefits from agentic adoption

SaaS firms

- Sell the platform layer an agent runs on
- Benefit twice, own agentic products complement existing applications, other firms' agents run on their platforms
- Agentforce near \$800M ARR, ServiceNow AI contracts past \$1B




Non-SaaS firms

- Sell curated data, ratings, reports, instruction, and brokered expertise
- What they sell is the output Gen-AI also produces, not the platform layer an agent runs on
- Do not benefit and stay repriced



Agentic AI launchers outperform by sample end



- Launches of firms' own agentic AI products from LSEG news archive
- 310 by 183 firms in the eight months after Cowork, 244 in the three years before

Findings.

- Launchers converge to peers by +120 and finish +7.7 p.p. above ($t = 2.55$) by $[-1, +160]$

Launch announcements pay off only after Cowork

Window	Before Cowork	After Cowork
$[-1, +1]$	-0.0008 (-0.17)	0.0136** (2.01)
$[-1, +10]$	-0.0110 (-0.93)	0.0331** (2.07)
N	124	102

- Mean CAR around each firm's first agentic AI launch

Findings.

- After Cowork, the first launch earns +3.3 p.p. over $[-1, +10]$ ($t = 2.07$)

Beliefs and fundamentals

Analysts mark down forecasts

	(1) EPS revision	(2) Revenue revision	(3) LTG change	(4) Target Price revision
AI-PDR	-0.029*** (-2.66)	-0.009* (-1.85)	-1.541** (-2.06)	-0.105*** (-5.32)
FE	GICS-2	GICS-2	GICS-2	GICS-2
<i>N</i>	1,960	2,533	1,028	2,655

- Pre = last IBES consensus before Jan 30
- Post = last consensus through Jun 30, 2026
- EPS, revenue, and target price, revision = (post – pre)/pre
- LTG change = post – pre in percentage points

Markdowns extend to later fiscal years

	(1) EPS revision	(2) Revenue revision
December 2026	−0.029*** (−2.66)	−0.009* (−1.85)
<i>N</i>	1,960	2,533
December 2027	−0.030*** (−3.04)	−0.022*** (−5.22)
<i>N</i>	1,943	2,416
December 2028	−0.039*** (−2.92)	−0.020*** (−3.54)
<i>N</i>	1,258	1,637
FE	GICS-2	GICS-2

Findings.

- EPS markdowns at every horizon, growing to −0.039 by December 2028
- The revenue markdown strengthens beyond 2026, −0.022 and −0.020 at 1%

Example endorsement messages

- 7.5M StockTwits messages, four weeks each side of each event
- Endorsements of Gen-AI displacing the firm's products

Examples.

Adobe (\$ADBE)

"\$ADBE What no one want to talk about: AI and Vibe coding is killing Blg Software"


Duolingo (\$DUOL)

"\$DUOL I vibe coded a Duolingo replacement overnight"

Microsoft (\$MSFT)

"\$MSFT Not low enough. Claude AI replaces most MSFT apps. For FREE."

Displacement talk jumps only at Cowork Plugins



Findings.

- Number of endorsing messages jumps 3.3-fold at Cowork Plugins

The belief jump loads on AI-PDR only at Cowork

- Change in firm-level endorsement on AI-PDR, GICS 2-digit FE

Event	<i>N</i>	Level (p.p.)	Log count	Calibrated log
ChatGPT Launch	2,434	0.014 (1.04)	0.004 (1.52)	0.005 (1.50)
DeepSeek R1	1,384	0.002 (0.70)	0.005 (1.42)	0.008 (1.54)
Claude Opus 4.5	1,484	−0.033 (−1.40)	−0.015 (−1.02)	−0.020 (−1.04)
Cowork Plugins	1,547	0.085*** (3.64)	0.091*** (4.69)	0.126*** (4.87)
FE		GICS-2	GICS-2	GICS-2

Fundamentals deteriorate


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

- Δ Revenue growth, change in the revenue growth rate from one year earlier
- Δ Gross margin, change in gross margin from one year earlier

Findings (2026 Q2).

- Revenue growth -3.7 p.p. and gross margin -0.9 p.p. per one-SD AI-PDR
- Both deteriorate on both sides of the SaaS divide

Fundamentals deteriorate




- 2026 Q2 is the first cohort in which both deteriorate significantly

Generality and robustness

Breadth and depth of product overlap is unmatched

- To what extent Cowork Plugins functionally overlap with Compustat firm products?



Other AI product releases


	Business Oriented (1)	Consumer Oriented (2)
$[-1, +1]$	-0.0018*** (-3.29)	0.0000 (0.12)
$[-1, +3]$	-0.0023*** (-3.31)	-0.0002 (-0.46)
$[-1, +5]$	-0.0026*** (-3.04)	-0.0011 (-1.35)
$[-1, +10]$	-0.0014 (-0.99)	-0.0015 (-1.30)
$[-1, +20]$	-0.0032 (-1.63)	-0.0024* (-1.70)
FE $\times$ event	GICS-6	GICS-6

- Business-oriented ones are M365 Copilot GA, Claude for Financial Services, Claude Managed Agents, Claude 3.5 computer use, and Claude Code

Findings.

- Business-oriented group, -0.18 p.p. per SD over $[-1, +1]$ ($t = -3.29$), gone within twenty trading days
- Cowork Plugins, nearly six times stronger per SD of overlap, repriced firms for months

No comparable effects across all LLM model releases



- 25 flagship LLM releases late 2022 to July 2026
- CAR $[-1, +10]$ on AI-PDR, GICS 6-digit FE, clustered 95% CIs

AI-PDR is not a software-exposure proxy

- Horse race at Cowork Plugins

	Software Beta (1)	Beta + AI-PDR (2)	Keyword Share (3)	Keyword + AI-PDR (4)
AI-PDR		−0.0415*** (−7.50)		−0.0412*** (−7.08)
Software Beta (IGV)	−0.0288*** (−6.39)	−0.0242*** (−5.41)		
Digital Keyword Share			−0.0280*** (−3.40)	−0.0099 (−1.29)
FE	GICS-6	GICS-6	GICS-6	GICS-6
<i>N</i>	2,845	2,845	2,845	2,845

Robustness battery

- Characteristics-matched abnormal returns, -0.0441

▶ matched table

▶ matched figure

- Excluding in-window earnings reporters, -0.0435

▶ ex-earners table

- Interest-rate-sensitivity control, -0.0466

▶ rate-control table

- Fama-French five-factor CARs, -0.0443

▶ FF5 table

- ESZT labor exposure as a joint regressor

▶ ESZT table

- Within-firm average score across filing years 2021 to 2026, -0.0530

▶ within-firm table

- GPT-3.5 scorer with a September 2021 training cutoff, -0.0356 ($t = -6.46$)

▶ GPT-3.5 table

- Added controls, alternative FE granularities

▶ controls table

▶ FE table

Contribution

- **A product-side measure of Gen-AI exposure.** AI-PDR scores what a firm sells, complementing labor-exposure indices that score whom it employs
- **When displacement is priced.** At broad entry into incumbent product space, not at model releases
- **Beyond prices.** Analyst forecasts, investor beliefs, and realized fundamentals move the same way
- **Agentic AI adoption.** First to document agentic AI adoption and its positive effect on firm value


Conclusions

1. Displacement is priced at product entry, not at model releases, -4.7 p.p. per SD
2. Shields attenuate it, supply side most
3. Forecasts, beliefs, and fundamentals move the same way
4. SaaS recovers, non-SaaS stays repriced

Open questions.

- Can high-AI-PDR non-SaaS firms race, or must they retreat?
- Who displaces, a few large Gen-AI providers or a fringe of Gen-AI-native startups?
- How far does the displacement frontier extend beyond digital products?

Scores are stable across filing years




Findings.

- Adjacent filing years correlate 0.94 to 0.95
- 2021 against 2026, about 0.89
- AI-PDR captures durable product characteristics, not filing-year noise

[◀ back](#)

No pre-trend, the slip begins at the Jan 12 preview






- Backward windows $[d, -1]$ plotted as negated slopes

Findings.

- Flat through late 2025, high-AI-PDR firms slip only about twelve days before launch, at the Cowork desktop preview
- The preview decline is small next to the launch drop, Jan 30 is the event date

The split-slope regression behind the SaaS divide

$$CAR_i^{[-1,+d]} = \alpha_{g(i)} + \delta SaaS_i + \beta_S SaaS_i AI-PDR_i + \beta_N (1 - SaaS_i) AI-PDR_i + \gamma' X_i + \varepsilon_i$$



- AI-PDR in full-sample SD units, GICS 6-digit FE, main-spec controls, clustered 95% bands

Matched abnormal returns give the same contrast

Event	CAR	Char.-matched	<i>N</i>
ChatGPT Launch	0.0068* (1.82)	0.0063* (1.70)	2,832
DeepSeek R1	0.0075** (2.01)	0.0075** (2.20)	2,887
Claude Opus 4.5	-0.0054 (-1.58)	-0.0042 (-0.99)	2,763
Cowork Plugins	-0.0466*** (-8.34)	-0.0441*** (-7.55)	2,845
FE	GICS-6	GICS-6	


- $[-1, +10]$ return net of the size $\times$ momentum cell
- 25 value-weighted cells formed at day -11

Findings.

- Cowork Plugins -0.0441 ($t = -7.55$) against -0.0466 with CAPM CARs
- Model releases at most 0.0075 in absolute value

◀ back

The SaaS divide holds under the matched benchmark



- Characteristics-matched abnormal returns within each group, GICS 6-digit FE
- At $[-1, +160]$, SaaS 0.005 ($t = 0.16$), non-SaaS -0.043 ($t = -1.94$)
- Non-SaaS significant at 5% through $[-1, +120]$, at 10% at +160

[◀ back](#)

Excluding in-window earnings reporters

	GICS 2-digit	GICS 6-digit
Full Sample	-0.0705^{***} (-13.68) $N = 2,865$	-0.0466^{***} (-8.34) $N = 2,845$
Excluding Earners	-0.0646^{***} (-11.42) $N = 1,979$	-0.0435^{***} (-6.96) $N = 1,951$

- Firms with an earnings announcement inside the Cowork Plugins window dropped, N falls to 1,951

Findings.

- -0.0435 ($t = -6.96$) at GICS 6-digit, against -0.0466 in the full sample

◀ back

Interest-rate sensitivity does not explain it

	GICS 2-digit	GICS 6-digit
Baseline	-0.0705*** (-13.68) $N = 2,865$	-0.0466*** (-8.34) $N = 2,845$
+ Rate Control	-0.0705*** (-13.63) $N = 2,865$	-0.0466*** (-8.28) $N = 2,845$

- Adds each firm's pre-event beta on daily 10-year Treasury yield changes as a control at Cowork Plugins

Findings.

- AI-PDR slope unchanged, -0.0466 ($t = -8.28$)

◀ back

FF5-adjusted CARs give the same contrast

	GICS 2-digit	GICS 6-digit
ChatGPT Launch	0.0023 (0.76) $N = 2,850$	0.0037 (0.96) $N = 2,832$
DeepSeek R1	0.0113*** (3.67) $N = 2,907$	0.0065 (1.27) $N = 2,887$
Claude Opus 4.5	-0.0147*** (-4.77) $N = 2,781$	-0.0086** (-2.39) $N = 2,763$
Cowork Plugins	-0.0648*** (-10.04) $N = 2,865$	-0.0443*** (-5.92) $N = 2,845$

- Abnormal returns from the Fama-French five-factor model, no intercept, in place of CAPM

Findings.

- Cowork Plugins -0.0443 ($t = -5.92$) at GICS 6-digit
- Model releases at most -0.0147

[< back](#)

AI-PDR survives the ESZT labor-exposure control

Event	AI-PDR	ESZT	<i>N</i>
ChatGPT Launch	−0.0068 (−0.99)	0.0191*** (3.32)	2,218
DeepSeek R1	0.0017 (0.26)	−0.0015 (−0.26)	2,049
Claude Opus 4.5	−0.0007 (−0.11)	−0.0105* (−1.85)	1,942
Cowork Plugins	−0.0451*** (−2.72)	−0.0125* (−1.75)	1,924
FE	NAICS-3	NAICS-3	

- ESZT-style specification, market-cap-weighted with NAICS 3-digit FE

Findings.

- ESZT carries the positive ChatGPT response, AI-PDR carries the Cowork repricing, distinct cross-sections at every event

◀ back

Averaging scores across filing years changes nothing

	GICS 2-digit	GICS 6-digit
ChatGPT Launch	0.0077** (2.24) $N = 3,220$	0.0089 (1.54) $N = 3,201$
DeepSeek R1	0.0188*** (5.73) $N = 2,963$	0.0118** (2.34) $N = 2,942$
Claude Opus 4.5	-0.0135*** (-4.36) $N = 2,856$	-0.0076** (-2.00) $N = 2,837$
Cowork Plugins	-0.0735*** (-15.92) $N = 2,901$	-0.0530*** (-9.58) $N = 2,881$

- AI-PDR averaged across each firm's 2021 to 2026 filings in place of the last pre-event filing year

Findings.

- Cowork Plugins -0.0530 ($t = -9.58$) at GICS 6-digit
- Model releases stay small

[< back](#)

A September 2021 scorer finds the same effect

Event	AI-PDR	<i>N</i>
ChatGPT Launch	0.0029 (0.79)	2,832
DeepSeek R1	0.0074** (2.08)	2,887
Claude Opus 4.5	0.0021 (0.53)	2,763
Cowork Plugins	-0.0356*** (-6.46)	2,845
FE	GICS-6	

- AI-PDR rebuilt with GPT-3.5-Turbo, training cutoff September 2021, before every event

Findings.

- Correlation 0.91 with the baseline score
- Alleviates look-ahead bias from the LLM scorer

◀ back

Richer controls leave the slope intact

Specification	AI-PDR	<i>N</i>
Baseline	−0.0466*** (−8.34)	2,845
+ Firm characteristics	−0.0423*** (−7.33)	2,659
+ Leverage, cash	−0.0415*** (−7.30)	2,650
+ Momentum, volatility	−0.0381*** (−6.39)	2,650
+ Firm age	−0.0381*** (−6.38)	2,650
+ R&D intensity	−0.0477*** (−4.19)	1,347
FE	GICS-6	

- Cowork Plugins $[-1, +10]$, GICS 6-digit FE, controls added cumulatively

Findings.

- Between -0.0381 and -0.0466 , t from -6.38 to -8.34
- R&D subsample -0.0477 on 1,347 firms

◀ back

The contrast survives every industry-FE granularity

	No FE	GLCS 2-digit	GLCS 4-digit	GLCS 6-digit	GLCS 8-digit
ChatGPT Launch	0.0040 (1.07)	0.0057* (1.92)	0.0062** (2.04)	0.0068* (1.82)	0.0072* (1.88)
DeepSeek R1	0.0141*** (4.86)	0.0164*** (4.94)	0.0098** (2.44)	0.0075** (2.01)	0.0096** (2.37)
Claude Opus 4.5	-0.0013 (-0.53)	-0.0116*** (-3.76)	-0.0074*** (-2.60)	-0.0054 (-1.58)	-0.0042 (-1.10)
Cowork Plugins	-0.0557*** (-11.28)	-0.0705*** (-13.68)	-0.0556*** (-9.33)	-0.0466*** (-8.34)	-0.0421*** (-7.20)

[◀ back](#)

AI input exposure is the winner side of Gen-AI

- AI Input Intensity, whether scaling Gen-AI requires more of the firm's product, 23-category taxonomy
- Value-weighted buy-and-hold returns, Jan 2022 through Sep 2026
 - Top-half AI-input firms 256.7%
 - Bottom-half 103.8%, remaining firms 41.9%, market 60.7%
- DeepSeek R1 hit AI-input firms (-0.053 at $[-1, +10]$), Cowork lifted them ($+0.060$ by $[-1, +10]$)
- The displacement and input channels are distinct and separable

We add the product dimension to three literatures

- **Gen-AI and labor**, occupation exposure indices (Felten, Webb, Eloundou, Eisfeldt et al.)
 - ESZT capture replaceable *workers*, we capture replaceable *products*
- **Gen-AI and firm value**, adoption and growth (Babina et al.), belief jumps (Andrews)
 - We identify which firms reprice, and when
- **Technological displacement**, creative destruction priced (Kogan et al.), patent-based threat (Kakhbod et al.)
 - Method builds on Hoberg-Phillips 10-K product text