Discussion of

Industry-University Collaboration and Innovation under Scientific Retractions

Jiao, Lyandres, and Ren (2026)

Elvira Sojli

UNSW Sydney

ABFER — May 13, 2026

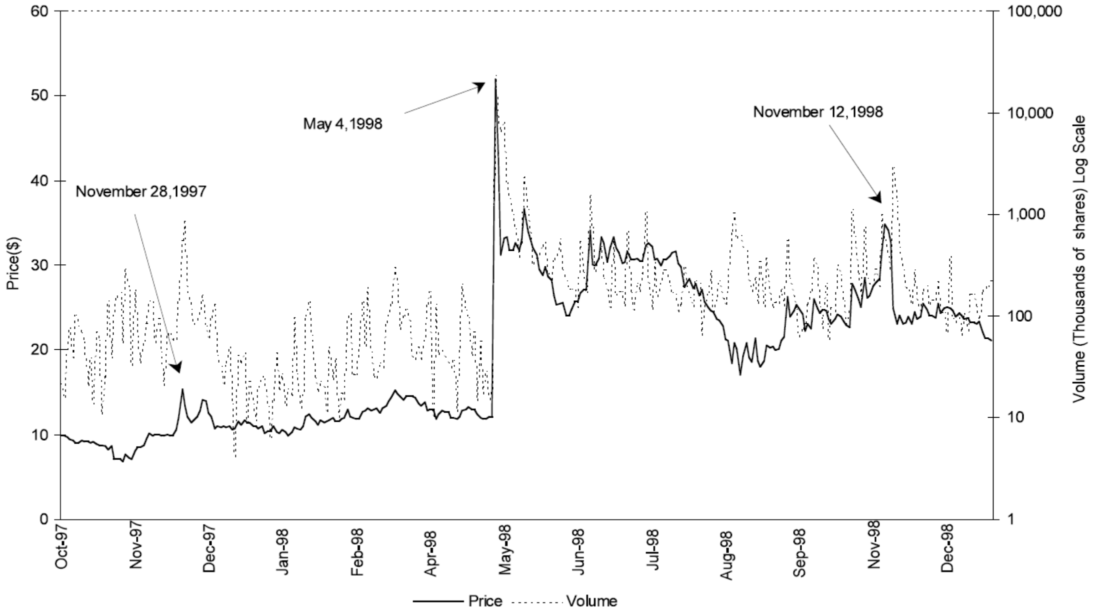

Figure 1. ENMD closing prices and trading volume, October 1, 1997, to December 30, 1998.

Science feeds technology

- Science and technology form a directed knowledge network. Technology cites science far more than reverse
  - Chen-Liu-Ma (2024), Bloom et al. (2023)
- Major innovations originate in science, propagate to technology with multi-year lag. AI: science leads patents by 20+ years
- Real-economy channel may bite even when equity prices miss the underlying science

Science feeds technology

- Science and technology form a directed knowledge network. Technology cites science far more than reverse
  - Chen-Liu-Ma (2024), Bloom et al. (2023)
- Major innovations originate in science, propagate to technology with multi-year lag. AI: science leads patents by 20+ years
- Real-economy channel may bite even when equity prices miss the underlying science

Conventional view: scientific credibility problems are an academic concern. Industry consumes outputs, not vulnerabilities.

This paper

When papers at a firm's university partners are retracted, what happens to firm innovation and stock price?

This paper

When papers at a firm's university partners are retracted, what happens to firm innovation and stock price?

Closest predecessors:

- Azoulay-Furman-Krieger-Murray (2015): academic spillovers from retractions
- Hussinger-Pellens (2019): retractions hurt prior academic co-authors
- Lu-Jin-Uzzi-Jones (2013): post-retraction citation declines

This paper

When papers at a firm's university partners are retracted, what happens to firm innovation and stock price?

Closest predecessors:

- Azoulay-Furman-Krieger-Murray (2015): academic spillovers from retractions
- Hussinger-Pellens (2019): retractions hurt prior academic co-authors
- Lu-Jin-Uzzi-Jones (2013): post-retraction citation declines

What is new here:

- Industry side of spillover, not academic side
- Equity-market valuation response in a retail-dominated market
- Decomposition into R&D spending vs R&D productivity

What is measured and how

Retraction Watch (treatment side)

- China-related retractions, 2007–2022, $\approx$15k events
- Time series: spike 2010–11 (paper-mill cleanups), trough 2012–19, spike 2020–22 (COVID, "rectification" campaigns)
- Each event flagged for reason (misconduct, error, methodology) and subject code
- Date of retraction $\neq$ date of concern

Firm-university link (paper's construction)

- Pair "linked" if joint patent in 5-year window
- Treatment = log retractions at all linked universities, untargeted by field
- Outcomes: firm patents, joint patents, CARs, R&D
- Panels: firm-year and firm-day

Key consequence. Treatment fires on a bioethics retraction at Tsinghua against a Tsinghua-firm lithography patent. Untargeted by construction.

What paper does

Main specifications.

- Panel OLS, firm and year fixed effects
- Treatment = log(retractions at linked universities) in year t
- Outcomes at $t+1, t+2, t+3$
- Standard errors clustered at firm

Event study (equity).

- CARs over $[-3, +3]$ and $[-5, +5]$ around retraction announcement date
- Market-model and Fama-French benchmarks
- Linked firms vs unlinked control

Mechanism. Splits R&D into spending (input) and patent-per-R&D (productivity). Latter falls. Former does not.

What paper finds (per 1 SD in retraction exposure)

Outcome	Reported effect	Horizon
Firm total patents (log)	−8% to −12%	t+1 to t+3
Joint patents w/ linked universities	−34% to −54%	t+1 to t+3
Citations to joint patents	−14% to −83%	t+1 to t+3
New collaborations	−85%	t+1
CAR around retraction (7-day)	−30 bp	event window
R&D spending	no significant effect	t+1 to t+3

What paper finds (per 1 SD in retraction exposure)

Outcome	Reported effect	Horizon
Firm total patents (log)	−8% to −12%	$t+1$ to $t+3$
Joint patents w/ linked universities	−34% to −54%	$t+1$ to $t+3$
Citations to joint patents	−14% to −83%	$t+1$ to $t+3$
New collaborations	−85%	$t+1$
CAR around retraction (7-day)	−30 bp	event window
R&D spending	no significant effect	$t+1$ to $t+3$

The paper's puzzle

- Outputs fall. Inputs do not.
- Paper attributes wedge to a "knowledge-productivity" channel: contaminated science makes ongoing R&D yield less

Four things stand out

- **Industry-side margin.** Almost all prior work on retractions studies academic spillovers, industry-side propagation is largely uncharted.
- **Capital-markets evidence.** CAR exercise puts a price on credibility in a market where the channel is unclear (more on this)
- **Productivity decomposition.** Separating R&D spending from R&D yield is the right move. Spending adjusts slowly, yield can move on impact. Lets paper say something about *which* margin is hit.
- **Setting.** China's 2020–22 retraction wave is one of the largest credibility events in recent science. Dataset is non-trivial to assemble.

So where to from here?

1. Who is pricing this? The equity-side channel
2. Are these “shocks”?
3. Mechanism — where is the magnitude link?
4. Missing innovation in the efficiency tests

Who is pricing this?

-30 bp CAR over 7 days around the retraction announcement date.

Who is pricing this?

-30 bp CAR over 7 days around the retraction announcement date.

For that to be a behavioural reaction, somebody has to *read the retraction and trade on it*

Who is pricing this?

-30 bp CAR over 7 days around the retraction announcement date.

For that to be a behavioural reaction, somebody has to *read the retraction and trade on it*

Three reasons that channel is unclear

Who is pricing this?

−30 bp CAR over 7 days around the retraction announcement date.

For that to be a behavioural reaction, somebody has to *read the retraction and trade on it*

Three reasons that channel is unclear

- **Chinese listed-firm trading is retail-dominated.** Retail investors are not reading Retraction Watch, *Nature*, or partner-university press releases. Even institutional desks rarely track retractions at universities one step removed from portfolio companies.

Who is pricing this?

-30 bp CAR over 7 days around the retraction announcement date.

For that to be a behavioural reaction, somebody has to *read the retraction and trade on it*

Three reasons that channel is unclear

- **Chinese listed-firm trading is retail-dominated.** Retail investors are not reading Retraction Watch, *Nature*, or partner-university press releases. Even institutional desks rarely track retractions at universities one step removed from portfolio companies.
- **Retraction date may not be the information date.** For high-concern retractions, PubPeer threads, journal expressions of concern, and Chinese-language media coverage often precede the formal retraction by weeks or months.

Who is pricing this?

-30 bp CAR over 7 days around the retraction announcement date.

For that to be a behavioural reaction, somebody has to *read the retraction and trade on it*

Three reasons that channel is unclear

- **Chinese listed-firm trading is retail-dominated.** Retail investors are not reading Retraction Watch, *Nature*, or partner-university press releases. Even institutional desks rarely track retractions at universities one step removed from portfolio companies.
- **Retraction date may not be the information date.** For high-concern retractions, PubPeer threads, journal expressions of concern, and Chinese-language media coverage often precede the formal retraction by weeks or months.
- **Huberman-Regev (2001) precedent.** Even publication of a cancer-cure breakthrough in *Nature* moved EntreMed by only 28%. A clean refutation by other labs in *WSJ* only walked the publicity back partway. Scientific information *by itself* barely registers without amplification.

Who is pricing this?

−30 bp CAR over 7 days around the retraction announcement date.

For that to be a behavioural reaction, somebody has to *read the retraction and trade on it*

Three reasons that channel is unclear

- **Chinese listed-firm trading is retail-dominated.** Retail investors are not reading Retraction Watch, *Nature*, or partner-university press releases. Even institutional desks rarely track retractions at universities one step removed from portfolio companies.
- **Retraction date may not be the information date.** For high-concern retractions, PubPeer threads, journal expressions of concern, and Chinese-language media coverage often precede the formal retraction by weeks or months.
- **Huberman-Regev (2001) precedent.** Even publication of a cancer-cure breakthrough in *Nature* moved EntreMed by only 28%. A clean refutation by other labs in *WSJ* only walked the publicity back partway. Scientific information *by itself* barely registers without amplification.

−30 bp may be real. But the path from a retraction at Tsinghua to a Shenzhen retail investor's sell order is not documented. Without it, the CAR is a correlation around an announcement date, not evidence on a credibility channel.

Suggestions

Easy

- Report institutional vs retail ownership of treated firms. If CAR is concentrated in most retail-held firms may need to rethink story
- Sample-balance test: do treated firms differ in analyst coverage, foreign institutional holding, QFII inclusion?

Suggestions

Easy

- Report institutional vs retail ownership of treated firms. If CAR is concentrated in most retail-held firms may need to rethink story
- Sample-balance test: do treated firms differ in analyst coverage, foreign institutional holding, QFII inclusion?

Medium

- Alternative event dates: first PubPeer mention, first Chinese-language news coverage of partner university misconduct, first expression of concern in the journal.
- Baidu search trends for partner-university name or retracted-paper topic around the event. If retail traders are responding, there should be a search-volume footprint.
- CAR around *retraction announcement* vs around journal's first *expression of concern* on same firms.

Suggestions

Easy

- Report institutional vs retail ownership of treated firms. If CAR is concentrated in most retail-held firms may need to rethink story
- Sample-balance test: do treated firms differ in analyst coverage, foreign institutional holding, QFII inclusion?

Medium

- Alternative event dates: first PubPeer mention, first Chinese-language news coverage of partner university misconduct, first expression of concern in the journal.
- Baidu search trends for partner-university name or retracted-paper topic around the event. If retail traders are responding, there should be a search-volume footprint.
- CAR around *retraction announcement* vs around journal's first *expression of concern* on same firms.

Time-consuming

- Text-mining of state media and finance portals (Eastmoney, Sina Finance) for partner-university mentions in 30 days around retractions.

Are these "shocks"?

Paper uses word "shock" repeatedly. Estimator is panel OLS with firm and year FE.

Within-firm variation may not be clean:

Universities self-select. Weak oversight cultures produce more retractions and plausibly different downstream effects on partners.

Firms choose partners. Treated firm set differs from never-treated set. Selection into linkage is part of the variation being identified.

Detection waves drive the time series.

- 2010–11: paper-mill exposures, >3,000/yr
- 2012–19: trough, <350/yr
- 2020–22: COVID, "rectification" campaigns, >2,300/yr

Overlap with macro shocks to Chinese R&D

Within-firm variation, conditional on year FE, may not be clean enough to carry causal language

Suggestions

Easy

- Event-study pre-trends at firm level. Four years before, four years after first treated retraction, normalized
- Sub-period split: 2010–11, 2012–19, 2020–22. Show results are not riding on one wave alone
- Soften language: replace "shock" with "exposure event"

Suggestions

Easy

- Event-study pre-trends at firm level. Four years before, four years after first treated retraction, normalized
- Sub-period split: 2010–11, 2012–19, 2020–22. Show results are not riding on one wave alone
- Soften language: replace “shock” with “exposure event”

Medium

- Placebo: non-linked firm-university pairs matched on industry, region, prior patenting where there should be no effect
- Drop 2020–22 entirely. If 2010–11 wave alone produces result, identification gets sharper

Suggestions

Easy

- Event-study pre-trends at firm level. Four years before, four years after first treated retraction, normalized
- Sub-period split: 2010–11, 2012–19, 2020–22. Show results are not riding on one wave alone
- Soften language: replace "shock" with "exposure event"

Medium

- Placebo: non-linked firm-university pairs matched on industry, region, prior patenting where there should be no effect
- Drop 2020–22 entirely. If 2010–11 wave alone produces result, identification gets sharper

More difficult

- IV from journal-level cleanup events. Some journals adopt image-forensics screening, triggering retraction batches at universities for reasons unrelated to firm exposure. Retraction Watch metadata supports the construction.

Mechanism magnitude link

Headline: 1 SD in partner retractions $\Rightarrow$ 8–12% drop in firm *total* patents

For knowledge-productivity channel to be sharp, retracted papers must occupy meaningful share of firm's knowledge stock.

Mechanism magnitude link

Headline: 1 SD in partner retractions $\Rightarrow$ 8–12% drop in firm *total* patents

For knowledge-productivity channel to be sharp, retracted papers must occupy meaningful share of firm's knowledge stock.

Treatment is untargeted. Bioethics retraction at Tsinghua counts identically against Tsinghua-firm lithography patent. Large diversified universities have more retractions *and* noisier exposure to any given partner.

Magnitudes are too large for direct contamination. Partner university produces thousands of papers per year. Even at extreme exposure, retracted papers are a small fraction of upstream knowledge. A 12% drop in firm *total* patents is a very large knowledge multiplier.

Mechanism magnitude link

Headline: 1 SD in partner retractions $\Rightarrow$ 8–12% drop in firm *total* patents

For knowledge-productivity channel to be sharp, retracted papers must occupy meaningful share of firm's knowledge stock.

Treatment is untargeted. Bioethics retraction at Tsinghua counts identically against Tsinghua-firm lithography patent. Large diversified universities have more retractions *and* noisier exposure to any given partner.

Magnitudes are too large for direct contamination. Partner university produces thousands of papers per year. Even at extreme exposure, retracted papers are a small fraction of upstream knowledge. A 12% drop in firm *total* patents is a very large knowledge multiplier.

Mechanism magnitude link

Headline: 1 SD in partner retractions $\Rightarrow$ 8–12% drop in firm *total* patents

For knowledge-productivity channel to be sharp, retracted papers must occupy meaningful share of firm's knowledge stock.

Treatment is untargeted. Bioethics retraction at Tsinghua counts identically against Tsinghua-firm lithography patent. Large diversified universities have more retractions *and* noisier exposure to any given partner.

Magnitudes are too large for direct contamination. Partner university produces thousands of papers per year. Even at extreme exposure, retracted papers are a small fraction of upstream knowledge. A 12% drop in firm *total* patents is a very large knowledge multiplier.

Mechanism magnitude link

Headline: 1 SD in partner retractions $\Rightarrow$ 8–12% drop in firm *total* patents

For knowledge-productivity channel to be sharp, retracted papers must occupy meaningful share of firm's knowledge stock.

Treatment is untargeted. Bioethics retraction at Tsinghua counts identically against Tsinghua-firm lithography patent. Large diversified universities have more retractions *and* noisier exposure to any given partner.

Magnitudes are too large for direct contamination. Partner university produces thousands of papers per year. Even at extreme exposure, retracted papers are a small fraction of upstream knowledge. A 12% drop in firm *total* patents is a very large knowledge multiplier.

Mechanism magnitude link

Headline: 1 SD in partner retractions $\Rightarrow$ 8–12% drop in firm *total* patents

For knowledge-productivity channel to be sharp, retracted papers must occupy meaningful share of firm's knowledge stock.

Treatment is untargeted. Bioethics retraction at Tsinghua counts identically against Tsinghua-firm lithography patent. Large diversified universities have more retractions *and* noisier exposure to any given partner.

Magnitudes are too large for direct contamination. Partner university produces thousands of papers per year. Even at extreme exposure, retracted papers are a small fraction of upstream knowledge. A 12% drop in firm *total* patents is a very large knowledge multiplier.

For such a large effect, channel is plausibly reputation/signaling, not knowledge devaluation. Paper's preferred mechanism needs a magnitude story.

Suggestions

Easy

- Report effect per retraction, not per SD of log retractions. Easier to assess against firm research portfolio.
- Back-of-envelope: average retracted paper accounts for what share of a partner university's citations? What share of firm's knowledge stock?
- Heterogeneity by retraction reason. Reputation channel should bite harder for misconduct; knowledge channel should bite harder for methodology.

Suggestions

Easy

- Report effect per retraction, not per SD of log retractions. Easier to assess against firm research portfolio.
- Back-of-envelope: average retracted paper accounts for what share of a partner university's citations? What share of firm's knowledge stock?
- Heterogeneity by retraction reason. Reputation channel should bite harder for misconduct; knowledge channel should bite harder for methodology.

Other options

- Topical proximity: Retraction Watch subject codes against CNIPA IPC classes on joint patents. Main result on topically proximate subset; topically distant subset is a built-in placebo. Helps separate channels: knowledge channel needs proximity, reputation channel does not.
- Decompose: effect on firms with no current joint research vs firms with active joint research at retracted PI. Reputation hits both; knowledge hits only the latter.

Missing innovation in efficiency tests

Paper's preferred mechanism: patents/R&D falls; R&D spending does not

That ratio requires *both* numerator and denominator to exist.

Two non-random selections

Missing innovation in efficiency tests

Paper's preferred mechanism: patents/R&D falls; R&D spending does not

That ratio requires *both* numerator and denominator to exist.

Two non-random selections

- **Missing R&D (denominator).** Firms with missing R&D drop from efficiency sample
- **Zero patents (numerator).** Mass point at zero efficiency. Firms with no patents in $t+k$ contribute a zero, indistinguishable from genuinely productive firms whose R&D simply has not paid off yet.

Missing innovation in efficiency tests

Paper's preferred mechanism: patents/R&D falls; R&D spending does not

That ratio requires *both* numerator and denominator to exist.

Two non-random selections

- **Missing R&D (denominator).** Firms with missing R&D drop from efficiency sample
- **Zero patents (numerator).** Mass point at zero efficiency. Firms with no patents in $t+k$ contribute a zero, indistinguishable from genuinely productive firms whose R&D simply has not paid off yet.

Koh et al. (2020): deletion and zero-imputation of unreported innovation produce **biased coefficients and false positives**. Effect is worst for patent-based metrics: 60%+ of positive-R&D firms never patent in the US sample. Chinese sample is unlikely to be cleaner.

Missing innovation in efficiency tests

Paper's preferred mechanism: patents/R&D falls; R&D spending does not

That ratio requires *both* numerator and denominator to exist.

Two non-random selections

- **Missing R&D (denominator).** Firms with missing R&D drop from efficiency sample
- **Zero patents (numerator).** Mass point at zero efficiency. Firms with no patents in $t+k$ contribute a zero, indistinguishable from genuinely productive firms whose R&D simply has not paid off yet.

Koh et al. (2020): deletion and zero-imputation of unreported innovation produce **biased coefficients and false positives**. Effect is worst for patent-based metrics: 60%+ of positive-R&D firms never patent in the US sample. Chinese sample is unlikely to be cleaner.

Missing innovation in efficiency tests

Paper's preferred mechanism: patents/R&D falls; R&D spending does not

That ratio requires *both* numerator and denominator to exist.

Two non-random selections

- **Missing R&D (denominator).** Firms with missing R&D drop from efficiency sample
- **Zero patents (numerator).** Mass point at zero efficiency. Firms with no patents in $t+k$ contribute a zero, indistinguishable from genuinely productive firms whose R&D simply has not paid off yet.

Koh et al. (2020): deletion and zero-imputation of unreported innovation produce **biased coefficients and false positives**. Effect is worst for patent-based metrics: 60%+ of positive-R&D firms never patent in the US sample. Chinese sample is unlikely to be cleaner.

Missing innovation in efficiency tests

Paper's preferred mechanism: patents/R&D falls; R&D spending does not

That ratio requires *both* numerator and denominator to exist.

Two non-random selections

- **Missing R&D (denominator).** Firms with missing R&D drop from efficiency sample
- **Zero patents (numerator).** Mass point at zero efficiency. Firms with no patents in $t+k$ contribute a zero, indistinguishable from genuinely productive firms whose R&D simply has not paid off yet.

Koh et al. (2020): deletion and zero-imputation of unreported innovation produce **biased coefficients and false positives**. Effect is worst for patent-based metrics: 60%+ of positive-R&D firms never patent in the US sample. Chinese sample is unlikely to be cleaner.

Missing innovation in efficiency tests

Paper's preferred mechanism: patents/R&D falls; R&D spending does not

That ratio requires *both* numerator and denominator to exist.

Two non-random selections

- **Missing R&D (denominator).** Firms with missing R&D drop from efficiency sample
- **Zero patents (numerator).** Mass point at zero efficiency. Firms with no patents in $t+k$ contribute a zero, indistinguishable from genuinely productive firms whose R&D simply has not paid off yet.

Koh et al. (2020): deletion and zero-imputation of unreported innovation produce **biased coefficients and false positives**. Effect is worst for patent-based metrics: 60%+ of positive-R&D firms never patent in the US sample. Chinese sample is unlikely to be cleaner.

> Productivity-channel result is the paper's headline mechanism. It is also the result most sensitive to how missing innovation is handled.

Suggestions

Easy

- Report share of firms with missing R&D, zero patents, or both, pre- and post-retraction. Sample-balance test.
- Show whether firms exit the efficiency sample more often after retractions. Selective attrition would mechanically generate the headline.
- Acknowledge the 637-obs gap between input and efficiency tables and explain it.

Suggestions

Easy

- Report share of firms with missing R&D, zero patents, or both, pre- and post-retraction. Sample-balance test.
- Show whether firms exit the efficiency sample more often after retractions. Selective attrition would mechanically generate the headline.
- Acknowledge the 637-obs gap between input and efficiency tables and explain it.

Medium

- Multiple imputation for missing R&D, then re-estimate efficiency regressions. Koh et al. (2020) provide the procedure; STATA `mi` command implements it.
- Two-part model on the efficiency variable: extensive margin (any patents at all) and intensive margin (patents/R&D conditional on patenting).
- Include non-USPTO patents. For Chinese setting, EPO and CNIPA filings are obvious extensions; reduces missing-patent rate substantially.

Suggestions

Easy

- Report share of firms with missing R&D, zero patents, or both, pre- and post-retraction. Sample-balance test.
- Show whether firms exit the efficiency sample more often after retractions. Selective attrition would mechanically generate the headline.
- Acknowledge the 637-obs gap between input and efficiency tables and explain it.

Medium

- Multiple imputation for missing R&D, then re-estimate efficiency regressions. Koh et al. (2020) provide the procedure; STATA `mi` command implements it.
- Two-part model on the efficiency variable: extensive margin (any patents at all) and intensive margin (patents/R&D conditional on patenting).
- Include non-USPTO patents. For Chinese setting, EPO and CNIPA filings are obvious extensions; reduces missing-patent rate substantially.

More work

- Report MI and IV estimates jointly. Different assumptions, different identification of the missing mechanism.

Other comments, briefly

- **Headline percentages, fragile by construction.** "54% drop in joint patents" uses $\beta\sigma_x/\bar{y}$ with $\bar{y} = 0.108$; really a few percentage points off the probability of any filing. "85% drop in new collaborations" is $3.7\% \rightarrow 0.6\%$. Conversely "8% patent decline" is $\exp(\beta\sigma_x) - 1 \approx -12\%$, which **understates** a real effect and sharpens Point 3.

Other comments, briefly

- **Headline percentages, fragile by construction.** "54% drop in joint patents" uses $\beta\sigma_x/\bar{y}$ with $\bar{y} = 0.108$; really a few percentage points off the probability of any filing. "85% drop in new collaborations" is $3.7\% \to 0.6\%$. Conversely "8% patent decline" is $\exp(\beta\sigma_x){-}1 \approx -12\%$, which **understates** a real effect and sharpens Point 3.
- **Intro contradicts body.** p. 4 says retractions *increase* probability of continued collaboration. Section 3.3 and Table 10 say *decrease*. Abstract is correct; intro sentence is the error.

Other comments, briefly

- **Headline percentages, fragile by construction.** "54% drop in joint patents" uses $\beta\sigma_x/\bar{y}$ with $\bar{y} = 0.108$; really a few percentage points off the probability of any filing. "85% drop in new collaborations" is 3.7% $\rightarrow$ 0.6%. Conversely "8% patent decline" is $\exp(\beta\sigma_x){-}1 \approx -12\%$, which **understates** a real effect and sharpens Point 3.
- **Intro contradicts body.** p. 4 says retractions *increase* probability of continued collaboration. Section 3.3 and Table 10 say *decrease*. Abstract is correct; intro sentence is the error.
- **Table 4 Panel C col (3).** Coefficients off by 100$\times$. Likely a units error.

Other comments, briefly

- **Headline percentages, fragile by construction.** "54% drop in joint patents" uses $\beta\sigma_x/\bar{y}$ with $\bar{y} = 0.108$; really a few percentage points off the probability of any filing. "85% drop in new collaborations" is $3.7\% \rightarrow 0.6\%$. Conversely "8% patent decline" is $\exp(\beta\sigma_x){-}1 \approx -12\%$, which **understates** a real effect and sharpens Point 3.
- **Intro contradicts body.** p. 4 says retractions *increase* probability of continued collaboration. Section 3.3 and Table 10 say *decrease*. Abstract is correct; intro sentence is the error.
- **Table 4 Panel C col (3).** Coefficients off by $100\times$. Likely a units error.
- **Endogenous sample split in Table 9.** Cols (4)–(6) split on *realized* CAR around the retraction. Conditions on outcome.

Other comments, briefly

- **Headline percentages, fragile by construction.** "54% drop in joint patents" uses $\beta\sigma_x/\bar{y}$ with $\bar{y} = 0.108$; really a few percentage points off the probability of any filing. "85% drop in new collaborations" is $3.7\% \rightarrow 0.6\%$. Conversely "8% patent decline" is $\exp(\beta\sigma_x){-}1 \approx -12\%$, which **understates** a real effect and sharpens Point 3.
- **Intro contradicts body.** p. 4 says retractions *increase* probability of continued collaboration. Section 3.3 and Table 10 say *decrease*. Abstract is correct; intro sentence is the error.
- **Table 4 Panel C col (3).** Coefficients off by $100\times$. Likely a units error.
- **Endogenous sample split in Table 9.** Cols (4)–(6) split on *realized* CAR around the retraction. Conditions on outcome.
- **Clustering.** Treatment varies at university-year level; clustering is at firm. Two-way at firm and university would be standard.

Other comments, briefly

- **Headline percentages, fragile by construction.** "54% drop in joint patents" uses $\beta\sigma_x/\bar{y}$ with $\bar{y} = 0.108$; really a few percentage points off the probability of any filing. "85% drop in new collaborations" is $3.7\% \to 0.6\%$. Conversely "8% patent decline" is $\exp(\beta\sigma_x) - 1 \approx -12\%$, which **understates** a real effect and sharpens Point 3.
- **Intro contradicts body.** p. 4 says retractions *increase* probability of continued collaboration. Section 3.3 and Table 10 say *decrease*. Abstract is correct; intro sentence is the error.
- **Table 4 Panel C col (3).** Coefficients off by $100\times$. Likely a units error.
- **Endogenous sample split in Table 9.** Cols (4)–(6) split on *realized* CAR around the retraction. Conditions on outcome.
- **Clustering.** Treatment varies at university-year level; clustering is at firm. Two-way at firm and university would be standard.
- **Two channels not separated.** "Knowledge devaluation" and "reputation/signaling" arrive on the same date. Null on R&D spending is consistent with productivity channel but also with slow input adjustment.

Other comments, briefly

- **Headline percentages, fragile by construction.** "54% drop in joint patents" uses $\beta\sigma_x/\bar{y}$ with $\bar{y} = 0.108$; really a few percentage points off the probability of any filing. "85% drop in new collaborations" is $3.7\% \to 0.6\%$. Conversely "8% patent decline" is $\exp(\beta\sigma_x) - 1 \approx -12\%$, which **understates** a real effect and sharpens Point 3.
- **Intro contradicts body.** p. 4 says retractions *increase* probability of continued collaboration. Section 3.3 and Table 10 say *decrease*. Abstract is correct; intro sentence is the error.
- **Table 4 Panel C col (3).** Coefficients off by $100\times$. Likely a units error.
- **Endogenous sample split in Table 9.** Cols (4)–(6) split on *realized* CAR around the retraction. Conditions on outcome.
- **Clustering.** Treatment varies at university-year level; clustering is at firm. Two-way at firm and university would be standard.
- **Two channels not separated.** "Knowledge devaluation" and "reputation/signaling" arrive on the same date. Null on R&D spending is consistent with productivity channel but also with slow input adjustment.
- **Collaboration = joint patents only.** Most industry-university transfer happens through sponsored research, talent flows, advisory boards. Reframe as "co-patenting collaborations."

Takeaway

Excellent question, impressive dataset, a real contribution in the making

Where the field is heading

- Science-technology spillovers are now a quantitative empirical agenda. Chen-Liu-Ma (2024), Bloom et al. (2023), Bryan-Williams (2021). Upstream knowledge network is becoming measurable.
- Papers that combine that network with credibly identified disruption events, a defensible equity-side channel, and disciplined handling of missing innovation will set the bar.
- This paper is among the earliest evidence on industry-side propagation of credibility shocks. Strengthening on these three margins puts it on the frontier.