

When AI-Washing Meets Initial Public Offerings

RUI DAI (UNIVERSITY OF SOUTH FLORIDA), GANGXING (GARY) GUO AND BOHUI ZHANG (CUHK-SHENZHEN)

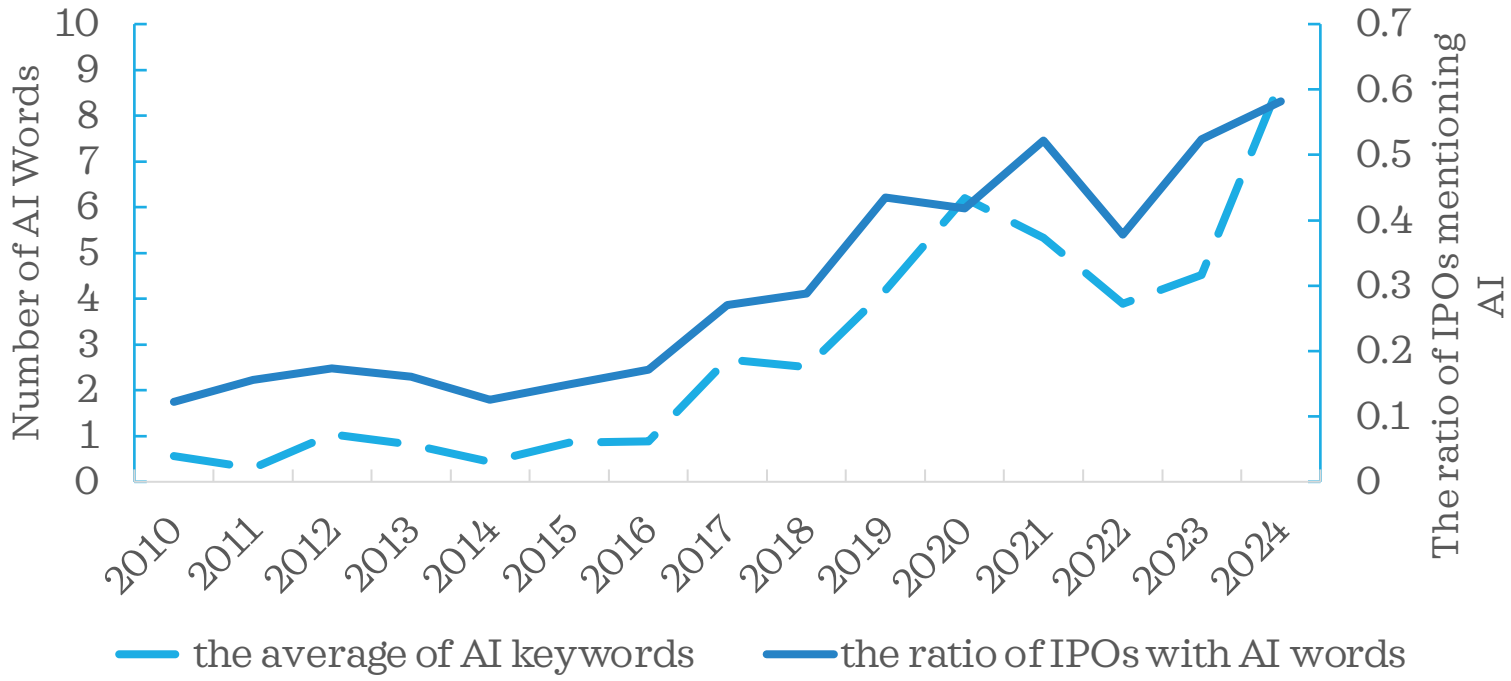
PRESENTED BY RUI DAI | 2026 FMA ANNUAL MEETING | OCTOBER 15, 2026

Background: Emerging AI disclosure in IPO prospectuses

AI is increasingly referenced by firms in the U.S. IPO market

- Mentions of AI in S-1 prospectuses have risen sharply over time

Difficulties in valuing AI create scope for potential overstatement



Background: SEC enforcement on AI washing

AI-washing (as defined by the SEC)

Deceptive practices in which firms embellish or exaggerate their use of AI and its capabilities to attract investors by creating a facade of innovation and reliability

SEC Charges Two Investment Advisers With Making False, Misleading AI Claims

By AARON MCDADE Published March 18, 2024 12:52 PM EDT



Credit: Graeme Sloan / Bloomberg via Getty Images



Research questions

Why IPOs? AI inputs are hard to observe, and issuers have no public track record

How prevalent is AI-washing in the IPO market?

How can we distinguish AI talk from observable AI investment?

- S-1 disclosures vs. pre-IPO AI employees and AI patent filings

How do investors price this disclosure–input mismatch?

- First-day returns and long-run stock performance

What explains the mispricing, and what curbs AI-washing?

Preview of findings

Prevalence: 107 of 1,671 IPOs (6.4%) are AI-washing; 14.3% of IPOs in 2023

Pricing: AI-washing IPOs earn 7.2 pp higher first-day returns, then underperform

- −8.8 pp (2 years) and −13.9 pp (3 years) vs. non-AI IPOs
- AI-investment IPOs: +10.2 pp and +22.0 pp; gap of 19.0 and 36.0 pp

Validation: AI talk fades after year 1; AI job postings (Lightcast) drop after lockup

Channels: valuation premium, verification difficulty, litigation risk

Discipline: less AI-washing under SEC scrutiny; AI-savvy intermediaries do not help

Dataset

1,671 U.S. IPOs from 2010 to 2024

- Exclude offer prices below $5, ADRs, REITs, closed-end funds, and issuers without first-day CRSP trading records (Lowry et al., 2017)

Main datasets

- SDC Platinum and Jay Ritter: IPO characteristics and controls
- CRSP/Compustat: stock returns, financials, and industry classifications
- SEC EDGAR: S-1 filings, 10-K filings, and other disclosures
- Revelio Labs: LinkedIn employee profiles (April 2025 snapshot)
- Lightcast (formerly Burning Glass): AI job postings
- USPTO AIPD: AI patent documents through 2023

Measuring AI talk and AI walk

AI talk: AI keyword mentions in the S-1 registration statement
- 97-term dictionary from prior studies; 57 terms appear in S-1s

AI walk 1 – human capital: AI-skilled employees at the listing date
- Revelio Labs (LinkedIn); Babina et al. (2024) skills; STEM degree required

AI walk 2 – patents: AI patent filings in the three years before listing
- USPTO AI Patent Dataset (0.93 threshold); KPSS links plus name matching

All inputs are dated no later than the listing date

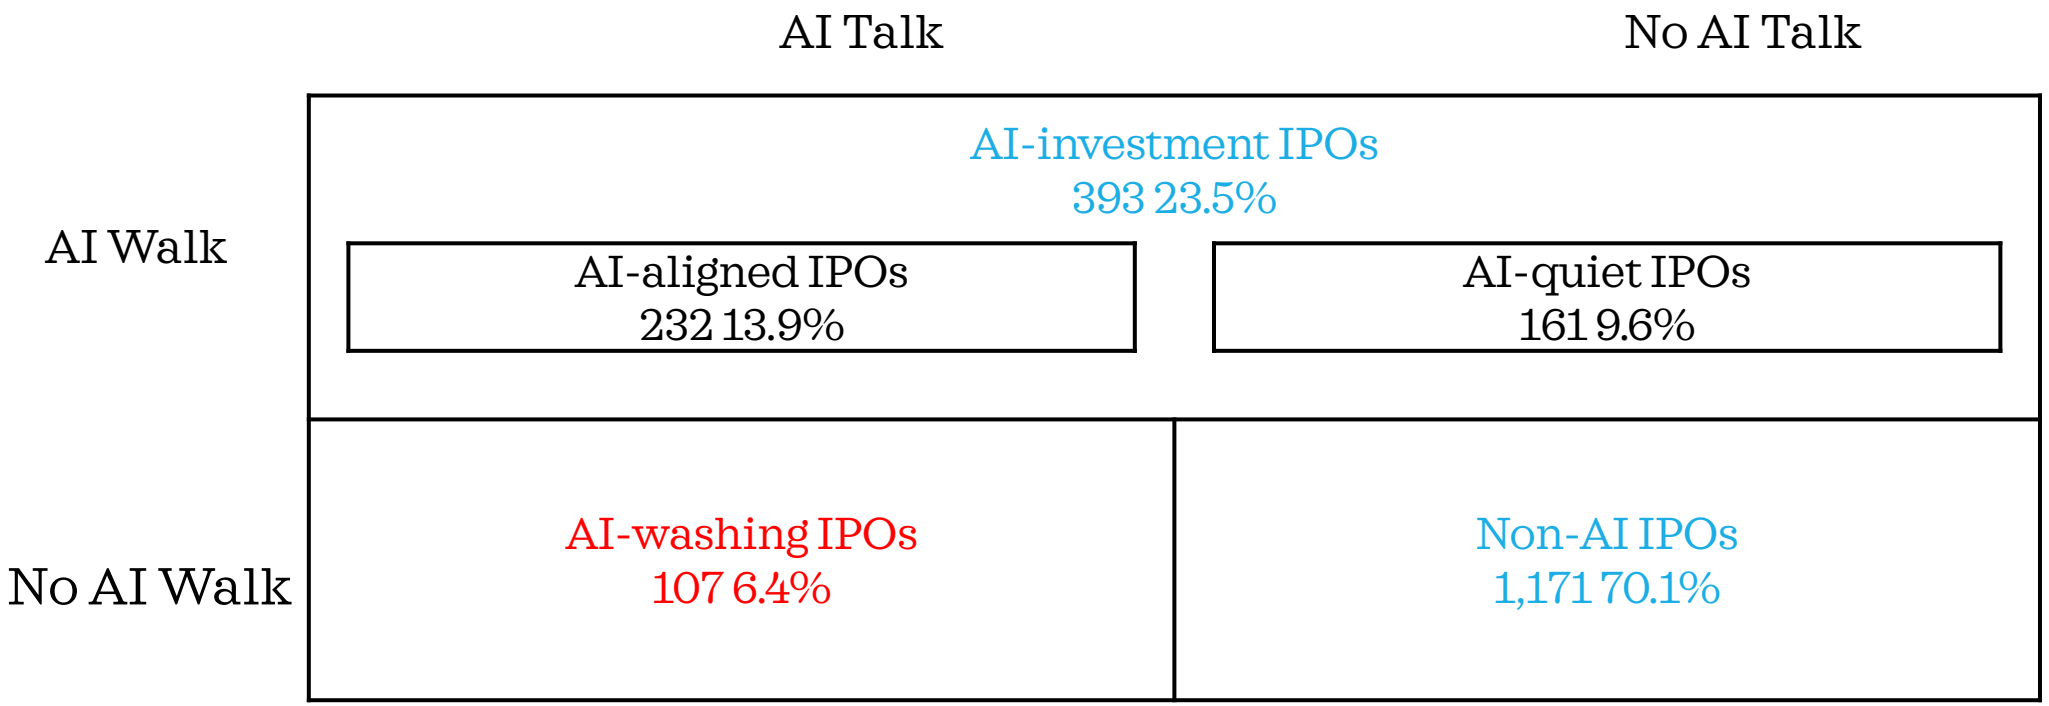
AI disclosure and AI inputs

Poisson regressions relate AI disclosures to pre-IPO AI employees and patents

$$E[AIWord \mid X] = \exp(\beta_1 AIEmployee + \beta_2 AIPatent + \text{Industry FE} + \text{Year FE})$$

Dep. Variable=			***AIWord***			
Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
AIEmployee	0.093***	0.046***			**0.077*** **	**0.042*** **
	(6.24)	(5.20)			**(7.98)**	**(6.34)**
AIPatent			0.031***	0.012***	**0.012*** **	**0.005**
			(4.21)	(2.94)	**(3.64)**	**(1.49)**
Fixed Effects	No	Ind & Year	No	Ind & Year	No	Ind & Year
Observations	1,671	1,660	1,671	1,660	1,671	1,660
Pseudo R2	8.1%	32.1%	4.6%	31.0%	8.6%	32.2%

Classifying IPOs: AI talk vs. AI walk

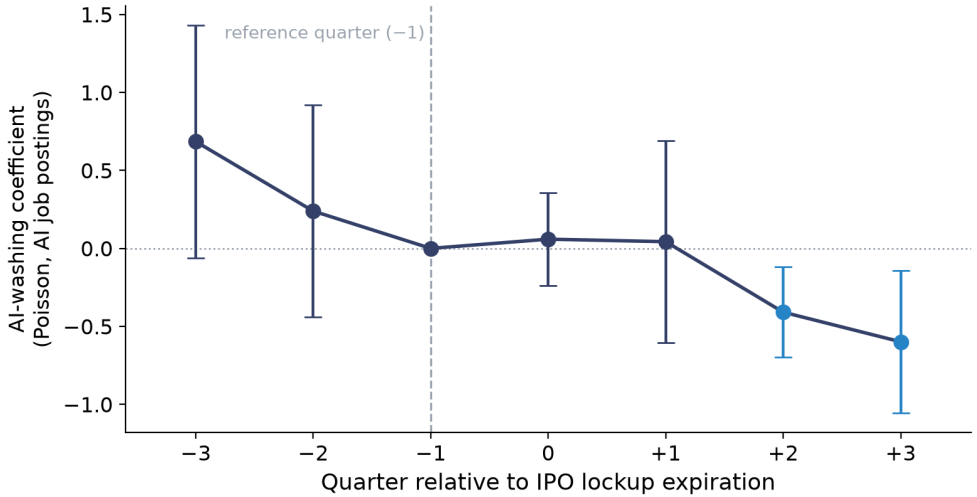


AI-washing: AI talk above the FF12 industry's historical 95th percentile (2007 to t−1), with zero AI employees and zero AI patents. **AI-investment:** at least one AI employee or AI patent.

AI washing or AI failures?

We estimate a dynamic regression using an event window of [−3, +3] quarters around the expiration day, focusing on the subsample of AI-washing and AI-investment IPOs:

$$AIJobPost_{i,t} = \sum_{\substack{k=-3 \\ k\neq -1}}^{3} \beta_k \left(AIWashing_i \times 1(rel_quarter = k) \right) + AIWashing_i + CONTROLS_i + FE_{i,t} + \varepsilon_{i,t}$$



AI-washing firms post fewer AI jobs (Lightcast) in quarters +2 and +3.

A firm genuinely building AI has no reason to stop hiring just when insiders can start selling.

Baseline test

OLS regressions of long-run *BHAR* on *AIWashing* and *AIInvestment*
Benchmark: Fama–French 5×5 size and book-to-market portfolios

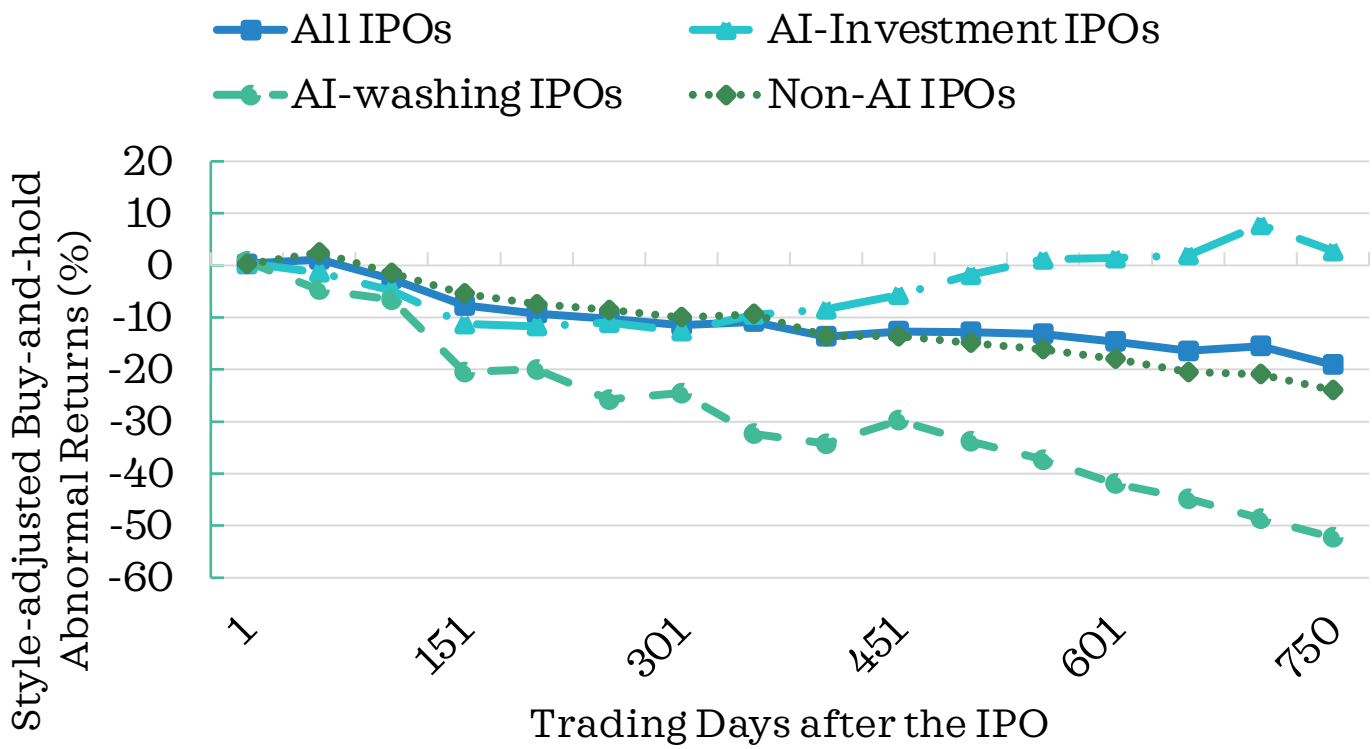
$$BHAR_{i,t} = \beta_1 * AIWashing_i + \beta_2 * AIInvestment_i + CONTROLS_i + FE_{i,t} + \varepsilon_{i,t}$$

Dep. Variable=	***BHAR[0,1yr]***	***BHAR[0,2yr]***	***BHAR[0,3yr]***	***InitialReturn***
Variable	(1)	(2)	(3)	(4)
AIWashing	-0.077	-0.088**	-0.139***	0.072**
	(-1.09)	(-2.35)	(-2.73)	(2.04)
AIInvestment	-0.022	0.102**	0.220***	0.014
	(-0.74)	(2.31)	(4.85)	(0.59)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Ind & Year	Ind & Year	Ind & Year	Ind & Year
Observations	1,581	1,538	1,498	1,652
Adjusted R2	14.5%	10.0%	7.6%	22.7%

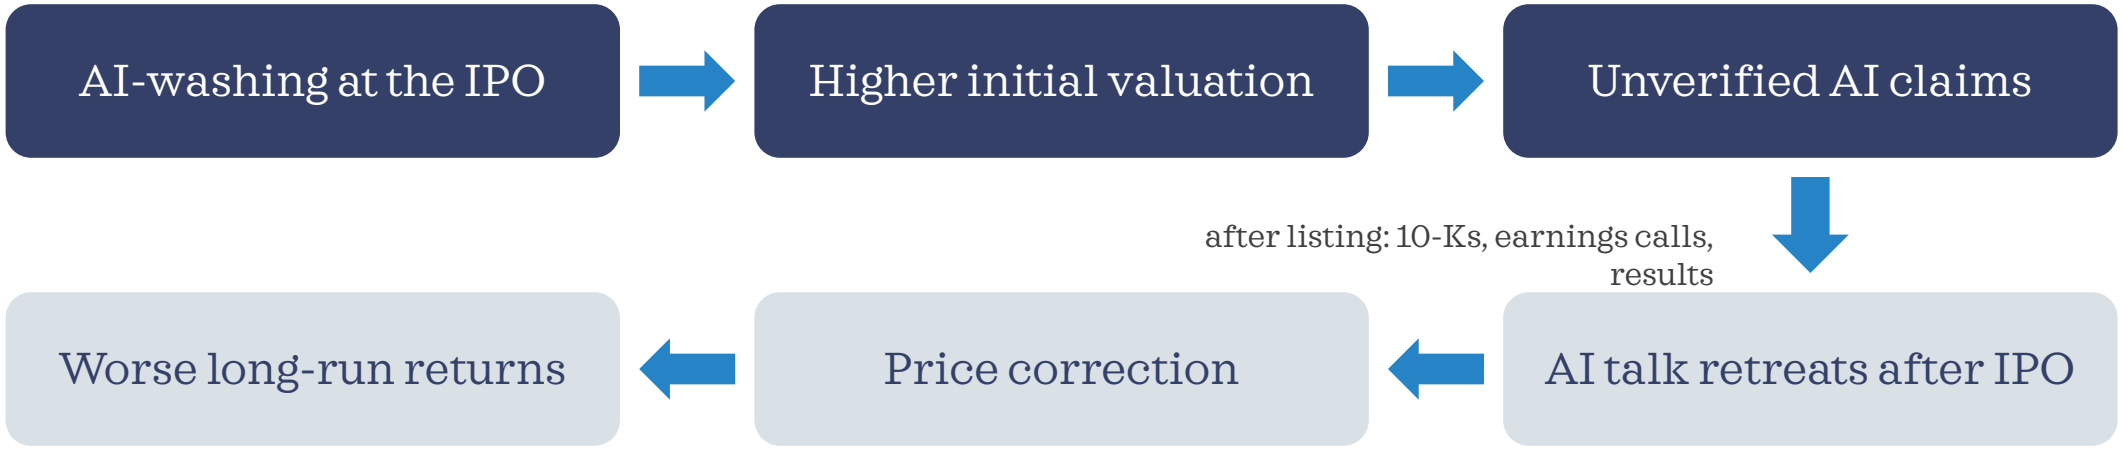
Robust to calendar-time alphas (FF5: −1.49%/month over 24 months), propensity-score matching, entropy balancing, and excluding firms with third-party AI tools, suppliers or partners.

Post-IPO stock performance

Style-adjusted buy-and-hold abnormal returns for 1,671 U.S. IPOs, 2010–2024



Mechanisms: why do AI-washing IPOs underperform?



After listing, the AI story fades

- **Earnings calls:** less AI talk after year 1 (Fig. 3A)
- **First 10-K:** fewer AI words than the S-1 (IA9)
- **Year-2 calls:** fewer AI words than the S-1 (IA9)

Mechanisms: the evidence

1 Valuation premium (Table 6)

Measure: AIVal, how much the market rewards AI patents at listed peers in the same industry (3 years before the IPO)

Test: AI-washing × AIVal

P/S at the offer: +16.1*** (t = 2.87)

2-year BHAR: −6.05*** (t = −3.16)

3-year BHAR: −5.41* (t = −2.00)

Reading: AI-washers borrow the AI premium where AI is prized, then give it back

2 Verification (Table 7)

Measure: AIDistance, 1 − similarity of the S-1 Business section to AI patents (sentence embeddings)

Test: AI-washing × AIDistance

2-year BHAR: −1.31** (t = −2.19)

3-year BHAR: −1.43* (t = −1.98)

Distance alone: no effect

Reading: claims far from the business are hardest to check, so mispricing lasts longer

3 Litigation risk (Table 8, IA9)

Measure: first-year litigation risk: Kim–Skinner (firm) and industry class actions (Donelson et al.)

Why IPOs: no safe harbor for forward-looking statements

2nd-year BHAR: −0.034*** (KS), −0.048** (DFY)

AI words, first 10-K vs S-1: −1.33*, −0.55*

AI words, calls vs S-1: −0.19***, −0.12*

Reading: legal exposure makes the AI story unwind

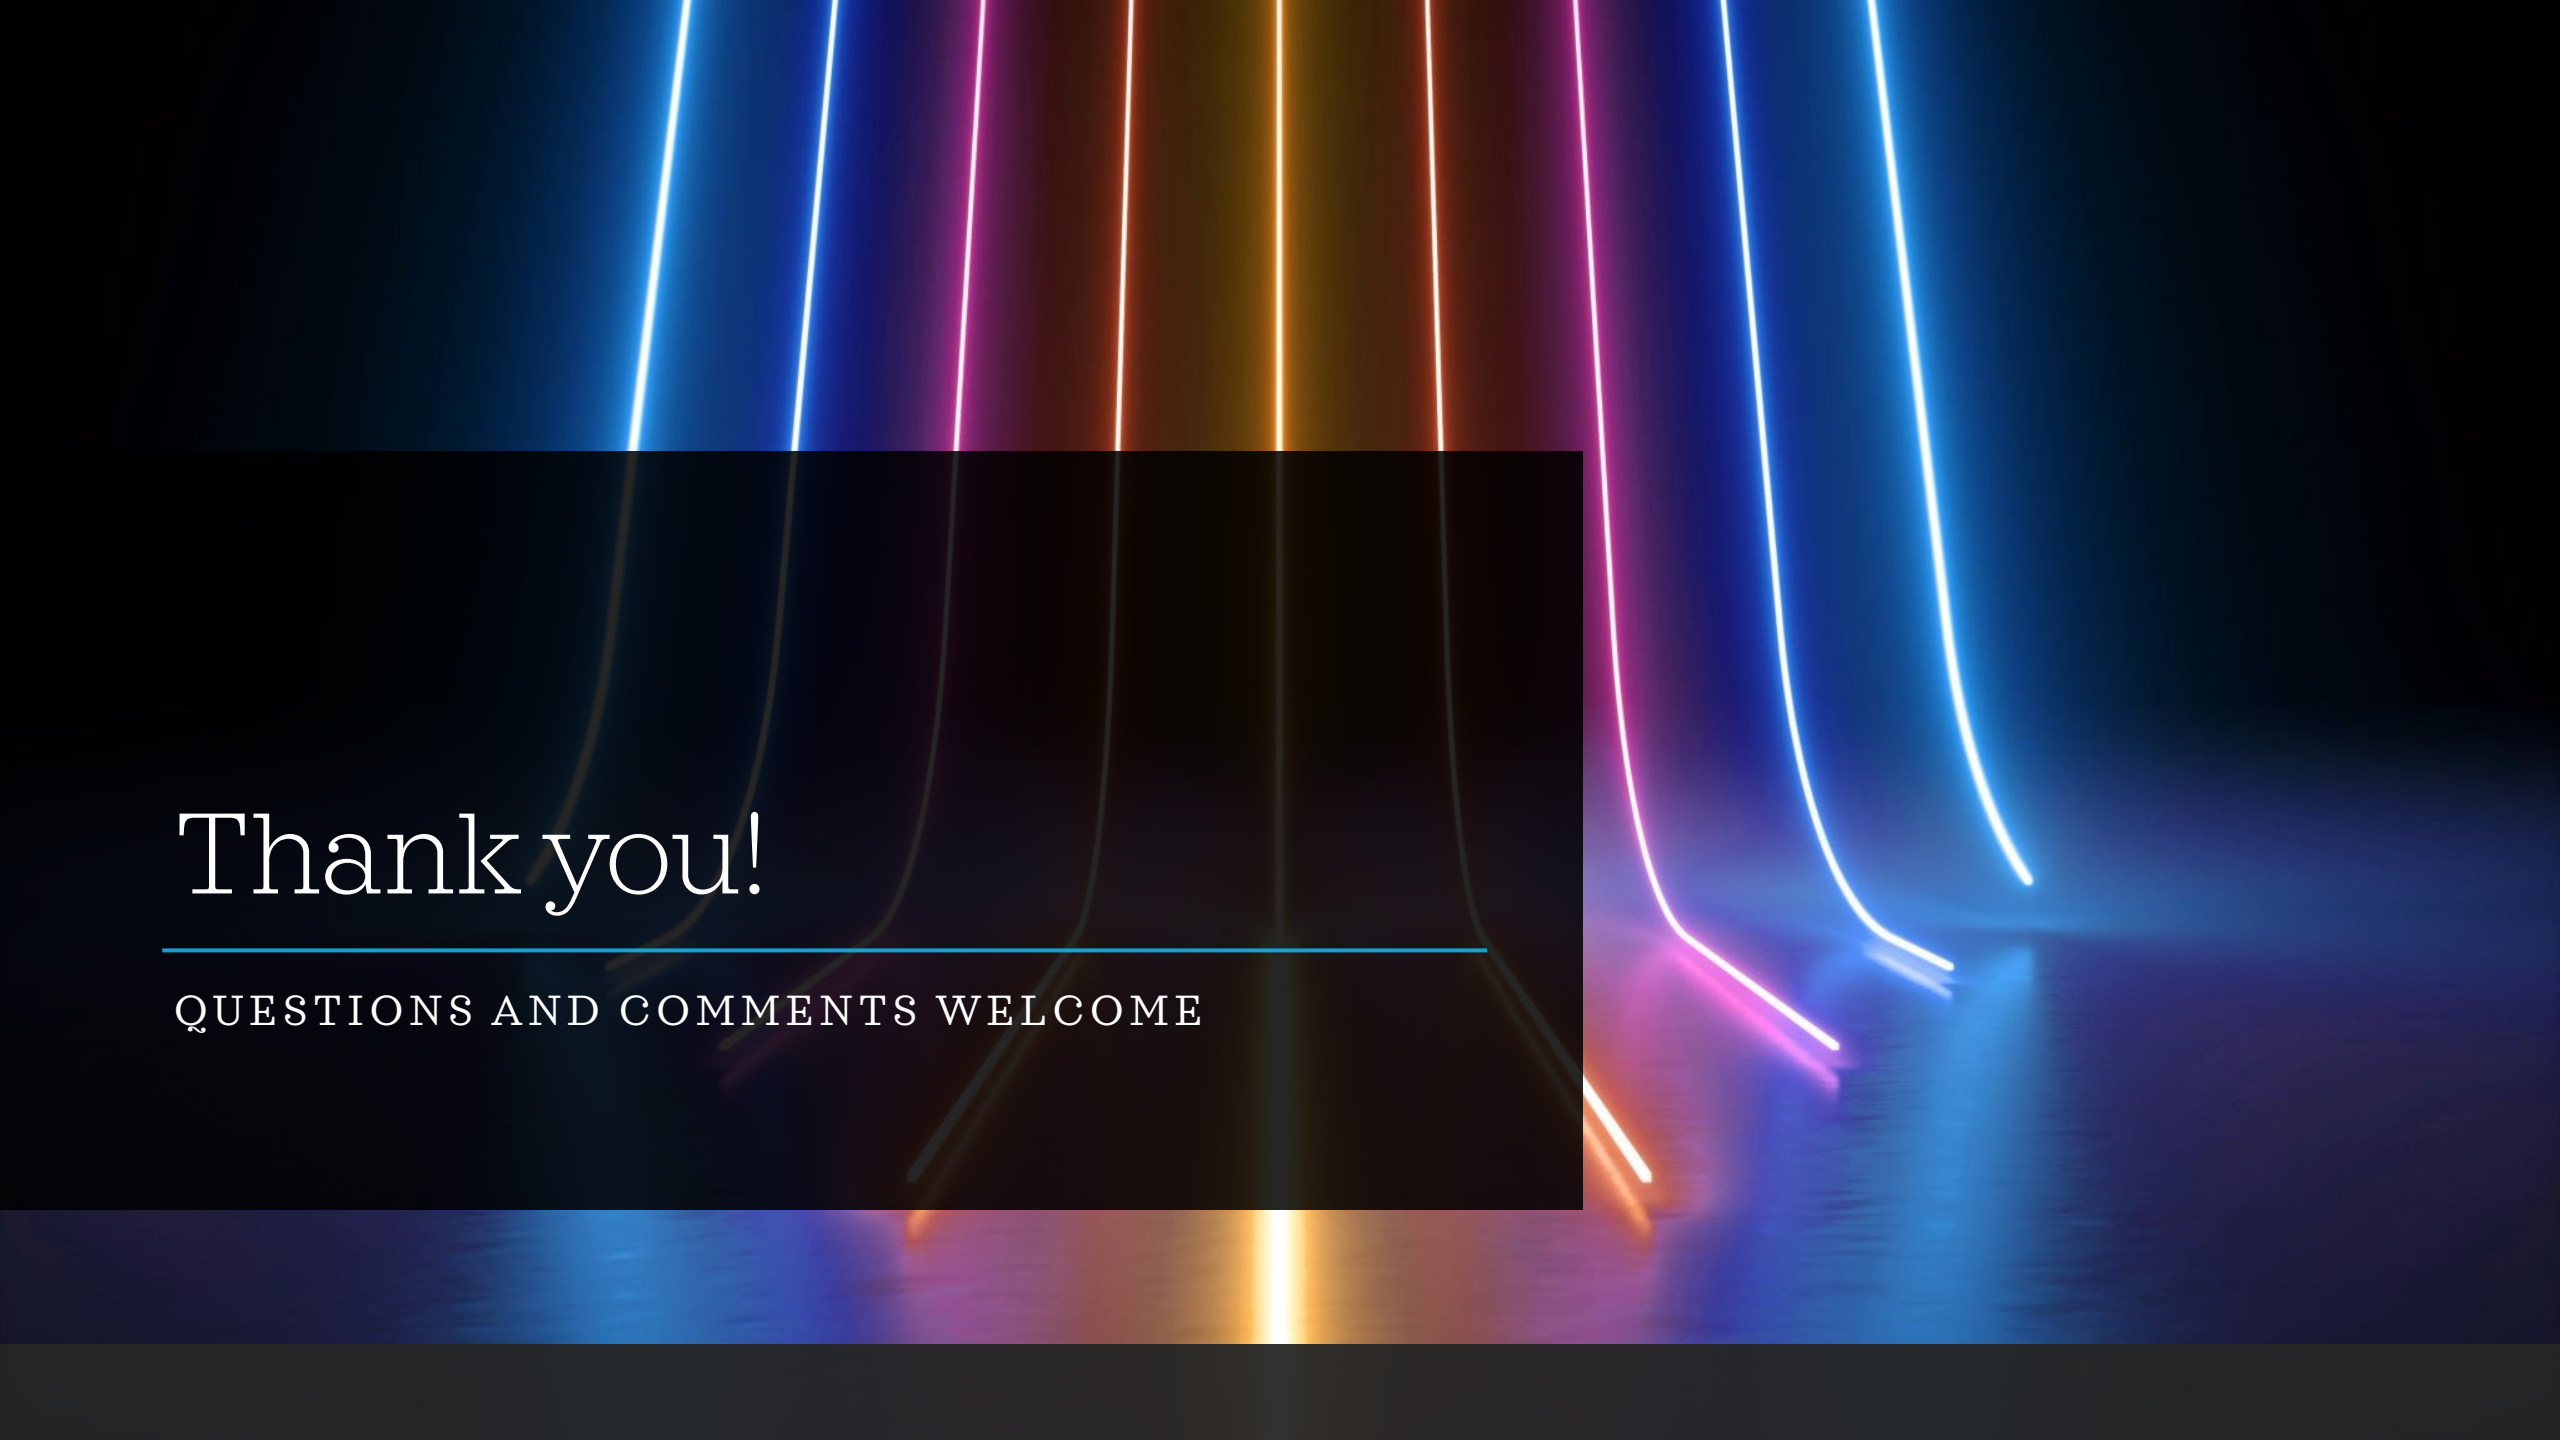
Who curbs AI-washing? Takeaways

Who curbs AI-washing?

- Intermediaries: AI expertise of VCs, underwriters and auditors shows no reliable effect
- SEC comment letters: more letters market-wide → less AI-washing, not less AI investment
- SEC enforcement (March 18, 2024): AI-washing falls more where it was common

Takeaways

- AI talk at the IPO is a weak signal unless backed by AI walk
- The talk–walk mismatch predicts long-run underperformance
- Regulators can check AI claims against observable inputs


Thank you!
QUESTIONS AND COMMENTS WELCOME



Appendix

BACKUP SLIDES

Why is AI washing special in IPOs?

Why AI capability is difficult to value at issuance

- Algorithms, proprietary data, and specialized talent are difficult to observe.
- AI investment can support growth and product innovation.
- IPO research has limited granular evidence on pre-IPO technological inputs.

Why disclosures may diverge from observable investment

- Limited public histories make AI claims difficult to verify at issuance.
- Capable firms may also limit disclosure to protect proprietary information.

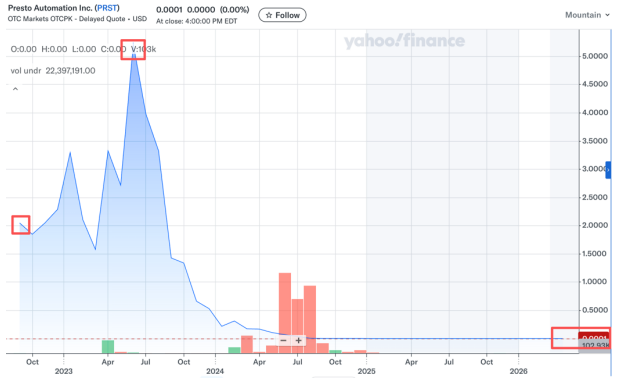
Anecdotal evidence: Presto Voice

Presto listed in September 2022 and promoted AI-powered restaurant technology



- *Vision* — Presto's AI-powered computer vision application provides real-time error detection, visual order tracking capabilities, and repeat guest identification. Presto Vision uses strategically located cameras and powerful artificial intelligence ("AI") algorithms to offer unique real-time insights to restaurant operations and analytics. Presto's next generation technology quickly identifies dropouts in drive-through lines, repeat customers, errors during food preparation, and missing items in the packing process. Vision's smart analytics promotes items with short preparation times when queues are long and notifies managers to deploy line-busting tablets when drive-through wait times or abandonment rates are high. The Vision product also helps boost order accuracy in the restaurant hospitality industry by deploying its technology to kitchens, which allows management to monitor whether orders going out the door are substantially accurate.

Its stock price peaked in **mid-2023**, then steadily declined and remained near zero



AP SUMMARY

SEC Charges Restaurant-Technology Company Presto Automation for Misleading Statements About AI Product

ADMINISTRATIVE PROCEEDING
File No. 3-22413

Contributions

We propose one of the first systematic frameworks to identify **AI-washing** in the IPO process

- ESG washing (Duchin et al., 2024 JF; Baker et al., 2024 JAR)

We extend traditional signaling theory by showing that **AI labels are not reliable signals** of firm quality or growth potential

- Real AI investment is associated with high value growth (Babina et al., 2024 JFE)
- AI-related narratives can be temporarily priced as signals (Beatty and Ritter, 1986 JFE; Rock, 1986 JFE; Ritter, 1991 JF)

We provide evidence supporting the need for **ex-ante regulation** of AI-related claims during IPOs

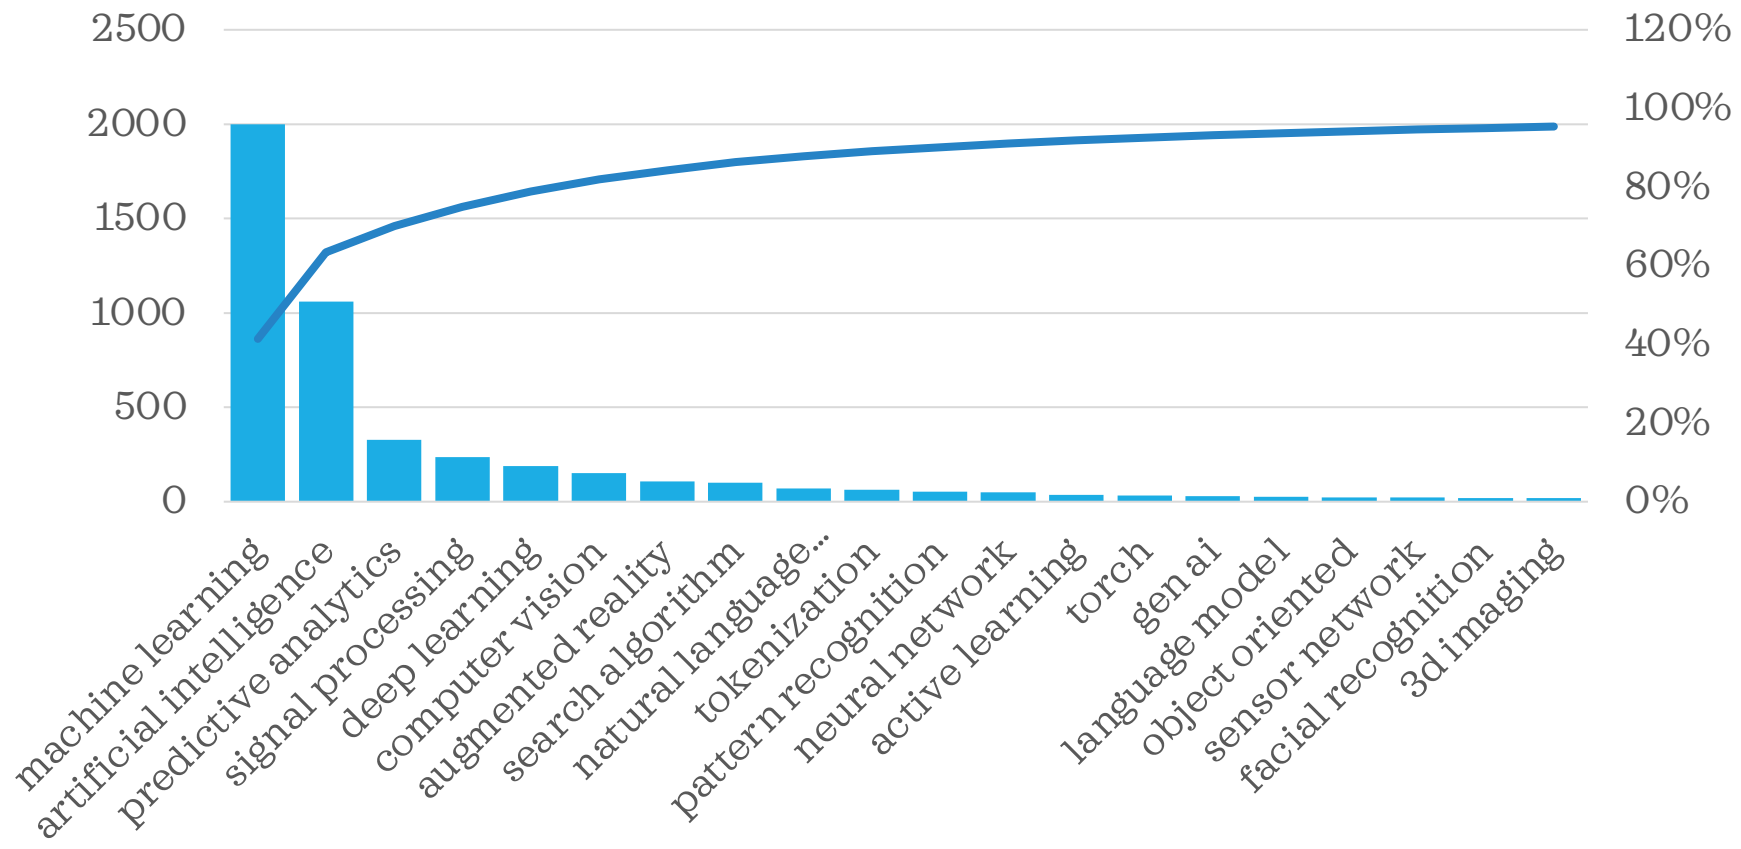
- Disclosure litigation risk and legal exposure in IPO prospectuses (Lowry and Shu, 2002 JFE; Hanley and Hoberg, 2012 JFE)

AI talk: AI keyword-based algorithm

Our dictionary contains 97 AI-related terms, 57 of which appear in S-1 filings.

KEYWORD	KEYWORD	KEYWORD	KEYWORD
3d imaging	computer vision	machine learning	search algorithm
acoustic model	data fusion	maximum likelihood	semantics
active learning	deep learning	model selection	semi-supervised learning
ai strategy	dimensionality reduction	motion capture	sensor network
ai training	facial recognition	natural language processing	sentiment analysis
aiops	feature selection	network analysis	signal processing
artificial intelligence	gen ai	neural network	supervised learning
augmented reality	generative artificial intelligence	object oriented	support vector machine
automatic speech recognition	gesture recognition	object recognition	tensor flow
autonomous system	h2o.ai	optimal search	tokenization
azure ml	hierarchical clustering	pattern recognition	torch
bayesian optimization	image matching	predictive analytics	virtual agent
classification algorithm	image recognition	random forest	
cluster analysis	intelligent agent	recommender system	
collaborative system	language model	regression model	

AI talk: Top 20 most frequent AI keywords



AI walk 1: AI human capital

Use employment histories from Revelio Labs' LinkedIn data

Count employees who remain at the firm on the IPO listing date

Exclude interns who have not been formally hired

Identify AI skills using the taxonomy in Babina et al. (2024, JFE)

Reduce noise in self-reported profiles by requiring:

At least one STEM degree

A minimum share of AI-related skills among all reported skills

AI walk 2: AI patents

Use the USPTO Artificial Intelligence Patent Dataset, which applies algorithms to identify AI patents

- Define a patent as AI-related if its score exceeds the threshold of 0.93

Match patents to listing firms using both match-keys from Kogan et al. (2017 QJE) and fuzzy-name matching to cover both granted and pre-granted filings

Use the filing date (instead of the granted date) to determine innovation timing (Luong et al., 2017 JFQA)

Count AI patent filings within the three-year window prior to the listing date

Our measures

AI talk: AI keyword mentions in S-1 registration statements

AI walk: AI employees at listing and AI patent filings in the prior three years

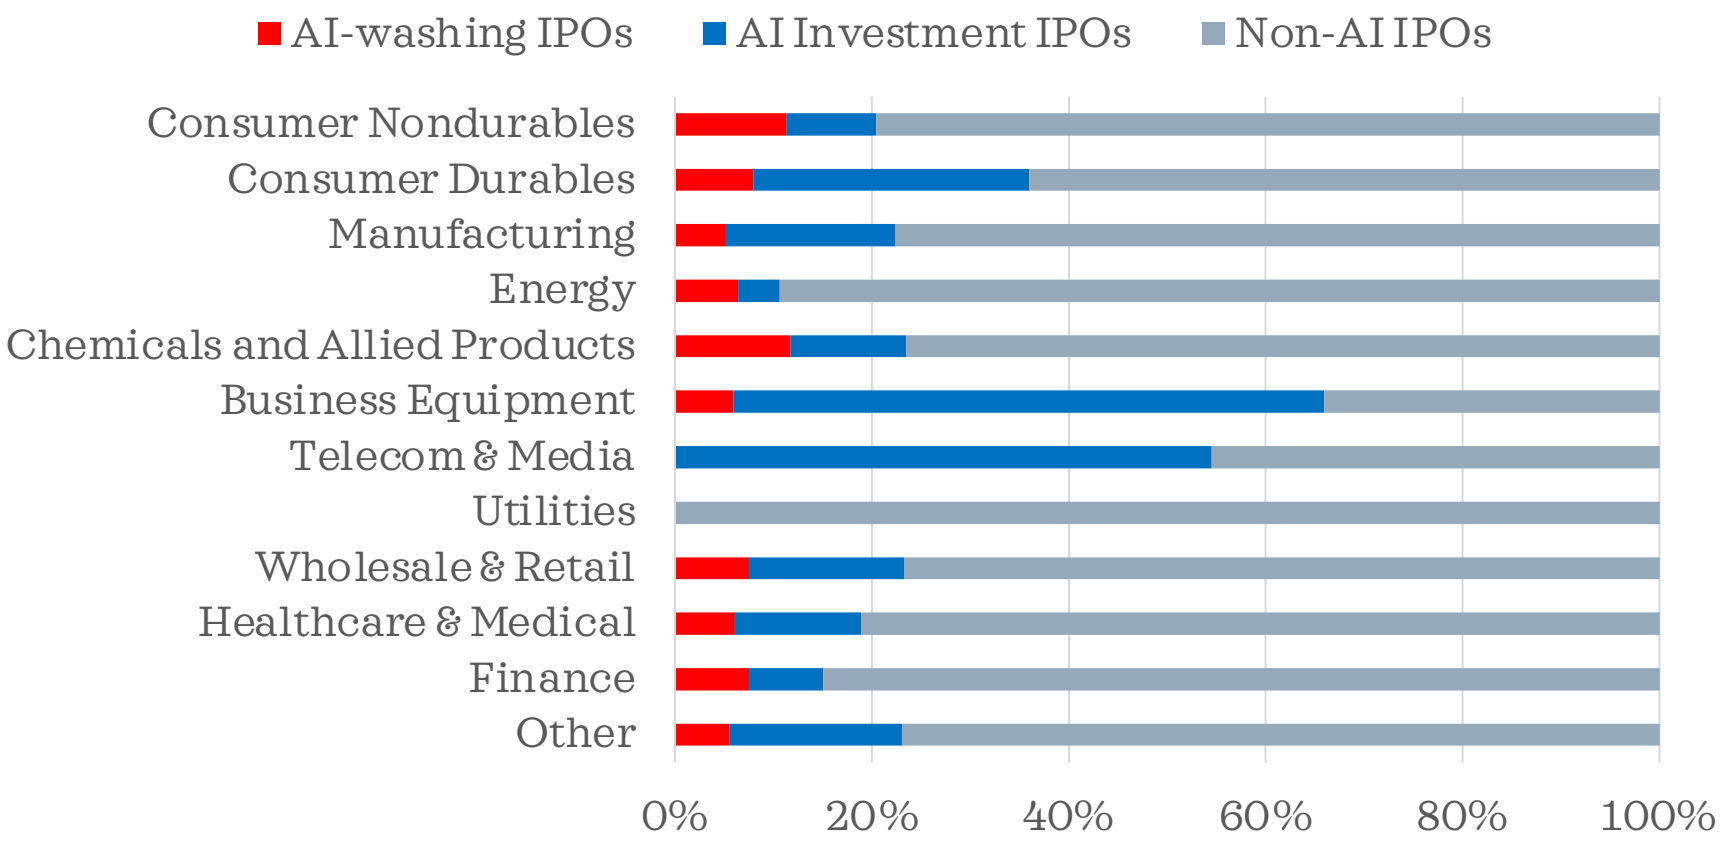
An AI-washing IPO meets all three conditions:

- AI talk exceeds its FF12 industry's historical 95th percentile (2007 to t – 1)
- No observable AI employees at listing
- No AI patent filings in the three years through listing

An AI-investment IPO has at least one AI employee or AI patent filing

Remaining issuers are non-AI IPOs

Industry distribution of AI IPOs



Summary statistics

Panel A1. Summary Statistics for First Day and Buy-and-hold Returns

	N	Mean	STD	10%	25%	Median	75%	90%
BHAR[0,1yr]	1,600	-0.110	0.658	-0.803	-0.545	-0.196	0.194	0.677
BHAR[1,2yr]	1,491	-0.094	0.626	-0.768	-0.489	-0.179	0.170	0.562
BHAR[0,2yr]	1,556	-0.137	1.005	-1.058	-0.754	-0.351	0.225	0.948
BHAR[0,3yr]	1,515	-0.172	1.297	-1.278	-0.946	-0.490	0.235	1.255
InitialReturn	1,671	0.197	0.327	-0.083	0.000	0.109	0.308	0.603

Panel A2. Summary Statistics for AI-related Characteristics

	N	Mean	STD	10%	25%	Median	75%	90%
AIWashing	1,671	0.064	0.245	0.000	0.000	0.000	0.000	0.000
AIInvestment	1,671	0.235	0.424	0.000	0.000	0.000	0.000	1.000
AIWord	1,671	2.841	8.603	0.000	0.000	0.000	1.000	8.000
AIPatent	1,671	1.530	6.050	0.000	0.000	0.000	0.000	3.000
AIEmployee	1,671	0.500	2.280	0.000	0.000	0.000	0.000	1.000

Summary statistics

Panel A3. Summary Statistics for Other Firm-level Characteristics

	N	Mean	STD	10%	25%	Median	75%	90%
HighUW	1,670	0.513	0.500	0.000	0.000	1.000	1.000	1.000
VC	1,671	0.591	0.492	0.000	0.000	1.000	1.000	1.000
Big4	1,671	0.393	0.489	0.000	0.000	0.000	1.000	1.000
Proceeds (in millions)	1,671	248.756	383.826	30.000	67.200	120.000	250.000	576.000
Age	1,665	17.920	22.697	3.000	6.000	10.000	19.000	40.000
Assets (in millions)	1,669	1414.601	4109.190	16.701	63.757	173.568	876.255	3414.723
Sales (in millions)	1,668	652.544	1764.007	0.000	0.000	85.929	395.431	1596.681
Carve	1,671	0.178	0.382	0.000	0.000	0.000	0.000	1.000
HighTech	1,671	0.550	0.498	0.000	0.000	1.000	1.000	1.000
Biotech	1,671	0.323	0.468	0.000	0.000	0.000	1.000	1.000
PrcRev	1,671	-0.013	0.139	-0.214	-0.077	0.000	0.071	0.143
Internet	1,671	0.095	0.293	0.000	0.000	0.000	0.000	0.000
Overhang	1,671	4.615	4.688	1.059	2.175	3.420	5.590	8.560
Primary	1,671	0.328	0.470	0.000	0.000	0.000	1.000	1.000
PositiveEPS	1,671	0.271	0.445	0.000	0.000	0.000	1.000	1.000
Second	1,671	0.097	0.231	0.000	0.000	0.000	0.000	0.368

Univariate test

Panel A. Non-AI versus AI-washing (N = 1,278)				Panel B. AI-investment versus AI-washing (N = 500)				
Variable	**Variable**	**N**	**Mean**	**N**	**Mean**	**Diff.**	**t-stat**	**stat**
HighUW	*HighUW*	393	0.697	106	0.575	0.122	2.376**	89***
VC	*VC*	393	0.814	107	0.701	0.113	2.067**	.619
Age	*Age*	393	2.553	107	2.513	0.040	0.525	.270
Proceeds	*Proceeds*	393	5.340	107	4.847	0.493	3.975***	.515
Sales	*Sales*	393	4.975	107	4.182	0.793	3.002***	07***
Assets	*Assets*	393	5.719	107	5.398	0.321	1.611	.794
HighTec	*HighTech*	393	0.728	107	0.495	0.232	4.650***	000
Biotech	*Biotech*	393	0.115	107	0.271	-0.157	-4.102***	35***
Internet	*Internet*	393	0.257	107	0.150	0.107	2.336**	55***
Second	*Second*	393	0.119	107	0.086	0.033	1.273	195
Big4	*Big4*	393	0.455	107	0.234	0.222	4.204***	28***
Positivel	*PositiveEPS*	393	0.211	107	0.299	-0.088	-1.918*	.227
PrcRev	*PrcRev*	393	0.037	107	-0.011	0.048	3.141***	.332
Carve	*Carve*	393	0.115	107	0.121	-0.007	-0.200	58**
Overhan	*Overhang*	393	5.799	107	6.841	-1.042	-1.893*	58***
Primary	*Primary*	393	0.293	107	0.187	0.106	2.190**	32***

Mechanism 1: AI value creation

We first run regressions for each of the Fama–French 12 industries and each year to estimate the long-run market valuation of AI investment

$$AIStock_{i,t-1} = \sum_{k=1}^{10} (1 - \delta)^k \, AIPatNum_{i,t-k}$$

$$lnMV_{i,t} = \beta_{AIValue} * AIStock_{i,t-1} + CONTROLS_{i,t-1} + FE_{i,t} + \varepsilon_{i,t}$$

After obtaining the industry-specific effects, we estimate the following equation:

$$Y_{i,t} = \beta_1 * AI_washing \times Beta_AIValue_{i,t} + \beta_2 * AI_washing_i + \beta_3 * AI_invest_i + \beta_4 * Beta_AIValue_{i,t} + CONTROLS_i + FE_{i,t} + \varepsilon_{i,t}$$

Mechanism 1: AI value premium

Dep. Variable=	***P/S***		***BHAR[0,2yr]***		***BHAR[0,3yr]***	
Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
AIWashing×AIVal		16.105***		-6.053***		-5.407*
		(2.87)		(-3.16)		(-2.00)
AIVal	-0.921	-2.104	1.516	1.842	0.675	0.974
	(-0.28)	(-0.64)	(0.95)	(1.10)	(0.17)	(0.25)
AIWashing	-0.356**	-0.465***	-0.086**	-0.043	-0.138**	-0.099
	(-2.55)	(-2.98)	(-2.34)	(-1.30)	(-2.59)	(-1.58)
AIInvestment	-0.403***	-0.412***	0.102**	0.106**	0.221***	0.224***
	(-3.62)	(-3.72)	(2.33)	(2.39)	(4.91)	(5.05)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Ind & Year	Ind & Year	Ind & Year	Ind & Year	Ind & Year	Ind & Year
Observations	605	605	1,538	1,538	1,498	1,498
Adjusted R2	24.0%	24.2%	10.0%	10.0%	7.4%	7.4%

Mechanism 2: Verification difficulty

Verification is harder when a firm's business is distant from AI

- Compare the issuer's S-1 Business section with the USPTO AI patent corpus.
- Embed both corpora using a sentence-transformer model and compute 1 − cosine similarity.
- Mean-center AIDistance, scale by its sample mean, and interact with AIWashing.

$$BHAR_{i,t} = \beta_1 AIWashing \times AIDistance_{i,t} + \beta_2 AIDistance_{i,t} + \beta_3 AIWashing_i + \beta_4 AIInvestment_i + Controls_i + FE_{i,t} + \varepsilon_{i,t}$$

Interpretation

- Greater technological distance is associated with more negative long-run returns for AI-washing IPOs.
- Distance captures verification difficulty. It may also reflect low technological fit.

Mechanism 2: Verification difficulty

Dep. Variable=	***BHAR[0,2yr]***		***BHAR[0,3yr]***	
Variable	**(1)**	**(2)**	**(3)**	**(4)**
AIWashing×AIDistance		-1.313**		-1.425*
		(-2.19)		(-1.98)
AIDistance	-0.099	0.003	-0.310	-0.200
	(-0.70)	(0.03)	(-1.25)	(-0.82)
AIWashing	-0.085**	-0.121**	-0.166***	-0.209***
	(-2.12)	(-2.37)	(-4.31)	(-4.69)
AIInvestment	0.062	0.072	0.156***	0.167***
	(1.12)	(1.38)	(3.23)	(3.44)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Ind & Year	Ind & Year	Ind & Year	Ind & Year
Observations	1,410	1,410	1,375	1,375
Adjusted R2	8.7%	8.7%	6.2%	6.2%

Mechanism 3: Litigation risk

Litigation risk measured over the first post-IPO year

LitRiskKS: Kim–Skinner (2012) predicted litigation exposure using firm characteristics, returns, volatility, skewness, and turnover.

LitRiskDFY: Disclosure-related class actions in the issuer's FF12 industry, following Donelson et al. (2022).

$$BHAR[1, 2yr]_{i,t} = \beta 1 AIWashing \times LitRisk_{i,t} + \beta 2 LitRisk_{i,t} + \beta 3 AIWashing_{i} + \beta 4 AIInvestment_{i} + Controls_{i} + FE_{i,t} + \varepsilon_{i,t}$$

Further supportive evidence: Disclosure gap

$$AIGap_{i,t} = \beta 1 AIWashing \times LitRisk_{i,t} + \beta 2 LitRisk_{i,t} + \beta 3 AIWashing_{i} + \beta 4 AIInvestment_{i} + Controls_{i} + FE_{i,t} + \varepsilon_{i,t}$$

Mechanism 3: Litigation risk

Dep. Variable=		***BHAR[1,2yr]***		
Variable	**(1)**	**(2)**	**(3)**	**(4)**
AIWashing×LitRiskKS		-0.034***		
		(-3.10)		
AIWashing×LitRiskDFY				-0.048**
				(-2.57)
LitRiskKS	-0.003	-0.001		
	(-0.32)	(-0.08)		
LitRiskDFY			0.023**	0.025***
			(2.46)	(2.81)
AIWashing	0.002	-0.097	0.001	0.080
	(0.05)	(-1.35)	(0.01)	(1.33)
AIInvestment	0.075*	0.074*	0.069*	0.063
	(1.76)	(1.73)	(1.69)	(1.55)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Ind & Year	Ind & Year	Ind & Year	Ind & Year
Observations	1,472	1,472	1,472	1,472
Adjusted R2	5.1%	5.1%	5.3%	5.4%

Mechanism 3: Post-IPO AI disclosure gaps

AI disclosure changes relative to the S-1, comparing AI-washing with AI-investment IPOs 10-K: matched Business, Risk Factors, and MD&A sections. Calls: quarters 5–8, scaled by words.

Dep. Variable=	***AIGapTenK***	***AIGapTenK***	***AIGapCall***	***AIGapCall***
Risk proxy	**KS**	**DFY**	**KS**	**DFY**
AIWashing × LitRisk	-1.328*	-0.547*	-0.188***	-0.118*
	(-1.76)	(-1.92)	(-3.28)	(-1.83)
LitRisk	0.146	0.005	0.044	0.078**
	(1.43)	(0.04)	(1.21)	(2.21)
AIWashing	-3.926	0.906	-0.698***	0.131
	(-1.35)	(1.31)	(-3.57)	(0.48)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Ind & Year	Ind & Year	Ind & Year	Ind & Year
Observations	212	214	373	374
Adjusted R^2	3.6%	1.7%	12.8%	13.4%

Regulatory oversight of AI-washing

Pre-filing review environment

ComLetter = ln(1 + comment letters issued market-wide from 12 to 6 months before filing).

This measures broader SEC review activity, not review of the focal issuer.

Logit specifications control for issuer characteristics, industry, and year effects.

$$\text{logit } \Pr(Y_i = 1) = \alpha + \beta_1 \, \text{ComLetter}_i + \gamma' X_i + \text{FE}i,t + \varepsilon i,t$$

SEC enforcement on misleading AI claims: March 18, 2024

Compare post-enforcement changes across industries with different 2010–2017 AI-washing exposure (Post × WashRatio).

Results indicate an association with regulatory scrutiny; they do not establish causality.

$$\text{logit } \Pr(\text{AIWashing}_i = 1) = \alpha + \beta_1(\text{Post}_i \times \text{WashRatio}_j) + \beta_2 \, \text{Post}_i + \beta_3 \, \text{WashRatio}_j + \gamma' X_i + \text{FE} + \varepsilon i,t$$

Regulatory oversight of AI-washing

Dep. Variable=	***AIWashing***		***AIInvestment***	
Variable	**(1)**	**(2)**	**(3)**	**(4)**
ComLetter	-0.264***	-0.249**	0.082	-0.044
	(-3.57)	(-2.44)	(0.84)	(-0.58)
Controls	No	Yes	No	Yes
Fixed Effects	Ind & Year	Ind & Year	Ind & Year	Ind & Year
Observations	1,445	1,436	1,524	1,515
Pseudo R^2	8.5%	11.4%	23.5%	33.9%

SEC enforcement on March 18, 2024

Dep. Variable=		***AIWashing***		
Variable	**(1)**	**(2)**	**(3)**	**(4)**
Post×WashRatio		-2.614***	-3.000***	-2.184**
		(-3.19)	(-2.75)	(-2.26)
Post	0.701**	1.405***	1.686***	2.472
	(2.06)	(3.79)	(3.23)	(1.62)
WashRatio	0.548	0.621*	0.520	0.364
	(1.55)	(1.74)	(1.52)	(1.05)
Controls	No	No	Yes	Yes
Fixed Effects	No	No	No	Filing quarter
Observations	748	748	735	692
Pseudo R^2	0.9%	1.2%	7.1%	10.1%