Stockholm School of Economics

Master of Science Thesis in Finance

Spring 2026

The PFAS Premium

Equity Market Pricing of Latent Environmental Liabilities

Marko Ivanovic

Supervisor: Professor Magnus Dahlquist

Department of Finance, Stockholm School of Economics

Abstract

Equity markets apply an annualized PFAS premium of −5.28% (block bootstrap $p = 0.005$) to PFAS-implicated firms over 1998–2026, a realized cumulative shareholder-value differential of approximately −USD 2.1 trillion relative to the non-PFAS benchmark (95% CI [−3.5, −0.6] trillion). The premium is the risk-adjusted alpha of a calendar-time long-short strategy that holds a value-weighted portfolio of 74 PFAS-implicated firms against the broad non-PFAS universe, with each firm entering the long leg only once it has been publicly identified as PFAS-implicated. Portfolios are formed within each Fama-French factor region and pooled by market capitalization, and the long-short return is benchmarked against the Fama and French (2015) five-factor and Carhart (1997) momentum model.

The average per-event differential is −0.24pp over the PFAS event window ($p = 0.024$) across 134 PFAS-related scientific, regulatory, and litigation events, accounting for only about 22% of the long-term premium.

The premium has a three-regime structure identified through data-driven Bai and Perron (2003) structural-break detection, not imposed ex ante. This decomposition shows heterogeneous PFAS-event reactions. The pre-2015 and post-2023 regime reactions are insignificant, while the middle regime (2015–2022) shows weakly negative reactions at announcement (−0.27pp, $p = 0.059$). Both the middle and post-2023 regimes exhibit long-horizon drift, although in opposite directions.

Three findings frame the PFAS premium. First, it accumulates continuously across the sample, with the majority falling outside announcement windows. Second, the premium is a cost-of-capital differential relative to the non-PFAS portfolio, not absolute underperformance. The PFAS-implicated firm portfolio earns a cumulative return of 135% in excess of the risk-free rate, against 243% for the non-PFAS portfolio. Capitalized as a perpetuity, the −5.28% premium implies a forward-looking cost-of-capital discount equivalent to approximately 47% of PFAS-implicated equity value, or roughly USD 1.7 trillion. This is about eight times the intangible-liability discount documented by Konar and Cohen (2001), on a comparable basis. Third, the sign of the PFAS premium is opposite to other environmental premia documented in the literature, which can be reconciled through Pástor, Stambaugh, and Taylor's (2022) expected-versus-realized-returns framework, or through van der Beck's (2021) flow-based pricing framework.

Keywords: PFAS, environmental premium, event study, calendar-time long-short portfolio, institutional cycle, latent liability

Author: Marko Ivanovic

Supervisor: Professor Magnus Dahlquist, Department of Finance, Stockholm School of Economics

Examiner: Associate Professor Anders Anderson, Department of Finance, Stockholm School of Economics

Acknowledgement: I would like to thank my supervisor Professor Magnus Dahlquist for the guidance throughout this thesis.

Contents

1 INTRODUCTION **1**

2 INSTITUTIONAL BACKGROUND **2**

3 LITERATURE REVIEW **4**
3.1 PFAS AS AN ENVIRONMENTAL LIABILITY WITH ECONOMIC COSTS 4
3.2 PRICING OF ENVIRONMENTAL LIABILITIES 5
3.3 THE TIMING OF PFAS LIABILITY REPRICING 8

4 HYPOTHESES DEVELOPMENT **9**
4.1 LONG-TERM PRICING 9
4.2 SHORT-TERM PRICING 10
4.3 RETURN DRIFT: EFFICIENT MARKETS, PERSISTENCE, OVERREACTION, OR REGIME DEPENDENCE 10
4.4 CROSS-SECTIONAL HETEROGENEITY 11

5 DATA **11**
5.1 PFAS-IMPLICATED FIRM SAMPLE 11
5.2 EVENT DATABASE 12
5.3 EQUITY RETURNS 13
5.4 RISK FACTORS 14
5.5 AUXILIARY DATA 14

6 METHODOLOGY **15**
6.1 LONG-TERM PRICING DESIGN (H1) 15
6.2 SHORT-TERM PRICING TEST AND REGIME DECOMPOSITION (H2) 17
6.3 PERSISTENCE ACROSS SUB-REGIMES (H3) 23
6.4 CROSS-SECTIONAL HETEROGENEITY (H4) 24

7 RESULTS **25**
7.1 LONG-TERM PRICING: THE PFAS PREMIUM (H1) 25
7.2 SHORT-TERM PRICING: REGIME DECOMPOSITION AND AGGREGATE MAGNITUDES (H2) 30
7.3 MULTI-HORIZON PERSISTENCE (H3) 34
7.4 CROSS-SECTIONAL HETEROGENEITY (H4) 35

8 DISCUSSION **38**
8.1 THE PFAS PREMIUM IS CONTINUOUSLY PRICED, NOT CONCENTRATED AT EVENTS 38
8.2 PRICING INTENSITY IS NON-STATIONARY ACROSS REGIMES 41
8.3 THE REGIME STRUCTURE IDENTIFIES INSTITUTIONAL PRICING INTENSITY, NOT FIRM OR EVENT COMPOSITION 44

9 LIMITATIONS **44**
9.1 PRIMARY METHODOLOGICAL CAVEATS 45
9.2 SAMPLE COMPOSITION AND COVERAGE 46
9.3 EVENT DATABASE CODING AND FIY IDENTIFICATION 47
9.4 GENERALIZABILITY 48

10 CONCLUSION **48**

11 REFERENCES **51**

APPENDIX **61**
A. SAMPLE AND DATA CONSTRUCTION 62
B. PROPENSITY-SCORE MATCHING 68
C. DESCRIPTIVE STATISTICS 74
D. LONG-TERM PRICING (H1) ROBUSTNESS 78
E. SHORT-TERM PRICING (H2) ROBUSTNESS 83
F. STRUCTURAL-BREAK IDENTIFICATION 87
G. AI DISCLOSURE 89

List of tables

Table I: Pooled and region long-short regression models 26
Table II: Per-event differential decomposition by Bai and Perron (2003) regimes 30
Table III: Multi-horizon persistence of post-event abnormal returns 34
Table IV: Nested OLS regression of firm-event CARs on firm and event characteristics 36
Appendix Table 1: Hypothesis resolutions 61
Appendix Table A1: IID selection rules 63
Appendix Table A2: IID switches per year 64
Appendix Table A3: Identifier switches for the seven largest PFAS-implicated firms 65
Appendix Table A4: Daily-return cleaning 67
Appendix Table B1: Propensity model and common support 69
Appendix Table B2: Covariate balance under the baseline PSM 70
Appendix Table B3: Matched-sample levels and industry composition 71
Appendix Table B4: Control-pool composition 72
Appendix Table B5: Region composition of matched pairs 73
Appendix Table B6: Robustness to the matching design 74
Appendix Table C1: PFAS versus non-PFAS descriptive statistics 75
Appendix Table C2: PFAS-implicated firms' descriptive statistics 76
Appendix Table C3: Event database 77
Appendix Table C4: Cross-sectional descriptive statistics of firm-events used for H4 77
Appendix Table C5: H4 pairwise Pearson correlation matrix 78
Appendix Table D1: Alternative factor models 79
Appendix Table D2: Jensen, Kelly, and Pedersen (*2023*) factor themes 80
Appendix Table D3: H1 long-short alpha under alternative portfolio construction 82
Appendix Table D4: Chow (1960) structural-break test on the pooled long-short alpha 83
Appendix Table E1: Baseline differential by sample definition 83
Appendix Table E2: Parameter-choice sensitivity 84
Appendix Table E3: Matched-control falsification 84
Appendix Table E4: Leave-one-out and restructuring sensitivity 85
Appendix Table E5: Cross-jurisdiction matrix 86
Appendix Table E6: Event-dating legitimacy 86
Appendix Table F1: Single-break diagnostics on the annual PFAS-event count, 2001–2026 87
Appendix Table F2: Pettitt non-parametric change-point test by event-type subset 88
Appendix Table F3: Bai and Perron (2003) BIC profile across 0–7 breaks 89

List of figures

Figure 1: Cumulative return of portfolios and rolling annualized 36-month alpha 28
Figure 2: Per-event CAR scatter plot across Bai and Perron (2003) regimes 31
Figure 3: Multi-horizon drift profiles across Bai and Perron (2003) regimes 35
Figure 4: Cross-sectional determinants of CARs 37
Figure 5: Events timeline and decomposition 42
Appendix Figure A1: IID selection algorithm performance 65
Appendix Figure D1: Alternative-factor-model robustness of the H1 long-short alpha 81
Appendix Figure F1: Chow (1960) *F*-statistic profile across candidate break years, by event-type subset 88

List of abbreviations

Abbreviation	Definition
AFFF	Aqueous film-forming foam
AQR	AQR Data Library (factor-data source)
BAB	Betting against beta
BIC	Bayesian Information Criterion
BM	Book-to-market ratio (book equity/market equity)
BMP	Boehmer, Musumeci, and Poulsen (1991) test
CAPM	Capital Asset Pricing Model
CAR	Cumulative abnormal return
CFC	Chlorofluorocarbon
CI	Confidence interval
CMA	Conservative minus aggressive (investment factor)
CRSP	Center for Research in Security Prices
ECHA	European Chemicals Agency
EMH	Efficient market hypothesis
EPA	United States Environmental Protection Agency
ESG	Environmental, social, and governance
EU	European Union
FDR	False discovery rate
FF3	Fama and French (1993) three-factor model
FF3 + MOM	Fama and French (1993) three-factor model and Carhart (1997) momentum model
FF5	Fama and French (2015) five-factor model
FF5 + MOM	Fama and French (2015) five-factor and Carhart (1997) momentum model
FIY	First identified year
GMB	Green minus brown
GVKEY	Compustat firm identifier
HML	High minus low (value factor)
IID	Issue identifier
MDL	Multidistrict litigation
MOM	Momentum factor
MTBE	Methyl tert-butyl ether
OLS	Ordinary least squares
PCB	Polychlorinated biphenyl
PFAS	Per- and polyfluoroalkyl substances
PFCA	Perfluorocarboxylic acid
PFHxS	Perfluorohexane sulfonic acid
PFOA	Perfluorooctanoic acid
PFOS	Perfluorooctane sulfonic acid

PSM	Propensity score matching
QMJ	Quality minus junk
REACH	Registration, Evaluation, Authorisation and Restriction of Chemicals (EU regulation)
RMW	Robust minus weak (profitability factor)
ROA	Return on assets (net income/total assets)
SAR	Standardized abnormal return
SIC	Standard Industrial Classification
SMB	Small minus big (size factor)
SMD	Standardized mean difference
VIF	Variance inflation factor
WRDS	Wharton Research Data Services

1 Introduction

Have equity markets repriced PFAS-implicated firms? Yes. This thesis identifies a risk-adjusted annualized PFAS premium of −5.28% (block bootstrap $p = 0.005$) over 1998–2026 across a sample of 74 PFAS-implicated firms, representing a realized cumulative shareholder-value differential of approximately −USD 2.1 trillion relative to the non-PFAS benchmark (95% CI [−3.5, −0.6] trillion). The pooled per-event differential between PFAS-implicated firms and matched controls is −0.24pp ($p = 0.024$) across 134 scientific, regulatory, and litigation events. An event-study lens detects the effect but understates its magnitude by roughly fivefold, as these event windows account for only about 22% of the long-term premium, with the remaining 78% accumulating continuously between events. The premium operates as a cost-of-capital differential, not absolute underperformance. The PFAS-implicated firm portfolio earns 135% cumulative excess returns against 243% for the non-PFAS portfolio. Capitalized as a perpetuity, the premium implies a forward-looking cost-of-capital discount equivalent to approximately 47% of PFAS-implicated equity value, or roughly USD 1.7 trillion.

The premium is non-stationary and operates through three regimes (pre-2015, middle, and post-2023) that align with institutional cycles and show sharply heterogeneous pricing. Per-event announcement differentials are insignificant in the pre-2015 and post-2023 regimes and marginally negative in the middle regime. Long-horizon drift appears in the latter two regimes, consistent with the slow-information-diffusion literature (Hong and Stein, 1999), but operates in opposite directions and nets out across the pooled sample. The premium is not explained by firm or event characteristics across the pooled cross-section, with the sole exception of two mega-settlement events whose positive reactions are based on a small sample and cannot be generalized.

This thesis tests four hypotheses: a long-term pricing effect (H1), a per-event reaction differential (H2), multi-horizon persistence (H3), and cross-sectional heterogeneity (H4). The empirical strategy is two-pronged. The PFAS premium is identified from a calendar-time, value-weighted long-short portfolio of 74 PFAS-implicated firms against the broad non-PFAS universe. The long-short portfolio is formed within each Fama-French factor region, pooled by market capitalization, and benchmarked against the Fama and French (2015) five-factor and Carhart (1997) momentum model (FF5 + MOM), using a stationary block bootstrap for inference. Per-event reactions are computed across 134 PFAS-related scientific, regulatory, and litigation events. Each PFAS firm-event is differenced against propensity-score-matched non-PFAS control firms on a three-day event window. Regime breaks are identified through Bai

and Perron (2003) data-driven structural-break detection (not imposed ex ante) and used in the event study. Persistence is tested at five horizons up to 365 days, for the full sample and within each regime, using cluster-bootstrap inference.

The contributions span four strands of literature. First, the thesis documents a three-regime institutional cycle in the PFAS-event setting, using the Bai and Perron (2003) structural breaks. Second, it extends the latent-liability event-study literature. The per-event reactions sit at the smaller end of the magnitudes that Karpoff, Lott, and Wehrly (2005) document for environmental-violation announcements, yet the capitalized premium is roughly eight times the intangible-liability estimate of Konar and Cohen (2001), on a comparable basis. The thesis also tests the litigation commencement-versus-conclusion asymmetry of Karpoff, Lott, and Wehrly (2005) and finds it only in the two mega-settlement events, not in the pooled sample. Third, it extends the environmental premia literature. The PFAS premium is opposite in sign to the premia that Hsu, Li, and Tsou (2023) and Bolton and Kacperczyk (2021, 2023) find in other environmental settings. This opposite sign has two possible explanations. *(i) Risk-based pricing:* the negative news that has characterized the PFAS institutional environment created negative shocks that outweigh the positive expected return that investors require as compensation for bearing the risk, in the style of the expected-versus-realized-returns framework of Pástor, Stambaugh, and Taylor (2022). *(ii) Flow-based pricing:* there has been an institutional investor outflow from PFAS-implicated assets, which created downward pressure on the returns, in the style of van der Beck (2021). In addition, the thesis sidesteps the Aswani, Raghunandan, and Rajgopal (2024) critique of Bolton and Kacperczyk's (2021, 2023) methodology by using the public, categorical, and non-vendor-estimated ChemSec (2025) database. Fourth, it extends the announcement-premium decomposition literature, with findings that contrast with the results of Savor and Wilson (2014) for the aggregate equity premium. Together, these four contributions amount to a multi-method decomposition of the PFAS premium, showing that the premium is continuous rather than event-driven, regime-dependent rather than stationary, and a cost-of-capital differential rather than absolute underperformance.

2 Institutional background

Per- and polyfluoroalkyl substances, or PFAS, is a collective name for a group of approximately 12,000 chemical compounds (Salvatore et al., 2022). There are several acronyms that fall under the PFAS umbrella, with major sub-classes including PFOA, PFOS, and PFHxS (Buck et al.,

2011). The chemicals are valued for their resistance to water, grease, and dirt, and have been used since the 1940s in applications such as non-stick cookware, Gore-Tex clothing, cosmetics, and firefighting foams (Glüge et al., 2020; Kurwadkar et al., 2022; Wee and Aris, 2023). However, scientists have warned of significant environmental and health risks. A growing body of studies shows that the accumulation of these chemicals in the human body is associated with immune suppression, liver and kidney disease, impaired reproductive function, and cancer (Fenton et al., 2021). PFAS is notoriously known as "forever chemicals" because it resists environmental degradation (Cousins et al., 2020), and it has been found in remote areas and water sources worldwide (Kurwadkar et al., 2022; Miner et al., 2021).

The environment surrounding PFAS has developed in parallel through scientific, regulatory, and litigation channels. Scientific research relating to health and environmental effects has accumulated over several decades. The C8 Science Panel, which ran from 2005 to 2013 and drew on a health study of over 69,000 participants, reported a probable link between PFOA, a type of PFAS, and several negative health effects (Frisbee et al., 2009; Steenland et al., 2020). Other research has mapped the environmental diffusion of PFAS, including its presence in drinking water (Andrews and Naidenko, 2020; Domingo and Nadal, 2019). There is an abundance of other literature relating primarily to health effects (Pelch et al., 2019).

The documented health effects of PFAS, together with litigation pressure, have resulted in new regulatory attempts (Klingelhöfer et al., 2024; Suffill et al., 2024), though their success is debated (Nevitt and Percival, 2022). In the United States, the EPA PFOA Stewardship Program began an industry phase-out in 2006, with several major companies voluntarily signing on (U.S. EPA, 2006). This was followed in 2024 by maximum contaminant levels for PFOA and PFOS in drinking water (U.S. EPA, 2024). The European Parliament and Council of the European Union (2006) introduced REACH in 2007, the framework under which PFOA and related substances were later restricted. In 2023, several European countries went further, proposing a broad PFAS ban (European Chemicals Agency, 2023). On a global level, the Stockholm Convention listed PFOS as a persistent organic pollutant in 2009, with PFOA, PFHxS, and PFCA subsequently added in 2019, 2022, and 2025, respectively (Stockholm Convention, n.d.).

The litigation channel began with the 2001 class-action filing of *Leach v. DuPont*, which led to a 2005 settlement that funded the C8 Science Panel (Frisbee et al., 2009). Various lawsuits relating to aqueous film-forming foam (AFFF) were consolidated into the Multidistrict Litigation (MDL) 2873 in 2018, which by May 2026 had grown to approximately 15,000

pending cases (U.S. JPML, 2026). A wave of settlements followed, the largest being 3M's USD 10.3 billion settlement (Reuters, 2023a) and the DuPont, Chemours, and Corteva settlement of USD 1.185 billion (Reuters, 2023b). This institutional development is examined through a theoretical lens in the following literature review, to assess whether, how, and when PFAS exposure is priced into asset values.

3 Literature review

This section synthesizes prior research to build a conceptual economic mechanism that links PFAS to asset pricing.

3.1 PFAS as an environmental liability with economic costs

PFAS has three characteristics that guide the economic interpretation and shape the asset-pricing predictions in this literature review. First, PFAS is a latent, slow-moving contaminant that accumulates as a persistent stock of pollutants (Cousins et al., 2020). The production of PFAS began in the 1940s, while understanding of its health effects emerged only from the 1990s onward, through studies such as the C8 Science Panel in 2005 (Frisbee et al., 2009; Steenland, Fletcher, and Savitz, 2010; Wang et al., 2017). This trend is analogous to other latent environmental contaminants such as asbestos, PCBs, MTBE, benzene, and CFCs, which the European Environment Agency reports all had a lag of 30 to 100 years between their introduction and the policy actions limiting them (Harremoës et al., 2001).

Second, firm-level exposure is heterogeneous and arises from the firm's own decision-making. Firms differ in the magnitude of their production volumes, application scope, and remediation responsibilities, as PFAS consists of approximately 12,000 different chemical compounds with varying toxicological and regulatory profiles (Cousins et al., 2020; Salvatore et al., 2022).

Third, given the heterogeneous exposure among firms, the liability distribution is also heterogeneous. Potential clean-up costs, litigation settlements, monitoring and compliance costs, and product and supply-chain redesign costs vary in size across firms (Konar and Cohen, 2001). Most firms face small future cash outflows, while a small number face enormous economic costs, with a tail risk extending far into the future as the institutional environment develops and new claims emerge. This skew is already evident empirically, with 3M's USD 10.3 billion settlement, the DuPont, Chemours, and Corteva settlement of USD 1.185 billion, the Tyco Fire Products settlement of USD 750 million, Carrier Global's USD 730 million settlement for its bankrupt Kidde-Fenwal subsidiary, and BASF's settlement of USD 316.5

million, all standing in contrast to the thousands of individual claims in MDL 2873 (BloombergNEF, 2024; Zarghamee and Bruno, 2023; Zarghamee, Brenner, and Fink, 2024). This same fat-tail pattern has been documented historically in asbestos litigation, where Carroll et al. (2005) report that 73 of the 8,400 named defendants filed for bankruptcy, some of them firms valued at more than USD 1 billion.

3.2 Pricing of environmental liabilities

Environmental liabilities are, in the simplest case, priced through Campbell's (1991) present-value identity, which decomposes into adjustments to expected future cash flows and adjustments to the discount rate applied to those cash flows. When new information about future PFAS-related liabilities arrives, such as clean-up costs, settlement payments, and lost sales from reputational damage, the market adjusts expected future cash flows accordingly. The same news may also change the perceived riskiness of the firm, which the market incorporates through a change in the discount rate. These two channels jointly produce immediate and full incorporation of any new information on the announcement day, assuming frictionless, efficient markets.

Multiple theories describe how the information environment creates frictions that impair the functioning of these two channels, reducing valuation precision. Below, these frictions are grouped into three categories: size uncertainty, distributional uncertainty, and regulatory uncertainty.

Size uncertainty pertains to a missing-information problem that primarily affects investors' ability to forecast the magnitude of expected future cash flows. PFAS, being latent and persistent and having costs that are difficult to measure (Cordner et al., 2021; Wang et al., 2017), adds to this size uncertainty. Literature on information economics and hidden-state models (Veronesi, 1999, 2000) explains this as a missing-input problem (Merton, 1987), in which investors must use limited information to infer expected cash flows (Grossman and Stiglitz, 1980). Investors are additionally hampered by incomplete disclosure of future obligations (Healy and Palepu, 2001) and by the high cost of acquiring precise information (Grossman and Stiglitz, 1980), which complicates cash-flow forecasting. The size-uncertainty friction therefore operates exclusively on the expected-cash-flow channel, not on the discount-rate channel.

Distributional uncertainty arises when the magnitude of the economic cost is relatively well understood but the distribution around it is not, leaving investors to face Knightian uncertainty rather than measurable risk (Knight, 1921). The highly variable nature of

environmental liabilities such as PFAS contributes to this uncertainty. Even if investors have reasonable knowledge of the approximate magnitude, the distribution around it remains poorly understood. This connects to well-established theories of ambiguity and uncertainty about probability distributions. Ambiguity theories show that investors often overweight negative scenarios when a distribution is poorly understood, primarily reducing expected cash flows but also raising the discount rate (Drechsler, 2013). Similarly, long-run risk models show that investors who do not understand the distribution assume a more fat-tailed one, which decreases expected cash flows and increases discount rates (Bansal and Yaron, 2004). Literature shows that when a liability arises from several sources with heterogeneous distributions, the combined distribution becomes wider, with a fatter tail than any of its components (Cooke, Nieboer, and Misiewicz, 2014), amplifying the ambiguity-aversion effect. This ambiguity aversion translates primarily into lower expected cash flows, as more negative scenarios receive higher weighting (Ju and Miao, 2012; Li, Tiwari, and Tong, 2017), and into higher discount rates, as investors add an ambiguity premium (Brenner and Izhakian, 2011; Rieger and Wang, 2012). In the PFAS setting, this operates through multiple heterogeneous distributions across litigation costs, clean-up costs, and operational costs, producing a liability distribution that is wider and fatter-tailed than any of its components.

The third friction to the information environment is regulatory uncertainty. Because investors do not know which regulatory regime will ultimately govern the treatment of PFAS liabilities, further uncertainty surrounds the valuation. Literature on political-regime uncertainty shows that shifts between lenient and strict policy regimes are priced through higher discount rates (Belo, Gala, and Li, 2013; Kelly, Pástor, and Veronesi, 2016; Pástor and Veronesi, 2012, 2013). Incomplete-markets theory similarly shows that regulatory uncertainty cannot be hedged, as it is a non-tradable risk, meaning that investors must adjust the discount rate to incorporate it in equilibrium (Merton, 1987; Pástor and Veronesi, 2012, 2013).

Several precedent papers provide a theoretical scaffolding for the asset-pricing prediction, in terms of both the sign and the magnitude of the premium for PFAS-implicated firms. First, the sin-stock literature of Hong and Kacperczyk (2009) shows that firms in sin industries earn risk-adjusted excess returns of approximately 3.5% per year (about 30 basis points per month) relative to comparable firms in the United States, for two reasons. The first is that, because they are screened out by norm-constrained institutional investors, the investors who do hold these stocks take larger, less diversified positions and demand higher compensation for the idiosyncratic risk. This aligns with Merton's (1987) segmented-markets

model, in which a larger investor base lowers an asset's expected return. The second is that investors demand higher compensation for the elevated litigation risk these firms face. Both channels apply directly to the PFAS setting. The ChemSec (2025) database of PFAS-implicated firms creates a similar norm-driven institutional avoidance, and the current PFAS litigation environment is analogous to the elevated litigation risk that Hong and Kacperczyk (2009) describe.

Second, the carbon-premium literature documents a more environmentally specific parallel. Bolton and Kacperczyk (2021) show that U.S. firms with higher carbon emissions over 2005–2017 earned higher returns, approximately 1.8–4.0% (depending on the emission scope) annualized per standard-deviation increase in scope 1, 2, or 3 emissions, while controlling for size, book-to-market, momentum, and other return predictors. They interpret this as compensation for bearing multiple sources of risk, including regulatory, commodity-price, and green-technology-shift risks. Bolton and Kacperczyk (2023) document the same positive association between carbon emissions and stock returns in a global cross-section of approximately 14,400 firms across 77 countries, holding across nearly all sectors and most countries. This establishes the carbon premium as a global phenomenon, which is the relevant precedent for the global PFAS sample studied here. The PFAS setting represents the same kind of latent environmental liability that drives the carbon premium described by Bolton and Kacperczyk (2021).

Third, Hsu, Li, and Tsou (2023) studied firms in the context of toxic emissions. They constructed a long-short portfolio sorted on toxic-emission intensity within industries and found an annual return of approximately 4.4% unexplained by standard risk factors. They attribute this to a pollution premium arising from environmental-policy uncertainty. By analogy to Bolton and Kacperczyk (2021, 2023) and Hong, Li, and Xu (2019), this thesis hypothesizes a non-zero risk-adjusted return for PFAS-implicated firms and tests for it empirically through Hypothesis 1 developed in Section 4.1.

Fourth, Pástor, Stambaugh, and Taylor (2021) show that investors' preferences may shift over time in relation to ESG, affecting stock performance. The non-zero risk-adjusted return for PFAS-implicated firms therefore need not remain constant across the sample period. This time-shifting notion is tested empirically over 1998–2026.

Fifth, the effect of litigation on stock returns is documented in multiple studies. Karpoff, Lott, and Wehrly (2005) studied 478 press announcements of environmental violations between 1980 and 2000. They found that initial allegations and settlements produced return reactions of

approximately −1.69% (significant at the 1% level) and −0.35% (significant at the 10% level), respectively. Furthermore, they found that the loss in equity market value is of similar magnitude to the legal penalties. They concluded that environmental violations are disciplined primarily through legal penalties, rather than reputational losses. In contrast, Karpoff, Lee, and Martin's (2008) study of 585 cases of corporate financial misrepresentation over 1978–2002 found that the reputational loss component was approximately 7.5 times the legal-penalty loss component. These two papers imply that the consequences of litigation in the PFAS setting have two components, a legal-penalty loss and a reputational loss, whose relative weights depend on the type of event (environmental versus corporate financial misrepresentation). Furthermore, Donelson, Flam, and Yust (2022) and Wang and Zhang (2023) document spillover effects through industry-wide reputational channels. These event-level litigation decompositions are tested empirically in the PFAS setting through the cross-sectional analysis of Hypothesis 4 developed in Section 4.4.

3.3 The timing of PFAS liability repricing

Fama's (1970, 1991) semi-strong-form efficient market hypothesis (EMH) predicts that public information is incorporated into asset prices immediately and without any exploitable post-event drift. The empirical record is mixed across decades and event types, but the EMH still provides a natural null hypothesis for a persistence test. In the absence of frictions, the announcement-window cumulative abnormal return (CAR) should capture the entire pricing reaction immediately. This announcement-window reaction is empirically tested in the event study of Hypothesis 2 developed in Section 4.2.

A body of literature documents market frictions that contradict the EMH. These frictions can be decomposed into short-term and long-term components. Some literature shows that information is only partially incorporated into prices in the short term (Bernard and Thomas, 1989; Jegadeesh and Titman, 1993), possibly because complex, technical information and limited investor attention cause investors to fail to incorporate the available information on the announcement day (Barber and Odean, 2008; Hirshleifer and Teoh, 2003; Hong and Stein, 1999; Lim and Teoh, 2010). In contrast, De Bondt and Thaler (1985) show that investors sometimes overreact to news and that these initial overreactions subsequently reverse. The PFAS setting's complexity, in terms of thousands of chemical compounds, technical scientific news that may require epidemiological knowledge to translate into firm liabilities, and a multi-geography regulatory environment, invites this kind of short-term market inefficiency. This thesis empirically tests for this drift through Hypothesis 3 developed in Section 4.3.

From a long-term perspective, the friction lies not in digesting the information on the announcement day but in the gradual formation of investor consensus around the asset class over many years. As news is released over multiple years, scientific, regulatory, and legal institutions converge around a new understanding of an environmental risk class, leading to a paradigmatic shift (Kuhn, 1962). Long-run risk models and climate-transition models argue that investors update their views gradually, in line with the evolution of the regulatory framework, over extended horizons (Bansal and Ochoa, 2011; Bolton and Kacperczyk, 2023). The regulatory-learning literature shows that regulatory and legal frameworks are updated gradually as new scientific evidence becomes available (Jasanoff, 1990). As described in the institutional background (Section 2), the regulatory environment has shifted substantively across the sample period, which may alter the cost-of-capital implications for PFAS-implicated firms. For PFAS, this framework suggests that the cumulative incorporation of stepwise information may be more important than the announcement-window reaction of any single event. The implication for the empirical design is that the assumption of stationary pricing across the full sample horizon of 2001–2026 may not hold, as the underlying institutional environment and investors' understanding of risk continue to evolve. This notion is tested with Bai and Perron's (2003) structural-break framework as part of Hypothesis 2 developed in Section 4.2.

4 Hypotheses development

Four hypotheses are developed, each empirically testing a theoretical prediction from the literature-review framework, with the purpose of understanding the relationship between PFAS and asset pricing.

4.1 Long-term pricing

H1: The long-short portfolio of PFAS-implicated firms minus a broad non-PFAS universe earns a non-zero risk-adjusted return over the period 1998–2026, after controlling for the standard Fama and French (2015) five factors and the Carhart (1997) momentum factor.

The liability-pricing framework described in Section 3.2 predicts that PFAS-implicated firms earn a non-zero long-term risk-adjusted return. By analogy, both Hong and Kacperczyk's (2009) sin-stock literature and Bolton and Kacperczyk's (2021, 2023) carbon-premium literature predict a positive PFAS premium. Hong and Kacperczyk's (2009) arguments about norm-constrained institutional investors screening out firms and about the heightened litigation

risk, as well as Bolton and Kacperczyk's (2021, 2023) argument that investors demand additional compensation for risk borne (including regulatory, transition, and green-technology-shift risk), apply directly to the PFAS setting described in Section 2. Furthermore, Pástor, Stambaugh, and Taylor's (2021) findings predict that this PFAS premium need not be constant across the sample period, in terms of both sign and magnitude.

4.2 Short-term pricing

H2: The pooled CAR across 134 events, for PFAS-implicated firms minus matched control firms, is on average negative in the three-day announcement window, with a sign and magnitude that may vary across structural regimes.

The theory outlined in Section 3.2 predicts that PFAS-related announcements conveying new information that crystallizes the environmental liability should generate negative abnormal returns in the three-day announcement window. An empirical precedent was established by Karpoff, Lott, and Wehrly (2005), who found negative abnormal returns around announcements of environmental violations, ranging from approximately −0.35% for settlements to −1.69% for allegations. Given that the institutional environment has shifted across 2001–2026, the sign and magnitude may vary across regimes.

4.3 Return drift: efficient markets, persistence, overreaction, or regime dependence

H3: The CAR profile over five post-event horizons (5, 30, 90, 180, and 365 trading days) is non-zero and may differ in both sign and magnitude across structural regimes.

The semi-strong form of Fama's (1970, 1991) EMH, discussed in Section 3.3, provides the natural null hypothesis, under which the announcement-day reaction is fully and immediately incorporated into asset prices within the announcement window and generates no predictable post-event return drift. Underreaction in short-term returns (Bernard and Thomas, 1989; Jegadeesh and Titman, 1993), potentially due to limited investor attention and slow diffusion of complex information (Hong and Stein, 1999), predicts a long-term return drift in the same direction as the initial reaction. De Bondt and Thaler (1985) and Daniel, Hirshleifer, and Subrahmanyam (1998) instead predict a post-event drift in the opposite direction, as the initial overreaction is subsequently corrected. Finally, Bai and Perron's (2003) structural-break framework (Section 3.3) raises a third possibility, in which return persistence is itself regime-dependent.

4.4 Cross-sectional heterogeneity

H4: The magnitude, and potentially the sign, of the announcement-window CAR differ across the three event types (scientific, litigation, and regulatory), with litigation and regulatory events producing more negative returns than scientific events, and across firm-level characteristics that proxy for the directness of cash flows.

The three event types differ in how directly they inform investors about cash flow effects. Litigation events often carry a penalty amount in the press release, and Karpoff, Lott, and Wehrly (2005) showed that the announcement-day loss in equity market value approximately equals the legal penalty. One can also distinguish between litigation commencement and conclusion, since case conclusions reduce uncertainty and thereby lower the discount rate. Regulatory events can have a direct effect on an industry, but their effects can be difficult to attribute to a specific firm, especially in a setting as complex as PFAS. Scientific events are the least direct, as any findings must pass through a further layer (regulatory or litigation) before they affect a specific firm's cash flows. In addition, limited investor attention further affects the speed with which investors update their understanding (Barber and Odean, 2008; Hirshleifer and Teoh, 2003; Lim and Teoh, 2010).

5 Data

The data used for this thesis consists of a ChemSec-provided (2025) database of PFAS-implicated firms, a manually compiled event database, S&P Global Market Intelligence Compustat (hereafter Compustat) (S&P Global Market Intelligence, 2026) equity returns accessed through Wharton Research Data Services (WRDS, 2026), French (2026) risk factors, and various auxiliary data.

5.1 PFAS-implicated firm sample

The cleaned dataset of PFAS-implicated firms consists of 74 publicly listed firms identified by the non-profit organization ChemSec (2025), which tracks PFAS-implicated firms globally. The firms are sourced from three annual lists compiled by ChemSec (2025) and published in 2016 (29 firms, after cleaning), 2020 (37 firms, after cleaning), and 2025 (74 firms, after cleaning). The ChemSec (2025) data includes both public and private firms, as well as listed private-equity firms with PFAS-implicated portfolios. The latter two categories are dropped from the list in the cleaning process. The private firms are dropped because they do not have observable daily returns. The listed private-equity firms are dropped for two reasons. First, the

PFAS implication operates at the portfolio-company level, not at the parent level, and the diversified nature of private-equity firms' portfolios means that any PFAS-related cash-flow effects are difficult to isolate, as they pertain only to a small fraction of their total portfolio. Second, listed private-equity firms are classified as financial firms, which are also screened out due to their different fundamental structures.

By the Standard Industrial Classification (SIC) of the U.S. Office of Management and Budget (1987), the sample is concentrated in manufacturing (63 of 74 firms) and, within manufacturing, in chemicals and allied products (SIC 28, 46 firms). These chemical firms span the value chain, from industrial organic and inorganic chemicals through plastics and synthetic resins, paints, and agricultural chemicals to pharmaceutical preparations. The remaining firms form a thin manufacturing tail in machinery, electronics, glass, petroleum refining, and instruments, alongside six wholesale chemical distributors (SIC 50 and 51), two life-science service firms (SIC 80 and 87), two diversified industrial conglomerates, and one agricultural producer. By country of incorporation, the sample spans 22 North American firms, 21 European firms, 14 Japanese firms, and 2 Asia-Pacific firms (Australia and Hong Kong). A further 15 emerging-market firms are pooled into an emerging market bucket, following Fama and French (2012) and Jensen, Kelly, and Pedersen (2023).

In addition to the publication years provided by ChemSec (2025) annual publications in 2016, 2020, and 2025, a first identified year (FIY) is established for each PFAS-implicated firm as the earliest date at which the market could publicly have known that the firm was PFAS-implicated. The FIY is defined as the earliest of: *i)* the earliest appearance of the firm in the ChemSec (2025) annual lists released in 2016, 2020, and 2025; *ii)* a manually identified earlier disclosure year obtained from 10-K and 10-Q filings, sustainability reports, news archives, corporate websites, and scientific literature naming the firm or any subsidiary of the parent. The FIY is also floored at the 1998 panel start, set by data availability with the event data beginning in 2001, and for firms incorporated after 1998, at the date their pricing data begins.

5.2 Event database

The event database consists of 144 manually compiled PFAS-related events from 2001 to 2026, classified into three event categories that correspond to the institutional channels described in Section 2. The database consists of 53 scientific events compiled from peer-reviewed studies, 45 regulatory events compiled from government rule proposals, suggestions, recommendations, and enforcements, and 46 litigation events compiled from case filings, settlements, and verdicts. The scientific events are sourced from peer-reviewed journals such

as *Environmental Health Perspectives*, *Environmental Science & Technology*, and *Environment International*, as well as the C8 Science Panel. The regulatory events are sourced from various levels of regulators, such as the U.S. EPA, U.S. state environmental agencies, the European Chemicals Agency (ECHA), the European Food Safety Authority, the European Commission and Parliament, the Stockholm Convention, and other national ministries. The litigation events are sourced from news agencies (Reuters, the New York Times, and the Washington Post), law-firm publications, and legal databases. Each event is assigned a date, the earliest publication date found. Out of these events, 41 name a firm from the ChemSec (2025) list of PFAS-implicated firms and are mapped to those firms. The other 103 events are non-firm-specific and are mapped to all firms. For the litigation events specifically, the events are coded as case commencement or case conclusion, and the penalty or settlement amount is logged where available. The final event dataset covers 57 U.S. events, 14 EU-level events, 11 events specific to individual European countries, 5 Canadian events, 5 Australian events, and 52 events coded as global. Of the 144 events, 10 are dropped in the filtering process described in Section 6.2, leaving 134 events in the database for the analyses.

5.3 Equity returns

Daily equity-price data is sourced for the whole firm universe in the period 1998 through April 2026 from Compustat (2026) Global and North America, both accessed through Wharton Research Data Services (WRDS) (WRDS, 2026). A filter is applied to all firms to include only common and ordinary shares with a main-exchange listing, following the literature (Ince and Porter, 2006; Jensen, Kelly, and Pedersen, 2023), and exclude financial firms, following Fama and French (1992), due to their different leverage profile. All firms are merged on GVKEY. In the few cases of duplicate GVKEYs across the two databases (Compustat (2026) Global and North America), the North American observation is preferred, with both files read beforehand to ensure that cross-listed firms retain full coverage. As a single GVKEY can have multiple issue identifiers (IID), whether from multiple listed securities, dual listings, or a change of IID over time, a rule-based algorithm constructs a chain of the single best IID for each GVKEY, working backwards from April 2026 and switching IID when one disappears. Constructing the IID chain backward avoids look-ahead bias (Banz and Breen, 1986). The algorithm is described in Appendix A.1. From each GVKEY's IID chain, returns are calculated from day-to-day changes in the split-adjusted closing price. The returns are cleaned using a 13-step rule that identifies breaks such as changes in IID, currency switches, stale prices, or implausible moves (Ince and Porter, 2006). The returns are winsorized at the 0.1st and 99.9th percentiles (Jensen,

Kelly, and Pedersen, 2023). Delisted firms are corrected by implementing a −30% return (Shumway, 1997). Refer to Appendix A.2 for the full list of cleaning rules. The cleaned panel contains 87,229 unique GVKEYs, of which 74 are PFAS-implicated firms and the rest constitute the pool of available control firms.

5.4 Risk factors

Monthly FF5 + MOM factor returns are sourced by region from the French Data Library (2026). The five regional datasets are North America, Europe, Japan, Asia-Pacific (excluding Japan), and emerging markets. Each PFAS-implicated firm and matched control is paired with the regional factor set corresponding to the firm's primary listing jurisdiction, based on the country-to-region pairing definitions of the French Data Library (2026). The risk-free rate is the one-month U.S. Treasury-bill rate from the French Data Library (2026), applied in common across regions, with all returns denominated in USD.

As part of the H1 robustness analysis, several alternative monthly factor models are obtained to test whether the PFAS premium reflects a misspecified benchmark. For the North American long-short portfolio, the augmented q-factor model, from the Global-q Data Library (2026), and the mispricing-factor model, from Stambaugh's website (2026), are used as stand-alone benchmarks. Four further factors are each added to FF5 + MOM in turn. These are the quality-minus-junk (AQR Data Library, 2026a) and betting-against-beta factors (AQR Data Library, 2026b), and the traded-liquidity and green-minus-brown factors, both from Pástor's website (2026). For the developed and emerging buckets, the factor themes, from Jensen, Kelly, and Pedersen's Global Factor Data (2026), provide an independently constructed benchmark.

5.5 Auxiliary data

Further auxiliary data is sourced from WRDS (2026). These data include annual and quarterly fundamentals from Compustat (2026) Global and North America, daily currency exchange rates from Compustat (2026) Global, the Center for Research in Security Prices (CRSP) delisting file, and the CRSP-to-GVKEY linking table (CRSP, 2026). The annual fundamentals are used to match firms on several characteristics in the propensity score matching (PSM) in Section 6.2 and as inputs to the cross-sectional return regressions. The quarterly fundamentals files provide the earnings-announcement dates, which are used to identify events falling close to an earnings announcement. The daily currency exchange rates are used to translate market capitalization into USD for comparability across markets, and the currency file is also used as a check in the returns-data cleaning. Finally, the CRSP delisting file is used for the Shumway

(1997) delisting correction, with the CRSP-to-GVKEY linking table used to link the delistings to the respective GVKEY.

6 Methodology

All four hypotheses must solve the same identification problem. PFAS-implicated firms are not a random sample of firms. They are often larger, are clustered primarily in sectors such as chemicals and electronics, and therefore face non-PFAS-related, time-varying risk premia relating to cyclicality and the regulatory environment. The methodology for the long-term performance in H1 solves this by adjusting for risk factors. The methodology for the short-term performance tested in H2, H3, and H4 solves this by using PSM to isolate the PFAS component in the CAR analyses.

All four hypotheses include each PFAS-implicated firm only from its FIY onward, to avoid look-ahead bias (Banz and Breen, 1986; Fama and French, 1992). A firm therefore enters the analyses only in the year the market could have publicly known that it was PFAS-implicated.

6.1 Long-term pricing design (H1)

To test for a PFAS premium under H1, a calendar-time portfolio analysis is conducted within each of the five Fama-French factor regions (North America, Europe, Japan, Asia-Pacific excluding Japan, and emerging markets) over 1998–2026 (Fama, 1998; Mitchell and Stafford, 2000). Calendar-time portfolios are preferred over an event-time long-horizon design because they are less susceptible to bad-model problems at long horizons, accommodate cross-sectional dependence across overlapping event-firm returns automatically, and produce test statistics closer to a normal distribution, better aligning with the *t*-test assumptions (Mitchell and Stafford, 2000).

The five region-specific long portfolios together hold the 74 PFAS-implicated firms, each of which enters its regional portfolio in its FIY and remains until delisting. The long portfolios are financed by five short portfolios that hold the whole universe of non-PFAS-implicated firms in the same regions, as available in Compustat (2026) Global and North America. Firms must have at least 24 months of return history at each rebalancing date, which limits the impact of new-listing data contamination and creates a more reasonable proxy for established firms, similar to the PFAS-implicated firm pool (Hou, Karolyi, and Kho, 2011). Firms domiciled in jurisdictions not covered by any of the five Fama-French factor files are excluded from the H1 test, because no clean benchmark applies. The five region-specific long-

short portfolio alphas are pooled by weighting each region's alpha by its PFAS leg's average market-capitalization share over the common sample period[1].

Both the long and short portfolios are value-weighted and rebalanced monthly on the beginning-of-month USD market capitalization, as described by Mitchell and Stafford (2000). Beginning-of-month market capitalization is used instead of end-of-month to avoid the bias that end-of-month weights introduce, since those weights would over-weight firms with positive returns over the month and under-weight firms with negative returns (Asparouhova, Bessembinder, and Kalcheva, 2013). Monthly returns are bounded at [−100%, +200%] before value-weighting, to limit the influence of microcap data errors that would otherwise compound over the month. This bound is slightly more conservative than those of Hou, Karolyi, and Kho (2011) and Ince and Porter (2006).

Using the cleaned monthly returns data, the long-short returns are regressed[2] on the standard FF5 + MOM model:

$$r_{\text{Long–Short},t} - r_{f,t} = \alpha + \beta_1(r_{m,t} - r_{f,t}) + \beta_2\text{SMB}_t + \beta_3\text{HML}_t + \beta_4\text{RMW}_t + \beta_5\text{CMA}_t + \beta_6\text{MOM}_t + \varepsilon_t \quad (1)$$

Where $r_{Long-Short,t}$ is the return of the long minus short portfolio, $r_{f,t}$ is the risk-free rate, SMB_t is small minus big, HML_t is high minus low, RMW_t is robust minus weak, CMA_t is conservative minus aggressive, MOM_t is the momentum factor, and ε_t is the residual term, with the t subscript denoting time. The intercept α is the monthly risk-adjusted return (the PFAS premium), the return attributed to the PFAS component after adjusting for market, size, value, profitability, investment, and the momentum factors. The monthly α is then compounded to an annualized rate. The pooled alpha is the market-capitalization-weighted average of the five region-specific FF5 + MOM intercepts. Standard errors, *t*-statistics, *p*-values, and 95% confidence intervals are computed using a stationary block bootstrap (Politis and Romano, 1994), with a six-month block length and 10,000 resamples. The bootstrap handles within-region autocorrelation and cross-region dependency jointly, without any parametric adjustment (Gonçalves, 2011). The six-month block length is set conservatively, guided by the Ljung-Box (1978) *Q* test, which rejects the no-autocorrelation null at the 1% level at lag 1 and at the 5% level at lag 3. This preserves the cross-region monthly correlation structure. Conventional

[1] The pooling weights are North America: 0.471, Europe: 0.265, Japan 0.103, Asia-Pacific 0.000, emerging market: 0.161. The Asia-Pacific weight is near zero, because the region has only two PFAS-implicated firms.

[2] Throughout this thesis, Greek letters without hats denote population parameters, and those with hats denote their fitted estimates.

Newey and West (1987) standard errors with a six-month lag[3], which correct for residual autocorrelation and conditional heteroskedasticity, are reported per region for the factor loadings. The short portfolio holding the non-PFAS-implicated firms is not propensity-score matched, because the FF5 + MOM model already captures that variation, providing risk adjustment equivalent to PSM.

Several robustness checks are conducted, grouped into three families, with full results in Appendix D. *(i) Alternative risk adjustments*: these checks ask whether the alpha survives a different factor model. The baseline FF5 + MOM is compared against a nested ladder of Fama-French benchmarks: the CAPM (Lintner, 1965; Sharpe, 1964), the Fama and French (1993) three-factor model (FF3), that three-factor model augmented with Carhart's (1997) momentum factor (FF3 + MOM), and the Fama and French (2015) five-factor model (FF5). The long-short alpha is then re-estimated against the alternative factor models described in Section 5.4. The augmented q-factor (Hou et al., 2021) and the mispricing-factor model (Stambaugh and Yuan, 2017) are used as stand-alone benchmarks. The quality-minus-junk (Asness, Frazzini, and Pedersen, 2019), betting-against-beta (Frazzini and Pedersen, 2014), traded-liquidity (Pástor and Stambaugh, 2003), and green-minus-brown (Pástor, Stambaugh, and Taylor, 2022) factors are each added to FF5 + MOM individually and jointly. The factor themes for the developed and emerging buckets are also applied (Jensen, Kelly, and Pedersen, 2023).

(ii) Portfolio construction: these checks vary how the portfolio is built. The portfolio is re-estimated with equal-weighting in place of value-weighting. The short leg is restricted to chemicals-industry firms, namely non-PFAS firms in the PFAS firms' two-digit SIC set, and the narrower SIC 28, 30, and 35, to test the premium against a tighter, industry-matched benchmark. An FIY-agnostic specification is estimated that includes all PFAS-implicated firms regardless of FIY.

(iii) Temporal stability: this check evaluates the time profile structural-break decomposition estimated using a Chow (1960) test.

6.2 Short-term pricing test and regime decomposition (H2)

The short-term pricing test uses PSM to find the closest non-PFAS control firms for the PFAS-implicated firms before each event date (Rosenbaum and Rubin, 1983). Taking the difference between a PFAS-implicated firm's abnormal return and its matched controls' abnormal returns

[3] The choice of six lags comes from Wooldridge's (2016) description of the bandwidth rule of thumb of $4 \times (n/100)^{(2/9)}$, where n equals the number of monthly time series observations in the sample. In this study $n = 337$ months, and the lag of 5.2 is rounded up to 6.

cancels out the pricing factors common to both, isolating the PFAS-specific abnormal return in the event announcement window.

The propensity score is estimated for each event date, re-matching the PFAS-implicated firms before each event instead of matching once and keeping the match throughout the panel, for two reasons. First, because firm characteristics change over time, a control matched at an earlier event might differ substantially for an event several years later. Second, the control pool changes constantly in terms of which firms are available (data availability, new listings, delistings, mergers and acquisitions, and restructurings), and re-matching ensures the control firms are available in the event window[4] (Heckman, Ichimura, and Todd, 1997).

Matching is conducted on the monthly data, and three eligibility filters are applied beforehand. First, the control firm must have at least 12 months of total monthly returns before the event, which excludes newly listed firms (Hou, Karolyi, and Kho, 2011) and aligns with the standard estimation-window practice of MacKinlay (1997) and Brown and Warner (1985). Second, the control firm must have at least 6 months of monthly returns immediately before the event, to ensure it has not been delisted, since performance-related delistings are often not announced (Shumway, 1997). Third, the firm fundamentals must be dated at least 90 days before the event (a defensive buffer beyond the 60-to-75-day filing deadlines in the United States and the longer international deadlines) and no more than 6 months before it, to ensure the matching covariates reflect what was publicly known at the time (Banz and Breen, 1986; Fama and French, 1992).

The propensity score is estimated by a logistic regression of a binary PFAS-treatment indicator on five firm-characteristic covariates, based on the latest filing date before the event date:

$$\log\left(\frac{\Pr(\mathrm{PFAS}_i = 1 \mid X_i)}{\Pr(\mathrm{PFAS}_i = 0 \mid X_i)}\right) = \gamma_0 + \gamma_1 \mathrm{logMktCap}_i + \gamma_2 \mathrm{BM}_i + \gamma_3 \mathrm{Lev}_i + \gamma_4 \mathrm{ROA}_i + \gamma_5 r_{i,-12:-1} \quad (2)$$

Where the left-hand side is the log-odds, $\mathrm{logMktCap}_i$ is the log of market capitalization in USD, $\mathrm{BM_i}$ is the book-to-market ratio (book equity/market equity), $\mathrm{Lev_i}$ is leverage (total liabilities/total assets), $\mathrm{ROA_i}$ is the return on assets (net income/total assets), and $r_{i,-12:-1}$ is the prior 12-month return excluding the most recent month, with the i subscript

[4] The alternative of a one-time matching at the start of the panel would require filtering the control pool with returns data availability across all events (or potentially lose statistical power through dropped events). This would introduce survivorship bias or ex-post selection bias, because which firm eventually survives is not knowable at the matching date (Banz and Breen, 1986).

denoting firm. Before fitting, the covariates are winsorized at the 1st and 99th percentiles: leverage and ROA over the full annual universe, and log market capitalization and book-to-market per event on the combined treated-and-control pool. The covariates are then z-scored within the matching pool, which makes the propensity-score coefficients γ comparable on a common basis. Standardization does not affect the matching outcome, which depends only on the rank order of the fitted scores.

The control pool contains all firms available in the Compustat (2026) Global and North America universe and excludes all PFAS-implicated firms, including any firm later discovered to be PFAS-implicated after the event, to avoid contamination between the PFAS-implicated firms and the control pool. Each PFAS-implicated firm is matched with three nearest-neighbor control firms, using the MatchIt package in R (Ho et al., 2011), with a caliper of 0.1 standard deviations of the logit of the propensity score, tighter than the 0.2 recommended by Austin (2011), to exclude poorly matched controls and be more conservative. Geography is not imposed as a matching constraint. The procedure selects controls on propensity score alone, so a PFAS-implicated firm may be matched to a control in a different region when that control is its closest match on the firm characteristics. The Fama-French regional classification, namely North America (United States and Canada), Europe, Japan, Asia-Pacific (Australia, Hong Kong, New Zealand, and Singapore), and an emerging-market bucket, following the country-to-region pairing definitions of the French Data Library (2026), is used to report the within-region matching rate (covariate balance before and after matching, and the full PSM diagnostics, are reported in Appendix B; Section 9.1 discusses the limitations).

Abnormal returns are calculated for each PFAS-implicated firm around each event announcement date, as the difference between the firm's return and the equal-weighted mean return of its matched controls (Barber and Lyon, 1997; Daniel et al., 1997). The CAR is calculated over a [−1, +1] trading-day window around the event announcement day, to capture effects occurring one day before or after the event. This follows the general guidelines of event-study methodology (Brown and Warner, 1985).

Event dates falling on a non-trading day are rolled forward to the next trading day, given the international scope of the panel and the differing trading calendars across exchanges. This roll uses each firm's own trading calendar. If the roll lands on another PFAS event date, or spans more than 7 calendar days (indicating stale prices or exit from the panel), that firm-event observation is dropped.

Two filters are applied at the firm-event level to clean the announcement window. First, if multiple events fall within [−1, +1] days of each other, the least important event is dropped, based on a hierarchy from most to least important (litigation, regulatory, then scientific), in line with the event-type importance argument in the hypothesis development. If multiple events share the same importance, all of them within the [−1, +1] window are dropped. The rationale for dropping events close in time is to avoid confounding abnormal returns across the events (Kothari and Warner, 2007; MacKinlay, 1997). Second, firm-event pairs falling within [−4, +4] calendar days of a firm's quarterly or annual earnings-release date are dropped, using the RDQ variable from WRDS Compustat (2026) Quarterly Fundamentals, to avoid confounding event abnormal returns with the earnings release (Bernard and Thomas, 1989).

Two further filters are applied at the firm-event level to clean the [−250, −10] trading day estimation window[5], which is used to standardize firm-event CARs for the Boehmer, Musumeci, and Poulsen (1991) BMP test described below. The differential test (also described below) does not require firm-level standardization, but the two filters are applied for it as well, to keep the sample size consistent across the two tests. The first filter is a minimum observations check. It drops firm-events where the intersection of the PFAS-implicated firm's return data with any of its three controls' return data falls below 30 daily observations in the estimation window, which protects against unstable standardization from too few observations. The second filter applies to control firms only. It drops a control firm whose estimation window contains more than 30% zero-return days, which removes low-liquidity firms and protects against deflated standardization of abnormal returns. The 30% threshold is a proxy for low liquidity and high transaction costs (Lesmond, Ogden, and Trzcinka, 1999). The forward roll, the two announcement-window filters, and the two estimation window filters together drop 1,301 of the 9,156 firm-events, leaving 7,855 across 134 events (10 of the 144 events are dropped entirely). Because all hypotheses include each firm only from its FIY, the headline differential and BMP tests are computed on the 5,596 of these firm-events that fall on or after the firm's FIY, spanning the same 134 events and 73 firms. The full 7,855-observation, FIY-agnostic panel is retained as a robustness specification.

Two statistical tests assess the pooled mean per-event reaction against zero: a per-event differential test, and the Boehmer, Musumeci, and Poulsen (1991) BMP-*t* test. The per-event differential test works in three steps. First, for a single event, the matched-adjusted CARs of

[5] The estimation window of [−250, −10] drops any abnormal returns within [−10, +10] calendar days of any event to avoid contamination to the standardization of abnormal returns.

all PFAS-implicated firms in that event are averaged. This is repeated for each of the 134 events, giving a series of 134 average PFAS CARs. The second step does the same for the matched control firms, computing for each control its own abnormal return over the event window as its return minus its estimation-window mean return (the mean-adjusted return of Brown and Warner, 1985) and averaging these within each event, which gives a second series of 134 average control CARs. The third step subtracts, for each event, the average control CAR from the average PFAS CAR, giving a series of 134 differentials. A one-sample *t*-test then assesses whether these 134 differentials differ from zero. Collapsing each event to a single PFAS mean and a single control mean before testing makes the event, rather than the firm-event, the unit of observation, which accounts for the cross-sectional correlation among firms reacting to the same event. The measure is a relative differential by construction. The PFAS leg is already expressed net of its matched controls, and the control leg captures the controls' own event-window reaction relative to their normal performance, so the differential measures PFAS underperformance net of the reaction accruing to comparable non-PFAS peers.

The second test is the BMP-*t* test (Boehmer, Musumeci, and Poulsen, 1991). This test standardizes the matched-adjusted CARs using the [−250, −10] estimation window, to account for differences in CAR volatility across firm-events, and tests the cross-sectional mean of these standardized values against zero. The standardized abnormal return is calculated as:

$$\frac{\mathrm{CAR}_{i,j}}{\sigma_{\mathrm{AR},i,j} \times \sqrt{L}} \tag{3}$$

where $\mathrm{CAR}_{i,j}$ is the CAR for a firm-event (i is the firm GVKEY, j is the event index) over the [−1, +1] event window, $\sigma_{\mathrm{AR},i,j}$ is the firm-event abnormal return volatility estimated from the same firm's pre-event data in the [−250, −10] estimation window, and $\sqrt{L}$ is the square root of the number of trading days L in the announcement window (3 for all events), which converts the firm's daily volatility into the volatility of a cumulative three-day window. The procedure also produces a *t*-statistic within each event, creating 134 *p*-values, one per event. A false-discovery-rate (FDR) correction at the 10% level controls for events appearing significant by chance (Benjamini and Hochberg, 1995).

The pooled tests above assume the per-event reactions are drawn from the same distribution across the whole 2001–2026 sample. The Bai and Perron (2003) multiple-breakpoint test identifies whether the sample contains regimes with different means and locates

the breaks. The test[6] is applied to the number of PFAS events per year, because that count follows the institutional environment, which economically motivates a regime-conditional CAR. The Bai and Perron (2003) break placement is assessed against three auxiliary diagnostics: the Bayesian Information Criterion (BIC) profile across zero to seven breaks, a Poisson rate-change test at each candidate year, and the Pettitt non-parametric change-point test (see Appendix F) by event-type subset.

Several robustness checks are conducted, grouped into four families, with full results in Appendix E. *(i) Parameter-choice sensitivity*: these checks vary the discretionary parameter choices. The event window is widened to [−7, +7]. The estimation window is varied to [−180, −10] and [−365, −10] alongside the [−250, −10] baseline. Entropy balancing (Hainmueller, 2012) is tested as an alternative to PSM. The event-collision policy is tested by dropping both colliding events instead of applying the importance ranking.

(ii) PFAS-specificity: these ask whether the reaction is specific to PFAS rather than a broader chemicals-industry effect. A matched-control falsification computes the same event-window abnormal return for the matched control firms, to check whether non-PFAS firms react on PFAS event dates. A separate falsification measures whether non-PFAS chemicals firms react on PFAS event dates, defined as firms in the PFAS firms' two-digit SIC set but not on the PFAS-implicated firms list, react on PFAS event dates. Finally, the reaction is split by PFAS exposure subtype, namely industrial, PFAS-intermediary, and pharmaceutical.

(iii) Concentration: These ask whether the result is concentrated in particular firms, events, or subsamples. Leave-one-firm-out and leave-one-event-out resampling test whether any single firm or event drives the result. The first event for each firm is compared against its later events. DuPont, Chemours, and Corteva are excluded to remove any restructuring effect. The sample is split into U.S. and non-U.S. jurisdictions to test whether the result is U.S.-driven. The cross-jurisdiction transmission of the reaction is tabulated by event jurisdiction against firm region.

(iv) Event-dating legitimacy: these verify that the event dating is legitimate. Pre-event drift is tested over [−15, −2] trading days for information leakage. A placebo test randomizes the event dates. Calendar regularities are checked by day-of-week, month-of-year, and an omnibus test.

[6] The test is implemented using the strucchange package in R (Zeileis et al., 2002) and uses the default minimum segment length $h = 0.15$ (meaning 15% of the observations has to be in a regime for it be counted as a regime) and finds the BIC-optimal number of breaks (Bai and Perron, 2003).

6.3 Persistence across sub-regimes (H3)

The matched-control design in Section 6.2 is extended with a persistence test across five post-event horizons [5, 30, 90, 180, 365] trading days. The CAR is calculated for each firm-event starting one trading day after the event window, to exclude the announcement reaction (Bernard and Thomas, 1989). The persistence window is therefore [+2, +(1 + *H*)], where *H* is the horizon (for example, the 5-day horizon uses the window [+2, +6]). Consistent with Section 6.2, at each date in the persistence window the abnormal return is the clean daily return of the PFAS-implicated firm minus the mean clean daily return of its PSM control firms (the same control firms as in Section 6.2). At each horizon, the per-firm event CARs are aggregated and tested with a one-sample mean test.

To account for the non-independence of events (at longer horizons, abnormal returns overlap across firms and events), the baseline inference is an event-clustered pairs bootstrap (Cameron, Gelbach, and Miller, 2008) that resamples whole events with replacement, using 10,000 resamples and a two-sided percentile *p*-value. This bootstrap is chosen as the baseline over two-way-clustered-robust standard errors because the small regime sample size makes cluster-robust inference unreliable (MacKinnon and Webb, 2017). A two-way clustered one-sample *t*-test, clustered by event × GVKEY and following Petersen (2009), is also reported as a benchmark. It uses cluster-robust degrees of freedom[7] equal to $\mathrm{MIN}(\text{number of events} - 1, \text{number of firms} - 1)$, following Cameron, Gelbach, and Miller (2011).

To test whether the drift differs by event type, the post-event CAR is regressed on event-type dummies (litigation and regulatory, with scientific as the omitted baseline), and the joint significance of the two dummies is assessed with a Wald *F*-test. This channel regression is also estimated on the collision-free subsample and by weighted least squares with inverse-variance weights, and the litigation channel is further split into commencement-stage and conclusion-stage events.

Persistence is also tested within the Bai and Perron (2003) regimes identified in Section 6.2, to evaluate whether the non-stationarity of the institutional environment shapes the drift dynamics, as outlined in the literature review (Pástor and Veronesi, 2012). The same calculations are conducted within each regime, including the channel regression by regime.

[7] Using the OLS residual degrees of freedom would understate the true sampling variability (Cameron, Gelbach, and Miller, 2011).

6.4 Cross-sectional heterogeneity (H4)

This section examines whether the mean CAR tested in Section 6.2 exhibits any systematic cross-firm or cross-event heterogeneity, by regressing the firm-event CAR on a set of predictors:

$$\begin{aligned} \mathrm{CAR}_{i,e} = \alpha + \beta_1 \mathrm{logMktCap}_{i,e} + \beta_2 \mathrm{BM}_{i,e} + \beta_3 \mathrm{Lev}_{i,e} + \beta_4 \mathrm{ROA}_{i,e} + \gamma_1 \mathrm{Lit}_e + \gamma_2 \mathrm{Reg}_e \\ + \delta_1 \mathrm{Named}_{i,e} + \delta_2 \mathrm{HasPen}_e + \delta_3 (\mathrm{logPen}_e - \mathrm{logPen}) + \delta_4 \mathrm{Mega}_e + \varepsilon_{i,e} \end{aligned} \quad (4)$$

Where $\mathrm{CAR}_{i,e}$ on the left-hand side is the CAR for firm i over the [−1, +1] trading-day window at event e. On the right-hand side, α is the intercept. The first block is four firm characteristics based on the most recent fundamentals. $\mathrm{logMktCap}_{i,e}$ is the size effect, the natural log of the market capitalization of firm i immediately before event e. $\mathrm{BM}_{i,e}$ is the value effect, the book-to-market ratio (book equity/market equity), where high-BM firms are value-firms and low-BM firms are growth-firms. $\mathrm{Lev}_{i,e}$ is the leverage effect (total liabilities/total assets), which controls for the possibility that more highly levered firms are impacted more strongly at an event, given their smaller equity cushion. $\mathrm{ROA}_{i,e}$ is the profitability effect (net income/total assets).

The second block is two event-level dummies that test whether the event type affects the price reaction, with scientific events as the omitted baseline[8]. Lit_e equals 1 if event e is a litigation event, and Reg_e equals 1 if event e is a regulatory event. The joint significance is assessed with a Wald F-test of whether both coefficients are zero.

The third block tests if the CAR has a stronger reaction if the firm is named in the event compared to unnamed firms. $\mathrm{Named}_{i,e}$ equals 1 if firm i is named directly in the news item for event e (for example, as a defendant in litigation). HasPen_e equals 1 if event e carries a penalty, evaluated at the average penalty size. $(\mathrm{logPen}_e - \mathrm{logPen})$ is the log penalty announced in event e minus the mean log penalty across all the events, measuring how far the penalty deviates from the mean. Mega_e equals 1 if the settlement in event e exceeds USD 1 billion. $\varepsilon_{i,e}$ is the residual. The regression is fit by OLS on the post-FIY panel (N = 5,596 firm-events from 134 events and 73 firms). The standard errors are clustered on event × GVKEY, because one firm across multiple events, and several firms within one event, are not independent observations (Petersen, 2009). Since this section relies on the panel created in Section 6.2, the robustness checks established there also apply here.

[8] There is no dummy variable for scientific events, to avoid perfect multicollinearity.

7 Results

The results section's structure mirrors the four hypotheses and four methodology sections, with the first four subsections each answering one hypothesis.

7.1 Long-term pricing: the PFAS premium (H1)

The compound-annualized[9] PFAS premium over 1998–2026 in the baseline model is −5.28% (monthly alpha of −0.451%), with stationary block bootstrap (Politis and Romano, 1994) t = −2.80 and p = 0.005 and a 95% confidence interval of [−8.79%, −1.62%], using a six-month block length and 10,000 resamples. This is the FF5 + MOM risk-adjusted, market capitalization-weighted alpha of the five region-specific long-short portfolios. Each portfolio uses monthly value-weighting, and the 74 PFAS-implicated firms enter post-FIY[10] over 337 months. H1 is fully supported. Translated into shareholder value, this premium corresponds to a realized cumulative differential of approximately −USD 2.1 trillion relative to the non-PFAS benchmark over the sample (path estimator, stationary block bootstrap p = 0.005, 95% CI [−3.5, −0.6] trillion), which is the dollar counterpart of the −5.28% annualized alpha. An equivalent forward-looking measure capitalizes the monthly alpha as a perpetuity at the panel's factor-implied cost of equity. The long-only PFAS portfolio earns a factor-implied expected return of 11.2% per annum[11] (pooled across regions by H1 weights) and dividing the −5.28% alpha by this rate gives an implied equity discount of USD 1.7 trillion or 47% of equity.

The region-specific results show that the pooled alpha is dominated by the North American region, which shows the strongest H1 rejection with an annualized alpha of −7.39% (bootstrap t = −3.80 and p < 0.001). It also carries the largest pool weight, 0.471. Emerging markets contribute a similar annualized alpha of −7.02%, which is not statistically significant (bootstrap t = −1.32, p = 0.17), and carry a pool weight of 0.161. Europe and Japan have statistically insignificant alphas of −2.99% and 1.54%, respectively. Asia-Pacific has a pool weight close to zero. Table I Panel A shows the full nested pooled panel, with Panel B showing the per-region FF5 + MOM regressions.

---

[9] The annualized α is calculated as $\left(1 + \alpha_{\text{monthly}}\right)^{12} - 1$.

[10] Post-FIY means that the PFAS-implicated firms only enter the long portfolio at the year that they were first publicly identified as PFAS-implicated, to avoid look-ahead bias. See Section 9 *(iv)* for limitations.

[11] Which can be compared to *Damodaran's* (2026) 10.5%–11.0% cost of equity for the global chemicals industry (basic, diversified, specialty chemicals).

Table I: Pooled and region long-short regression models

Nested risk-factor models for the value-weighted long-short PFAS-minus-non-PFAS portfolio over 1998–2026. Columns apply (1) CAPM, (2) FF3, (3) FF3 + MOM, (4) FF5, and (5, baseline) FF5 + MOM. Alpha standard errors are from a six-month block bootstrap with 10,000 resamples; factor-loading standard errors are Newey-West (1987) (six-month lag), in parentheses. Pooling is by each region's PFAS-implicated-firm market-capitalization share. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Panel A: Pooled models, weighted across regions by PFAS-implicated-firm market capitalization

Variable	(1) CAPM	(2) FF3	(3) FF3 + MOM	(4) FF5	(5) FF5 + MOM
α (monthly, %)	−0.173	−0.272	−0.291*	−0.456***	−0.451***
	(0.185)	(0.167)	(0.164)	(0.162)	(0.161)
$R_m - R_f$	−0.162***	−0.113***	−0.116***	−0.044	−0.054
	(0.042)	(0.032)	(0.038)	(0.035)	(0.037)
SMB		−0.299***	−0.291***	−0.228**	−0.215**
		(0.094)	(0.094)	(0.092)	(0.091)
HML		0.392***	0.387***	0.243***	0.215**
		(0.066)	(0.067)	(0.087)	(0.092)
RMW				0.265***	0.257**
				(0.101)	(0.100)
CMA				0.330***	0.341**
				(0.127)	(0.134)
MOM			0.005		−0.025
			(0.058)		(0.058)
α (annualized, %)	−2.06	−3.21	−3.44	−5.34	−5.28
Bootstrap t	−0.94	−1.62	−1.78	−2.82	−2.80
Bootstrap p	0.340	0.100	0.064	0.007	0.005
95% CI (annual., %)	[−6.33, +2.17]	[−6.93, +0.62]	[−7.25, +0.21]	[−8.73, −1.55]	[−8.79, −1.62]
Observations	337	337	337	337	337
Weighted R^2	0.033	0.172	0.176	0.203	0.210
Weighted Adj. R^2	0.030	0.164	0.166	0.191	0.195

Panel B: Per-region FF5 + MOM estimates

Variable	North America	Europe	Japan	Asia-Pacific	Emerging
α (monthly, %)	−0.637***	−0.253	0.127	0.078	−0.605
	(0.168)	(0.366)	(0.165)	(0.207)	(0.458)
$R_m - R_f$	−0.027	−0.042	0.006	−0.712***	−0.189
	(0.045)	(0.075)	(0.049)	(0.022)	(0.137)
SMB	−0.188**	−0.613***	−0.265***	0.011	0.395
	(0.080)	(0.233)	(0.084)	(0.038)	(0.342)
HML	0.142*	0.081	0.203*	0.047	0.656**
	(0.086)	(0.245)	(0.105)	(0.063)	(0.309)
RMW	0.187**	0.204	0.039	0.149**	0.689*
	(0.082)	(0.245)	(0.183)	(0.076)	(0.385)
CMA	0.479***	0.705**	−0.237	0.070	−0.291
	(0.130)	(0.358)	(0.152)	(0.077)	(0.443)
MOM	−0.094	−0.002	−0.084	−0.046	0.179
	(0.067)	(0.099)	(0.054)	(0.030)	(0.250)
α (annualized, %)	−7.39	−2.99	1.54	0.93	−7.02
Bootstrap t	−3.80	−0.69	0.77	0.38	−1.32
Bootstrap p	<0.001	0.504	0.407	0.749	0.173
95% CI (annual., %)	[−11.03, −3.68]	[−11.37, +5.50]	[−2.28, +5.56]	[−3.67, +6.26]	[−16.60, +3.63]
Observations	337	337	337	337	337
R^2	0.334	0.123	0.077	0.727	0.073
Pool weight	0.471	0.265	0.103	0.000	0.161

The annualized alpha (PFAS premium) ranges from −2.06% (CAPM) to −5.34% (FF5), with the most controlled specification, FF5 + MOM, used as the baseline and showing an annualized alpha of −5.28%. All the alphas are negative, implying that the sign of the estimate is stable across specifications. The RMW and CMA factors added in the FF5 model are significant, and the model shows a large jump of 1.90pp in the annualized alpha (from −3.44% to −5.34%). This means that the PFAS-implicated firms behave like profitable and conservative firms, and adding these factors as controls helps uncover the positive factor exposure[12] such firms had in the sample and makes the alpha more negative. This intuitively makes sense, given that several of the PFAS-implicated firms are large, established, profitable companies that do not have large increases in their assets year over year. The bootstrap standard errors for the alpha are approximately the same across the models, which shows that additional factors do not reduce residual noise but instead reallocate the systematic risk component from the alpha correctly to the factors that explain the variance. The alpha is not significant in the CAPM model ($p = 0.34$) or FF3 ($p = 0.100$), is marginally significant in FF3 + MOM ($p = 0.064$), and is significant at the 1% level in FF5 and FF5 + MOM. The specifications with fewer factors included, such as the CAPM regression, bundle the negative PFAS-specific effect with the positive RMW and CMA that partially cancel out the alpha, which leaves an alpha not distinguishable from zero in these lower specification models.

Looking at the FF5 + MOM model in more detail, the profitability factor (weighted RMW = 0.257) and the investment factor (weighted CMA = 0.341) show that PFAS-implicated firms look more profitable and invest more conservatively than the non-PFAS stocks. The size factor (weighted SMB = −0.215) and value factor (weighted HML = 0.215) indicate that the PFAS-implicated firms behave more like large-cap value stocks than the non-PFAS-implicated firms. The momentum effect (MOM = −0.025) is close to zero and shows that the PFAS premium is not driven by any hidden momentum. The PFAS premium is the alpha after netting out all these factor coefficients. Figure 1 visualizes the cumulative returns of various portfolios.

The regional decomposition shows that North America and emerging markets drive a large part of the annualized pooled PFAS premium of −5.28%. North America, which is 47% of the pool, has an annualized alpha of −7.39% (bootstrap $t = -3.80$, $p < 0.001$, CI [−11.03%, −3.68%]) and is the only region whose 95% CI lies wholly below zero. Emerging markets,

[12] Over the 1998–2026 sample period, the RMW and CMA returned on average 0.302% and 0.210% per month, respectively. Multiplying the mean monthly factor returns with the factor loading coefficients in model 4, gives $0.265 \times 0.302\% + 0.330 \times 0.210\% = 0.149\%$, which explains ≈ 90% of the −0.165pp per month shift between model 3 and 4 (the residual of 0.016pp per month comes from a shift in the weighting on the other factors).

which are 16% of the pool, have an annualized alpha of −7.02% (bootstrap $t = -1.32$, $p = 0.173$, CI [−16.60%, +3.63%]) and contribute a sizeable negative point estimate, but the CI crosses zero. Europe (27% of the pool, annualized alpha of −2.99%, bootstrap $t = -0.69$, $p = 0.504$, CI [−11.37%, 5.50%]) and Japan (10% of the pool, annualized alpha of 1.54%, bootstrap $t = 0.77$, $p = 0.407$, CI [−2.28%, +5.56%]) are non-significant. Asia-Pacific contributes approximately zero to the pool weighting and is included for completeness. Figure 1 visualizes this.

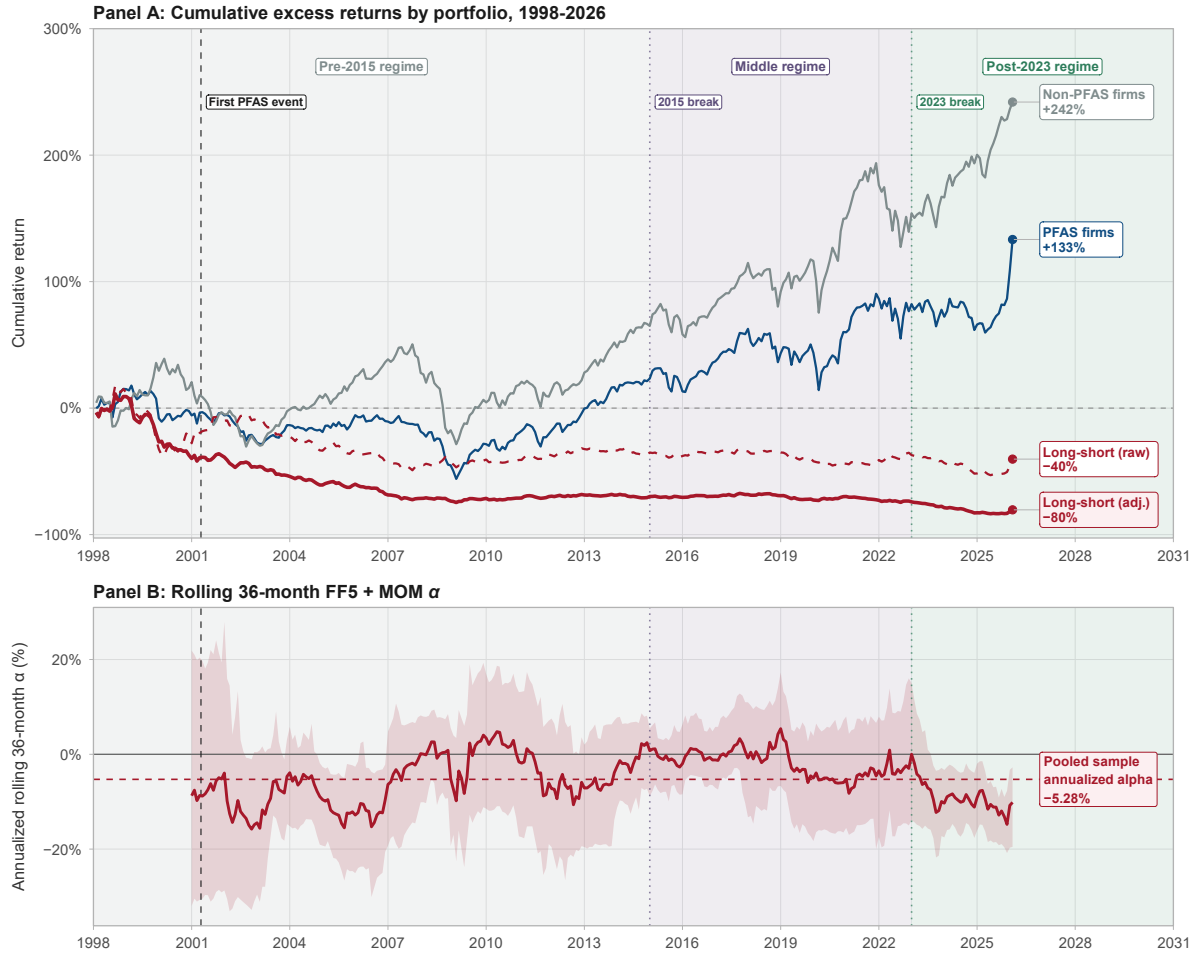


Figure 1: Cumulative return of portfolios and rolling annualized 36-month alpha

This figure decomposes the long-short PFAS strategy into its portfolio legs (Panel A) and its risk-adjusted performance over time (Panel B). Panel A plots the cumulative excess return of four series over 1998–2026: a long position in the PFAS-implicated firms (PFAS-firms), the broad non-PFAS benchmark (Non-PFAS firms), and the long-short spread shown both as raw returns (Long-short (raw)) and after risk adjustment for FF5 + MOM (Long-short (adj.)). Being long only the PFAS-implicated firms earns a positive absolute excess return, meaning that owning these stocks implied an opportunity cost, not an absolute underperformance. Panel B plots the rolling 36-month annualized FF5 + MOM alpha of the pooled long-short, with a shaded 95% bootstrap confidence band. The dashed line at 2001 marks the first PFAS-related event, and the dotted lines at 2015 and 2023 mark the Bai and Perron (2003) regime breaks discussed in the following sections. The 1998–2000 period precedes the first event. The relatively flat cumulative-return lines over this window support the interpretation that the PFAS premium is event driven, not a generic incumbent risk factor unrelated to PFAS exposure.

The R^2 and adjusted-R^2 across models are small, by design of the value-weighted long-short portfolio, which cancels out market exposure and leaves little variation for the models to explain. The $R_m - R_f$ coefficient in the CAPM model (Lintner, 1965; Sharpe, 1964) is close to zero for this reason, as the model has no systematic variance to explain. The increase in the R^2 for models two to five shows that the additional factor loadings contribute slightly to explaining the variance between the long-short portfolios. The R^2 reported here is in line with that reported by Pástor, Stambaugh, and Taylor's (2022) paper.

The PFAS premium is robust to the checks in Section 6.1, with the full results in Appendix D and a forest-plot summary across augmented and alternative complete factor models in Appendix Figure D1. *(i) Alternative risk adjustments*: the negative alpha survives every benchmark. Augmenting FF5 + MOM with the quality-minus-junk, betting-against-beta, and traded-liquidity factors leaves the annualized alpha between −7.5% and −8.2% (bootstrap $p < 0.001$), and the stand-alone q-factor and mispricing-factor models give −7.9% (bootstrap $p < 0.001$) and −5.0% (bootstrap $p = 0.037$). Adding the green-minus-brown factor over its shorter 2009–2023 window reduces the alpha to −4.1% but keeps it significant (bootstrap $p = 0.038$) Under the developed-and-emerging factor themes, the developed long-short alpha is −5.3% ($p = 0.001$) while the emerging leg is positive and insignificant (4.0%, $p = 0.47$). The market capitalization-weighted pool is −3.9%. The positive emerging-market point estimate is consistent with the headline decomposition, where the emerging region also carries a wide, insignificant alpha, reflecting the small number of emerging-market PFAS firms and the resulting estimation noise rather than a reversal of the premium.

(ii) Portfolio construction: equal-weighting widens the alpha to −8.54% (bootstrap $t = -3.77$, $p < 0.001$), which indicates that the premium is not concentrated in large-cap firms but is spread across the broader portfolio. Restricting the short leg to chemicals-industry firms leaves the alpha negative and significant across the same factor models, implying that the premium is not an artifact of the broad non-PFAS benchmark. The FIY-agnostic specification, which includes all PFAS-implicated firms regardless of FIY, reduces the alpha to −3.04% (bootstrap $t = -1.96$, $p = 0.051$) (See Appendix Table D3). This weakening is itself confirmatory, since diluting the portfolio with firm-years from before the market could have known of a firm's PFAS implication should weaken a premium genuinely tied to that information, consistent with EMH.

(iii) Temporal stability: The Chow (1960) test shows no significant structural break on the pooled sample after 2009, which can be considered as the relevant year for PFAS setting, as the institutional environment had not developed by 2009.

7.2 Short-term pricing: regime decomposition and aggregate magnitudes (H2)

The pooled per-event differential between PFAS-implicated firms and their matched controls, across all 134 events is −0.24pp ($t = -2.28$, $p = 0.024$). The PFAS-implicated firm component is −0.08%, while the matched controls earn 0.16% on event days. The differential of −0.24pp is therefore more negative than it would have been in a one-sample test (the PFAS-implicated firms' abnormal return alone, without netting against controls). Any market-wide effect cancels out, because it affects both the PFAS-implicated firms and the control firms symmetrically.

Table II decomposes the pooled estimate into three regimes identified through the Bai and Perron (2003) test, based on the annual count of PFAS events. Both the pre-2015 and post-2023 regimes show effects that are not distinguishable from zero. The middle regime of 2015–2022 shows a marginally significant negative differential. The middle regime coincides with the restructuring of DuPont and the Chemours spin-off, which created an entity that separated the PFAS liability from DuPont. The post-2023 regime coincides with the conclusion of several large litigations, including the 3M USD 10.3 billion settlement and the DuPont, Chemours, and Corteva USD 1.185 billion settlement.

Table II: Per-event differential decomposition by Bai and Perron (2003) regimes

PFAS-implicated-firm-minus-matched-control differentials across the Bai and Perron (2003) regime decomposition and the pooled sample. The per-event differential t and the firm-event-level BMP t are reported below (with p-values in parentheses), followed by number of events. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Regime (Years)	Pre-2015 (2001–2014)	Middle (2015–2022)	Post-2023 (2023–2026)	Pooled (2001–2026)
Differential (pp)	−0.38	−0.27	0.00	−0.24
Differential t	−1.64	−1.92*	−0.02	−2.28**
(p)	(0.113)	(0.059)	(0.982)	(0.024)
BMP t	−2.07**	−2.19**	1.31	−1.70*
(p)	(0.039)	(0.029)	(0.192)	(0.090)
N events	29	77	28	134

The −0.24pp per-event differential across the 134 event windows explains only 22% of the −5.28% FF5 + MOM-adjusted long-term annualized alpha established in Section 7.1. Pooling the 134 events at the −0.24pp per-event differential gives −32.1pp, and dividing by the approximately 28.1-year sample period gives −1.14pp per year, which compares to the −5.28% in the long-short FF5 + MOM portfolio. The remaining approximately 78% of the long-term annualized alpha is explained by effects outside the announcement windows.

The three regimes contribute unequally to the cumulative event-window differentials. The pre-2015 regime shows a per-event differential of −0.38pp (p = 0.113), which adds up to −11.0pp, or −0.39pp per year, and explains about 7% of long-short alpha. The middle regime, with a per-event differential of −0.27pp (p = 0.059) across 77 events, accumulates to −21.0pp, or −0.75pp per year, and explains about 14% of the long-short alpha, the largest single driver by regime. The post-2023 regime has a per-event differential indistinguishable from zero (p = 0.98) and contributes negligibly.

The BMP test reproduces the baseline differential results in direction but more strongly supports the pre-2015 and middle regimes than the differential test does. The convergence of statistical significance at both the pooled and middle-regime levels support H2. There is a divergence between the differential test and the BMP test in the post-2023 regime in Table II. The differential is approximately zero, while the BMP statistic is slightly positive (t = 1.31), but insignificant. At the two mega-settlement events, the PFAS-implicated firms react positively, driving the BMP point estimate in a positive direction (BMP measures the standardized firm-event reaction without netting against the controls at the event-level). The matched control firms reacted positively in the same direction as the PFAS-implicated firms (possibly because of the litigation conclusion, although the sample size is too small to support any generalizable finding), and because the differential test nets against the controls at the event level, this drives the differential to zero. Figure 2 shows the events across the three regimes.

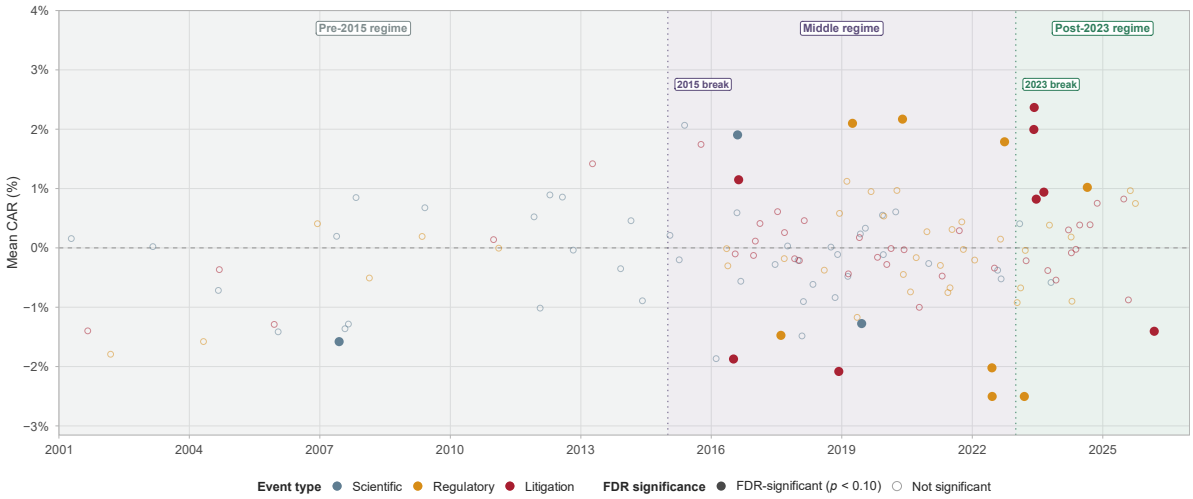


Figure 2: Per-event CAR scatter plot across Bai and Perron (2003) regimes

The figure shows per-event CARs across 2001–2026, with the background shading showing the three regimes. Each circle represents one of the 134 events. Color encoding shows event type (scientific blue, regulatory amber, and litigation red). Filled circles are significant at 10%, showing a CAR different from zero, after the Benjamini and Hochberg (1995) FDR adjustment across all 134 events. Most events are not significant, but this does not contradict the pooled results. Increasing the number of events increases the statistical power of the pooled test, while at the same time making the FDR adjustment stricter for classifying a single event as significant. This

explains why most events are not significant individually. The dense negative cluster of events in 2018–2022 matches the negative CARs seen in the middle regime.

The approximately zero differential in the post-2023 regime in Table II has two opposing patterns that cancel out but are visible in the Figure 2 event scatter plot. The two mega-settlement events triggered positive reactions, while ongoing litigation and regulatory events continued to trigger negative or no reactions. This implies that PFAS-implicated firms did not stop reacting to events, but that there is heterogeneity in reactions across event types.

Translating the CAR into dollar terms across event types shows heterogeneity. The litigation channel aggregates to approximately USD 47 billion (mean of USD 1.07 billion per event, $N = 44$), driven by the conclusion reactions in the post-2023 regime. The regulatory channel aggregates to approximately −USD 45 billion (mean of −USD 1.03 billion per event, $N = 44$). The scientific channel aggregates to approximately −USD 39 billion (mean of −USD 0.85 billion per event, $N = 46$). The bootstrap 95% confidence intervals, from 10,000 resamples of the per-event impacts within each event type, are wide and include zero at every channel. Summed across all 134 events, the three channels net to approximately −USD 37 billion of event-window shareholder value, but this aggregate is not statistically distinguishable from zero (bootstrap 95% approximately [−369, +295] billion, $p = 0.82$), because a few large-cap firms at high-magnitude events dominate the dollar sum and inflate its variance. The statistical insignificance is specific to the dollar aggregate and does not affect the per-event differential, which is significant (−0.24pp, $p = 0.024$), because the differential is computed on percentage abnormal returns and is not subject to the same market-capitalization concentration.

The baseline differential is supported by the robustness checks in Section 6.2, with full results in Appendix E. *(i) Parameter-choice sensitivity*: the differential is stable to the discretionary analysis choices. Varying the estimation window across [−180, −10], [−250, −10], and [−365, −10] leaves the standardized reaction essentially unchanged (BMP $t \approx -1.70$ throughout). The collision-free subsample, which drops both colliding events instead of applying the importance ranking, reproduces the result, with the mean matched-adjusted reaction shifting by only 0.015pp. Entropy balancing as an alternative to PSM achieves near-perfect covariate balance, with all standardized mean differences below 0.001. Widening the event window to [−7, +7] weakens the reaction to statistical insignificance (BMP $t = -0.34$, $p = 0.73$, versus −1.70, $p = 0.09$ at [−1, +1]), confirming that the reaction is concentrated in the three-day announcement window.

(ii) PFAS-specificity: these checks confirm the reaction is specific to PFAS rather than a broader chemicals-industry effect. The matched-control falsification computes the same

event-window abnormal return for the matched control firms on PFAS event dates. The controls earn a small positive mean reaction of approximately 0.21pp, which the differential test removes by construction, since the per-event differential nets the PFAS reaction against the control reaction at the event level. A separate falsification measures whether non-PFAS chemicals firms, defined as the full universe of firms in the PFAS firms' two-digit SIC set but not on the PFAS-implicated firms list, react on PFAS event dates. Across 142 PFAS event dates and more than one million firm-event observations, the pooled event-level reaction is statistically indistinguishable from zero (−0.05pp, $t = -0.74$, $p = 0.46$; BMP $t = -1.10$, $p = 0.27$), which indicates that the chemicals industry as a whole does not move on PFAS news and that the PFAS-firm reaction is information-specific (see Appendix Table E3). Splitting the reaction by PFAS exposure subtype shows no significant differential in any subtype (industrial −0.04pp, $p = 0.14$, PFAS-intermediary −0.09pp, $p = 0.52$, pharmaceutical −0.11pp, $p = 0.68$), confirming that the result is not concentrated in one operational category.

(iii) Concentration: the differential is not driven by any single firm, event, or subsample. Leave-one-firm-out and leave-one-event-out resampling move the differential *t*-statistic only marginally, with a maximum change of 0.17 and 0.31, respectively. Excluding DuPont, Chemours, and Corteva leaves the differential unchanged (−0.24pp, $p = 0.025$, versus −0.24pp, $p = 0.024$). Splitting first events from later events shows that the result is not a first-event repricing artifact, with first events insignificant ($N = 26$, $p = 0.55$) while later events carry the marginal pooled reaction ($p = 0.10$). The reaction is, however, concentrated outside the United States. The U.S. subsample is insignificant (differential −0.04pp, $t = -0.27$, $p = 0.79$), while the non-U.S. subsample is significant (−0.38pp, $t = -2.60$, $p = 0.011$). This contrasts with the North-America-driven long-run alpha in Section 7.1 and suggests that the announcement-window reaction is driven primarily by European events. The cross-jurisdiction matrix is consistent with weak transmission. Of the four event-jurisdiction-by-firm-region cells, only non-U.S. events affecting North American firms show a marginally significant negative reaction (−0.50%, $p = 0.054$), with no significant within-jurisdiction or U.S.-to-rest-of-world transmission (see Appendix Table E4 and Appendix Table E5).

(iv) Event-dating legitimacy: the event dating shows no sign of information leakage. Pre-event drift over [−15, −2] is insignificant (BMP $t = 0.75$, $p = 0.45$), and a formal pre-trends test confirms no drift into the window (slope −0.15 basis points per day, $p = 0.60$). A placebo test that randomly shifts all event dates never reproduces the baseline reaction. Across 763 draws, the per-event differential *t*-statistic averages 0.39 with a 95% interval of [−1.82, +1.99],

and not one draw reaches the baseline value of −2.28 (empirical $p = 0.001$), confirming the reaction is specific to the true event dates. Calendar checks show no day-of-week effect ($F = 1.55$, $p = 0.16$), though a mild month-of-year pattern is present ($F = 2.36$, $p = 0.007$) that does not align with the event distribution and does not invalidate the result.

7.3 Multi-horizon persistence (H3)

The pooled-sample multi-horizon profile shows statistically insignificant mean CARs at every horizon (5, 30, 90, 180, and 365 days) under both event-cluster bootstrap and cluster-robust standard errors (see Table III). This rules out aggregate post-event drift on the pooled sample. However, the Bai and Perron (2003) regime decomposition reveals persistence-drift patterns in opposing directions, which cancel out in the pooled sample.

Table III: Multi-horizon persistence of post-event abnormal returns

The top pooled-sample panel reports the result across all 134 events, while the subsequent three panels show the results across the Bai and Perron (2003) regimes. CARs are measured over the post-event window [+2, (1 + *H*)] trading days at five horizons. Two inference tests are reported, the baseline being the event-cluster bootstrap, a two-sided *p*-value from a 10,000-resample event-cluster pairs bootstrap (Cameron, Gelbach, and Miller, 2008). Cluster-*t* and cluster-*p* are the two-way clustered one-sample *t*-statistic and *p*-value on the per-firm-event CARs with clusters on firm × event, reported as a benchmark to the baseline cluster bootstrap. Firm-event counts differ slightly between the five-day horizons and the rest in each panel, due to a two-trading day minimum filter in the persistence window that excludes some firm-events for the shortest horizon. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Sample	*N* firm-events	*N* events	Horizon	Mean CAR (%)	Cluster bootstrap *p*	Cluster *t*	Cluster *p*
Pooled sample	5,555	134	5	0.10	0.217	1.25	0.216
	5,596	134	30	0.09	0.659	0.40	0.694
	5,596	134	90	−0.01	0.974	−0.02	0.983
	5,596	134	180	0.27	0.557	0.38	0.703
	5,596	134	365	0.25	0.723	0.20	0.841
Pre-2015 (2001–2014)	600	29	5	0.08	0.485	0.53	0.603
	616	29	30	−0.21	0.560	−0.53	0.601
	616	29	90	0.60	0.518	0.60	0.550
	616	29	180	−0.66	0.550	−0.52	0.606
	616	29	365	0.32	0.864	0.20	0.847
Middle regime 2015–2022	3,413	77	5	0.08	0.475	0.70	0.485
	3,428	77	30	0.02	0.946	0.07	0.948
	3,428	77	90	0.30	0.480	0.55	0.583
	3,428	77	180	1.10	0.058*	1.32	0.193
	3,428	77	365	2.34	0.006***	1.61	0.114
Post-2023 (2023–2026)	1,542	28	5	0.14	0.243	1.21	0.236
	1,552	28	30	0.36	0.245	1.01	0.320
	1,552	28	90	−0.93	0.044**	−1.13	0.266
	1,552	28	180	−1.19	0.143	−0.84	0.408
	1,552	28	365	−4.38	0.001***	−1.73	0.095 *

Decomposed into the three Bai and Perron (2003) regimes, the post-2023 regime ($N = 28$ events) and the middle regime ($N = 77$ events) show statistically significant long-horizon drifts in opposite directions. The middle regime exhibits positive drift of 1.10% at $H = 180$ (cluster-

bootstrap p = 0.058) and 2.34% at H = 365 (cluster-bootstrap p = 0.006) after a matched-control-adjusted abnormal return that is not statistically distinguishable from zero (−0.07pp, cluster-bootstrap p = 0.50). The post-2023 regime drifts −0.93% at H = 90 (cluster-bootstrap p = 0.044) and −4.38% at H = 365 (cluster-bootstrap p = 0.001) after an also statistically insignificant announcement-window abnormal return (0.12pp, cluster bootstrap p = 0.51). Both regimes show signs of post-event drift, and the opposite-signed directions offset in the pooled sample. The pre-2015 regime (N = 29 events) shows no consistent sign pattern across horizons under either cluster-bootstrap or cluster-robust standard errors.

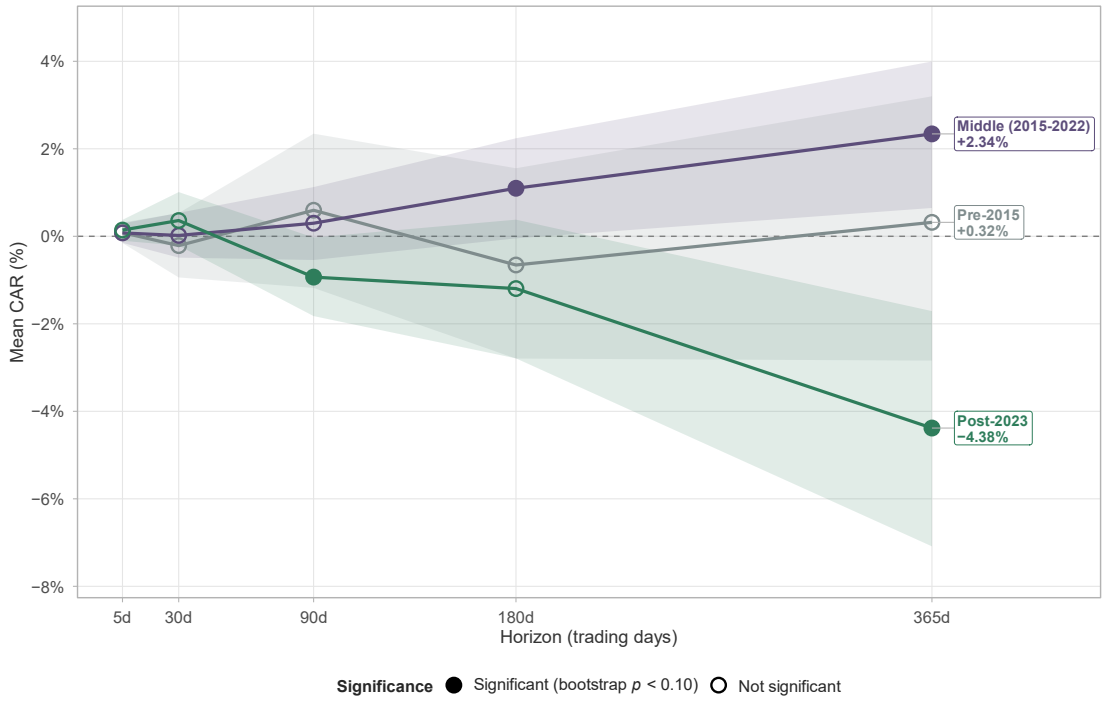


Figure 3: Multi-horizon drift profiles across Bai and Perron (2003) regimes

The three lines represent the three Bai and Perron (2003) regimes (pre-2015 regime grey, middle regime purple, post-2023 regime green). Ribbons are 95% confidence intervals from the event-cluster bootstrap with 10,000 resamples, with filled dots indicating $p < 0.10$.

7.4 Cross-sectional heterogeneity (H4)

The dominant cross-sectional driver of the event-window CARs is the Bai and Perron (2003) regime in which the event occurred, not firm or event-channel characteristics (except for mega-settlement dummy).

Table IV: Nested OLS regression of firm-event CARs on firm and event characteristics

Each column adds variables to the regression model: (1) firm characteristics only, (2) adds the event-channel type dummies (scientific as baseline), (3) adds the named-defendant dummy, (4) adds the penalty control, and (5) adds the mega-settlement dummy. Coefficients are reported in percentage points. Standard errors are two-way clustered on firm × event date (Petersen, 2009) and reported in parentheses below each coefficient. The log penalty variable is mean-centered, to be interpretable as a level shift away from the geometric mean penalty size. The mega settlement dummy equals 1 if the event had a reported settlement of over USD 1 billion. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Variable	(1) Firm chars.	(2) + Channels	(3) + Named	(4) + Penalty	(5) Full
Intercept	0.264	0.202	0.218	0.239	0.277
	(1.233)	(1.220)	(1.218)	(1.210)	(1.211)
Panel A. *Firm characteristics*					
Log market cap.	−0.012	−0.015	−0.015	−0.017	−0.017
	(0.042)	(0.043)	(0.043)	(0.042)	(0.042)
Book-to-market	0.168	0.158	0.158	0.153	0.144
	(0.177)	(0.177)	(0.177)	(0.176)	(0.176)
Leverage	−0.279	−0.274	−0.282	−0.264	−0.275
	(0.406)	(0.403)	(0.396)	(0.396)	(0.398)
ROA	0.447	0.529	0.512	0.518	0.398
	(0.981)	(0.985)	(0.979)	(0.978)	(0.968)
Panel B. *Event channels*					
Litigation (vs. sci.)		0.239	0.235	0.231	0.232
		(0.177)	(0.173)	(0.206)	(0.207)
Regulatory (vs. sci.)		0.102	0.102	0.103	0.102
		(0.203)	(0.203)	(0.203)	(0.203)
Panel C. *Exposure and severity*					
Named defendant			0.230	0.237	0.178
			(0.854)	(0.864)	(0.818)
Has Penalty				0.011	−0.248
				(0.270)	(0.240)
Log penalty				0.108	−0.022
				(0.112)	(0.117)
Mega settlement					1.633***
					(0.471)
Observations	5,596	5,596	5,596	5,596	5,596
R^2	0.001	0.002	0.002	0.002	0.005
Adjusted R^2	0.000	0.000	0.000	0.001	0.003
Max VIF	1.51	1.52	1.52	1.82	1.82

The pooled regression on 5,596 post-FIY firm-events across 73 PFAS-implicated firms[13] returns non-significant coefficients on log market capitalization, book-to-market, leverage, return on assets, the event type dummies, named-defendant status, whether the event had a monetary penalty, and the log magnitude of the penalty (cluster robust $p > 0.10$ on all of these variables). The only statistically significant variable is the mega-settlement dummy (1.633pp,

[13] One firm is dropped from the sample, due to it being listed in 2025 and has no firm-event that passes the multi-year fundamentals data filter explained in Section 6.2.

$p = 0.0005$). However, this variable consists of only two events[14] in the sample and should be considered a direction, not a generalizable finding.

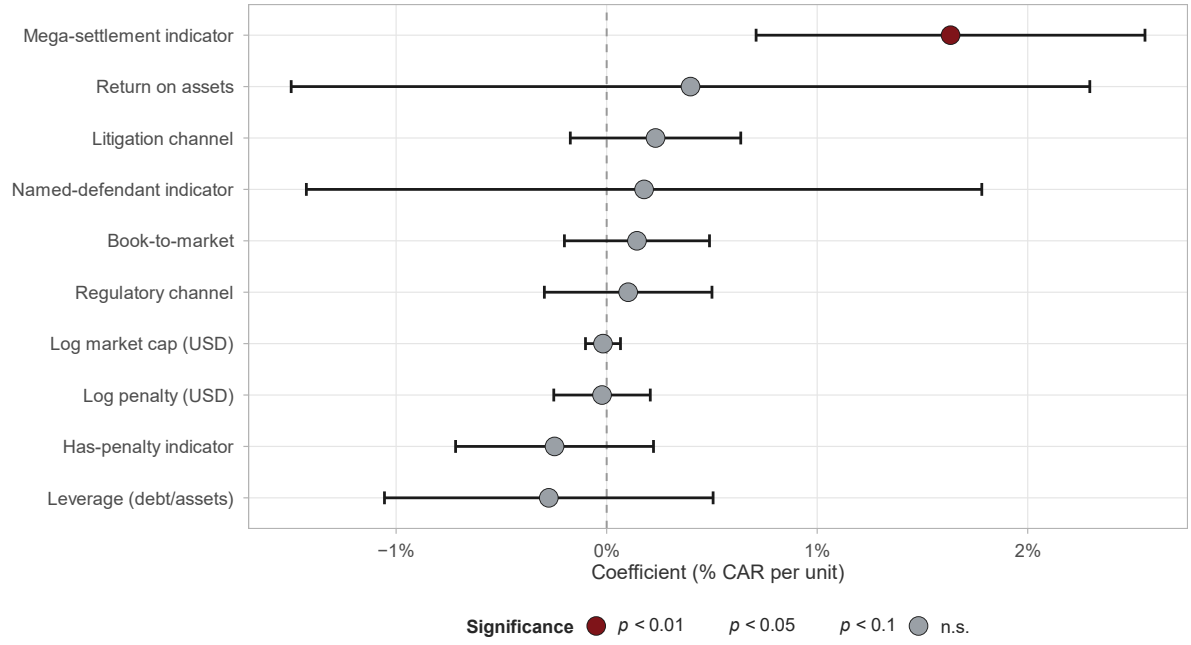


Figure 4: Cross-sectional determinants of CARs

Regression coefficients on the $N = 5{,}596$ firm-events, with two-way clustered standard errors on firm × event date. Whiskers show 95% confidence interval. Scientific event type is the baseline category.

The Wald joint test of the litigation coefficient = regulatory coefficient = 0 gives $F = 0.98$ and $p = 0.38$, meaning that the three event types (litigation, regulatory and scientific) are statistically indistinguishable in their effect on the CAR. The mean CARs across the three event types are 0.08%, −0.07%, and −0.17% for litigation, regulatory, and scientific, respectively, a spread of only 0.25pp. This fails to support the hypothesis that regulatory and litigation events have a larger effect than scientific events. The within-litigation event test of commencement ($N = 20$) versus conclusion ($N = 24$) announcement sub-types is also null (BMP $p = 0.789$ and $p = 0.248$, respectively), which rejects the part of the hypothesis that commencement events and conclusion events have different reactions.

Because the H4 regression operates on the same firm-event panel as the differential test in Section 7.2, the panel-construction checks validating that panel apply equally to the data underlying this regression. The cross-sectional nulls are otherwise stable across the nested specifications in Table IV, with the mega-settlement dummy the sole significant coefficient.

[14] This leaves only two clusters, below the recommended number, which biases the standard errors downward and inflates the *t*-statistics (Bertrand, Duflo, and Mullainathan, 2004; Cameron and Miller, 2015).

8 Discussion

This section integrates the results from Section 7 and discusses them in the context of their economic implications through three themes.

8.1 The PFAS premium is continuously priced, not concentrated at events

Approximately 78% of cumulative alpha accumulates outside the 134 PFAS event windows. The pooled H3 null finding at every horizon, under both cluster-bootstrap and cluster-robust standard errors, argues against post-event drift as an explanation. This bears on two strands of the literature. The latent-liability event-study literature (Karpoff, Lott, and Wehrly, 2005; Karpoff, Lee, and Martin, 2008) reviewed in Section 3.2 pools announcement-window CARs as its primary object of interest, yet the findings here show that the announcement-window channel explains a minority (22%) of the PFAS premium. The results also depart strongly from the announcement-premium patterns documented by Savor and Wilson (2014), who show that news days carry 60% of the aggregate equity premium, which can be put in contrast to the 78% of the PFAS premium that accumulates outside event windows. The rest of this section places that 78% in the context of the broader literature on news announcements and returns, and discusses candidate and ruled-out mechanisms behind the continuous alpha.

Two details sharpen the comparison between the PFAS non-announcement-day accumulation and Savor and Wilson's (2014) inverse. This thesis uses a three-day event window, whereas Savor and Wilson (2014) use a one-day window (the news day itself). Comparing each day's share of the premium to its share of trading days, the Savor and Wilson (2014) study's announcement days are about 13% of trading days and carry about 60% of the premium, a ratio of about 4.6 times. In contrast, for the PFAS event windows, this ratio is about 3.8 times, meaning that the PFAS premium clusters on event windows to a lesser extent than what Savor and Wilson (2014) found.

Two explanations follow, one about the nature of the risk and one about the portfolio construction. First, the type of risk measured differs. Macroeconomic risk arrives at scheduled learning events. Wachter and Zhu (2022) formalize this as investors receiving macroeconomic news through regular scheduled channels (e.g., monetary policy decisions, employment reports, inflation releases) that update their beliefs at those points. The PFAS information environment has no equivalent. Regulations evolve over multi-year rulemakings, with proposals, revisions, and votes that arrive incrementally as news. Scientific findings are dispersed across many institutions and develop continuously through working papers, publications, and multi-year studies. The litigation environment generates a high, continuous

flow of mostly low-impact news before culminating in higher-impact verdicts. As such, the PFAS premium accrues continuously because information arrives continuously and the 134 events studied here are the most impactful events, not every negligible update. Second, the long-short risk-adjusted portfolio hedges macroeconomic risk by construction. This means that even if PFAS-implicated firms carried a macroeconomic premium it would cancel out, leaving a PFAS-specific premium and not a market-wide premium as in the Savor and Wilson (2014) paper.

Two mechanisms can explain the negative sign of the PFAS premium: risk-based pricing of a PFAS-specific risk factor, and flow-based pricing through institutional ESG divestment. Under the risk-based mechanism, the premium is positive in expectation but negative as realized. Following the equilibrium intuition of Pástor, Stambaugh, and Taylor (2021), an environmental characteristic is priced, and PFAS-implicated firms earn a positive expected return as compensation for bearing PFAS exposure, which appears as alpha because FF5 + MOM contains no PFAS factor. Pástor, Stambaugh, and Taylor (2022) reconcile such a positive expected return with a negative realized one through the expected-versus-realized-returns framework, in which realized return has two components: the equilibrium expected return (the long-term positive risk premium investors require to bear the risk) and the cumulative effect of unanticipated shocks that arrive during the sample period (positive, negative, or zero if no shocks). When shocks are mild or symmetric, the realized returns track the expected return. When shocks are large and systematically one-sided, realized returns diverge sharply from expected. They use this to explain why brown stocks underperformed green stocks over 2012–2020, as investors developed stronger green preferences than anticipated, dragging realized brown returns below their positive expected-return baseline. The PFAS setting maps onto this directly. Investors require a positive expected return for bearing PFAS risk in equilibrium, but the continuously developing institutional environment delivered unanticipated negative shocks that heavily outweighed that positive baseline, producing a negative realized premium.

Hsu, Li, and Tsou's (2023) positive-signed 4.4% pollution premium does not invalidate the negative PFAS premium under the Pástor, Stambaugh, and Taylor (2022) framework. It simply indicates that their study's sample experienced milder, less one-sided regime conditions than the PFAS setting. Bolton and Kacperczyk's (2021, 2023) positive carbon premium also fits the same Pástor, Stambaugh, and Taylor (2022) framework, but differs in one respect unrelated to sign. Aswani, Raghunandan, and Rajgopal (2024) argue that Bolton and

Kacperczyk's (2021, 2023) use of vendor-estimated, unscaled emissions, instead of firm-disclosed emissions, introduces measurement error. This thesis sidesteps the methodology critiqued by Aswani, Raghunandan, and Rajgopal (2024) by using a public, categorical, and non-vendor-estimated database from the non-profit organization ChemSec (2025).

The second mechanism is flow-based pricing through institutional ESG divestment, and it operates independently of the first risk-based one. Van der Beck (2021) describes this mechanism as the cumulative price impact of ESG-mandated portfolios rebalancing away from an asset. In the PFAS setting, this would imply negative realized returns for the portfolio as investors move away from the asset. The two mechanisms have heterogeneous impact channels. The risk-based mechanism scales with how PFAS-intensive a firm's operations are, whereas the flow-based mechanism depends on the share of the firm owned by ESG-mandated funds and on when that capital exits, concentrating its price impact in the periods when divestment actually occurs. Whichever mechanism applies, the −5.28% annualized alpha can be interpreted as a cost-of-capital differential rather than an absolute loss for holding the PFAS portfolio, since the pooled long-only PFAS portfolio earns a cumulative excess return of 135% over the sample period.

The premium also carries a magnitude that can be set against the latent-liability literature. Realized over 1998–2026, it represents a cumulative shareholder-value differential of approximately −USD 2.1 trillion relative to the non-PFAS benchmark. Capitalized as a perpetuity using the long-only PFAS portfolio factor-implied expected return of 11.2% per annum (pooled across regions by H1 weights) and dividing, it implies a forward-looking cost-of-capital discount of approximately 47% of PFAS-implicated equity value, or USD 1.7 trillion. This forward figure capitalizes the alpha as a perpetuity because it is a per-period cost-of-capital differential the long-short Chow (1960) test shows to be stable across the PFAS era (no break at either Bai and Perron (2003) regime boundary, see Appendix D.3), whereas the USD 2.1 trillion is the assumption-free realized path-sum that requires no such stationarity. Konar and Cohen (2001) found that S&P 500 firms with poor environmental performance carried an average intangible liability of approximately USD 380 million, which they report as roughly 9% of the replacement value of tangible assets. This amount translates to 4.7% of market value (they define market value as equity plus debt plus preferred stock, and their reported means imply a mean of approximately USD 8.1 billion). The forward-looking discount for the PFAS-implicated firms amounts to 39.5% on a comparable market-value basis and is roughly eight times the value found by Konar and Cohen (2001), showing that PFAS represents a

categorically more severe market-priced liability. However, the two numbers are not strictly comparable. Konar and Cohen (2001) estimate an intangible liability already embedded in prices through a one-year market-value decomposition, whereas the figure here capitalizes a continuing cost-of-capital differential into a perpetuity. Both nonetheless express an environmental discount as a share of market value.

The forward-looking implied discount can be sized against two reference points. As a fraction of the panel's current value-weighted total equity market capitalization (a 12-month rolling mean of USD 3.5 trillion ending April 2026), the forward-looking discount represents approximately 47% of that equity value. Benchmarking against the range of published external estimates for total PFAS-related corporate liability, which range from approximately USD 30 billion (ChemSec, 2022), to USD 120–175 billion for U.S. water district remediation specifically (Milliman, 2024), and figures exceeding USD 1 trillion (Bloomberg Law, 2023). The implied discount is at the high end of this range. These estimates are not strictly comparable, as they measure different liability concepts, but together they indicate the market is pricing something beyond the realized-litigation channel. One explanation lies in the stranded-assets literature from climate finance (Caldecott et al., 2016; Carbon Tracker Initiative, 2013). Caldecott et al. (2016) characterize stranded assets as those whose value is written down, lost, or turned into liabilities earlier or more abruptly than anticipated. PFAS-implicated product lines fit this characterization: regulatory and litigation exposure may shorten their effective economic lives well below depreciation schedules, converting expected future cash flows into impairments and liabilities before those assets' productive lives are complete. On this reading, the forward-looking implied discount can be interpreted as a market-derived stranded-asset valuation for PFAS-implicated firms.

8.2 Pricing intensity is non-stationary across regimes

The three Bai and Perron (2003) regimes show materially different per-event differentials, with the 2015 and 2023 breaks aligning with institutional milestones described in Section 2 (the break placement is corroborated by the diagnostics in Appendix F). The regime-conditional segmentation extends the latent-liability event-study literature (Karpoff, Lott, and Wehrly, 2005; Karpoff, Lee, and Martin, 2008), which focuses on pooled estimates, and reframes a single pooled differential into three that align with institutional cycles.

The pre-2015 regime is dominated by scientific news and establishes PFAS as an environmental liability, with litigation focused mainly on DuPont. The 2015 break coincides with the Chemours spin-off, which carved the PFAS liability out of DuPont into a separately

listed entity. The middle regime shows a fivefold increase in events per year, rising from approximately 2 to approximately 10 events per year, with a balanced mix of event types, alongside industry-wide regulatory initiatives (EPA advisory, EU REACH, ECHA proposals) and the consolidation of litigation into MDL 2873. The 2023 break coincides with the conclusion of the largest cases, the 3M settlement (USD 10.3 billion) and the DuPont, Chemours, and Corteva settlement (USD 1.185 billion). The post-2023 regime is dominated by litigation-type events. The event-mix evolution across the three regimes, from scientific-dominated to balanced to litigation-dominated, maps to the liability-uncertainty frictions established in Section 3.2: scientific events characterize hazard (distributional uncertainty), regulatory events codify scope (regulatory uncertainty), and litigation conclusions bound magnitude (size uncertainty). The regime structure tracks which friction is being crystallized most actively in each period. The post-2023 differential shows the disappearance of a PFAS-specific differential alongside substantial event-level dispersion. Two mega-settlement events carry a 1.41pp event-level differential, while the remaining 26 of 28 events show a slightly negative differential.

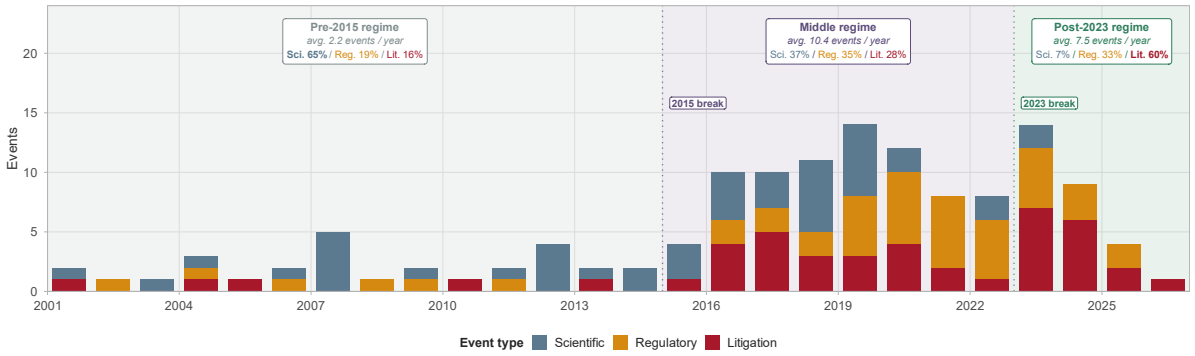


Figure 5: Events timeline and decomposition

The 144 events in the raw database are decomposed into the three event types (scientific, regulatory, and litigation) across the sample period and Bai and Perron (2003) regimes. The event study itself uses the 134 events remaining after the exclusion criteria (Section 6.2). The timeline shows event-type dominance shifting from scientific in the pre-2015 regime (65% of the regime's events), to an approximately balanced mix in the transitional middle regime (37% scientific, 35% regulatory, 28% litigation), followed by a dominance of litigation events in the post-2023 regime (60% of that regime's events). The pre-2015 regime establishes PFAS in a health context. The middle regime shows an uptick in both regulatory and litigation events. The post-2023 regime is dominated by litigation, mainly conclusions of processes that began in the earlier regimes.

The baseline contribution relative to the latent-liability literature is the regime structure itself. The pooled differential of −0.24pp sits below the ranges that Karpoff, Lott, and Wehrly (2005) and Karpoff, Lee, and Martin (2008) report. The regime structure conforms with McQueen and Roley (1993), who show that the relationship between news and stock prices depends on the

business-cycle state. This thesis extends that to the scientific, regulatory, and litigation cycle at which PFAS-specific news arrives.

The regime-conditional persistence pattern seen in the middle regime and post-2023 regime reinforces this. The two regime-specific drifts run in opposite directions and offset at the pooled level. To compare announcement-window reactions with multi-horizon drift on a consistent basis, both use the firm-event level, matched-control-adjusted CAR instead of the per-event differential. The middle regime shows a statistically insignificant initial reaction (−0.07pp, cluster-bootstrap $p = 0.50$) followed by positive drift of 1.10pp at $H = 180$ (cluster-bootstrap $p = 0.058$) and 2.34pp at $H = 365$ (cluster-bootstrap $p < 0.01$), consistent with the slow-information-diffusion mechanism (Hong and Stein, 1999) applied to positive news content over a long horizon. One potential explanation is that the complexity of scattered news from three different channels makes the cash-flow impact difficult to interpret at the announcement window. The H3 drift test measures information flow in specific 365-day post-event windows rather than calendar-time alpha, and the middle-regime positive drift therefore does not negate the claim that the net 2001–2026 news stream was an aggregate negative shock, as discussed in Section 8.1.

The post-2023 regime shows a statistically insignificant reaction at the announcement window (0.12pp, cluster-bootstrap $p = 0.51$), followed by a significant negative drift of −0.93pp at $H = 90$ (cluster-bootstrap $p = 0.044$) and −4.38pp at $H = 365$ (cluster-bootstrap $p < 0.01$). This is again consistent with the slow-information-diffusion mechanism (Hong and Stein, 1999), but in the opposite direction to the middle regime. A plausible reading is that the market underreacted to the litigation conclusions because the settlement amounts were already priced in, while the negative drift reflects cash-flow consequences not fully priced at the announcement window, namely lost revenue, large capital expenditures to shift operations away from PFAS, and damage to the firm's reputation. This finding suggests that the institutional environment in which a PFAS event arrives evolves across regimes and shapes not only the announcement window but also the post-event window.

The pre-2015 regime, in contrast, shows no drift in either direction. This is notable, because the slow-information-diffusion mechanism outlined in Section 3.3 would predict negative drift in a regime dominated by complex scientific news. The PFAS sample does not show it, indicating that the slow-information-diffusion mechanism does not hold uniformly across all three regimes.

8.3 The regime structure identifies institutional pricing intensity, not firm or event composition

The three Bai and Perron (2003) regimes contextualized in Section 8.2 reflect the institutional environment at each phase of the cycle of scientific recognition, regulatory codification, and litigation resolution, and are not a function of firm or event characteristics. This strengthens generalizability to other latent-liability settings with similar institutional phasing, such as sin industries and climate-transition liability.

The H4 null finding on firm characteristics and event-channel composition rules out the alternative interpretation of the Bai and Perron (2003) regime decomposition, namely that the three regimes look different because firms differ (log market capitalization, book-to-market, leverage, and ROA) or because event-type channels (scientific, regulatory, and litigation) differ across regimes. None of these variables is statistically significant in the Section 7.4 regressions. The only significant variable is the mega-settlement indicator, which supports this interpretation, because the two events the variable captures are the very milestones that mark the post-2023 regime break. The mega-settlement variable thus functions as a regime-defining marker, not a generic event characteristic. The regime structure must therefore carry something the firm- and event-level variables do not capture. The alignment of the regime breaks with the separately identified institutional milestones of Section 8.2 supports the interpretation that this missing factor is institutional intensity, rather than any firm- or event-level characteristic the H4 specification omitted.

The litigation-type evidence points the same way. There are no statistically significant commencement-versus-conclusion differentials, and the litigation-type heterogeneity reported by Karpoff, Lott, and Wehrly (2005) does not appear in the pooled PFAS sample. The two mega-settlement events, the 3M settlement (USD 10.3 billion) and the DuPont, Chemours, and Corteva settlement (USD 1.185 billion), show a positive per-event differential, descriptively consistent with the uncertainty-reduction mechanism (Karpoff, Lott, and Wehrly, 2005) and the three-friction framework of Section 3, though the small sample restricts generalizability. This reinforces the H4 reading that event characteristics do not drive the cross-sectional differences in PFAS-event reactions. The limits of this interpretation are addressed in Section 9.

9 Limitations

This section consolidates the methodology, sample composition, event database coding, statistical power and generalizability dimensions that are limitations.

9.1 Primary methodological caveats

Seven limitations warrant acknowledgement, because each could weaken the inferential basis of the study.

(i) Matched control bias on PFAS event days: the matched control firms earn a positive mean CAR of approximately 0.21pp per event on the event days. The per-event differential test cancels out this bias by construction because the bias affects the PFAS-implicated firms and control firm-events symmetrically, under the assumption that there is symmetric exposure of the common shock. If the matching procedure introduces an asymmetric common shock, the cancellation is imperfect. The leave one event out resampling and the entropy balancing alternative confirm the baseline results, but do not diagnose the mechanism that generates the positive CAR of approximately 0.21pp for the control firms, which is left as an empirical anomaly that the differential design is robust to but does not explain.

(ii) Limited sample size: several of the tests have small underlying sample data. The baseline per-event differential is identified by $N = 134$ events. The regime decomposition segments this into $N = 29$, 77, and 28 per respective regime. The middle-regime results are at the boundary of conventional significance, while the power to detect the post-2023 regime is low and the pre-2015 regime is similarly underpowered. The H4 mega settlement coefficient is identified by only two event clusters and should be interpreted as a case-study description, not a generalizable result.

(iii) Cross-jurisdictional matching: approximately 50% of PFAS-control matched pairs are matched across geographic regions. This implies that region-specific common shocks are not controlled for. The robustness check in Section 7.2 reports a jurisdiction split. The reaction is concentrated in non-U.S. events, which indicates geography matters for where the effect appears but does not, on its own, invalidate the pooled differential.

(iv) Long-horizon overlap in the H3 persistence test: A firm-event's $H = 365$ horizon CAR mechanically incorporates the announcement reactions of firm-events that fall within the 365-trading day horizon window, and the data shows that approximately 97% of the firm-events across the pooled sample overlap with another firm-event. The numbers at the other horizons are 16%, 62%, 90%, 96%, and 98% respectively at 5 days, 30 days, 90 days, 180, and 365 days horizon. The two-way clustering on firm × event date adjusts for the variance contribution but cannot separate the effect of post-event drift after the event in question and subsequent events arriving during the horizon window. A calendar-time portfolio approach (Mitchell and Stafford, 2000) would be methodologically cleaner in handling event-time

clustering automatically. However, with many months in the sample containing at least one event, this would approach something that looks more like a permanently long PFAS-implicated portfolio, which is why this design was not used. Thus, the long-horizon drifts across regimes are robust to event overlap in direction but not in magnitude.

(v) Bai and Perron (2003) break identification design: the structural-break composition is based on the annual count of PFAS events, not on the per-event CAR series itself (see Section 6.2). This design choice rests on the point that identifying breaks from the same data (CARs) used to test the segment differences would be a circularity.

(vi) Temporal alignment of the Bai and Perron (2003) breaks with institutional milestones: the 2015 and 2023 identified Bai and Perron (2003) breaks align independently with institutional milestones, such as the Chemours spin-off in 2015 and the mega settlements in 2023. The discussion in Section 8.2 interprets this as a result of institutional pricing intensity. However, this could also be attributable to unusually high event clustering in the sample data.

(vii) Categorical PFAS-implicated treatment design: the treatment indicator is based on the ChemSec (2025) data and is binary. Either a firm is a PFAS-implicated firm, or it is not. This design does not distinguish between the PFAS intensity. Some firms, like Chemours, where fluorochemicals are a large share of revenue, are treated the same as diversified conglomerates. This symmetrical treatment of PFAS-implicated firms does not allow the cross-firm heterogeneity models that test H4 to explain if the PFAS intensity explains any of the variation. This limitation also does not allow for a quantitative decomposition of the risk versus flow mechanisms discussed in Section 8.1.

9.2 Sample composition and coverage

The firm sample includes 74 listed firms (described in Section 5.1) and excludes private firms, listed private-equity firms with portfolio level exposure, and firms with insufficient data to pass the filters. This creates a bias towards larger producers with more transparent corporate structures (given exchange listing reporting requirements). Smaller private specialty chemical firms that are PFAS-implicated are absent from the analysis, even though they might have a higher liability concentration as a percentage of their revenue. As such, the PFAS premium estimated here is not necessarily representative for the whole PFAS-implicated firm universe. The geographic distribution of the sample is also non-uniform (see Appendix C), reflecting a pattern of public listings and ChemSec (2025) geographic coverage. This might underweight geographies where PFAS reporting coverage is lower and where firms are more often private than listed.

9.3 Event database coding and FIY identification

The Bai and Perron (2003) analysis thus also depends on the author's manual identification of events and the judgements used to categorize which events are sufficiently important to be included and which are not to be included. The individual events in each category are not of equal value, with some events having a larger effect on the PFAS liability crystallization than others. Furthermore, as the data was manually identified, some events may have been overlooked. The event database is manually constructed and involves two layers of author decisions. First, which events to include as PFAS events. This requires a decision whether the event carries sufficient information to warrant inclusion. As a general rule, the scientific events were included if they included a major health or environmental discovery, were peer-reviewed and published in a journal. The litigation events included major commencements and conclusions events that assumingly signaled a crystallization of the latent liability. Finally, the regulatory events included both proposals and introductions of rules, advisories, and recommendations. A different author might construct an event database that differs based on inclusion decision on the same event universe. Second, the classification within events also involves decisions, although the coding rule is clearer here than in the first layer. Regulatory events were coded as such if they were sourced from regulatory sources or if the proposal or decision was regulator created. The scientific events were coded based on if the source was a scientific paper in a journal. Finally, the litigation events were coded as such if the information was formally in a court. The leave one event out resampling robustness test shows that the baseline results are not driven by any single event or small subset but does not address systematic inclusion or coding biases in the coding process.

Manual identification of the first PFAS-implication news may be incomplete. For 45 of the 74 PFAS-implicated firms, the FIY is manually identified based on 10-K and 10-Q filings, sustainability reports, news archives and scientific literature that names the firm as PFAS-implicated. The median number of years backtracked using this rule is 15 years before the earliest ChemSec (2025) data publication, which is non-trivial and affects the entry of the PFAS-implicated firm in the long-short portfolio. This filter depends on the author's manual identification, and if the author overlooked earlier disclosures, those firms' FIY would be set too late. Additionally, investors might know about the PFAS implication through other sources than public knowledge, e.g., if they work in the industry and have specialized knowledge on the topic. The post-FIY specification is essential for the baseline statistical significance, as

dropping it and including all PFAS-implicated firms from their first trading day post-1998 (panel start) provides non-significant results.

9.4 Generalizability

The findings in this thesis are specific to the 2001–2026 PFAS pricing setting and generalizability should be interpreted conservatively. PFAS shares structural features with other environmental liabilities, such as asbestos, PCBs, and leaded gasoline, as outlined in Section 3.2. But there are nuances in the institutional development of these various settings and their regulatory frameworks.

10 Conclusion

This thesis identifies a PFAS premium in global equity markets and characterizes its structure across three institutional regimes. The long-short portfolio of PFAS-implicated firms minus non-PFAS-implicated firms earns a negative risk-adjusted return over 1998–2026, reflecting the equity market's pricing of the latent environmental liability. The premium accumulates continuously across the sample, with the majority falling outside the announcement windows, and it operates as a relative cost-of-capital differential to the non-PFAS universe rather than as an absolute negative cumulative return for holding the PFAS-implicated portfolio. At the per-event level, reactions vary sharply across the three regimes. Within the post-2023 regime, the two mega-settlement resolution events deliver a positive differential, consistent with an uncertainty-reduction mechanism, although the sample size is limited. The substantive contribution of this thesis is the decomposition, rather than a single estimate. See Appendix Table 1 for a summary of the hypothesis resolutions.

Two candidate mechanisms generate the continuous premium. The first is the expected-versus-realized-returns framework of Pástor, Stambaugh, and Taylor (2022). The PFAS institutional environment has been continuously negative in its news flow, with three channels of scientific evidence, regulatory action, and escalating litigation dragging PFAS returns below the expected-return baseline. The second is flow-based price pressure (van der Beck, 2021), as institutional ESG divestment rebalances away from PFAS-implicated firms. Both mechanisms are out of scope for this thesis' methodology and represent promising directions for future research. The multi-horizon persistence drift is null at the full-sample level under both cluster-bootstrap and cluster-robust standard errors, which argues against systematic mispricing in the aggregate sample, since there is no overreaction-and-correction pattern (De Bondt and Thaler, 1985). The regime-conditional drift, however, is significant and opposite-signed across the

middle and post-2023 regimes, pointing to mechanisms that depend on the institutional environment and its development across regimes and that offset at the pooled level.

This thesis contributes to four strands of the literature. First, it documents a three-regime institutional cycle in PFAS-event pricing: the scientific-dominated pre-2015 regime, the balanced middle regime, and the litigation-dominated post-2023 regime. This regime structure is identified through Bai and Perron (2003) data-driven structural-break detection, not imposed ex ante, which separates testing from identification.

Second, it extends the latent-liability event-study literature in several aspects. The per-event differential reactions are somewhat smaller than those documented by Karpoff, Lott, and Wehrly (2005) for environmental-violation announcements. Capitalized as a perpetuity, the premium implies a forward-looking discount of roughly 47% of PFAS-implicated equity value, about eight times the intangible-liability estimate of Konar and Cohen (2001), on a comparable market-value basis. The PFAS sample shows no evidence for commencement-versus-conclusion heterogeneity in event differentials (Karpoff, Lott, and Wehrly, 2005). The post-2023 mega-settlement events show a positive differential, though on a limited sample size, and should be treated as a case study, not a generalizable finding.

Third, it extends the environmental premia literature through a long-short calendar-time alpha that is opposite in sign to the environmental premia found by Bolton and Kacperczyk (2021, 2023) and Hsu, Li, and Tsou (2023). The opposite sign has two potential explanations. *(i)* Risk-based pricing: explained by the expected-versus-realized-returns framework of Pástor, Stambaugh, and Taylor (2022). *(ii)* Flow-based pricing: through norm-constrained institutional investors exiting PFAS assets (van der Beck, 2021). This thesis sidesteps the Aswani, Raghunandan, and Rajgopal (2024) critique by using the public, categorical, and non-vendor-estimated ChemSec (2025) database.

Fourth, it contrasts with the announcement-window premium literature. The PFAS event window accounts for only 22% of the premium, while the remaining 78% accrues continuously outside event windows. This departs from the announcement-day concentration found by Savor and Wilson (2014) for the aggregate equity market.

Three implications for adjacent literature and policy follow. First, the magnitude of the priced PFAS premium sets it apart from generalized environmental risk. On a comparable market-value basis, the capitalized discount borne by PFAS-implicated firms is approximately eight times the intangible liability that Konar and Cohen (2001) attribute to poor corporate environmental performance. Together with a long-short alpha that survives an extensive

robustness battery, this supports treating PFAS exposure as a distinct priced equity-risk factor rather than as one instance of generalized environmental risk. Where major commercial ESG-rating providers capture PFAS, they score it within broad, generic categories. MSCI (2025) files PFAS-related cases under its "Toxic Emissions & Waste" controversy theme, which is one of 28 themes that aggregate into a composite score. Morningstar Sustainalytics' ESG Risk Rating is built from roughly 22 broad material ESG issues, none of them directly PFAS related (Morningstar Sustainalytics, 2026). Because a priced risk of this magnitude is absorbed into composite scores, it is not visible in the headline ratings most investors rely on, which implies that a continuously scored standalone PFAS measure would surface risk-relevant information that current ratings do not. Second, the outside-of-event-window component explaining 78% of the premium implies that conventional event-study methodology understates the equity-market-priced magnitude of slow-moving environmental liabilities. For PFAS specifically, regulators, policymakers, and ESG-fund constructors who rely on event-study evidence to anchor expected costs are observing only the discretely announced component of a continuous repricing process. Third, the magnitude of the forward-looking implied discount (approximately USD 1.7 trillion based on the 12-month rolling mean panel market capitalization as of April 2026) is at the upper end of industry analyst estimates (ranging from USD 30 billion to above USD 1 trillion) for total realized plus expected PFAS-related liability. Consistent with Pástor and Veronesi's (2012) discount-rate channel for policy-uncertainty resolution, the market is not underpricing eventual realized liability. It is pricing ambient regulatory and litigation uncertainty across the entire panel, in continuous response to slow-moving information flows that arrive outside discrete event windows. PFAS is therefore not a story about market inefficiency or about single litigation. It is a story about how equity markets price categorical, multi-decade environmental liabilities, and why the conventional toolkit for measuring this pricing systematically misses the dominant component.

What investors holding PFAS-implicated firms going forward should expect to earn is a different question from what investors holding the portfolio over 1998–2026 actually earned. The future is an equilibrium expected-return question that this thesis cannot answer. The history is the realized pricing decomposition that this thesis does answer. The realized decomposition is the contribution, and the next test of the framework is whether the regime decomposition appears in other latent-liability settings.

11 References

Andrews, David Q., and Olga V. Naidenko, 2020, Population-wide exposure to per- and polyfluoroalkyl substances from drinking water in the United States, *Environmental Science & Technology Letters* 7, 931–936.

AQR Data Library, 2026a, Quality Minus Junk: Factors, Monthly, available at https://www.aqr.com/Insights/Datasets/Quality-Minus-Junk-Factors-Monthly, accessed May 29, 2026.

AQR Data Library, 2026b, Betting Against Beta: Equity Factors Data, Monthly, available at https://www.aqr.com/Insights/Datasets/Betting-Against-Beta-Equity-Factors-Monthly, accessed May 29, 2026.

Asness, Clifford S., Andrea Frazzini, and Lasse Heje Pedersen, 2019, Quality minus junk, *Review of Accounting Studies* 24, 34–112.

Asparouhova, Elena, Hendrik Bessembinder, and Ivalina Kalcheva, 2013, Noisy prices and inference regarding returns, *Journal of Finance* 68, 665–714.

Aswani, Jitendra, Aneesh Raghunandan, and Shiva Rajgopal, 2024, Are carbon emissions associated with stock returns?, *Review of Finance* 28, 75–106.

Austin, Peter C., 2011, Optimal caliper widths for propensity-score matching when estimating differences in means and differences in proportions in observational studies, *Pharmaceutical Statistics* 10, 150–161.

Bai, Jushan, and Pierre Perron, 2003, Computation and analysis of multiple structural change models, *Journal of Applied Econometrics* 18, 1–22.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Bansal, Ravi, and Marcelo Ochoa, 2011, Temperature, aggregate risk, and expected returns, NBER Working Paper 17575, National Bureau of Economic Research.

Banz, Rolf W., and William J. Breen, 1986, Sample-dependent results using accounting and market data: Some evidence, *Journal of Finance* 41, 779–793.

Barber, Brad M., and John D. Lyon, 1997, Detecting long-run abnormal stock returns: The empirical power and specification of test statistics, *Journal of Financial Economics* 43, 341–372.

Barber, Brad M., and Terrance Odean, 2008, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *Review of Financial Studies* 21, 785–818.

Belo, Frederico, Vito D. Gala, and Jun Li, 2013, Government spending, political cycles, and the cross section of stock returns, *Journal of Financial Economics* 107, 305–324.

Benjamini, Yoav, and Yosef Hochberg, 1995, Controlling the false discovery rate: A practical and powerful approach to multiple testing, *Journal of the Royal Statistical Society: Series B (Methodological)* 57, 289–300.

Bernard, Victor L., and Jacob K. Thomas, 1989, Post-earnings-announcement drift: Delayed price response or risk premium?, *Journal of Accounting Research* 27, 1–36.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, 2004, How much should we trust differences-in-differences estimates?, *Quarterly Journal of Economics* 119, 249–275.

Bloomberg Law, 2023, Trillions in PFAS liabilities threaten corporate bankruptcy wave, available at https://news.bloomberglaw.com/bankruptcy-law/trillions-in-pfas-liabilities-threaten-corporate-bankruptcy-wave, accessed June 2, 2026.

BloombergNEF, 2024, 3M incurs $10.5 billion liability for polluting waterways with 'forever chemicals,' BloombergNEF Insight Report (October), available at https://tnfd.global/wp-content/uploads/2024/10/BNEF_When-the-Bee-Stings_3M.pdf, accessed May 3, 2026.

Boehmer, Ekkehart, Jim Musumeci, and Annette B. Poulsen, 1991, Event-study methodology under conditions of event-induced variance, *Journal of Financial Economics* 30, 253–272.

Bolton, Patrick, and Marcin Kacperczyk, 2021, Do investors care about carbon risk?, *Journal of Financial Economics* 142, 517–549.

Bolton, Patrick, and Marcin Kacperczyk, 2023, Global pricing of carbon-transition risk, *Journal of Finance* 78, 3677–3754.

Brenner, Menachem, and Yehuda Izhakian, 2011, Asset prices and ambiguity, Working Paper EC-11-10, Stern School of Business, New York University.

Brown, Stephen J., and Jerold B. Warner, 1985, Using daily stock returns: The case of event studies, *Journal of Financial Economics* 14, 3–31.

Buck, Robert C., James Franklin, Urs Berger, Jason M. Conder, Ian T. Cousins, Pim de Voogt, Allan Astrup Jensen, Kurunthachalam Kannan, Scott A. Mabury, and Stefan P. J. van Leeuwen, 2011, Perfluoroalkyl and polyfluoroalkyl substances in the environment: Terminology, classification, and origins, *Integrated Environmental Assessment and Management* 7, 513–541.

Caldecott, Ben, Elizabeth Harnett, Theodor Cojoianu, Irem Kok, and Alexander Pfeiffer, 2016, Stranded assets: A climate risk challenge (Inter-American Development Bank), available at https://publications.iadb.org/publications/english/document/Stranded-Assets-A-Climate-Risk-Challenge.pdf, accessed June 2, 2026.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2008, Bootstrap-based improvements for inference with clustered errors, *Review of Economics and Statistics* 90, 414–427.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2011, Robust inference with multiway clustering, *Journal of Business & Economic Statistics* 29, 238–249.

Cameron, A. Colin, and Douglas L. Miller, 2015, A practitioner's guide to cluster-robust inference, *Journal of Human Resources* 50, 317–372.

Campbell, John Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.

Carbon Tracker Initiative, 2013, Unburnable carbon 2013: Wasted capital and stranded assets, available at https://www.lse.ac.uk/granthaminstitute/wp-content/uploads/2014/02/PB-unburnable-carbon-2013-wasted-capital-stranded-assets.pdf, accessed June 2, 2026.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.

Carroll, Stephen J., Deborah R. Hensler, Jennifer Gross, Elizabeth M. Sloss, Matthias Schonlau, Allan Abrahamse, and J. Scott Ashwood, 2005, *Asbestos Litigation* (RAND Corporation, Santa Monica, CA).

Center for Research in Security Prices, 2026, CRSP US stock database, University of Chicago Booth School of Business, accessed via Wharton Research Data Services (WRDS), https://WRDS-www.wharton.upenn.edu/, accessed April 28, 2026.

ChemSec, 2022, Investors with $8 trillion call for phase-out of dangerous “forever chemicals”, available at https://chemsec.org/investors-with-8-trillion-call-for-phase-out-of-dangerous-forever-chemicals/, accessed June 2, 2026.

ChemSec, 2025, SIN Producers, available at https://sinproducers.chemsec.org/, accessed January 24, 2025.

Chow, Gregory C., 1960, Tests of equality between sets of coefficients in two linear regressions, *Econometrica* 28, 591–605.

Cooke, Roger M., Daan Nieboer, and Jolanta Misiewicz, 2014, *Fat-Tailed Distributions: Data, Diagnostics and Dependence*, Vol. 1 (ISTE Ltd, London, and John Wiley & Sons, Hoboken, NJ).

Cordner, Alissa, Gretta Goldenman, Linda S. Birnbaum, Phil Brown, Mark F. Miller, Rosie Mueller, Sharyle Patton, Derrick H. Salvatore, and Leonardo Trasande, 2021, The true cost of PFAS and the benefits of acting now, *Environmental Science & Technology* 55, 9630–9633.

Cousins, Ian T., Jamie C. DeWitt, Juliane Glüge, Gretta Goldenman, Dorte Herzke, Rainer Lohmann, Mark Miller, Carla A. Ng, Martin Scheringer, Lena Vierke, and Zhanyun Wang, 2020, Strategies for grouping per- and polyfluoroalkyl substances (PFAS) to protect human and environmental health, *Environmental Science: Processes & Impacts* 22, 1444–1460.

Damodaran, Aswath, 2026, Cost of equity by sector (U.S. and global), available at https://pages.stern.nyu.edu/~adamodar/pc/datasets/waccGlobal.xls, accessed June 2, 2026.

Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* 52, 1035–1058.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839–1886.

De Bondt, Werner F. M., and Richard Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 40, 793–805.

Domingo, José L., and Martí Nadal, 2019, Human exposure to per- and polyfluoroalkyl substances (PFAS) through drinking water: A review of the recent scientific literature, *Environmental Research* 177, 108648.

Donelson, Dain C., Rachel W. Flam, and Christopher G. Yust, 2022, Spillover effects in disclosure-related securities litigation, *Accounting Review* 97, 275–299.

Drechsler, Itamar, 2013, Uncertainty, time-varying fear, and asset prices, *Journal of Finance* 68, 1843–1889.

European Chemicals Agency, 2023, ECHA receives PFASs restriction proposal from five national authorities, news release, February 13, available at https://echa.europa.eu/web/guest/-/echa-receives-pfass-restriction-proposal-from-five-national-authorities, accessed May 3, 2026.

European Parliament and Council of the European Union, 2006, REACH Regulation, available at https://environment.ec.europa.eu/topics/chemicals/reach-regulation_en, accessed May 3, 2026.

Fama, Eugene F., 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25, 383–417.

Fama, Eugene F., 1991, Efficient capital markets: II, *Journal of Finance* 46, 1575–1617.

Fama, Eugene F., 1998, Market efficiency, long-term returns, and behavioral finance, *Journal of Financial Economics* 49, 283–306.

Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fama, Eugene F., and Kenneth R. French, 2012, Size, value, and momentum in international stock returns, *Journal of Financial Economics* 105, 457–472.

Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.

Fenton, Suzanne E., Alan Ducatman, Alan Boobis, Jamie C. DeWitt, Christopher Lau, Carla Ng, James S. Smith, and Stephen M. Roberts, 2021, Per- and polyfluoroalkyl substance toxicity and human health review: Current state of knowledge and strategies for informing future research, *Environmental Toxicology and Chemistry* 40, 606–630.

Frazzini, Andrea, and Lasse Heje Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.

French, Kenneth R., 2026, Data Library, available at https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html, accessed April 28, 2026.

Frisbee, Stephanie J., A. Paul Brooks, Jr., Arthur Maher, Patsy Flensborg, Susan Arnold, Tony Fletcher, Kyle Steenland, Anoop Shankar, Sarah S. Knox, Cecil Pollard, Joel A. Halverson, Verónica M. Vieira, Chuanfang Jin, Kevin M. Leyden, and Alan M. Ducatman, 2009, The C8 Health Project: Design, methods, and participants, *Environmental Health Perspectives* 117, 1873–1882.

Global Factor Data, 2026, Factor return series, available at https://jkpfactors.com/data, accessed May 29, 2026.

Global-q Data Library, 2026, q-factor model data, available at https://global-q.org/factors.html, accessed May 29, 2026.

Glüge, Juliane, Martin Scheringer, Ian T. Cousins, Jamie C. DeWitt, Gretta Goldenman, Dorte Herzke, Rainer Lohmann, Carla A. Ng, Xenia Trier, and Zhanyun Wang, 2020, An overview of the uses of per- and polyfluoroalkyl substances (PFAS), *Environmental Science: Processes & Impacts* 22, 2345–2373.

Gonçalves, Sílvia, 2011, The moving blocks bootstrap for panel linear regression models with individual fixed effects, *Econometric Theory* 27, 1048–1082.

Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70, 393–408.

Hainmueller, Jens, 2012, Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies, *Political Analysis* 20, 25–46.

Harremoës, Poul, David Gee, Malcolm MacGarvin, Andy Stirling, Jane Keys, Brian Wynne, and Sofia Guedes Vaz, eds., 2001, *Late Lessons from Early Warnings: The Precautionary Principle 1896–2000*, Environmental Issue Report No. 22 (European Environment Agency, Office for Official Publications of the European Communities, Luxembourg).

Healy, Paul M., and Krishna G. Palepu, 2001, Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature, *Journal of Accounting and Economics* 31, 405–440.

Heckman, James J., Hidehiko Ichimura, and Petra E. Todd, 1997, Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme, *Review of Economic Studies* 64, 605–654.

Hirshleifer, David, and Siew Hong Teoh, 2003, Limited attention, information disclosure, and financial reporting, *Journal of Accounting and Economics* 36, 337–386.

Ho, Daniel E., Kosuke Imai, Gary King, and Elizabeth A. Stuart, 2011, MatchIt: Nonparametric preprocessing for parametric causal inference, *Journal of Statistical Software* 42, 1–28.

Hong, Harrison, and Jeremy C. Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets, *Journal of Finance* 54, 2143–2184.

Hong, Harrison, and Marcin Kacperczyk, 2009, The price of sin: The effects of social norms on markets, *Journal of Financial Economics* 93, 15–36.

Hong, Harrison, Frank Weikai Li, and Jiangmin Xu, 2019, Climate risks and market efficiency, *Journal of Econometrics* 208, 265–281.

Hou, Kewei, G. Andrew Karolyi, and Bong-Chan Kho, 2011, What factors drive global stock returns?, *Review of Financial Studies* 24, 2527–2574.

Hou, Kewei, Haitao Mo, Chen Xue, and Lu Zhang, 2021, An augmented q-factor model with expected growth, *Review of Finance* 25, 1–41.

Hsu, Po-Hsuan, Kai Li, and Chi-Yang Tsou, 2023, The pollution premium, *Journal of Finance* 78, 1343–1392.

Ince, Ozgur S., and R. Burt Porter, 2006, Individual equity return data from Thomson Datastream: Handle with care!, *Journal of Financial Research* 29, 463–479.

Jasanoff, Sheila, 1990, *The Fifth Branch: Science Advisers as Policymakers* (Harvard University Press, Cambridge, MA).

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Jensen, Theis Ingerslev, Bryan T. Kelly, and Lasse Heje Pedersen, 2023, Is there a replication crisis in finance?, *Journal of Finance* 78, 2465–2518.

Ju, Nengjiu, and Jianjun Miao, 2012, Ambiguity, learning, and asset returns, *Econometrica* 80, 559–591.

Karpoff, Jonathan M., John R. Lott, Jr., and Eric W. Wehrly, 2005, The reputational penalties for environmental violations: Empirical evidence, *Journal of Law and Economics* 48, 653–675.

Karpoff, Jonathan M., D. Scott Lee, and Gerald S. Martin, 2008, The cost to firms of cooking the books, *Journal of Financial and Quantitative Analysis* 43, 581–611.

Kelly, Bryan T., Ľuboš Pástor, and Pietro Veronesi, 2016, The price of political uncertainty: Theory and evidence from the option market, *Journal of Finance* 71, 2417–2480.

Klingelhöfer, Doris, Markus Braun, David A. Groneberg, and Dörthe Brüggmann, 2024, The "forever" per- and polyfluoroalkyl substances (PFAS): A critical accounting of global research on a major threat under changing regulations, *Chemosphere* 354, 141694.

Knight, Frank H., 1921, *Risk, Uncertainty and Profit* (Houghton Mifflin, Boston, MA).

Konar, Shameek, and Mark A. Cohen, 2001, Does the market value environmental performance?, *Review of Economics and Statistics* 83, 281–289.

Kothari, S. P., and Jerold B. Warner, 2007, Econometrics of event studies, in B. Espen Eckbo, ed.: *Handbook of Corporate Finance: Empirical Corporate Finance*, Vol. 1 (Elsevier/North-Holland, Amsterdam), 3–36.

Kuhn, Thomas S., 1962, *The Structure of Scientific Revolutions* (University of Chicago Press, Chicago, IL).

Kurwadkar, Sudarshan, Jason Dane, Sushil R. Kanel, Mallikarjuna N. Nadagouda, Ryan W. Cawdrey, Balram Ambade, Garrett C. Struckhoff, and Richard Wilkin, 2022, Per- and polyfluoroalkyl substances in water and wastewater: A critical review of their global occurrence and distribution, *Science of the Total Environment* 809, 151003.

Lesmond, David A., Joseph P. Ogden, and Charles A. Trzcinka, 1999, A new estimate of transaction costs, *Review of Financial Studies* 12, 1113–1141.

Li, Cindy Z., Ashish Tiwari, and Lin Tong, 2017, Investment decisions under ambiguity: Evidence from mutual fund investor behavior, *Management Science* 63, 2509–2528.

Lim, Sonya S., and Siew Hong Teoh, 2010, Limited attention, in H. Kent Baker and John R. Nofsinger, eds.: *Behavioral Finance: Investors, Corporations, and Markets* (John Wiley & Sons, Hoboken, NJ), 295–312.

Lintner, John, 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* 47, 13–37.

Ljung, Greta M., and George E. P. Box, 1978, On a measure of lack of fit in time series models, *Biometrika* 65, 297–303.

MacKinlay, A. Craig, 1997, Event studies in economics and finance, *Journal of Economic Literature* 35, 13–39.

MacKinnon, James G., and Matthew D. Webb, 2017, Wild bootstrap inference for wildly different cluster sizes, *Journal of Applied Econometrics* 32, 233–254.

McQueen, Grant, and V. Vance Roley, 1993, Stock prices, news, and business conditions, *Review of Financial Studies* 6, 683–707.

Merton, Robert C., 1987, A simple model of capital market equilibrium with incomplete information, *Journal of Finance* 42, 483–510.

Milliman, 2024, Milliman PFAS liability estimate, pt. 1: Water district remediation, available at https://www.milliman.com/en/insight/pfas-liability-estimate-water-district-remediation, accessed June 2, 2026.

Miner, Kimberley R., Heather M. Clifford, Todd Taruscio, Mariusz Potocki, Gene Solomon, Mark Ritari, Imogen E. Napper, Ananta P. Gajurel, and Paul A. Mayewski, 2021, Deposition of PFAS 'forever chemicals' on Mt. Everest, *Science of the Total Environment* 759, 144421.

Mitchell, Mark L., and Erik Stafford, 2000, Managerial decisions and long-term stock price performance, *Journal of Business* 73, 287–329.

Morningstar Sustainalytics, 2026, ESG Risk Ratings Methodology, version 3.1, April 7, available at https://mstar-sustops-cdn-mainwebsite-s3.s3.amazonaws.com/docs/default-source/transparency-documents/methodology-disclosures/esg-risk-ratings/morningstar_sustainalytics_esg_risk_ratings_methodology_3-1_07-04-2026.pdf?sfvrsn=58cacc53_1, accessed June 2, 2026.

MSCI, 2025, MSCI ESG Controversies and Global Norms Methodology, available at https://www.msci.com/documents/1296102/14524248/ESG-Research-Controversies-Methodology.pdf, accessed June 2, 2026.

Nevitt, Mark P., and Robert V. Percival, 2022, Can environmental law solve the "forever chemical" problem?, *Wake Forest Law Review* 57, 239.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Pástor, Ľuboš, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.

Pástor, Ľuboš, and Pietro Veronesi, 2012, Uncertainty about government policy and stock prices, *Journal of Finance* 67, 1219–1264.

Pástor, Ľuboš, and Pietro Veronesi, 2013, Political uncertainty and risk premia, *Journal of Financial Economics* 110, 520–545.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2021, Sustainable investing in equilibrium, *Journal of Financial Economics* 142, 550–571.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2022, Dissecting green returns, *Journal of Financial Economics* 146, 403–424.

Pástor, Ľuboš, 2026, Liquidity factors and green-minus-brown factor data, available at https://faculty.chicagobooth.edu/lubos-pastor/data, accessed May 29, 2026.

Pelch, Katherine E., Anna Reade, Taylor A. M. Wolffe, and Carol F. Kwiatkowski, 2019, PFAS health effects database: Protocol for a systematic evidence map, *Environment International* 130, 104851.

Petersen, Mitchell A., 2009, Estimating standard errors in finance panel data sets: Comparing approaches, *Review of Financial Studies* 22, 435–480.

Politis, Dimitris N., and Joseph P. Romano, 1994, The stationary bootstrap, *Journal of the American Statistical Association* 89, 1303–1313.

Reuters, 2023a, 3M reaches tentative $10.3 billion deal over US ‘forever chemicals’ claims, news release, June 22, available at https://www.reuters.com/legal/3m-reaches-tentative-103-billion-deal-over-us-forever-chemicals-claims-2023-06-22/, accessed May 3, 2026.

Reuters, 2023b, Chemical makers reach PFAS-related settlement with U.S., news release, June 2, available at https://www.reuters.com/business/chemical-makers-reach-pfas-related-settlement-us-2023-06-02/, accessed May 3, 2026.

Rieger, Marc Oliver, and Mei Wang, 2012, Can ambiguity aversion solve the equity premium puzzle? Survey evidence from international data, *Finance Research Letters* 9, 63–72.

Rosenbaum, Paul R., and Donald B. Rubin, 1983, The central role of the propensity score in observational studies for causal effects, *Biometrika* 70, 41–55.

S&P Global Market Intelligence, 2026, Compustat North America and Compustat Global, accessed via Wharton Research Data Services (WRDS), https://WRDS-www.wharton.upenn.edu/, accessed April 28, 2026.

Salvatore, Derrick, Kira Mok, Kimberly K. Garrett, Grace Poudrier, Phil Brown, Linda S. Birnbaum, Gretta Goldenman, Mark F. Miller, Sharyle Patton, Maddy Poehlein, Julia Varshavsky, and Alissa Cordner, 2022, Presumptive contamination: A new approach to PFAS contamination based on likely sources, *Environmental Science & Technology Letters* 9, 983–990.

Savor, Pavel, and Mungo Wilson, 2014, Asset pricing: A tale of two days, *Journal of Financial Economics* 113, 171–201.

Sharpe, William F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425–442.

Shumway, Tyler, 1997, The delisting bias in *CRSP* data, *Journal of Finance* 52, 327–340.

Stambaugh, Robert F., and Yu Yuan, 2017, Mispricing factors, *Review of Financial Studies* 30, 1270–1315.

Stambaugh, Robert F., 2026, Mispricing factor data, available at https://finance.wharton.upenn.edu/~stambaug/, accessed May 29, 2026.

Steenland, Kyle, Tony Fletcher, and David A. Savitz, 2010, Epidemiologic evidence on the health effects of perfluorooctanoic acid (PFOA), *Environmental Health Perspectives* 118, 1100–1108.

Steenland, Kyle, Tony Fletcher, Cheryl R. Stein, Scott M. Bartell, Lyndsey Darrow, María-José López-Espinosa, P. Barry Ryan, and David A. Savitz, 2020, Evolution of evidence on PFOA and health following the assessments of the C8 Science Panel, *Environment International* 145, 106125.

Stockholm Convention on Persistent Organic Pollutants, n.d., Industrial POPs: per- and polyfluoroalkyl substances (PFAS) overview, Secretariat of the Stockholm Convention, available at https://www.pops.int/Implementation/IndustrialPOPs/PFAS/Overview/tabid/5221/Default.aspx, accessed May 3, 2026.

Suffill, Ellise, Mathew P. White, Sarah Hale, and Sabine Pahl, 2024, Regulating “forever chemicals”: Social data are necessary for the successful implementation of the essential use concept, *Environmental Sciences Europe* 36, 111.

U.S. Environmental Protection Agency, 2006, EPA seeking PFOA reductions, news release, January 25, available at https://www.epa.gov/archive/epapages/newsroom_archive/newsreleases/fd1cb3a075697aa485257101006afbb9.html, accessed May 3, 2026.

U.S. Environmental Protection Agency, 2024, PFAS National Primary Drinking Water Regulation: General Questions and Answers, fact sheet, April 9, available at https://www.epa.gov/system/files/documents/2024-04/pfas-npdwr_qa_general_4.9.24v1.pdf, accessed May 3, 2026.

U.S. Judicial Panel on Multidistrict Litigation, 2026, Pending MDL Dockets by Actions Pending, May 1, available at https://www.jpml.uscourts.gov/sites/jpml/files/Pending_MDL_Dockets_By_Actions_Pending-May-1-2026.pdf, accessed May 3, 2026.

U.S. Office of Management and Budget, 1987, *Standard Industrial Classification Manual* (U.S. Government Printing Office, Washington, DC).

van der Beck, Philippe, 2021, Flow-driven ESG returns, Swiss Finance Institute Research Paper No. 21-71, Swiss Finance Institute.

Veronesi, Pietro, 1999, Stock market overreaction to bad news in good times: A rational expectations equilibrium model, *Review of Financial Studies* 12, 975–1007.

Veronesi, Pietro, 2000, How does information quality affect stock returns?, *Journal of Finance* 55, 807–837.

Wachter, Jessica A., and Yicheng Zhu, 2022, A model of two days: Discrete news and asset prices, *Review of Financial Studies* 35, 2246–2307.

Wang, Zhanyun, Jamie C. DeWitt, Christopher P. Higgins, and Ian T. Cousins, 2017, A never-ending story of per- and polyfluoroalkyl substances (PFASs)?, *Environmental Science & Technology* 51, 2508–2518.

Wang, Zhonghui, and Chenchen Zhang, 2023, Do shareholder litigations have spillover effects on peer companies from the perspective of financing constraints?, *Finance Research Letters* 58, 104401.

Wee, Sze Yee, and Ahmad Zaharin Aris, 2023, Environmental impacts, exposure pathways, and health effects of PFOA and PFOS, *Ecotoxicology and Environmental Safety* 267, 115663.

Wharton Research Data Services, 2026, WRDS, available at https://WRDS-www.wharton.upenn.edu/, accessed April 28, 2026.

Wooldridge, Jeffrey M., 2016, *Introductory Econometrics: A Modern Approach*, 6th ed. (South-Western Cengage Learning, Mason, OH).

Zarghamee, Reza, and Tamara D. Bruno, 2023, Recent PFAS settlements emphasize the scale of liability exposure for primary manufacturers, PFAS Observer, Pillsbury Winthrop Shaw Pittman LLP, July 25, available at https://pfas.pillsburylaw.com/pfas-settlements-liability-primary-manufacturers/, accessed May 3, 2026.

Zarghamee, Reza, Steve R. Brenner, and Macarena Ferreira Fink, 2024, PFAS fire protection products litigation liability, PFAS Observer, Pillsbury Winthrop Shaw Pittman LLP, December 9, available at https://pfas.pillsburylaw.com/pfas-fire-protection-products-litigation-liability/, accessed May 3, 2026.

Zeileis, Achim, Friedrich Leisch, Kurt Hornik, and Christian Kleiber, 2002, strucchange: An R package for testing for structural change in linear regression models, *Journal of Statistical Software* 7, 1–38.

Appendix

This appendix documents the construction of the sample and return panel (Appendix A), the PSM procedure (Appendix B), descriptive statistics (Appendix C), and the robustness of H1 and H2 tests (Appendix D and Appendix E), and the structural-break identification (Appendix F), followed by the AI disclosure (Appendix G). Appendix Table 1 is a summary of the hypothesis resolutions.

Appendix Table 1: Hypothesis resolutions

This table summarizes the resolution of the four hypotheses established in Section 4. Hypotheses are decomposed into their constituent sub-hypotheses for ease of overview. Resolutions follow the conventional statistical-significance thresholds.

Hypothesis	Sub-claim	Finding	Key statistic
H1	(a) Non-zero long-term PFAS premium, FF5 + MOM risk-adjusted over the sample period 1998–2026	Supported	−5.28% annualized (bootstrap $t = -2.80$, $p = 0.005$)
H2	(a) Pooled event-window CAR [−1, +1] is non-zero	Supported	Per-event differential −0.24pp ($t = -2.28$, $p = 0.024$)
	(b) Sign and magnitude vary across structural regimes	Supported for magnitude but not for sign	**Pre-2015:** −0.38pp ($p = 0.113$) **Middle:** −0.27pp ($p = 0.059$) **Post-2023:** 0.00pp ($p = 0.98$)
H3	(a) Post-event drift is non-zero at the pooled-sample level	Not supported	All horizons $p > 0.1$ under both cluster-bootstrap and cluster-robust standard errors
	(b) Sign and magnitude of the drift vary across Bai and Perron (2003) regimes	Supported	**Post-2023:** $H = 90$: −0.93pp (cluster-bootstrap $p = 0.044$) $H = 365$: −4.38pp (cluster-bootstrap $p = 0.001$) **Middle:** $H = 180$: 1.10pp (cluster-bootstrap $p = 0.058$) $H = 365$: 2.34pp (cluster-bootstrap $p = 0.006$) **Pre-2015:** null at every horizon under both cluster-bootstrap and cluster-robust standard errors
H4	(a) Firm-characteristic heterogeneity (log market capitalization, book-to-market, leverage, and ROA)	Not supported	All firm-characteristic covariates $p > 0.1$
	(b) Event-type heterogeneity (litigation, regulatory, and scientific)	Not supported	Wald joint $F = 0.98$, $p = 0.38$, mean event-type differentials span only 0.25pp
	(c) Exposure and severity heterogeneity (named defendant, penalty indicator, and log penalty magnitude)	Not supported	Both penalty indicator and log penalty magnitude (centered) $p > 0.1$
	(d) Mega-settlement direction	Supported in sample, generalizability limited	1.63pp, $p = 0.0005$, identified by only two event clusters

A. Sample and data construction

This appendix documents two steps in constructing the return data: the backwards IID selection algorithm that selects an IID per firm in each month, and the screens used to clean the daily returns.

A.1 Backwards IID selection algorithm

Firms in Compustat (2026) Global and North America carry several security identifiers (IID) under one firm identifier (GVKEY). Some firms are dual-listed, and a firm's IID can change over time. Constructing one coherent return chain per firm therefore requires an algorithm that *(i)* selects the best IID for each GVKEY and *(ii)* continues the chain onto a new IID when the current one disappears. The pipeline implements a backward-rolling, monthly IID selection with hysteresis, following the convention of Jensen, Kelly, and Pedersen (2023).

The algorithm works in three steps. Step 1 restricts the universe to common shares listed on a main exchange, applying the thirty-code exchange-exclusion list of Jensen, Kelly, and Pedersen's (2023), where every non-excluded exchange is treated as a main exchange. For each GVKEY, each calendar month, and each IID, the algorithm computes two statistics: the number of trading days in the month and the average market capitalization of the IID in local currency. Step 2 walks backward month by month, from the panel's most recent month (April 2026), retaining the current IID and changing it only when that IID disappears from the data, at which point a five-rule hierarchy elects a replacement. The backward direction avoids look-ahead bias (Banz and Breen, 1986). Step 3 filters the daily price panel to the elected IID for each firm-month, computes month-end observations, and in doing so resolves the duplicate GVKEY entries that arise across Compustat (2026) Global and North America.

Before the five rules are applied, the algorithm imposes a liquidity pre-filter. A candidate IID must clear a panel-wide floor of 20 trading days across the full 1998–2026 sample, or a within-month floor of five trading days. This screens out very lightly traded identifiers. If no candidate clears the pre-filter, the pre-filter is dropped and all candidates proceed through the hierarchy.

Appendix Table A1: IID selection rules

Five rules in decreasing priority. The first rule that resolves the selection ends the algorithm for that firm GVKEY-month.

Rule	Logic
(i) Standard IID	The modal value of Compustat's (2026) primary-security field (PRIUSA for U.S. firms, PRICAN for Canadian firms, PRIROW for the rest). The modal value, not the most recent, guards against transient Compustat (2026) coding errors.
(ii) Daily fallback	If the standard IID is missing from the current month's monthly statistics but the firm has at least two years (504 trading days) of daily price data on that IID across the panel, the standard IID is retained anyway. This handles the rare cases in which Compustat's (2026) monthly file drops the primary security at a corporate-action date while the daily file retains it. This rule was motivated by the Praxair-Linde merger of October 2018.
(iii) Home exchange suffix and source database	Prefer the IID whose suffix match the firm's home region (no suffix for North American listings, "W" elsewhere), combined with a database source-preference (Compustat (2026) North America preferred over Compustat (2026) Global for U.S. firms).
(iv) Higher average local currency market cap.	Select the IID with the larger average market capitalization that month.
(v) Higher trading day count	Select the IID that traded on more days that month.

Of the 1,590 switches in the panel, the standard-IID rule (rule *i*) resolves only 2.7% of cases. The remaining 97.3% resolve through the home-suffix, source, the market-capitalization or trading-day tiebreakers (rules *iii–v*). This is expected, as a switch only occurs when the current identifier disappears, and at that moment the primary security is usually unavailable as well. The primary-security map's larger contribution is upstream, in retaining the primary security through hysteresis, which prevents unnecessary switches.

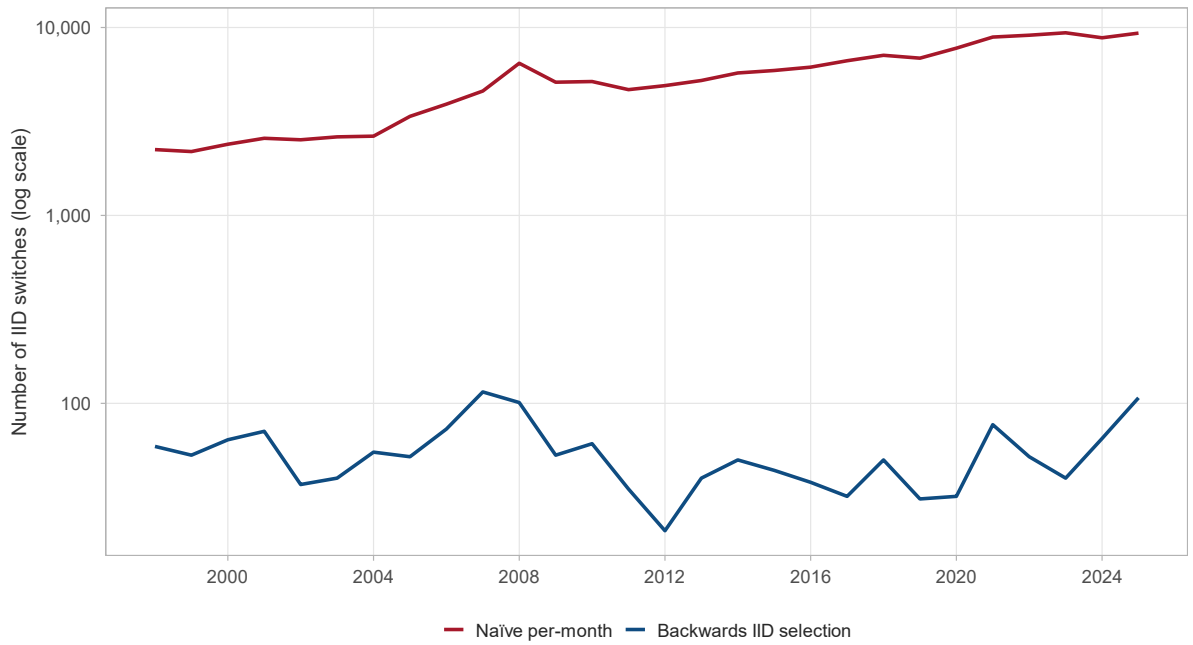
The backward selection yields 1,590 durable identifier switches across the full 1998–2026 panel, against roughly 155,000 under a naive rule that re-elects the best-covered identifier independently each month, resulting in a reduction of about 98 times (Appendix Figure A1). The effect is concentrated where it matters most for the value-weighted long leg for the PFAS-implicated firms. For example, Arkema's 100 naive switches collapse to zero, and the remaining large PFAS-implicated firms are essentially switch-free (Appendix Table A3).

Appendix Table A2: IID switches per year

Number of elected identifier changes per calendar year under naive selection (best-covered identifier re-elected independently each month) versus the backward-rolling IID selection algorithm with hysteresis used in the thesis. Reduction is the ratio of the two.

Year	*N* naive switches	*N* backwards IID selection switches	Reduction ratio (multiple)
1998	2,240	59	38
1999	2,187	53	41
2000	2,392	64	37
2001	2,573	71	36
2002	2,526	37	68
2003	2,618	40	65
2004	2,641	55	48
2005	3,366	52	65
2006	3,909	73	54
2007	4,591	115	40
2008	6,451	101	64
2009	5,116	53	97
2010	5,158	61	85
2011	4,670	35	133
2012	4,905	21	234
2013	5,224	40	131
2014	5,728	50	115
2015	5,904	44	134
2016	6,151	38	162
2017	6,654	32	208
2018	7,113	50	142
2019	6,867	31	222
2020	7,762	32	243
2021	8,894	77	116
2022	9,097	52	175
2023	9,372	40	234
2024	8,814	65	136
2025	9,333	107	87
2026	2,933	42	70
Total	**155,189**	**1,590**	**98**

The backward-rolling IID selection algorithm with hysteresis reduces the amount of IID switches by 98 times, turning the panel into continuous traded-return chains rather than identifier noise. The table is also visualized in Appendix Figure A1.



Appendix Figure A1: IID selection algorithm performance

The figure shows a comparison of the backwards IID selection algorithm versus a naive per-month selection that picks whichever IID has the most trading day that month (with ties broken by market cap., then IID lexicographical order).

Appendix Table A3: Identifier switches for the seven largest PFAS-implicated firms

The firms carrying the most weight in the value-weighted long leg. Columns count identifier switches under naive and the backwards IID selection algorithm.

Firm	*N* naive switches	*N* backwards IID selection switches
3M	0	0
Arkema	100	0
BASF	5	2
Chemours	0	0
Corteva	0	0
Daikin	0	0
DuPont	0	0

These seven PFAS-implicated firms carry the largest weight in the value-weighted long leg. A naive per-month rule would have switched Arkema's series 100 times and dropped those return data points in cleaning. Because these firms dominate the long-short portfolio by market capitalization, suppressing the switches is what makes the input data for the H1 test a sequence of traded returns and not an artifact of identifier churn.

A.2 Cleaning daily and monthly returns

The returns implied by Compustat (2026) price data is contaminated by identifier changes (Appendix A.1), currency redenominations, exchange-code changes, corporate actions, stale

prices, and implausible moves. The break-detection screen sets the affected daily or monthly return to missing whenever any rule triggers, instead of letting a contaminated value propagate. Delisting-return correction follows Shumway (1997). The screen applies thirteen flags grouped into seven buckets. The numeric bounds differ between the daily and monthly panels, with the monthly bounds wider to accommodate genuine month-over-month moves.

An observation may carry more than one flag. The final column is the share of the bucket's observation for which it is the only flag. The stale-quote bucket dominates and is caught cleanly, 99.9% of the time as the sole flag, whereas an implausible market-capitalization move is almost never isolated, only 4.3% of the time. A bad market-capitalization move typically co-occurs with a price or identifier anomaly, exactly as a genuine data error would, which is reassuring for the validity of the cleaning screen.

The price-ratio, market-capitalization, and FX bounds are sufficiently loose to retain genuine large moves while removing only mechanical artifacts. The quote quality rules target the two most common contaminants in a broad multi-exchange panel, namely carried-forward stale quotes and zero-after-positive prices. Quote-quality artifacts account for 99.66% of all flags and 12.6% of the approximately 258 million raw observations, almost entirely the stale-quote flag, which marks a price identical to the three preceding closes. This is expected in a global panel that spans thinly traded names and many exchange calendars, as carried-forward staleness would otherwise inject false zero-returns and understate volatility. The substantive corrective flags are rare. Every flag other than the stale marker touches well under one-twentieth of one percent of observations, and the calendar-discontinuity and negative-market-capitalization flags never trigger.

Appendix Table A4: Daily-return cleaning

The screen applies thirteen break flags grouped into seven buckets. A flagged observation has its return set to missing to avoid contaminating the returns series. Panel A summarizes incidence by bucket. Flagged observations is the count carrying any flag in the bucket. Share of all flags and Share of all observations express it against all flagged and all raw observations, respectively. Sole flag is the share of the bucket's observations for which it is the only flag triggered, which is a measure of how cleanly the contaminant is isolated. Panel B disaggregates to the individual flag.

Panel A: By bucket

Bucket	*N* flagged observations	Share of all flags (%)	Share of all observations (%)	Sole flag in bucket (%)
(i) Identifier or exchange change	46,918	0.14	0.018	47.3
(ii) Unflagged corporate action	82,537	0.25	0.032	88.6
(iii) Quote-quality artifact	32,521,869	99.66	12.584	99.9
(iv) Implausible price move	12,237	0.04	0.005	38.8
(v) Implausible market-cap. move	11,796	0.04	0.005	4.3
(vi) Currency-to-FX cross-check	619	0.002	0.000	0.0
(vii) Calendar discontinuity	0	0	0	0

Panel B: By the thirteen individual flags, mapped to their bucket

Flag	Bucket	*N* Observations	Share of all flags (%)	Share of all observations (%)
Repeated identical close	(*iii)* Quote-quality artifact	32,498,067	99.58	12.575
Shares-outstanding jump (>50%)	*(ii)* Unflagged corporate action	82,357	0.25	0.032
Reporting-currency change	*(i)* Identifier or exchange change	46,255	0.14	0.018
Zero or missing price after positive	*(iii)* Quote-quality artifact	23,802	0.07	0.009
Implausible market-cap. move	*(v)* Implausible market-cap. move	11,796	0.04	0.005
Implausible price drop	*(iv)* Implausible price move	8,442	0.03	0.003
Implausible price jump	*(iv)* Implausible price move	3,795	0.01	0.001
Identifier (IID) change	*(i)* Identifier or exchange change	1,197	0.004	0.000
Exchange-code change	*(i)* Identifier or exchange change	797	0.002	0.000
Currency-switch to FX mismatch	*(vi)* Currency-to-FX cross-check	619	0.002	0.000
ADR-ratio change (IID change with price jump)	*(ii)* Unflagged corporate action	227	0.001	0.000
Negative market cap.	*(v)* Implausible market-cap. move	0	0.000	0.000
Calendar gap (>365 days)	*(vii)* Calendar discontinuity	0	0.000	0.000

B. Propensity-score matching

Each PFAS firm-event is matched to non-PFAS control firms by propensity score, with the propensity model re-estimated at every event date on five covariates: log market capitalization, book-to-market, leverage, return on assets, and prior 12-month return. Matching is nearest-neighbor within the common-support region of the propensity score, using a caliper of 0.10 standard deviations of the logit propensity score and three matched controls per treated firm-event. A firm enters the control pool for an event only if it is not itself PFAS-implicated, has at least one year of pre-event return history (252 trading days), and has non-missing matching covariates. All characteristics are measured as of 90 days before the event date, meaning the matching does not use information from the event window itself. Matched pairs are required to share the same Compustat (2026) report database, North America or Global, which aligns accounting standard and reporting currency for the fundamentals. Matching is not constrained to the same Fama-French region or the same industry, both of which are examined as balance diagnostics below. This section documents the propensity model, the overlap it achieves, the covariate balance that results, the composition of the control, and the insensitivity of all of this to the matching design choices.

B.1 Propensity model and common support

Selection into the PFAS group is overwhelmingly due to size. Log market capitalization carries a large positive coefficient that is significant at every event (mean absolute z of 11.6), book-to-market is a distant second, and return on assets is irrelevant. PFAS-implicated firms are large, established producers, which is precisely the dimension the matching must accommodate. Common support is complete, as no treated firm-event lies outside the control propensity range. The caliper never discards a PFAS firm-event for lack of comparable controls. Treated units sit at a higher average propensity, 0.029 against 0.002, which is expected given size-based selection, but remain entirely within the control support. See Appendix Table B1.

Appendix Table B1: Propensity model and common support

Panel A reports the propensity-score logit, re-fit at every one of the 144 event dates on contemporaneously-available, *z*-scored characteristics. Mean coefficient is the cross-event average log-odds shift per one-standard-deviation increase. SD across events is the dispersion across the per-event fits. Share events significant is the share of events at which the covariate is significant at 5%. Mean |*z*| is the average absolute *z*-statistic. Panel B reports common-support diagnostics. The propensity range spanned by control firm-events, the count of treated firm-events falling outside it, and the mean propensity score in each group.

Panel A: Propensity-score logit

Covariate	Mean coefficient	SD across events	Share events significant (%)	Mean |*z*|
Intercept	−8.629	0.425	100.0	21.97
Log market capitalization	2.459	0.310	100.0	11.58
Book-to-market	0.273	0.253	42.4	1.82
Leverage	0.146	0.141	18.1	1.06
Return on assets	0.036	0.169	0.0	0.47
Prior 12-month return	−0.879	4.558	38.2	1.71

Panel B: Common support

Quantity	Value
Treated firm-events	9,275
Control firm-events	4,122,182
Control propensity-score range	0 to 0.551
Treated below and above control support	0 and 0
Share of treated outside control support	0%
Mean treated propensity score	0.029
Mean propensity score	0.002

B.2 Covariate balance

Balance of the covariates is good throughout. Average |SMD| ≤ 0.05 and none exceeds the 0.25 severe-imbalance threshold in more than 2.8% of events, meaning PFAS and control firm-events are comparable on observables.

Appendix Table B2: Covariate balance under the baseline PSM

Standardized mean differences (SMD) between PFAS and matched control firm-events. Panel A pool across all events, showing the average and worst-case absolute SMD per covariate and the share of events breaching the Austin (2011) severe-imbalance threshold of 0.25. Panel B shows the per-year average and maximum absolute SMD across the matching covariates. The convention is that |SMD| below 0.10 indicates good balance.

Panel A: Pooled, by covariate

Covariate	Mean |SMD|	Max |SMD|	Share of events |SMD| > 0.25 (%)
Log market capitalization	0.012	0.069	0.0
Book-to-market	0.049	0.178	0.0
Leverage	0.046	0.178	0.0
Return on assets	0.028	0.169	0.0
Prior 12-month return	0.046	0.373	2.8

Panel B: By calendar year

Year	Events	Mean |SMD|	Max |SMD|
2001	2	0.058	0.179
2002	1	0.052	0.165
2003	1	0.063	0.176
2004	3	0.041	0.156
2005	1	0.039	0.114
2006	2	0.035	0.125
2007	5	0.035	0.169
2008	1	0.024	0.070
2009	2	0.050	0.150
2010	1	0.043	0.127
2011	2	0.071	0.161
2012	4	0.039	0.137
2013	4	0.050	0.139
2014	2	0.071	0.373
2015	4	0.027	0.065
2016	11	0.032	0.204
2017	11	0.059	0.367
2018	11	0.026	0.177
2019	15	0.023	0.085
2020	13	0.028	0.092
2021	10	0.036	0.126
2022	8	0.032	0.178
2023	14	0.032	0.123
2024	11	0.036	0.125
2025	4	0.052	0.156
2026	1	0.022	0.059

B.3 Matched-sample levels and industry composition

On every covariate the PFAS and matched control groups sit at nearly identical central values. Log market capitalization 23.06 against 23.00, leverage 0.561 against 0.551, return on assets at 0.050 against 0.054, and prior 12-month return 0.103 against 0.100. PFAS-implicated events

concentrate in chemicals (62%), with a long tail across electronic equipment, stone, clay, and glass, industrial machinery, the wholesale trades, and petroleum refining. The matched controls are more dispersed: chemicals supply only 10% of controls, and the largest single control industries carry no PFAS firms at all (business services at 8.7% of controls, utilities at 7.1%, food at 5.6%, and communications at 5.5%). The result does not hinge on this, as the propensity model balances size and value under a tight caliper regardless of sector. See Appendix Table B3.

Appendix Table B3: Matched-sample levels and industry composition

Panel A shows the levels the treated and control groups sit at on each matching covariate, evaluated at the as-matched event-time observation. The reported values are the mean, median, and standard deviations, and this complements the differences in Appendix Table B2. Panel B shows the PFAS and control groups segmented by their two-digit SIC industry distribution. The twelve largest industries by PFAS firm-event count are shown. The remaining industries not shown are each below 1% of controls.

Panel A: Matching-covariate levels

Covariate	Group	Mean	Median	SD
Log market capitalization	PFAS	23.061	23.204	1.847
	Matched control	23.005	23.200	1.770
Book-to-market	PFAS	0.653	0.521	0.513
	Matched control	0.570	0.378	0.708
Leverage	PFAS	0.561	0.572	0.159
	Matched control	0.551	0.554	0.257
Return on assets	PFAS	0.050	0.047	0.051
	Matched control	0.054	0.052	0.157
Prior 12-month return	PFAS	0.103	0.043	0.471
	Matched control	0.100	0.035	0.595

Panel B: Industry (SIC2) composition, largest industries

SIC2 Industry	*N* PFAS firm-events	*N* Control firm-events	Share of all PFAS firm-events (%)	Share of all controls firm-events (%)
28 Chemicals	5,886	2,782	61.8	10.4
36 Electronic equipment	571	2,101	6.0	7.8
32 Stone, clay, and glass	425	355	4.5	1.3
35 Industrial machinery	418	1,368	4.4	5.1
51 Wholesale, nondurable goods	364	302	3.8	1.1
29 Petroleum refining	300	675	3.1	2.5
50 Wholesale, durable goods	288	399	3.0	1.5
38 Instruments	232	1,111	2.4	4.1
37 Transportation equipment	166	1,391	1.7	5.2
80 Health services	144	255	1.5	0.9
34 Fabricated metal products	144	248	1.5	0.9
87 Engineering and mgmt. services	140	293	1.5	1.1

B.4 Control-pool composition

Of the 74 PFAS-implicated firms, 73 are matched in the event study, with one firm never finding a control within the caliper and is thus dropped. The control pool is large and no firm is used more than 43 times out of 27,060 pairings. The median use of a control firm is two. Every event draws from a large control pool, with on average 28,626 eligible control firms per event, meaning no event is starved of comparable controls. Reuse does not bias the differential, as a reused control enters only as a subtraction, and at roughly 188 controls per event no single firm carries meaningful weight, meaning that the idiosyncratic return diversifies away while its common movement is differenced out. Each event is then reduced to a single PFAS-minus-control figure, and the test is run on those 134 event-level figures, meaning that firms reacting to the same event are not counted as if they were independent. See Appendix Table B4.

Appendix Table B4: Control-pool composition

Panel A shows the size and reuse structure of the control pool. A control firm may serve at more than one event provided the events are well separated. Panel B shows the number of controls and the eligible pool per event.

Panel A: Pool size and reuse	
Quantity	**Value**
Matched pairs (firm × event × control)	27,060
Distinct control firms used	7,385
Mean and Median uses per control firm	3.66 and 2
Most-reused control (max uses)	43
99th-percentile uses	28
Controls used only once	3,646
Controls used 10 or more times	694
Panel B: Common support	
Quantity	**Value**
PFAS events matched	144
Mean and median controls per event	187.9 and 193.5
Minimum and maximum controls per event	112 and 209
Mean PFAS-implicated firms per event	63.6
Mean eligible control pool per event	28,626

B.5 Region composition

Matching is exact on the Compustat (2026) reporting database (North America or Global), not on Fama-French regions. The global database pools Europe, Japan, Asia-Pacific, and the emerging markets, which makes about 50% of the pairs cross-region. North American firms, by contrast, stay within the North American database for the most part. This segmentation is acceptable because a control is a valid comparison when it matches the firm characteristics that drive expected returns, such as size and book-to-market, and the matched groups are closely

balanced on those characteristics (Appendix Table B2). Section 9.1 *(iii)* discussed the cross-region composition as a limitation for the event study.

Appendix Table B5: Region composition of matched pairs

Matched pairs share the same Compustat (2026) reporting database (North America or Global) by construction and matching is not exact on Fama-French region. Panel A shows the share of pairs that are same-region under the five-region Fama-French classification used in the H1 portfolios. Panel B shows the largest cross-region flows under that classification.

Panel A: Same-region share of matched pairs in the Fama-French classification		
Quantity	**Pairs**	**Share %**
Same Fama-French region	13,398	49.5
Different region	13,662	50.5
Panel B: Cross-region flows, largest		
PFAS region	**Control region**	**Pairs**
Europe	Emerging	3,149
Japan	Emerging	2,706
Japan	Europe	1,770
Emerging	Europe	1,231
Europe	Japan	1,126
Emerging	Japan	800
Europe	North America	785
Europe	Asia-Pacific	512
Emerging	Asia-Pacific	376
Japan	Asia-Pacific	371

B.6 Robustness to the matching design

The matched balance is insensitive to every design choice. Tightening or loosening the caliper, moving from one to five controls, and raising the minimum pre-event history to two years (the baseline is 252 trading days), all leave the mean |SMD| essentially unchanged, between 0.032 and 0.049. Using the alternative method of entropy balancing drives the covariate SMD to the order of one out of a hundred thousand. The equivalence of the matched groups is therefore a property of the data and the covariates and not an artifact of the nearest-neighbor algorithm or the tuning parameters. See Appendix Table B6.

Appendix Table B6: Robustness to the matching design

Each panel shows sensitivities to parameters for the matching. Panel A through C are parameters for the nearest-neighbor matching, while Panel D shows the alternative method of Entropy balancing. Panel A varies the caliper (maximum propensity-score distance). Panel B varies the number of matched controls. Panel C varies the minimum pre-event return history required for eligibility. Panel D shows the entropy balancing method (Hainmueller, 2012), which reweights controls to match the PFAS covariate moments exactly, applied to the four fundamental covariates (log market capitalization, book-to-market, leverage, and return on assets), with 12-month prior return excluded in the entropy-balancing constraint set (because it is the highest-dispersion covariate in the propensity model and only one requiring a complete pre-event return history, which unnecessarily limits the specification without materially improving the balance). Baseline settings are in bold.

Panel A: Caliper width (SD of logit propensity score)

Caliper	Events	Mean matched pairs	Mean |SMD|
0.05	15	170.3	0.042
0.10	**15**	**178.5**	**0.036**
0.20	15	181.0	0.034
0.50	15	181.9	0.034

Panel B: Number of matched controls

Controls per treated firm-event	Events	Mean matched pairs	Mean |SMD|
1	15	60.3	0.049
3	**15**	**178.5**	**0.036**
5	15	288.3	0.032

Panel C: Minimum pre-event return history (trading days)

Minimum history	Events	Mean matched pairs	Mean |SMD|
0	15	179.0	0.038
126	15	178.9	0.037
252	**15**	**178.5**	**0.036**
504	15	174.1	0.036

Panel D: Entropy-balancing benchmark (pooled)

Covariate	Mean |SMD|	Max |SMD|	% events |SMD| > 0.10
Log market capitalization	8.5×10^{-6}	5.1×10^{-5}	0
Book-to-market	3.9×10^{-6}	8.1×10^{-5}	0
Leverage	4.7×10^{-6}	3.2×10^{-5}	0
Return on assets	3.8×10^{-6}	2.2×10^{-5}	0

C. Descriptive statistics

This appendix shows the descriptive statistics of the PFAS-implicated firm sample and full non-PFAS firm universe, Appendix C.1 covers the monthly return panel underlying the H1 portfolio test, comparing returns, volatility, size, region, and regime. Appendix C.2 covers the event database underlying the matched-sample tests. Industry composition is documented in Section 5.1 and is not repeated here. Appendix C.3 shows descriptive statistics related to cross-sectional regression in H4.

C.1 Firm return and size descriptives

This appendix compares the return and size profiles of the 74 PFAS-implicated firms with the non-PFAS universe (Appendix Table C1) and describes how the PFAS sample is distributed across regions, exposure types, and time (Appendix Table C2).

Appendix Table C1: PFAS versus non-PFAS descriptive statistics

Return and size profiles of the 74 PFAS-implicated firms against the non-PFAS universe. Panel A shows monthly returns and size. Panel B shows daily returns. Panel C shows the market cap. distribution. Panel D shows a regional segmentation. The non-PFAS universe is slightly larger in the daily panel (Panel B) than in the monthly panel (Panel A), because 242 firms have daily return observations but never form a clean month-end return that clears the monthly screens.

Panel A: Monthly returns and size, PFAS versus non-PFAS

Group	*N* firms	*N* firm-months	Mean (%)	Median (%)	SD (%)	Ann. Mean (%)	Ann. SD (%)	Negative months (%)	Mean market cap. (USD bn)	Median market cap. (USD bn)
PFAS	74	19,899	0.95	0.61	10.32	11.44	35.76	46.1	32.00	9.12
Non-PFAS	86,601	11,868,050	1.12	0.00	18.16	13.45	62.92	49.3	2.58	0.10

Panel B: Daily returns, PFAS versus non-PFAS

Group	*N* Firms	*N* Firm-days	Mean (%/day)	SD (%/day)	Median (%/day)	Zero return days (%)
PFAS	74	412,240	0.05	2.37	0.00	5.5
Non-PFAS	86,843	224,261,506	0.09	4.87	0.00	16.4

Panel C: Market-capitalization distribution, PFAS versus non-PFAS

Group	Mean market cap. (USD bn)	Median market cap. (USD bn)	Share in largest universe quantile (%)	Mean market cap. within largest quantile (USD bn)
PFAS	32.0	9.1	86.5	37.0
Non-PFAS	2.6	0.1	19.9	12.4

Panel D: Descriptive statistics by Fama-French region, PFAS versus non-PFAS

Region	Group	*N* firms	*N* firm-months	Ann. mean (%)	Ann. SD (%)	Mean market cap. (USD bn)	Median market cap. (USD bn)
North America	PFAS	22	5,626	11.91	33.37	55.34	18.42
	Non-PFAS	11,737	1,490,314	15.42	62.93	5.34	0.43
Europe	PFAS	21	5,819	7.04	32.47	29.96	12.85
	Non-PFAS	10,682	1,421,701	5.75	59.68	2.29	0.10
Japan	PFAS	14	4,538	10.06	30.88	13.21	7.87
	Non-PFAS	4,955	1,067,935	9.24	44.17	1.23	0.12
Asia-Pacific	PFAS	2	278	20.49	38.05	8.54	4.88
	Non-PFAS	5,788	883,690	12.88	75.63	0.75	0.04
Emerging	PFAS	15	3,638	18.79	47.73	24.33	2.18
	Non-PFAS	53,439	7,004,410	15.31	64.20	2.50	0.08

Appendix Table C2: PFAS-implicated firms' descriptive statistics

Panel A through D show various segmentations of the PFAS-implicated firm sample.

Panel A: By geographic region		
Region	***N* firms**	**Share of total (%)**
North America	**22**	30
Europe	21	28
Emerging market	15	20
Japan	14	19
Asia-Pacific	**2**	3
Panel B: By exposure type		
Exposure type	***N* Firms**	**Share of total (%)**
Industrial	61	82
PFAS-intermediary	7	9
Pharmaceutical	6	8
Panel C: By annual ChemSec (2025) database appearance		
Year	***N* Firms**	**Share of total (%)**
2016	29	39
2020	8	11
2025	37	50
Panel D: By FIY, segmented by Bai and Perron (2003) regimes		
FIY	***N* Firms**	**Share of total (%)**
Pre-2015	33	44
Middle (2015–2022)	28	38
Post-2023	13	18

C.2 Event-database descriptives

The 144 events split roughly evenly across litigation, regulatory, and scientific channels and roughly 60%/40% between global and North-American jurisdictions. Only 16 events carry a quantified penalty, totaling USD 16.03 billion, and a single settlement (3M's USD 10.3 billion) dominates the total.

Appendix Table C3: Event database

Panel A cross-tabulates the 144-event database by jurisdiction region and primary channel. Panel B summarizes announced penalties among the subset of events carrying one. The penalty distribution is heavily right-skewed, with a handful of post-2023 mega-settlements dominating the total.

Panel A: Event distribution by jurisdiction and channel (144 raw events)				
Jurisdiction	**Litigation**	**Regulatory**	**Scientific**	**Total**
Global	17	21	49	87
North America	29	24	4	57
Total	46	45	53	144

Panel B: Distribution of announced monetary penalties, across events that carry one	
Statistic	**Value**
Events with a quantified penalty	16
Total announced penalties (USD bn)	16.03
Mean penalty (USD m)	1,001.9
Median penalty (USD m)	329.5
25th percentile (USD m)	15.0
75th percentile (USD m)	775.0
Largest single penalty (USD bn)	10.3

C.3 H4 regression descriptive statistics

These tables describe the firm-event panel that enters the cross-sectional H4 regression. CARs average near zero with fat tails (min −31.874, max 26.936). Penalties and mega-settlements are rare (12.4% and 2.0% of firm-events).

Appendix Table C4: Cross-sectional descriptive statistics of firm-events used for H4

Each row shows a variable in the H4 regression over the 5,596 post-FIY firm-events. CAR is the [−1, +1] cumulative abnormal return. Log market cap., book-to-market, leverage, and ROA are firm characteristics 90 days before the event. Named firm, has penalty, log-penalty (zero if none), and mega settlement (settlements larger than USD 1 billion) are event indicators. Q1 and Q3 are 25th and 75th percentiles.

Variable	***N***	**Mean**	**SD**	**Min**	**Q1**	**Median**	**Q3**	**Max**
CAR (pp)	5,596	−0.045	3.659	−31.874	−1.916	−0.128	1.689	26.936
Log market cap.	5,596	23.304	1.690	17.196	22.226	23.328	24.536	26.989
Book-to-market	5,596	0.667	0.472	0.013	0.307	0.563	0.955	2.739
Leverage	5,596	0.564	0.163	0.141	0.464	0.578	0.666	1.035
ROA	5,596	0.049	0.049	−0.291	0.022	0.045	0.074	0.240
Named firm (0/1)	5,596	0.009	0.093	0.000	0.000	0.000	0.000	1.000
Has penalty (0/1)	5,596	0.124	0.330	0.000	0.000	0.000	0.000	1.000
Log penalty	5,596	0.000	0.865	−4.664	0.000	0.000	0.000	4.106
Mega settlement (0/1)	5,596	0.020	0.139	0.000	0.000	0.000	0.000	1.000

All pairwise correlations are below 0.5 in absolute value, meaning that multicollinearity is not a concern. The strongest associations are mechanical (log penalty with mega settlement of 0.495) and the well-known size-value relationship (log market cap. with book-to-market, −0.390).

Appendix Table C5: H4 pairwise Pearson correlation matrix

Correlation matrix of the dependent variable CAR and the H4 regressors. All pairwise correlations are below 0.5 in absolute value, implying low correlation between variables.

Variable	CAR (pp)	Log market cap.	Book-to-market	Leverage	ROA	Named firm (0/1)	Has penalty (0/1)	Log penalty	Mega settlement (0/1)
CAR (pp)	1.000								
Log market cap.	−0.011	1.000							
Book-to-market	0.026	−0.390	1.000						
Leverage	−0.020	−0.108	−0.312	1.000					
ROA	0.001	0.292	−0.272	−0.226	1.000				
Named firm (0/1)	0.007	0.047	−0.044	0.044	0.037	1.000			
Has penalty (0/1)	0.013	0.005	0.009	0.004	0.008	0.058	1.000		
Log penalty	0.026	0.014	0.027	−0.041	0.003	−0.009	0.000	1.000	
Mega settlement (0/1)	0.056	0.016	0.023	−0.022	0.026	0.042	0.376	0.495	1.000

D. Long-term pricing (H1) robustness

The headline pooled value-weighted FF5 + MOM long-short annualized alpha is −5.28% per year (bootstrap t = −2.80, p = 0.005). This appendix shows it survives alternative risk adjustments (Appendix D.1), alternative portfolio construction (Appendix D.2), and is stable over time (Appendix D.3).

D.1 Alternative risk adjustments

The CAPM, FF3, FF3 + MOM, FF5, along with the baseline FF5 + MOM specification, are already shown in Table I and will not be repeated here. Appendix Table D1 therefore only shows the augmented q-factor (Hou et al., 2021), the mispricing-factor model (Stambaugh and Yuan, 2017), the quality-minus-junk (Asness, Frazzini, and Pedersen, 2019), betting-against-beta (Frazzini and Pedersen, 2014), traded-liquidity (Pástor and Stambaugh, 2003), green-minus-brown (Pástor, Stambaugh, and Taylor, 2022) and the Jensen, Kelly, and Pedersen (2023) factor themes. The discount is robust to every alternative factor (annual alphas −8.2% to −4.1%, all significant) in Appendix Table D1 Panel A. The negative green loadings confirm PFAS firms tilt brown, but adding the green-minus-brown factor does not absorb the alpha. Against an industry-matched chemicals short leg in Panel B, the discount persists near 7% per year under the standard alternative factors. It attenuates and loses significance only in the short 2009–2023 green-factor windows, where the sample is half as long and power is limited.

Appendix Table D1: Alternative factor models

α monthly and annualized is the regression intercept and its compounded equivalent. SE, *t*, and *p* are the parametric standard error, *t*-statistic, and asymptotic two-sided *p*-value. Bootstrap *p* and the 95% CI come from a stationary block bootstrap of the long-short series (robust to serial correlation and non-normality). The alternative factor libraries (QMJ, BAB, PS-liquidity, q^5, M4, and GMB) are U.S. series, so the specifications are only estimated on the North-America long-short panel. Green *β* (*t*) is the green-minus-brown loading, reported only for the two specifications that include it. The relevant baseline is the North America FF5 + MOM alpha of −7.39% (*t* = −3.80), not the pooled −5.28%. Panel A shows the specification against the baseline broad short-leg. Panel B restricts the short leg to the SIC2 of the PFAS.

Panel A: H1 long-short alpha under alternative factor models, broad-universe short leg, North America

Model	Sample	*N* months	*α* monthly (%)	*α* annual. (%)	SE (%)	*t* (*p*)	Boot-strap *p*	95% CI annualized (%)	R^2	Green *β* (*t*)
FF5 + MOM + QMJ	1998–2026	337	−0.709	−8.18	0.153	−4.62 (<0.001)	<0.001	[−11.62, −4.58]	0.345	–
FF5 + MOM + BAB	1998–2026	337	−0.649	−7.52	0.158	−4.12 (<0.001)	<0.001	[−11.18, −3.85]	0.337	–
FF5 + MOM + PS liq.	1998–2025	335	−0.644	−7.46	0.159	−4.06 (<0.001)	<0.001	[−10.99, −3.74]	0.359	–
FF5 + MOM + GMB	2009–2023	180	−0.352	−4.14	0.170	−2.06 (0.040)	0.038	[−8.30, −0.22]	0.362	−0.173 (−1.82)
QMJ + BAB + PS + GMB	2009–2023	180	−0.388	−4.56	0.171	−2.27 (0.024)	0.028	[−8.58, −0.53]	0.401	−0.237 (−2.60)
Hou et al. q^5 (stand-alone)	1998–2024	323	−0.687	−7.94	0.172	−4.00 (<0.001)	<0.001	[−11.87, −3.77]	0.294	–
M4 (stand-alone)	1998–2016	227	−0.429	−5.03	0.210	−2.04 (0.042)	0.037	[−9.98, −0.27]	0.312	–

Panel B: H1 long-short alpha under alternative factor models, industry neutral, chemicals-sector short leg, North America

Model	Sample	*N* months	*α* monthly (%)	*α* annual. (%)	SE (%)	*t* (*p*)	Boot-strap *p*	95% CI annualized (%)	R^2	Green *β* (*t*)
FF5 + MOM + QMJ	1998–2026	337	−0.656	−7.59	0.154	−4.25 (<0.001)	<0.001	[−11.10, −3.85]	0.326	–
FF5 + MOM + BAB	1998–2026	337	−0.639	−7.40	0.156	−4.10 (<0.001)	<0.001	[−11.05, −3.56]	0.327	–
FF5 + MOM + PS liq.	1998–2025	335	−0.621	−7.20	0.164	−3.80 (<0.001)	0.001	[−10.79, −3.28]	0.338	–
FF5 + MOM + GMB	2009–2023	180	−0.300	−3.54	0.168	−1.79 (0.076)	0.091	[−7.76, +0.60]	0.302	−0.107 (−0.97)
QMJ + BAB + PS + GMB	2009–2023	180	−0.294	−3.47	0.162	−1.81 (0.072)	0.091	[−7.42, +0.57]	0.331	−0.145 (−1.49)
Hou et al. q^5 (stand-alone)	1998–2024	323	−0.658	−7.62	0.176	−3.74 (<0.001)	<0.001	[−11.57, −3.22]	0.225	–
M4 (stand-alone)	1998–2016	227	−0.471	−5.50	0.220	−2.14 (0.033)	0.033	[−10.77, −0.49]	0.333	–

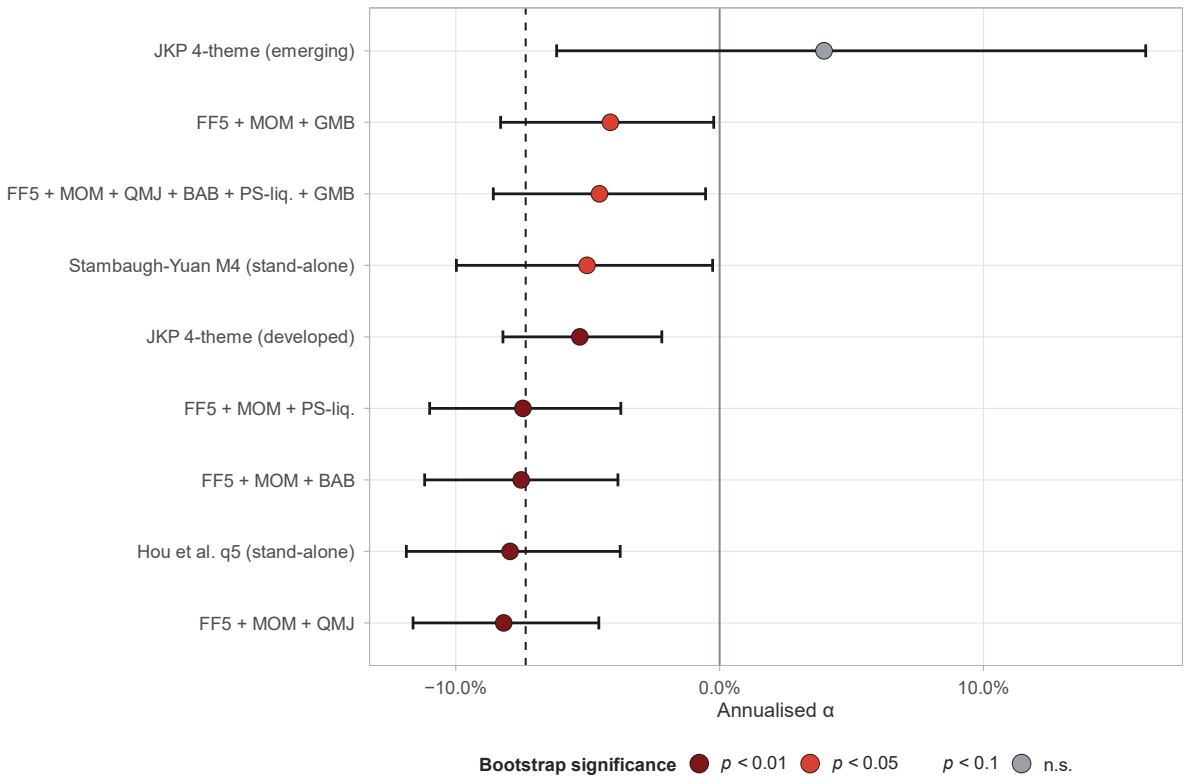
Appendix Table D2 shows the Jensen, Kelly, and Pedersen (2023) factor themes. In developed markets the discount survives the full JKP (2023) augmentation. The four-theme alpha −5.30% ($p < 0.001$), with the low-risk and investment themes absorbing the most individually (−4.02% and −3.40%, both significant). The premium partly overlaps those exposures but is not

explained away by them. The emerging-market alpha is positive and insignificant at every augmentation (15 firms and wide bootstrap intervals), confirming the discount is a developed-market phenomenon, consistent with regional H1 results.

Appendix Table D2: Jensen, Kelly, and Pedersen (*2023*) factor themes

Robustness of the H1 discount to the modern JKP (2023) factor set, estimated separately for the developed and emerging regions. Each row augments the market factor with one JKP (2023) theme cluster (quality, profitability, investment, low risk) or with all four at once as the final specification of each region. Region is the JKP (2023) developed or emerging aggregate. α is the augmented-model intercept (monthly and annualized), monthly and annualized is the regression intercept and its compounded equivalent. SE, t, and p are the parametric standard error, t-statistic, and two-sided p-value. Bootstrap p and the 95% CI come from a stationary block bootstrap of the long-short series (robust to serial correlation and non-normality). R^2 is the regression fit. Across regions, the cap-weighted four-theme annualized alpha is −3.87% (N = 335 months).

Region	Theme augmentation	*N* months	*α* month. (%)	*α* annual. (%)	SE (%)	*t* (*p*)	Boot-strap *p*	95% CI annual. (%)	R^2
Developed	mkt + quality	335	−0.073	−0.87	0.133	−0.55 (0.583)	0.592	[−4.20, +2.58]	0.048
	mkt + profitability	335	−0.193	−2.29	0.124	−1.55 (0.122)	0.130	[−5.27, +0.69]	0.046
	mkt + investment	335	−0.288	−3.40	0.113	−2.55 (0.011)	0.018	[−6.04, −0.56]	0.137
	mkt + low risk	335	−0.341	−4.02	0.118	−2.89 (0.004)	0.006	[−6.80, −1.09]	0.146
	mkt + all four themes	335	−0.453	−5.30	0.130	−3.49 (<0.001)	0.001	[−8.21, −2.19]	0.199
Emerging	mkt + quality	335	0.284	3.46	0.434	0.65 (0.513)	0.496	[−6.11, +14.10]	0.068
	mkt + profitability	335	0.222	2.70	0.440	0.50 (0.614)	0.615	[−6.94, +13.22]	0.065
	mkt + investment	335	0.014	0.17	0.432	0.03 (0.974)	0.971	[−9.28, +10.85]	0.062
	mkt + low risk	335	0.227	2.76	0.424	0.54 (0.592)	0.594	[−6.62, +13.32]	0.060
	mkt + all four themes	335	0.324	3.96	0.474	0.68 (0.495)	0.468	[−6.18, +16.15]	0.079



Appendix Figure D1: Alternative-factor-model robustness of the H1 long-short alpha

Annualized alpha of the value-weighted PFAS-minus-broad long-short strategy, re-estimated under alternative factor specifications on the North-America long-short with the broad short leg. Each point is the estimated alpha and each bar the 95% stationary-block-bootstrap interval. Specifications are ordered by alpha. The solid vertical line marks zero and the dashed line marks the −7.39% North-America FF5 + MOM baseline. The battery augments the FF5 + MOM model with Asness, Frazzini, and Pedersen's (2019) quality-minus-junk (QMJ), Frazzini and Pedersen's (2014) betting-against-beta (BAB), Pástor and Stambaugh's (2003) traded-liquidity, and Pástor, Stambaugh, and Taylor's (2022) green-minus-brown factors, individually and jointly, and replaces it with two complete alternative models, the q5 model (Hou et al., 2021) and Stambaugh and Yuan's (2017) four-factor mispricing model. Jensen, Kelly, and Pedersen's (2023) four-theme developed and emerging benchmarks are shown for comparison. The alpha remains negative and close to the headline magnitude under every augmentation and alternative complete model, and the developed-market four-theme benchmark agrees, while the emerging-market benchmark is positive and statistically insignificant. Market shading denotes bootstrap significance.

D.2 Portfolio construction

This section asks whether the discount depends on how the long-short portfolio is built or whether it is a genuine pricing effect, varying the weighting scheme, the composition of the short-leg, and the sample-timing rule.

Appendix Table D3: H1 long-short alpha under alternative portfolio construction

Pooled-sample robustness to how the long-short portfolio is built. Panel A contrasts equal and value-weighting of the baseline FF5 + MOM specification. Panel B restricts the short leg to chemicals-industry firms. The short leg is narrowed from all non-PFAS firms in the broad universe to non-PFAS chemicals peers, on two definitions: broad (all SIC2 codes the PFAS firms occupy) and narrow (SIC2 codes 28 chemicals, 30 plastic/rubber, 35 industrial machinery) Panel C drops the FIY timing rule, with the baseline re-estimated over all firm-months, ignoring the FIY rule.

Panel A: Equal- versus value-weighting

Statistic	Equal-weighted FF5 + MOM	Value-weighted (baseline) FF5 + MOM
α (monthly, %)	−0.741***	−0.451***
(SE)	(0.197)	(0.161)
α (annualized, %)	−8.54	−5.28
Bootstrap t	−3.77***	−2.80***
Bootstrap p	<0.001	0.005
95% CI (annual, %)	[−12.77, −4.34]	[−8.79, −1.62]
Observations	337	337

Panel B: Short leg restricted to chemicals-industry firms pooled, FF5 + MOM

Short leg	α annualized (%)	Bootstrap t	Bootstrap p	95% CI annualized (%)
Broad chemicals (PFAS SIC2 set)	−6.51	−3.11	0.002	[−10.47, −2.48]
Narrow chemicals (SIC2 28, 30, 35)	−6.04	−3.44	0.001	[−9.29, −2.56]

Panel C: FIY-agnostic specification

Statistic	FIY-agnostic FF5 + MOM	Post-FIY (baseline) FF5 + MOM
α (monthly, %)	−0.257*	−0.451***
(SE)	0.131	(0.161)
α (annualized, %)	−3.04	−5.28
Bootstrap t	−1.96*	−2.80***
Bootstrap p	0.051	0.005
95% CI (annual, %)	[−6.03, 0.01]	[−8.79, −1.62]
Observations	337	337

The premium is broader than a few large caps. Equal-weighting widens the alpha (FF5 + MOM equal-weighting −8.54% versus value-weighting −5.28%) in Panel A. Restricting the short leg to chemicals peers in Panel B widens the discount rather than shrinking it, meaning that the premium is specific to PFAS within the chemicals industry and not a generic chemical-sector effect. In other words, PFAS firms underperform even their own industry peers. Dropping the FIY rule in Panel C pulls the alpha toward zero (–3.04%, $p = 0.051$), exactly as expected if the rule correctly times the onset of pricing, in line with the EMH.

D.3 Temporal stability

This section stress tests the premium's stability over time, by assessing whether the discount is roughly constant across the sample. Neither Bai and Perron (2003) boundary marks a break in the pooled long-short alpha ($p = 0.299$ and $p = 0.425$). The post-2023 steepening (–8.71%

versus −4.84%) is visible but far from significant, and the 2015 boundary shows the alpha narrowing (−6.59% to −3.27%). The pooled discount is best read as a persistent and gradually-varying premium that does not switch at the regime dates. The 2015 and 2023 breaks result from an event-intensity phenomenon and are not a break relevant to the H1 alpha.

Appendix Table D4: Chow (1960) structural-break test on the pooled long-short alpha

The pooled −5.28% per year alpha is not a single regression but a market-cap-weighted average of five regional FF5 + MOM regressions, meaning that a standard break test has no single regression to split. Instead, the test is run on one combined monthly series whose average equals the pooled alpha, formed by removing the factor influence out of each regional long-short return and blending the five series with the headline weights. The break test therefore asks whether the average of this series shifts at each Bai and Perron (2003) break year. *F* is the Chow (1960) $F(1, 335)$. α pre and α post are the averages before and after the breaks.

Break year	*F*	*p*	α pre (%)	α post (%)	*N* pre	*N* post
2015	1.08	0.299	−6.59	−3.27	203	134
2023	0.64	0.425	−4.84	−8.71	299	38

E. Short-term pricing (H2) robustness

The headline H2 result is a per-event PFAS-minus-control differential of −0.240 (per-event p = 0.024, BMP t = −1.70, BMP p = 0.090). Appendix Table E1 first shows the differential under alternative sample definitions. Then the four families of robustness checks follow in four sections, parameter-choice sensitivity (E.1), PFAS-specificity (E.2), concentration (E.3), and event-dating legitimacy (E.4).

Appendix Table E1: Baseline differential by sample definition

This table shows how the event study's results differ across various definitions of the sample. Mean CAR is the pooled firm-event average. Per-event *t* and *p* are the paired test on event-level PFAS-minus-control differential. BMP *t* and *p* are the standardized cross-sectional test. Firm-specific events only keeps the events that name a specific firm, measuring all PFAS-implicated firms at those events. Narrow keeps only firms named in a firm-specific event. Broad keeps only events on the FIY-agnostic firm set.

Specification	*N* firms	*N* events	Mean CAR (%)	Differential (pp)	Per-event *t*	Per-event *p*	BMP *t*	BMP *p*
Post-FIY (baseline)	5,596	134	−0.045	−0.240	−2.28	0.024	−1.70	0.090
FIY-agnostic	7,855	134	0.012	−0.147	−1.38	0.171	−0.61	0.539
Firm-specific events only	2,297	39	0.096	0.099	0.59	0.562	0.75	0.454
Narrow	49	32	0.220	0.141	0.23	0.817	−0.32	0.748
Broad	7,221	124	0.046	−0.123	−1.15	0.253	−0.01	0.990

E.1 Parameter-choice sensitivity

The BMP statistic collapses from −1.70 to −0.34 once the window widens to two weeks, meaning the reaction is concentrated at the event window and not spread over the longer [−7, +7] window, as shown in Panel A. Panel B shows that the BMP statistic is invariant to the estimation-window length, with a BMP $t \approx -1.70$ throughout, meaning that the inference does not depend on the sigma estimate.

Appendix Table E2: Parameter-choice sensitivity

Panel A shows H2 recomputed using a wider [−7, +7] event window and compared to the baseline [−1, +1] window. Panel B varies the estimation-window length used to standardize abnormal returns in the H2 BMP test. Mean sigma is the average estimation-window return standard deviation.

Panel A: Event-window width sensitivity

Event window (days)	*N* Firm-events	Differential (pp)	Per-event *t*	Per-event *p*	BMP *t*	BMP *p*
[−1, +1]	5,596	−0.240	−2.28	0.024	−1.70	0.090
[−7, +7]	5,596	−0.135	−0.39	0.698	−0.34	0.735

Panel B: Sigma estimation window sensitivity

Estimation window (days)	*N* Firm-events	Mean sigma	Mean SAR	BMP *t*	BMP *p*
[−150, −10]	5,596	0.021	−0.022	−1.70	0.088
[−200, −10]	5,596	0.021	−0.022	−1.69	0.091
[−250, −10]	5,596	0.021	−0.022	−1.70	0.090

E.2 PFAS-specificity

The matched controls show a small positive mean CAR at the PFAS event dates (the anomaly described in Section 9.1). It does not undermine the event study's findings, which rest on the per-event PFAS-minus-control differential.

Appendix Table E3: Matched-control falsification

Three checks that assess whether the short-window reaction is specific to PFAS-implicated firms and not a broader chemical-industry effect. Panel A applies the H2 event-window abnormal return calculation to the matched control firms themselves at the PFAS event dates. Panel B assesses whether non-PFAS chemical-industry firms (those in the PFAS firms' two-digit SIC set but not PFAS-implicated firms themselves) react on the same dates. Panel C splits the PFAS firms' own reaction by exposure type (industrial, PFAS-intermediary, and pharmaceutical). All three panels report the BMP standardized cross-sectional test, and mean CAR is the firm-event average in percentage points.

Panel A: Matched-control falsification

Specification	*N* firm-events	*N* events	*N* control firms	Mean CAR (pp)	BMP *t*	BMP *p*
Matched controls at PFAS event dates	26,365	144	7,196	0.205	8.51	<0.001

Panel B: Non-PFAS chemicals falsification

Group	*N*	Mean CAR (pp)	*t*	*p*	BMP *t*	BMP *p*
Event-level	142 events	−0.045	−0.74	0.463	−1.10	0.274
PFAS baseline	134 events	−0.240	−2.28	0.024	−1.70	0.090

Panel C: Reaction by PFAS-implication type					
Subtype	*N* firm-events	*N* firms	Mean CAR (%)	BMP *t*	BMP *p*
Industrial	4,801	60	−0.037	−1.49	0.137
PFAS-intermediary	403	7	−0.085	−0.64	0.519
Pharmaceutical	392	6	−0.107	−0.42	0.676

E.3 Concentration

An event-study differential can be an artifact of a single dominant observation. If one firm supplies a large share of the firm-events, or one outlier event drives the average abnormal return, the pooled baseline number could reflect an effect that represents the outlier rather than the isolated PFAS effect. The checks below remove observations one at a time, and remove the firm-events around the sample's corporate restructuring, to confirm that the differential is diffuse and not concentrated.

Appendix Table E4: Leave-one-out and restructuring sensitivity

Two checks assess whether the baseline differential is driven by particular observations. Panel A recomputes the per-event differential (baseline −2.28) and BMP *t*-statistics (baseline −1.70), each time dropping a single PFAS firm, then a single event. Units is the number of such re-runs in the row, one for each firm dropped (73) or event dropped (134). Max |Δ differential *t*| is the largest absolute change in the differential *t* across those re-runs, and Max |Δ BMP *t*| is the same quantity for BMP *t*. A small value means that no single firm or event, when removed, moves the statistic by much, so the result is not concentrated in one observation. Panel B recomputes H2 after dropping only the PFAS firm-events that fall within a ±180-calendar-day window of corporate restructuring date: July 1, 2015 (the Chemours spin-off, two firm rows), and June 1, 2019 (the DuPont and Corteva split, one firm row). The three firms remain in the sample for all their other events. Only the four firm-events inside those two six-month (± 180 days) windows are removed.

Panel A: Leave-one-out influence			
Specification	*N*	Max |Δ differential *t*|	Max |Δ BMP *t*|
Leave-one-firm-out	73	0.17	0.35
Leave-one-event-out	134	0.31	0.52

Panel B: Restructuring sensitivity							
Specification	*N* Firm-events	Mean CAR (%)	Per-event differential	Per-event *t*	Per-event *p*	BMP *t*	BMP *p*
Baseline	5,596	−0.0454	−0.240	−2.279	0.024	−1.698	0.0896
Excluding DuPont, Corteva, and Chemours	5,592	−0.0452	−0.239	−2.263	0.025	−1.692	0.0907

The effect is concentrated outside the U.S., where the differential is −0.38pp (per-event p = 0.011) versus an insignificant −0.04pp for U.S. events, as shown in Appendix Table E5. In the raw-CAR matrix the one notable cell is North-America firms reacting negatively to global events (–0.50%, p = 0.054), consistent with global regulatory and scientific news pricing into the most exposed U.S. producers.

Appendix Table E5: Cross-jurisdiction matrix

Mean firm-level CAR by the event's jurisdiction and the reacting firm's region, plus the per-event differential split by jurisdiction. The matrix cells are raw firm CARs (not differentials).

Panel A: Per-event differential by jurisdiction

Specification	*N* firm-events	*N* events	Differential (pp)	Per-event *t*	Per-event *p*	BMP *t*	BMP *p*
U.S.	2,369	55	−0.039	−0.27	0.791	−0.31	0.758
Non-U.S.	3,227	79	−0.380	−2.60	0.011	−1.96	0.050

Panel B: Raw CAR by event jurisdiction × firm region

Event jurisdiction	Firm region	*N* events	Mean CAR (%)	*t*	*p*
Global	Global (non-North America)	7	0.40	1.45	0.147
Global	North America	7	−0.50	−1.93	0.054
North America	Global (non-North America)	32	0.06	0.58	0.564
North America	North America	32	0.16	1.23	0.219

E.4 Event-dating legitimacy

An event study is credible only if the measured reaction occurs at the event date and reflects the event itself. Several things could undermine that, and these are tested in Appendix Table E6.

Appendix Table E6: Event-dating legitimacy

Three checks that the measured reaction occurs at the event date and reflects the event itself. Panel A reports the abnormal return over a pre-event window [−15, −2] to detect information leakage. Panel B is a placebo test that re-runs the event study on randomly shifted dates. Each of the 763 valid draws applies a random calendar shift to every event date and re-estimates the per-event PFAS-minus-control differential *t*-statistic. The empirical *p* is the share of placebo draws whose differential *t* is at least as extreme as the baseline (−2.28). Panel C tests for calendar regularities, reporting day-of-week and month-of-year *F*-tests. BMP *t* and BMP *p* are the standardized cross-sectional test statistics.

Panel A: Pre-event drift

Window	*N*	Mean CAR (%)	BMP *t*	BMP *p*
Pre-event [−15, −2]	5,595	0.155	0.75	0.451

Panel B: Placebo event study on randomly shifted dates

Statistic	Value
Valid placebo draws	763
Baseline differential *t*	−2.28
Placebo differential *t*, SD	1.06
Placebo differential *t*, 95% interval	[−1.82, +1.99]
Empirical *p* (one-sided, $t \leq -2.28$)	0.001

Panel C: Calendar regularities

Check	*F*	*p*
Day-of-week	1.55	0.159
Month-of-year	2.36	0.007

F. Structural-break identification

The two Bai and Perron (2003) breaks at 2015 and 2023 define the regimes. Break placement is corroborated by three supporting diagnostics, which include a Chow (1960) level-break and Poisson rate-change test at each candidate year (Appendix F.1), the Pettitt non-parametric change-point test by event-type subset (Appendix F.2), and the BIC profile across zero to seven breaks (Appendix F.3).

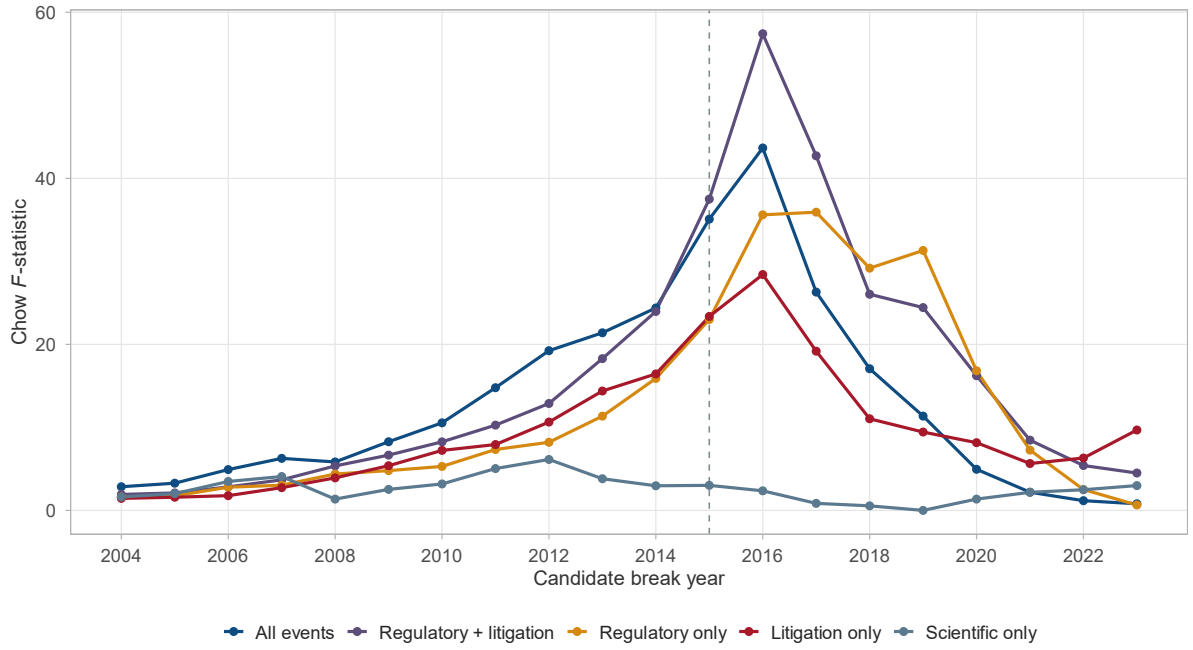
F.1 Single-break diagnostics on the annual event count

Both diagnostics peak in 2015–2016, where the arrival rate jumps from roughly two to ten events per year. The exhaustive Chow (1960) profile confirms the peak for every stream except the scientific-only stream.

Appendix Table F1: Single-break diagnostics on the annual PFAS-event count, 2001–2026

Chow *F*-test and two regime Poisson model at each candidate break year. Pre- and post-rates are events per year. Chow *F* and *p* test a level break at the candidate year. Poisson ΔBIC is the BIC improvement from a single-rate to a two-rate Poisson model. Candidate years are a-priori PFAS regulatory and litigation milestones from an event-density screening.

Candidate year	Pre-rate	Post-rate	Chow *F*	Chow *p*	Poisson ΔBIC
2016	2.33	9.91	43.65	<0.001	63.05
2015	2.21	9.42	35.09	<0.001	59.85
2017	2.88	9.80	26.30	<0.001	48.27
2018	3.35	9.67	17.07	<0.001	36.44
2019	3.78	9.50	11.36	0.003	26.73
2009	2.00	7.11	8.27	0.008	28.13
2006	1.60	6.48	4.92	0.036	19.42
2021	4.80	8.00	2.19	0.152	4.57



Appendix Figure F1: Chow (1960) *F*-statistic profile across candidate break years, by event-type subset

The Chow *F*-statistic for a single level break in the annual event count, computed at every candidate year from 2004 to 2023, with one line per event-type subset (all events, regulatory and litigation; regulatory only; litigation only; and scientific only). A higher *F* indicates stronger evidence of a structural break at that year. The four regulatory and litigation streams peak sharply in 2016–2017, placing the break in the immediate neighborhood of the first Bai and Perron (2003) break at 2015. The scientific-only stream peaks weakly at 2012 and never approaches the others, consistent with scientific evidence accumulating gradually rather than at a discrete date.

F.2 Non-parametric change-point test

The 2015 break is corroborated by all three diagnostics. Chow (1960) and Poisson peak at 2015–2016. Pettitt locates 2014–2015 across the full series and the regulatory and litigation subsets.

Appendix Table F2: Pettitt non-parametric change-point test by event-type subset

Change-point year is the year maximizing the Pettitt statistic. *p*-value is the approximate significance for no change-point.

Subset	Change point year	*p*-value
All events	2014	0.004
Regulatory + litigation	2015	0.001
Regulatory only	2015	0.002
Litigation only	2014	0.001
Scientific only	2011	0.124

F.3 BIC profile across break counts

BIC is minimized at two breaks, at the years 2015 and 2023. The 2015 break is a frequency jump, while the 2023 break is BIC-only and reflects a compositional shift toward litigation conclusions and mega-settlement events, not a rise in event intensity. The Chow (1960) and

Poisson candidate set spans 2006–2021 and the Pettitt test returns only the single dominant mid-2010s change point, meaning that neither addresses 2023. The second break therefore rests on the BIC-optimal two-break partition alone.

Appendix Table F3: Bai and Perron (2003) BIC profile across 0–7 breaks

BIC of the segmented event-intensity model for zero to seven breaks. BIC is minimized at 2 breaks, which are placed at the years 2015 and 2023 in the test.

N breaks	BIC
0	160.3
1	139.9
2	131.7
3	136.5
4	141.9
5	148.1
6	153.8
7	161.7

G. AI disclosure

What AI tools have been used and how?

Anthropic's Claude Opus 4.7 and 4.8 were used, with the supervisor's prior approval, to proofread the text and to assist with citation formatting, flagging inconsistencies across the document, and improving its style, structure, formatting, spelling, grammar, and language before submission.

In what ways have these tools contributed to increasing the quality of the thesis?

AI tools like Anthropic's Claude Opus 4.7 and 4.8 are effective at finding and suggesting corrections for proofreading items such as style, grammar, formatting, spelling, and inconsistencies across a long document, and can do so much faster and more consistently than a human can. For example, the tool could quickly identify whether the citations in the main text and the reference list aligned properly. The thesis has thereby gained improved clarity, style, and a higher degree of polish.

What potential risks were found using AI and what measures were taken to reduce these risks?

The tool sometimes generates information that is not correct. For example, it suggested that a reference was missing part of the full citation and proposed a correction that did not align with the actual source. It is therefore important to verify the tool's suggestions, so as not to implement something that is incorrect.

What are the insights gained from using AI tools in the thesis-writing process?

AI is a tool that can solve tasks consistently and quickly, but its suggestions should be verified before they are trusted.

Reference: Anthropic, 2026, Claude Opus 4.7 and 4.8, available at https://claude.ai, accessed June 5, 2026.