When System Harm Reaches Asset Prices

A four-test discovery study using Fama–French models

Erik Postnieks

Center for Decision Accounting

Salt Lake City, Utah, United States

erik@decisionaccounting.org

https://decisionaccounting.org

Working Paper, Version 1.1

July 28, 2026

Planned public paper and data page (reserved until staging):

https://decisionaccounting.org/study/when-system-harm-reaches-asset-prices/

Planned public code page (reserved until staging):

https://decisionaccounting.org/research/when-system-harm-reaches-asset-prices/code/

Abstract

Markets can price costs that companies expect to bear. Social harm is harder: it may remain outside company accounts until regulation, litigation, contracts, reputation, financing, or demand connects the harm to expected cash flows. This paper asks whether a retrospective industry-level ratio of attributed system-welfare loss to revenue is associated with U.S. security returns. The public build combines a current SEC-listed-security universe, Yahoo Finance monthly returns, filing-item metadata, and standard factors from 1996 through 2026. Across the reported ordinary-month models, high-minus-low alpha is statistically indistinguishable from zero. That failure to reject is conditional on equal weighting, current-list selection, an unaudited security crosswalk, and unvalidated return extremes. In selected exit-cost and impairment filing months, maximum and mean mapped exposure are positively associated with absolute excess-return magnitude. The coefficient disappears under median exposure, after removing opioids, and in a single-domain model with broader controls. The predicted negative portfolio-skew difference is unsupported. A separate covariance calculation uses the union of all months containing a selected filing anywhere in the sample and is descriptive because its current uncertainty estimates do not resample that common calendar. The analysis establishes neither historical omission nor causal repricing. It identifies the data and design changes required for a journal test: point-in-time issuer exposures, validated filing text, one common-equity security per issuer, conventional event windows, and licensed returns with delisting coverage.

Keywords: System Asset Pricing Model; Missing System Theorem; beta-W; event studies; market efficiency; SEC filings; system welfare; asset pricing.

JEL codes: G12, G14, D62, D53, K22, Q54.

Highlights.

- Tests a retrospective system-harm measure against public stock returns.
- Does not reject zero alpha in the reported ordinary-month models.
- Finds that the selected-filing association changes across defensible specifications.
- Finds no support for negative skew; treats covariance as descriptive.
- Reports the design, results, failures, and limits together.

Contents

Plain language summary

Suppose two industries are associated with very different costs to public health, the climate, a financial benchmark, or another shared system. Will their stock prices reveal that difference? This paper looks in two places. First, it asks whether securities in the higher-exposure group earned different returns outside their own selected filing months. The reported factor models cannot distinguish the estimated difference from zero. Second, it asks whether securities with greater measured exposure also had larger monthly moves when their issuers reported exit costs, material impairments, or one of four hand-entered public matters. Some exposure assignments produce a positive association. Other reasonable choices erase or reverse it. The event result is unsettled.

The study also asks two harder questions. Do high-exposure securities have more negatively skewed returns? Do securities assigned to the same system move together after ordinary market factors are removed? The skew estimate points in the opposite direction, although its interval crosses zero. The common-movement coefficient is positive, but the calculation uses a market-wide set of filing months and its uncertainty estimates do not account for that shared calendar. Exit costs and impairments often concern ordinary business events, and the exposure measure was calculated after most returns occurred. The evidence defines a sharper next test; it does not show that system-welfare information caused the price movements.

Summary of findings

- **Ordinary months:** The study does not detect a high-minus-low exposure premium under CAPM, the Fama–French three- and five-factor models, or a momentum-augmented model. The author-control model has an estimated 80-percent minimum detectable effect of about 0.50 percent per month; the independent model is less precise.
- **Selected-filing months:** The estimated association changes with the rule used to assign exposure to multi-domain securities. A stricter single-domain audit finds no incremental association.
- **Return tails:** High-exposure securities do not show the predicted negative skew. Practical outlier treatments leave the point estimate in the opposite direction, while the bootstrap intervals include zero.
- **Shared-system movement:** The calculation uses every available residual return for 915 event-exposed securities across the union of 233 months in which any sample security had a selected filing. The coefficient is descriptive. Current uncertainty calculations cover securities repeated across pairs but do not resample the shared calendar, carry first-stage residual-estimation uncertainty, or validate the assigned system from filing text.

Reader guide

Section 1 states the question and the paper's contribution. Section 2 compares the paper with the closest finance, externality, and event-study research. Sections 3 and 4 define the four predictions and the public-data design. Sections 5–8 report every result. Sections 9–12 explain what the evidence can support and the stronger replication design. Appendix A gives the estimating procedures in reproducible form.

1. The question

An oil producer can impose climate costs that never appear on its income statement. A pharmaceutical company can contribute to a public-health loss before litigation connects that loss to expected cash flows. A bank can profit from a benchmark whose integrity is deteriorating until an investigation turns the deterioration into fines, restitution, and lost business. In each case, the social fact exists before the financial consequence.

That timing creates a simple empirical question: does the market price measured system harm continuously, or does the harm become visible in returns mainly when an institution creates a connection to firm value?

In this paper, **alpha** is the average return left unexplained by the comparison portfolios in a regression. A **factor** is a broad return pattern used to separate the proposed exposure from familiar market, size, value, and momentum effects. A **loading** is the estimated sensitivity of one quantity to changes in another. These terms describe statistical relationships; they do not by themselves identify a cause.

The paper studies that question with a ratio called system beta, written βW. It divides estimated annual system-welfare loss by annual industry revenue measured over the same activity boundary and period. A value above one means the estimated loss exceeds the industry's revenue; a value below one means it is smaller. A negative value denotes an estimated welfare gain. The ratio is an accounting exposure measure in this study. It is not a stock-market beta, and this paper does not identify the causal welfare effect of an individual firm.

The companion Missing System Theorem puts the boundary plainly: if a pricing rule receives exactly the same information in two situations, it cannot distinguish them merely because system welfare differs (Postnieks, 2026). The formal paper calls everything the rule receives its complete declared input. This return study does not infer that market prices were blind merely because system beta lacked a name. Beliefs, order flow, expected regulation, liability, and investor preferences can transmit relevant information before a formal welfare measure exists. The empirical test asks whether the later constructed exposure proxy adds information beyond those observable channels and controls.

The return test compares each security's ordinary months with months containing selected SEC filing items. The paper calls the second group **selected-filing months**. That label describes the selection rule; it does not establish that each filing conveyed new system-welfare information. The paper asks four questions. Does the high-minus-low exposure portfolio have a detectable alpha outside each security's own selected-filing months? Is exposure associated with absolute return magnitude in those months? Do high-exposure returns have a more negative tail? Do securities mapped to the same system move together?

Prior research already shows that social norms, climate concerns, pollution exposure, regulation, litigation, public attention, and investor preferences can affect expected or realized returns. The candidate contribution here is narrower: a reproducible cross-domain system-welfare exposure proxy and an explicitly mixed four-part test. The implementation remains exploratory because the exposure measure is retrospective and domain-level, and the event proxy relies heavily on broad filing metadata.

2. Relation to existing research

The paper draws from standard factor models, incomplete-information and missing-market analysis, environmental asset pricing, and event studies. Table 1 states what each literature established and identifies the narrower question tested here.

Table 1. How this paper fits the literature

Author or work	What it established	Inputs and assumptions	What remains unresolved for this paper	What this paper adds
Fama (1970, 1991); Jensen (1978)	Market efficiency is defined relative to an information set and an expected-return model.	A specified information set, trading process, and pricing model.	Whether the proposed exposure adds information beyond variables already observed by investors.	Uses efficiency as a disciplined null framework without claiming that information requires an academic name before it can enter price.
Fama and French (1993, 2015); Carhart (1997)	Market, size, value, profitability, investment, and momentum factors explain important return patterns.	Tradable factor portfolios and historical return covariation.	Whether system-harm exposure contains incremental return information.	Uses nested factor models as controls in ordinary-month and selected-filing tests.
Merton (1973)	Investors may price assets that hedge changes in state variables affecting investment opportunities.	Intertemporal optimization and payoff-relevant state variables.	Whether the welfare-attribution ratio is a priced state-variable loading.	Keeps the accounting exposure measure separate from a return loading.
Merton (1987)	Limited investor recognition can reduce firm value and raise expected returns.	Investors know and hold only subsets of securities.	Recognition of a security differs from recognition of social harm associated with its activity.	Tests a domain-level exposure measure while treating investor recognition as a close classical precedent.
Arrow (1969); Lindahl (1919)	Arrow analyzes limits of market allocation and the role of nonmarket allocation. Lindahl develops benefit-based taxation for public goods.	Incomplete markets or a public good whose financing must be assigned across users.	Whether a retrospective social-cost ratio leaves a detectable signature in securities.	Tests that return question without treating either predecessor as an asset-pricing result.

Author or work	What it established	Inputs and assumptions	What remains unresolved for this paper	What this paper adds
Allcott, Montanari, Ozaltun, and Tan (2025 revision)	Defines corporate social impact as the welfare loss caused by a firm's exit and estimates it for 74 firms in 12 industries.	Firm-exit counterfactual, consumer and worker substitution surveys, and product- and labor-market models.	A causal firm-exit welfare quantity differs from an attributed domain loss divided by industry revenue.	Tests whether a separate cross-domain accounting ratio aligns with returns, while treating the firm-level welfare framework as the closest measurement precedent.
Hong and Kacperczyk (2009)	Socially contested firms had lower norm-constrained ownership and higher expected returns than comparables.	Industry classifications, investor norms, ownership, and return tests.	Norms, litigation, neglect, and correlated exposures can compete as explanations.	Adds a continuous exposure proxy and an independently constructed sin-industry control.
Pástor, Stambaugh, and Taylor (2021)	Investor tastes and climate-risk hedging can affect green and brown prices and expected returns.	Observable greenness, heterogeneous preferences, and equilibrium demand.	Does not measure total system-welfare attribution across domains.	Tests a broader retrospective exposure measure while recognizing taste and hedging channels as alternatives.
Bolton and Kacperczyk (2021)	Higher-emissions firms earned higher returns in their sample; some institutions screened direct emissions.	Firm emissions, equity returns, and holdings.	Risk, tastes, data timing, and firm information compete as explanations.	Extends the question beyond carbon and tests whether the proposed proxy adds information after controls.

Table 1 continued

Author or work	What it established	Inputs and assumptions	What remains unresolved for this paper	What this paper adds
Pástor, Stambaugh, and Taylor (2022)	Unexpected climate-concern changes help explain realized green outperformance even when expected green returns are lower.	Green-minus-brown portfolios and climate-concern shocks.	Separating expected returns from realized concern shocks remains difficult.	Treats concern-driven repricing as a direct predecessor to the selected-filing hypothesis.
Ardia et al. (2023)	Unexpected increases in news-based climate concern are associated with green gains and brown losses.	Daily concern index, green/brown classification, and daily returns.	Focuses on climate and depends on exposure classification.	Asks whether a cross-domain proxy adds information beyond an attention-based environmental channel.
Choi, Gao, and Jiang (2020)	Local abnormal warmth raises climate attention and carbon-intensive firms underperform in those months.	Temperature, search attention, emissions proxies, and monthly returns.	Attention may move price without new fundamental information.	Identifies attention as an alternative explanation requiring direct tests in the next design.
Engle et al. (2020)	Climate-news innovations can support portfolios that hedge climate news.	Textual news measures, ESG exposures, and portfolio construction.	Hedge performance does not identify social welfare or causality.	Provides a methodological benchmark for validating both news and exposure measures.
Krüger (2015)	Negative corporate-social-responsibility events produce stronger price reactions, especially when legal and economic content is greater.	Hand-classified news and standard event-study methods.	Reactions may reflect cash flows, liability, reputation, or preferences.	Provides a direct predecessor and raises the standard for validating the filing content used here.
Hsu, Li, and Tsou (2023)	Toxic-emission intensity is associated with a pollution premium and environmental-policy uncertainty is a proposed risk channel.	Toxic emissions, within-industry portfolios, factors, and litigation penalties.	The premium's mechanism and measurement remain contested.	Tests a broader exposure proxy and must compare it directly with pollution and litigation measures.
Garel, Romec, Sautner, and Wagner (2024)	A firm biodiversity footprint is unrelated to average returns before salient policy developments, then predicts returns and short-window valuation changes around biodiversity-policy events.	Point-in-time firm footprints, lagged monthly returns, and verified global policy dates.	The result is domain-specific and does not measure total welfare across systems.	Uses a broader retrospective ratio across domains; the Garel et al. design supplies the closest template for firm-level timing and event validation.
Zhang (2025)	Respecting point-in-time emissions availability changes estimated carbon-return relations.	Historical disclosure timing, emissions, and global returns.	Exposure timing can reverse results.	Makes the retrospective status of the present exposure measure explicit and defines the prospective replication needed.
MacKinlay (1997); Kothari and Warner (2007)	Event studies connect dated information to abnormal returns and require careful timing, benchmarks, and dependence corrections.	Verified event dates, estimation windows, abnormal returns, and statistical assumptions.	Broad monthly filing metadata do not meet the full conventional event-study standard.	Presents the current analysis as an event-associated return-magnitude design and specifies the required reconstruction.
Shumway (1997)	Omitting delisting returns can bias return studies.	Complete histories including delisted firms.	Missing failed firms may hide downside outcomes.	Treats public-data survivorship as a limitation and requires a licensed-data replication.

The proposed contribution is the combination of a retrospective cross-domain attributed-loss-to-revenue ratio with four jointly reported return diagnostics. Firm-level welfare-impact measurement, domain-specific footprint pricing, and event-driven environmental repricing already exist. The present proxy has not yet shown incremental return information beyond validated firm-level measures.

3. Predictions

The study uses system beta as an average accounting ratio for sorting and exposure measurement. A causal welfare loading is a different object and requires a separate identification design. A return loading would also be different: it would describe how returns covary with a priced factor. The public-data analysis does not treat these three quantities as interchangeable.

Why might the four tests differ? An attributed social loss can remain outside expected firm cash flows until a legal, regulatory, contractual, reputational, financing, or demand event makes some part of that loss consequential to the firm. Under that account, ordinary returns need not carry a stable premium even though returns move more when an institutional event creates a connection to firm value. Negative skew and common movement are stronger implications: they require those connections to arrive mainly as adverse, system-level shocks. The study tests all four claims separately because the first two can appear while the stronger two fail.

P1: no detected ordinary-month premium. In the selected ordinary-month sample, a high-minus-low system-beta portfolio should have an alpha statistically indistinguishable from zero after the stated factor controls (Fama & French, 1993; Carhart, 1997). Failure to reject zero is consistent with this prediction. It cannot establish historical exclusion because limited power, portfolio construction, and offsetting priced channels can produce the same result.

P2: selected-filing association. System beta should be positively associated with absolute monthly excess-return magnitude in months containing selected filing items. The implementation tests this association under several exposure assignments and control sets. It does not establish that every filing conveyed system-welfare information or that the filing caused the observed return.

P3: negative portfolio skew. The time-series skewness of the equal-weighted high-exposure portfolio should be lower than the time-series skewness of the equal-weighted low-exposure portfolio if adverse system-level shocks disproportionately affect the high-exposure leg.

P4: within-system covariance. If the exposure is systemic, returns for securities assigned to the same shared-system category may move together when the system becomes financially consequential. The present calculation is a preliminary proxy for that prediction: it estimates pairwise covariance over the union of all months containing a selected filing anywhere in the sample. It does not require either security in a pair to file in a given month, so it cannot establish the predicted event-specific movement.

The four predictions are distinct. The corrected estimates do not reject P1. P2 changes across defensible exposure and sample choices. P3 receives no support. P4 produces a positive descriptive coefficient whose uncertainty is unresolved because the securities share a market-wide event calendar. None of these results carries a causal interpretation.

Table 2. Predicted return signatures

Proposed quantity	Finance analogue	Public-data test	Result in this paper
βW exposure	excluded system-welfare coordinate	sort and event exposure	used throughout
Ordinary-month pricing	ordinary factor premium	P1 high-minus-low alpha	zero alpha is not rejected in a 363-month security-specific design
Selected-filing relation	filing-associated return magnitude	P2 absolute monthly excess return	specification-sensitive and unresolved
Tail-risk signature	crash-like left tail	P3 skew	fails
Systemic covariance	shared damaged-system risk	P4 covariance during selected filing months	positive descriptive coefficient; sampling uncertainty unresolved

Note: The present design does not establish historical exclusion or causal repricing. P1 is a conditional failure to reject zero alpha. P2 is unstable. P3 receives no support. P4 is descriptive.

4. Data and scope

This study uses a current-listed public-data build with documented source paths and exclusions. Return validation, security-type filtering, and historical survivorship reconstruction remain incomplete. A licensed-data replication could add historical issuer coverage, delisting returns, and point-in-time issuer characteristics.

The issuer universe comes from SEC company ticker and exchange metadata plus data.sec.gov submissions files for SIC metadata where available. Equity returns are Yahoo Finance adjusted monthly closes obtained through `yfinance`. The sample runs from January 1996 through May 2026. The factor set includes the Ken French monthly market, size, and value factors plus momentum. The size and value factors are conventionally labeled SMB, for small minus big, and HML, for high minus low book-to-market. I add three author-constructed controls: a sin-industry portfolio based on tobacco, alcohol, and gambling mappings (Hong & Kacperczyk, 2009); a litigation-exposure portfolio based on opioids, PFAS, tobacco, firearms, and private-prison mappings; and a green-minus-brown portfolio. The last subtracts the equal-weighted return of securities mapped to oil and gas, coal, cement, aviation, and shipping from the equal-weighted return of securities mapped to nuclear and water domains. It is a rough transition-exposure proxy. It has not been established as a traded risk factor or as a measure of corporate environmental performance.

Version 1.1 uses a file snapshot of 58 numeric domain estimates parsed from `paperData.js` on July 27, 2026 (SHA-256 `1211784a5d98f9271319612b224583dc62092956b039fc6b3956403aa3b07fa7`). The normalized CSV has SHA-256 `246642f2c4e9f43d49afb25620cdc4da620e2e8b1bd6539bf0b4411b7407045f`. The rerun does not verify per-row admission against a single machine-readable controlling receipt. These rows are frozen analysis inputs rather than a fully verified measurement authority.

The issuer-domain crosswalk is a 2026 heuristic security classification. A row enters when a current SEC SIC code or a company-name or SIC-description substring matches a domain rule. The current export does not reliably distinguish SIC matches from keyword-only matches, and the keyword rules have not received a row-level false-positive audit. The public build contains 4,922 ticker symbols rather than 4,922 distinct firms. The symbols correspond to at most 4,118 identified CIK issuers, and some issuers contribute multiple listed securities. Results that use this crosswalk are diagnostics of the classification algorithm and cannot yet be interpreted as firm-level system exposure.

Across the 58 domain estimates, the mean is 4.84, the median is 2.33, the interquartile range is 1.04 to 6.57, and the range is 0.30 to 21.98. Across mapped tickers, the maximum-domain measure has a mean of 6.06 and a median of 5.00; the mean-domain alternative has a mean of 3.45 and a median of 2.57. These distributions make the mapping and influence checks central to interpretation.

The corrected coded-event panel uses SEC 8-K Item 2.05, which reports costs associated with exit or disposal activities, and Item 2.06, which reports material impairments. These items establish financial-reporting content; they do not establish a connection to the system assigned by the crosswalk. A separate hand-entered list contains ten ticker-event observations associated with four public legal or regulatory matters: the 1998 tobacco Master Settlement Agreement, the 2015 Volkswagen emissions notice, a 2021 opioid settlement, and a 2023 PFAS water settlement. The present package does not contain a reproducible filing-text validation ledger. Generic Item 8.01 filings are excluded because that category overwhelmed the earlier panel without identifying event content.

The system-beta estimates were constructed after most sample returns occurred. This is therefore a retrospective exposure design: it tests whether a later measure aligns with earlier return patterns. It does not test whether investors used contemporaneously available system-beta data.

Reported t-statistics and intervals are unadjusted for the family of predictions, exposure constructions, domain deletions, lag horizons, outcome treatments, and control sets. The sensitivity specifications are diagnostic and do not constitute separate confirmatory tests. Version 1.1 makes no family-wise error-rate claim.

Table 3. Public-data sample

The resulting data sets contain:

Quantity	Count
Quantified domains	58
Non-null domain-security mapping rows	13,079
Mapped ticker symbols	4,922
Identified CIK issuers represented	at most 4,118
Tickers mapped to more than one domain	3,277
Maximum domains assigned to one ticker	10
Monthly return periods used in principal P1 series	363
Monthly periods in independent standard-factor audit	364
Raw Item 2.05, Item 2.06, or hand-entered event records before ticker-month collapse	5,210
Unique selected-filing ticker-months before exposure and market-data joins	1,841
Unique selected-filing ticker-months after exposure merge	1,840
Complete ticker-month observations in the principal P2 regression	1,703
Hand-entered public-matter ticker-months	10
Complete hand-entered observations	9

The universe begins with securities listed in the SEC current-company file in 2026 and adds eight hand-selected failed, acquired, or reorganized securities. It therefore excludes most issuers that delisted or disappeared earlier in the 1996–2026 sample. The resulting panel is survivor selected, is not a historical investable universe, and cannot support a complete downside-tail analysis. Delisting bias is a known issue in return studies (Shumway, 1997).

5. P1: Security-specific ordinary-month test

P1 forms equal-weighted high- and low-system-beta portfolios from the bottom and top 20 percent of mapped ticker symbols ranked by maximum mapped beta at the 2026 classification date. Membership is fixed through the historical sample. Each month, a ticker is omitted only if that ticker has a selected Item 2.05, Item 2.06, or hand-entered event in the same month. The intercept in each factor regression is the estimated alpha. Inference uses heteroskedasticity-and-autocorrelation-consistent standard errors with six monthly lags. Historical market capitalizations are unavailable in the public build.

Table 4. Main estimates

Prediction and estimate	Estimate	t-stat	Result
P1: author-control specification	-0.00222147	-1.26	Zero alpha not rejected
P1: FF5 plus momentum audit	-0.00312809	-1.04	Zero alpha not rejected
P2: maximum mapped beta	0.00642947	5.50	Positive in this specification
P2: single-domain audit with external controls	-0.00128600	-0.78	No association detected
P3: high-versus-low skew difference	13.0513	bootstrap interval crosses zero	Opposite predicted sign
P4: shared-system covariance coefficient	0.0007711	inferential statistic withdrawn	Descriptive only

The author-control specification estimates an alpha of -0.22 percent per month, with a t-statistic of -1.26 across 363 months. The author-control regression has $R^2 = 0.9986$. Its constructed controls use the same security returns and provisional classification map as the tested portfolio, so the resulting alpha precision and 0.50-percent MDE are treated as secondary. The independent FF5-plus-momentum estimate supplies the principal external-factor benchmark. A normal approximation based on the author-control model's standard error gives an 80-percent minimum detectable effect of about 0.50 percent per month, or 6.1 percent compounded annually. The independent audit, using only single-domain securities and the Fama–French five factors plus momentum, estimates -0.31 percent per month, with a t-statistic of -1.04 across 364 months; its corresponding monthly minimum detectable effect is about 0.84 percent. CAPM, Fama–French three-factor, and Carhart four-factor estimates also fail to reject zero. These failures to reject are conditional on a current-listed-security universe, equal weighting, an unaudited exposure crosswalk, and unvalidated Yahoo return extremes. They do not establish a zero premium.

6. P2: Association in selected-filing months

P2 asks whether system beta is associated with absolute monthly excess-return magnitude in selected-filing months. The corrected panel contains one observation per ticker-month and excludes generic Item 8.01 filings. The specification includes the Fama–French and momentum factors plus the author-constructed sin, litigation-exposure, and green-minus-brown controls. It remains a filing-associated monthly-return analysis rather than a conventional short-window event study based on independently verified announcements (MacKinlay, 1997; Kothari & Warner, 2007).

With the maximum mapped beta as exposure, the coefficient is 0.00643 with a two-way ticker- and month-clustered t-statistic of 5.50 across 1,703 complete observations. The dependent variable is `|R_it -`

`R_ft|`, so the coefficient corresponds to 0.643 percentage point of absolute monthly excess-return magnitude per one-unit increase in beta-W. Replacing the maximum with the security's mean mapped beta produces a coefficient of 0.00633 and a t-statistic of 2.75. A percentile-rank or logarithmic transform remains positive. The median mapped beta is negative and statistically indistinguishable from zero. Removing the opioid domain changes the coefficient to -0.00059 with a t-statistic of -0.49.

The 1,703 ticker-month rows represent 887 ticker symbols and 856 identified CIK issuers. Because the current regression clusters ticker symbols separately, its standard errors do not account for dependence across multiple securities of the same issuer. Issuer/month clustering requires a rebuilt one-common-equity-security-per-issuer panel.

An independent audit takes a stricter approach: it excludes every multi-domain ticker, adds the Fama–French profitability and investment factors, absolute market movement, prior 12-month volatility, event-type indicators, and year effects, and retains 539 complete events. In that sample, the coefficient is -0.00129 with a ticker- and month-clustered t-statistic of -0.78. Adding system beta lowers adjusted R-squared by 0.00089. The hand-entered sample contains only nine complete observations and yields no positive evidence on its own: its univariate coefficient is -0.00634 with a t-statistic of -1.03, and a controlled model cannot be estimated at that sample size.

P2 therefore does not survive the full specification audit. The positive estimate depends on how multi-domain securities are assigned exposure and on the opioid domain. Exit costs and impairments can concern ordinary restructuring, accounting recognition, or issuer-specific operating news. Several mechanisms can produce the same association, and the nine complete hand-entered observations cannot separate them.

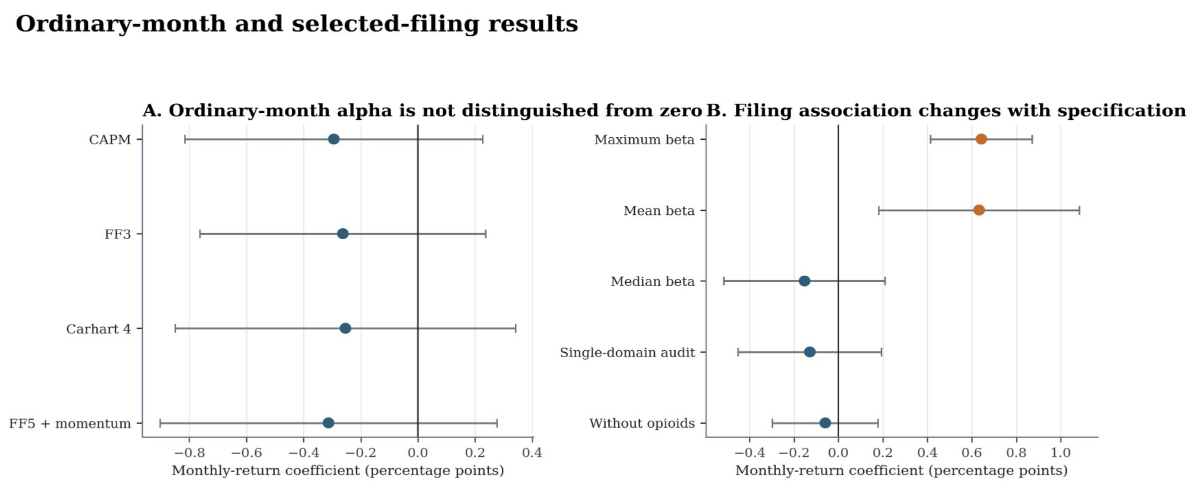

Figure 1. The ordinary-month failure to reject zero recurs across four standard factor models, while the selected-filing association changes across exposure and sample choices. Dots show monthly-return coefficients in percentage points; bars show approximate 95 percent intervals from reported standard errors.

6.1 Within-security lag diagnostic

One important alternative explanation is persistent security-level volatility: securities mapped to higher system beta may simply be more volatile. The current follow-on calculation compares coded-event observations with available observations for the same ticker 12 and 24 months earlier, then demeans the variables within ticker. This absorbs a time-invariant additive ticker component and the ticker-constant system-beta level. It does not remove time-varying volatility, differences in event importance, or selection into the event panel.

The current within-ticker calculation is unbalanced. At the 12-month horizon it combines 1,703 event rows with 1,665 available lagged rows; at the 24-month horizon it combines 1,703 event rows with 1,562 available lagged rows. Some lagged months are themselves selected event months and enter both arms. The dependent variable is `|R_it - R_ft|`. The regressors are system beta, an event indicator, their interaction, Fama–French and momentum controls, and the three constructed exposure controls. The reported interactions describe this existing panel and are not matched-pair estimates.

Table 5. Event-versus-placebo comparisons

Lag horizon	Observations	Securities	System beta × selected-filing indicator	t-stat
12 months	3,368	887 ticker symbols	0.00391153	3.01
24 months	3,265	887 ticker symbols	0.00255600	2.04

Notes: Both ticker-demeaned specifications include the stated factor and exposure controls. Standard errors are clustered by ticker and calendar month. At 12 months, 35 ticker-months occur in both arms; at 24 months, 36 do. A release-grade matched design must retain complete pairs, identify pair membership, and prespecify the treatment of lagged months that contain another selected event.

The interaction is positive at both lag horizons under the maximum-exposure mapping. Because the panel is incomplete and some months enter both arms, the calculation does not support a matched-event interpretation. It also retains the exposure-assignment sensitivity found in P2. The result motivates a rebuilt event study with complete pairs, verified announcements, and unique issuer-event observations.

7. P3: Skew fails

P3 predicts that high-system-beta activities should show more negative skew than low-system-beta activities. In the raw current-vintage sort, the high-minus-low skew difference is +13.0513, the opposite sign from the prediction (Table 4). The magnitude is driven by extreme public-return observations. A security bootstrap gives a 95 percent interval from -4.04 to 18.93; a 12-month block bootstrap gives -3.59 to 19.22.

Practical outlier treatments greatly reduce the magnitude but preserve the wrong sign. Clipping individual monthly returns to the range -100 to +100 percent yields +0.280; winsorizing at 0.5 and 99.5 percent yields +0.238; winsorizing at 1 and 99 percent yields +0.192; and excluding securities with any monthly return above 1,000 percent yields +0.444. The evidence does not support the predicted negative-skew signature.

This result means the public-data evidence does not support the claim that system beta is a broad hidden tail-risk factor. The point estimates have the opposite sign, and the intervals remain wide enough to include zero.

8. P4: Covariance on the global coded-event calendar

The current calculation first takes the union of all months containing at least one selected filing anywhere in the panel. It then retains the available factor-residual return of every event-exposed security in every month on that global calendar, whether or not that security filed in that month. For each pair, the dependent variable is the sample covariance of those residual returns over the pair's shared available months. The regression compares the covariance of same-system and other pairs while controlling for a shared SIC code. This is a global event-calendar covariance diagnostic. It is not a test of simultaneous security-specific system events.

The constructed panel contains 418,155 overlapping security pairs involving 915 ticker symbols. Pairwise covariance is estimated from 39 to 233 common monthly observations, with a median of 108. The same-system coefficient is +0.0007711 in decimal-return-squared units. An actor-incidence sandwich and a 500-draw actor bootstrap account for securities repeated across pairs. They condition on the common event-month calendar and the first-stage factor residuals; they do not resample calendar months or carry first-stage estimation uncertainty into the pair regression. Because disjoint pairs share calendar observations and the dependent variables are estimated covariances, the coefficient is descriptive and no t-statistic or confidence interval is used for inference. A rebuilt test must define system-specific event dates and resample both issuers and calendar events.

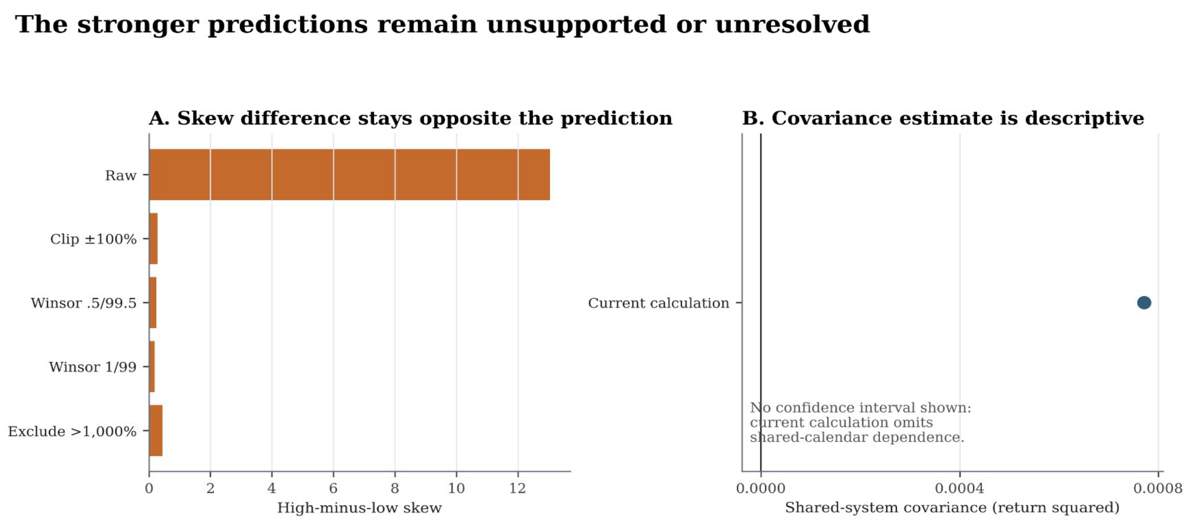

Figure 2. The negative-skew prediction receives no support across practical outlier treatments. The right panel shows only the exploratory covariance point estimate because the current uncertainty calculations do not cover common-calendar dependence.

9. Interpretation

The proposed signature has four parts: no ordinary-month premium, a selected-filing association, negative skew, and positive within-system covariance. Zero alpha is not rejected in any reported ordinary-month factor model. P2 is positive under maximum and mean multi-domain exposure assignments, then disappears under the median assignment, the opioid deletion, and the stricter single-domain audit. P3 points in the opposite direction across practical outlier treatments. P4 produces a positive descriptive coefficient from a design whose current uncertainty calculations are incomplete.

P1 is a failure to reject zero, not proof of informational exclusion. P2 is unresolved. The ticker-demeaned diagnostic preserves a positive association under one exposure rule while retaining an incomplete comparison panel and the event-selection, mapping, timing, and retrospective-measurement problems. P3 provides no support for a broad hidden tail-risk measure. P4 remains descriptive until a system-event design accounts jointly for issuer and calendar dependence.

10. A stronger follow-on test

The next test should reconstruct the design around verified, issuer-specific information events. Each event should have a dated announcement, checked source text, a stated connection to expected cash flows, and one unique issuer-event observation. Standard short-window abnormal returns should be estimated against a prespecified benchmark with dependence-appropriate inference. The current security-specific ordinary-month construction should be retained and compared with value-weighted portfolios.

Exposure timing also needs repair. A prospective design should use welfare or physical-exposure information available before each return window. A retrospective design remains useful for discovery, but its claim must remain retrospective. The strongest follow-on would compare system beta with established emissions, pollution, litigation, social-norm, attention, and transition-risk measures, then test whether the system measure adds out-of-sample information.

11. Current evidence boundary

The corrected run is a public-data discovery analysis with documented inputs, exclusions, model variants, influence tests, and an independent code audit. Its present constraints are binding: retrospective exposure measurement, an unaudited heuristic security crosswalk, repeated securities for some issuers, equal-weighted portfolios, selected filing items that need not concern system welfare, nine complete hand-entered observations, impossible public-return values, missing delisting returns, and incomplete calendar dependence in P4. A journal submission requires a point-in-time issuer and security map, filing-text validation, a conventional event reconstruction, cleaned licensed returns, and reproduction from a frozen public package.

12. Conclusion

This study reports a conditional failure to reject zero ordinary-month alpha, a specification-sensitive filing-month association, no support for the predicted negative skew, and a descriptive covariance coefficient. The ordinary-month result recurs across CAPM, Fama–French three-factor and five-factor models, and momentum. The filing-month association does not survive reasonable changes in exposure assignment and sample construction. P4's present uncertainty estimates do not cover its shared calendar.

These findings do not establish historical exclusion, causal repricing, or a new priced factor. They show which parts of a cross-domain return diagnostic survive the current public build and identify why the apparent positive results remain unsettled. The next version must use audited point-in-time issuer mappings, one common-equity security per issuer, checked announcement content, standard abnormal-return windows, licensed returns, and a frozen public replication package.

Data and code availability

The Version 1.1 paper, data documentation, code, frozen inputs, result files, software environment, and checksum manifest are assembled in the version-matched self-service release package distributed with this manuscript. The reserved public paper page is https://decisionaccounting.org/study/when-system-harm-reaches-asset-prices/. The reserved public code page is https://decisionaccounting.org/research/when-system-harm-reaches-asset-prices/code/. These URLs must not be described as live until the package is staged and the links are verified.

Statements and declarations

Funding. This research received no external funding.

Competing interests. The author founded the Center for Decision Accounting and developed the System Asset Pricing Model used in this paper. The author may benefit professionally, reputationally, or commercially from further research or adoption of this work. The Center is independent; no affiliation or endorsement by Microsoft, Anthropic, Google, NIST, or any other organization is implied.

Generative-AI disclosure. The author used generative artificial intelligence tools for research assistance, code review, source checking, editorial revision, and document production. The author directed the work, checked the reported claims and results, and accepts responsibility for the manuscript.

Data ethics. The study uses public company, filing, market, and factor data. It contains no human-subject data.

References

Allcott, H., Montanari, G., Ozaltun, B., & Tan, B. (2023, revised April 2025). An economic view of corporate social impact. *NBER Working Paper* No. 31803. https://doi.org/10.3386/w31803

Aronow, P. M., Samii, C., & Assenova, V. A. (2015). Cluster-robust variance estimation for dyadic data. *Political Analysis*, 23(4), 564–577. https://doi.org/10.1093/pan/mpv018

Arrow, K. J. (1969). The organization of economic activity: issues pertinent to the choice of market versus non-market allocation. In *The Analysis and Evaluation of Public Expenditure* (pp. 47-64). U.S. Government Printing Office.

Ardia, D., Bluteau, K., Boudt, K., & Inghelbrecht, K. (2023). Climate change concerns and the performance of green versus brown stocks. *Management Science*, 69(12), 7607–7632. https://doi.org/10.1287/mnsc.2022.4636

Bolton, P., & Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, 142(2), 517–549. https://doi.org/10.1016/j.jfineco.2021.05.008

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance*, 52(1), 57–82. https://doi.org/10.1111/j.1540-6261.1997.tb03808.x

Choi, D., Gao, Z., & Jiang, W. (2020). Attention to global warming. *The Review of Financial Studies*, 33(3), 1112–1145. https://doi.org/10.1093/rfs/hhz086

Engle, R. F., Giglio, S., Kelly, B., Lee, H., & Stroebel, J. (2020). Hedging climate change news. *The Review of Financial Studies*, 33(3), 1184–1216. https://doi.org/10.1093/rfs/hhz072

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2), 383–417. https://doi.org/10.1111/j.1540-6261.1970.tb00518.x

Fama, E. F. (1991). Efficient capital markets: II. *The Journal of Finance*, 46(5), 1575–1617. https://doi.org/10.1111/j.1540-6261.1991.tb04636.x

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56. https://doi.org/10.1016/0304-405X(93)90023-5

Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), 1–22. https://doi.org/10.1016/j.jfineco.2014.10.010

Garel, A., Romec, A., Sautner, Z., & Wagner, A. F. (2024). Do investors care about biodiversity? *Review of Finance*, 28(4), 1151–1186. https://doi.org/10.1093/rof/rfae010

Hong, H., & Kacperczyk, M. (2009). The price of sin: The effects of social norms on markets. *Journal of Financial Economics*, 93(1), 15-36.

Hsu, P.-H., Li, K., & Tsou, C.-Y. (2023). The pollution premium. *The Journal of Finance*, 78(3), 1343–1392. https://doi.org/10.1111/jofi.13217

Jensen, M. C. (1978). Some anomalous evidence regarding market efficiency. *Journal of Financial Economics*, 6(2-3), 95-101.

Kothari, S. P., & Warner, J. B. (2007). Econometrics of event studies. In B. E. Eckbo (Ed.), *Handbook of Empirical Corporate Finance* (Vol. 1, pp. 3–36). Elsevier. https://doi.org/10.1016/B978-0-444-53265-7.50015-9

Krüger, P. (2015). Corporate goodness and shareholder wealth. *Journal of Financial Economics*, 115(2), 304–329. https://doi.org/10.1016/j.jfineco.2014.09.008

Lindahl, E. (1919). Just taxation - a positive solution. In R. A. Musgrave & A. T. Peacock (Eds.), *Classics in the Theory of Public Finance* (pp. 168-176). International Economic Association.

MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature*, 35(1), 13–39. https://doi.org/10.2307/2729691

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica*, 41(5), 867-887.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *The Journal of Finance*, 42(3), 483–510. https://doi.org/10.1111/j.1540-6261.1987.tb04565.x

Pástor, Ľ., Stambaugh, R. F., & Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of Financial Economics*, 142(2), 550–571. https://doi.org/10.1016/j.jfineco.2020.12.011

Pástor, Ľ., Stambaugh, R. F., & Taylor, L. A. (2022). Dissecting green returns. *Journal of Financial Economics*, 146(2), 403–424. https://doi.org/10.1016/j.jfineco.2022.07.007

Postnieks, E. (2026). *The Missing System Theorem*. Center for Decision Accounting working paper. https://decisionaccounting.org/proof/

Shumway, T. (1997). The delisting bias in CRSP data. *Journal of Finance*, 52(1), 327-340.

Zhang, S. (2025). Carbon returns across the globe. *The Journal of Finance*, 80(2), 615–645. https://doi.org/10.1111/jofi.13402

Appendix A. Reproducible estimating procedures

This appendix describes the procedures used in the current build. A version-matched replication package accompanies the manuscript and contains the frozen inputs, outputs, code, environment specification, manifest, and checksums needed to reproduce the headline calculations without contacting the author. The reserved website routes will become the public distribution point after the package is staged and verified.

A.1 Exposure snapshot and issuer mapping

The numerical input is frozen before returns are joined. The current CSV contains 58 domain estimates and records a source-file path, snapshot date, and file hashes. It does not yet contain a row-level admission decision, evidence receipt, or version identifier. The paper therefore treats the file as a reproducible numerical snapshot rather than a complete measurement authority.

Security assignment starts with SEC ticker and SIC metadata. Prespecified SIC ranges and company-name or SIC-description substrings assign ticker symbols to domains. The current export does not preserve a reliable trigger field for every match: 27.1 percent of its rows are keyword-only, and audit examples show false positives caused by substrings. The 4,922 symbols correspond to at most 4,118 CIK issuers, so some issuers enter through more than one security. For a symbol mapped to several domains, the principal specification uses the maximum mapped estimate. Sensitivity specifications use the mean, median, percentile rank, logarithm of one plus the estimate, winsorized maximum, single-domain securities only, and leave-one-domain-out samples. No assignment rule uses returns. These mappings remain provisional until a row-level audit identifies the rule and supporting text for every assignment and selects one common-equity security per issuer.

A.2 Returns, factors, and portfolios

Adjusted monthly closes are converted to simple monthly returns. The Ken French Data Library supplies the market excess return, SMB, HML, RMW, CMA, momentum, and risk-free series. RMW is the return on robust-profitability firms minus weak-profitability firms. CMA is the return on conservatively investing firms minus aggressively investing firms. The independent model ladder estimates CAPM, Fama–French three-factor, Carhart four-factor, and Fama–French five-factor plus momentum specifications.

The frozen maximum-exposure sort places the highest 20 percent of mapped ticker symbols in the high group and the lowest 20 percent in the low group. Before monthly return availability and security-specific event exclusions, the sort contains 985 ticker symbols in each group. For month t, each leg is the equal-weighted return of group securities with an available return after removing only securities that have their own selected filing or hand-entered event in month t. The tested portfolio return is the high return minus the low return. The intercept in each factor model is P1's alpha. Standard errors allow serial correlation through six monthly lags. Because membership is based on securities listed in 2026, this is not a point-in-time investable portfolio and does not represent delisted issuers.

A.3 Selected-filing panel and P2

The selected-filing panel retains SEC Item 2.05, SEC Item 2.06, and the hand-entered observations described in Section 4. Generic Item 8.01 is excluded. Item 2.05 identifies costs associated with exit or disposal activities, and Item 2.06 identifies material impairments; neither item by itself establishes a system-welfare event. The event key is reduced to one ticker-month observation before estimation. The dependent variable is the absolute value of the security's monthly excess return, |R_it - R_ft|. The principal regression includes the exposure measure, market and style factors, momentum, and the three author-constructed controls. Standard errors are clustered by ticker and calendar month.

The independent audit applies a narrower sample and a broader external control set. It keeps only single-domain securities and adds the absolute market return, prior 12-month security volatility, event-type indicators, and year indicators to the Fama–French five factors and momentum. The paper reports both analyses because the difference between them determines the scientific conclusion.

A.4 Within-security lag diagnostic

For each selected ticker-month, the diagnostic seeks the same security's observation 12 months earlier and then 24 months earlier. The resulting panels are unbalanced: the 12-month calculation contains 1,703 selected observations and 1,665 lag observations, while the 24-month calculation contains 1,703 selected observations and 1,562 lag observations. Some ticker-months appear in both the selected and lag arms. Demeaning within security removes each security's time-invariant average. The coefficient of interest is the interaction between system beta and an indicator for the selected observation. Because the panels are incomplete and the arms overlap, this is not a matched-pair design. The calculation is a diagnostic pending reconstruction with complete nonoverlapping pairs.

A.5 Skew uncertainty

P3 computes sample skewness separately for the monthly equal-weighted high- and low-exposure portfolio returns and reports the high-minus-low difference. It does not estimate the cross-sectional distribution of security-level skewness. Sampling uncertainty is measured in two ways: 1,000 bootstrap samples of securities and 2,000 circular block-bootstrap samples using 12-month blocks. The outlier checks clip returns at -100 and +100 percent, winsorize individual returns at 0.5/99.5 and 1/99 percent, winsorize monthly portfolio-leg returns, and exclude securities with any monthly return above 1,000 percent. These checks diagnose sensitivity; they do not repair corporate-action errors or replace licensed return data with delisting coverage.

A.6 Exploratory covariance calculation

P4 begins with the union of all 233 calendar months in which any sample security has a selected filing. Every ticker symbol that appears at least once in the selected-filing universe contributes its available factor-residual return in every month on that global calendar. A paired security need not have a filing in the month used to estimate covariance. For each of 418,155 security pairs with at least six shared observations, the dependent variable is the sample covariance of those residual returns. Pairwise estimates use 39 to 233 common monthly observations, with a median of 108. The regression includes an indicator for a shared maximum-beta system label and an indicator for a shared SIC code. Its coefficient is expressed in decimal-return-squared units.

The current code calculates an actor-incidence sandwich estimator and a 500-draw Poisson actor bootstrap. Those procedures address repeated securities across pairs. They do not resample the shared calendar months, propagate uncertainty from the first-stage factor residuals, or account for the fact that the regression outcome is itself an estimated covariance. The reported coefficient is therefore descriptive; its calculated t-statistic and interval are retained in the computational record and are not used for inference. A release-grade P4 test requires prespecified system-specific event dates and resampling over both issuers and calendar events.

Acknowledgments

The author thanks the U.S. Securities and Exchange Commission and the Ken French Data Library for maintaining public data used in this study. Computational reviews tested the analysis from separate scripts and preserved conflicting results for comparison.

Version history

Version 1.1, July 28, 2026. Release-candidate revision. It reports the current-listed-security ordinary-month analysis, selected-filing diagnostic, sensitivity and influence checks, and global-calendar covariance calculation. The revision states the exact portfolio cutoffs, event-count stages, return-magnitude outcome, P2 economic scale, beta-input vintage, P4 estimand, and remaining uncertainty limits. It also supplies a version-matched self-service package. The issuer crosswalk, return cleaning, filing validation, lag comparison, point-in-time exposure, and covariance-inference limits continue to define the next empirical build.