Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

What Does Mandatory Portfolio Disclosure Reveal? A Placebo Test of the 13F Regime

Boon Chuan Lim

boonchuan@singapore.to

Research Article

Keywords: Mandatory disclosure, Form 13F, institutional holdings, event study, placebo test, sovereign wealth funds

Posted Date: August 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10533940/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

What Does Mandatory Portfolio Disclosure Reveal? A Placebo Test of the 13F Regime

Boon Chuan Lim

Independent Researcher, Singapore

Corresponding author: boonchuan@singapore.to

ORCID: 0009-0005-8477-9393

Abstract

Section 13(f) requires large institutional managers to disclose their US equity holdings quarterly. The rule is fifty years old, applies identically to filers differing by more than two orders of magnitude in portfolio breadth, and is defended partly on the ground that disclosure is costly. I test whether these disclosures move prices using a negative control supplied by the regime: Norges Bank Investment Management files quarterly reports containing no positions under confidential treatment, with the information table following about a year later. Across 29,421 security-events from 894 filings by 26 managers during 2013–2026, newly disclosed and materially increased positions earn a two-day abnormal return of −7.0 basis points relative to the control when security-events are weighted equally, and between −2 and −6 basis points when filings are weighted equally. The control returns 0.7 basis points, close to the zero an empty filing should produce. None of the alternative weighting estimates is significant at five per cent, although some six-factor fixed-effect specifications reach that threshold; all have the same sign. The estimates differ because control filings hold about 1,100 securities each and treatment filings about 33. There is no gradient in position size relative to trading volume, where execution costs should bind most strongly. Under the preferred six-factor specifications, confidence intervals generally exclude effects more negative than about fifteen to sixteen basis points, while the market model is less precise. Security-event and filing weighting answer different questions and are reported side by side.

Keywords: Mandatory disclosure; Form 13F; institutional holdings; event study; placebo test; sovereign wealth funds

JEL classification: G14, G18, G23, G28

1 Introduction

Section 13(f) of the Securities Exchange Act, added by the 1975 amendments, requires any institutional investment manager exercising discretion over at least USD 100 million in Section 13(f) securities to report those holdings quarterly on Form 13F, within 45 days of the quarter end. Congress framed the rule as a data-collection measure: it would create a central repository of institutional holdings for regulatory and analytical use at a time when the institutionalisation of equity markets was a live policy concern. The rule has not been materially recalibrated since.

Its reach is uneven. Form 13F reports long positions in securities appearing on the Commission's Official Section 13(f) List. That list covers US-exchange-traded stocks, exchange-traded funds and closed-end funds, and also certain convertible debt securities, held equity options and warrants; shares of foreign issuers traded on qualifying US exchanges may be reportable. Short positions and written options are excluded, as are holdings that do not appear on the list (Securities and Exchange Commission 2026a). The USD 100 million threshold has never been indexed to inflation, so the filer population is far larger than Congress contemplated; the Commission proposed raising it to USD 3.5 billion in 2020, which by its own estimate would have removed about 4,500 of the 5,089 managers then filing, 89.2 per cent of filers, while retaining 90.8 per cent of the reported dollar value; the proposal was not adopted (Securities and Exchange Commission 2020).

The burden the rule imposes is not uniform. For an index-like portfolio, publication is close to costless: there is no strategy to reverse-engineer. For a concentrated book with an ongoing accumulation programme, publication is a tax on execution. Shi (2017) shows that disclosure degrades hedge fund performance, and the confidential treatment provisions of Rule 24b-2 exist precisely because the Commission accepts that some managers face a real cost. It follows that managers should position themselves in relation to the regime in proportion to the information content of their books, and that the observed distribution of filing behaviour is itself informative.

This paper asks the empirical question underneath the calibration debate. Do these disclosures move prices, and if so, where? A substantial literature documents copycat trading and proprietary costs associated with portfolio disclosure (Frank et al. 2004; Shi 2017; Verbeek and Wang 2013), while direct evidence on filing-date price reactions is thinner and less settled. Kwon (2022) shows that the burden falls differentially on long-horizon investors. What is missing is a credible benchmark: an event that shares the statutory timing and market environment of a real disclosure but conveys no position-level information.

The regime supplies one. Norges Bank Investment Management, manager of the Norwegian Government Pension Fund Global, files Form 13F under a request for confidential treatment. Its quarterly originals for the first and third quarters are placeholders: a single row, zero reported value, no securities. The information table arrives roughly twelve months later as an amendment. I verify this pattern across the fund's complete electronic filing history: first-quarter and third-quarter originals are placeholders in all thirteen sample years without exception. Because a manager requesting confidential treatment is not relieved of the obligation to file publicly, these empty filings appear on the EDGAR dissemination schedule exactly as informative filings do.

I use them as a within-regime negative control. Carrying forward the fund's most recently disclosed holdings, each empty filing generates a set of security-events that share the disclosure date, the filer, the estimation machinery and the securities of a real disclosure, and differ only in containing no information. Whatever abnormal return the negative control generates is the offset the design carries. The quantity of interest is the difference.

That difference is what the design identifies. Pooled across 29,421 security-events, newly disclosed and materially increased positions earn a two-day abnormal return of −6.4 basis points under a six-factor model. The control over the same window returns 0.7 basis points, close to zero in the preferred specification. Net of it, the effect is −7.0 basis points with a t-statistic of −1.88, a 95 per cent confidence interval from −14.4 to 0.3 and a p-value of 0.06. The estimate survives calendar-quarter fixed effects, restriction to securities held in both arms, security fixed effects, clustering by filing or by event date, and exclusion of every mapping the identifier audit flags.

It does not survive two changes that have as much claim to being right. Under a market model the same comparison gives −6.9 basis points with a t-statistic of −0.73. And when each filing rather than each security-event is weighted equally, the estimate falls to between −2 and −6 basis points with t-statistics between −0.5 and −1.1. The reason is arithmetic: the control comprises 48 filings holding about 1,100 securities each, the treatment 894 holding about 33, so weighting security-events equally gives a single control filing roughly thirty times the influence of a typical treatment filing on the point estimate. Clustering corrects the standard error and leaves that untouched.

The honest summary is therefore neither an effect nor its absence. Every estimate is negative, none is significant at five per cent under the preferred weighting, and the range runs from about −2 to −7 basis points. Part of that range reflects two distinct estimands rather than uncertainty: security-event weighting measures the response of the average disclosed position, filing weighting the response associated with the average filing. The distinction is a policy one as much as a statistical one: a regulator asking about the average disclosure event may prefer the filing-weighted figure, while an investor asking whether the average disclosed stock is traded against wants the security-weighted one. What can be said is that any effect is small, bounded above by roughly fifteen basis points in the most favourable specification, and that its measurement is more sensitive to weighting and factor model than to anything about the disclosures themselves.

One methodological point deserves emphasis because it nearly cost this paper its result. The control classes are large: a placeholder filing carries forward up to 2,400 securities, and estimating all of them across 48 filings is expensive. An earlier version sampled them at 250 securities per filing under a fixed seed and obtained a control of −5.6 basis points and a difference of −0.9, from which it concluded that the raw effect was an artefact of the estimation. Pricing the full control population reverses that: the control is 0.7 basis points and the difference is −7.0. Across one hundred independently seeded draws of the same size, the difference has a median of −6.8 and a fifth-to-ninety-fifth percentile range of −9.9 to −3.4; the original draw lies above that interval, among the draws closest to zero. A sampled control is a control estimated with error, and at this ratio of control to treatment observations that error was large enough to change the substantive conclusion.

The paper makes three contributions. First, it provides a within-regime negative control for disclosure event studies, constructed from filings that are mandatory, scheduled and empty, together with the diagnostics needed to assess its comparability. Second, it bounds a small, consistently negative but model- and weighting-sensitive short-window association. The preferred security-event estimate is −7.0 basis points, while filing-level estimates range from about −2 to −6. Third, it documents the institutional architecture by which large managers relate to the

regime, including one sovereign fund that discloses comprehensively but manages the timing of its mandatory filings, and another that does not appear in the filer universe at all.

The evidence is suggestive but inconclusive, and I treat it as such. I report a confidence interval and a minimum detectable effect for every specification, so that the claim is bounded rather than asserted. I report the negative control at every window, including the one where it misbehaves. Where the sample cannot support a test, I say so rather than reporting a statistic that a small number of clusters cannot sustain.

2 Related literature

The paper sits at the intersection of four strands: the theory of mandatory disclosure, the empirical measurement of proprietary costs in portfolio disclosure, the literature on confidential 13F filings, and the smaller body of work on announcement effects around the filings themselves. I take each in turn and state where the present design differs.

2.1 Disclosure regulation and proprietary cost

The theoretical case for mandating disclosure rests on the gap between private and social value. Verrecchia (1983) shows that when disclosure is costly, informed parties withhold information over an interval, so voluntary regimes are systematically incomplete. Admati and Pfleiderer (2000) extend this to a setting where disclosures by one firm affect the valuation of others, and show that the Nash equilibrium disclosure policies are frequently socially inefficient, which supplies the externality argument for compulsion. Both models make the proprietary cost of disclosure central: it is the reason voluntary disclosure fails and the reason mandated disclosure has a welfare cost to set against its benefit.

That cost is the object of the calibration debate around Section 13(f), and it is what the confidential treatment provisions of Rule 24b-2 concede. The present paper does not test the theory. It asks how large the market-visible component of that cost is for the filers the regime actually captures.

2.2 Proprietary costs of portfolio disclosure

The empirical literature on portfolio disclosure has approached the cost from the disclosing manager's side. Musto (1999) shows that the prospect of disclosure alters what managers hold, so the reported portfolio is itself shaped by the requirement to report it. Frank, Poterba, Shackelford and Shoven (2004) construct copycat mutual funds that buy the assets of actively managed funds as soon as holdings are disclosed, and find that after expenses the copycats earn returns statistically indistinguishable from, and possibly higher than, the funds they imitate. Verbeek and Wang (2013) extend the exercise across the mutual fund universe and find that mimicking portfolios beat the funds they copy by a few basis points on average. Both results imply that disclosed holdings retain enough value to be worth replicating.

Shi (2017) provides the sharpest evidence on the cost side. Using a difference-in-differences design around the point at which a hedge fund begins filing Form 13F, he documents a drop in fund performance and an increase in return correlation with same-style funds, concentrated among funds that disclose a greater fraction of assets or hold more illiquid stocks. He concludes that the cost is economically large and that the current form of the requirement could be improved by less frequent disclosure or longer delays.

These studies measure the cost as it accrues to the manager over subsequent quarters, through degraded performance. The present paper measures a different object: whether the disclosure itself moves the price of the disclosed security when it becomes public. The two are complementary and need not agree. A manager can bear a substantial execution cost through worse fills on later trades without any detectable price movement on the filing date, and that possibility is stated explicitly in the conclusion rather than assumed away.

2.3 Confidential filings

The literature closest to this paper's mechanism concerns managers who use Rule 24b-2 to delay disclosure. Agarwal, Jiang, Tang and Yang (2013) study confidential holdings, which are disclosed with a delay through amendments to Form 13F and are therefore absent from standard databases. They find that funds managing large risky portfolios with non-conventional strategies seek confidentiality more frequently, that the stocks concerned are disproportionately associated with information-sensitive events, and that confidential holdings outperform for

up to twelve months. Aragon, Hertzel and Shi (2013) ask directly why hedge funds avoid disclosure and reach a compatible conclusion. Together these papers establish that confidentiality is sought by managers with something to protect, and that the protection has value.

This paper uses the same institutional provision for the opposite purpose. Where that literature treats confidential filings as evidence of private information, I exploit the fact that a confidentiality request does not suspend the public filing obligation. The manager must still file publicly, and when the request covers the entire portfolio the resulting public filing is empty. Norges Bank Investment Management files on precisely this pattern, and its empty filings supply an event with the timing, filer and securities of a real disclosure and none of the content. To my knowledge these filings have not previously been used as a falsification benchmark.

2.4 Announcement effects and the missing benchmark

Work on the market reaction to 13F releases themselves is thinner and less settled. Sias (2004) documents institutional herding in the holdings data. Griffin and Xu (2009) use 13F holdings to evaluate hedge fund stock selection and find limited evidence of superior ability. Brown and Schwarz (2020) study the market reaction to the releases themselves and report abnormal volume and returns around the filing date, though estimates across this small literature vary with the sample and the window. Kwon (2022) argues that the disclosure rule burdens long-horizon equity investors differentially, because their positions are the easiest to replicate.

Christoffersen, Danesh and Musto (2015) approach the same question from the filer's side, asking why institutions use the full 45-day window rather than reporting early. Examining fourteen years of filings, they find that front-running and the concealment of voting power motivate delay, while copycatting does not. That result is the filer-side counterpart to the one reported here. If copycat trading were the operative channel, disclosure should move prices when it occurs, and managers should time filings to manage that reaction. Neither appears: they find no copycat motive in the timing decision, and I find no detectable price reaction at the release. The two findings are consistent, and together they suggest that whatever cost mandatory disclosure imposes operates through subsequent execution rather than through an announcement effect in the disclosed security.

What this literature lacks, and what limits its interpretation, is a benchmark event. Abnormal returns around 13F releases are typically measured against zero. At the sample sizes available in holdings data — tens of thousands of security-events — a systematic offset of a few basis points arising from factor-model misspecification will be statistically significant and indistinguishable from an announcement effect. The results below bear on this in two ways. Raw abnormal returns in the treatment sample are significant at one per cent across three windows, and the complete control does not explain them away: filings disclosing nothing return 0.7 basis points at the two-day window against −6.4 for the treatment. An earlier version of this study estimated the control from a random subsample and obtained −5.6, which did appear to explain them away. A benchmark is necessary, and estimating it carelessly is its own hazard.

2.5 Sovereign wealth funds as institutional investors

The sovereign wealth fund literature supplies the filers that motivate the design. Bortolotti, Fotak and Megginson (2015) document a discount in the announcement returns to sovereign fund equity investments relative to comparable private-sector investments, and attribute it to passivity in monitoring. Dewenter, Han and Malatesta (2010) and Kotter and Lel (2011) examine target selection and announcement effects with related findings. This work uses investment announcements as the event and treats 13F filings only as a source of holdings data.

The present paper inverts that relationship: the filing is the event, and the sovereign funds are of interest because they occupy the extremes of the disclosure spectrum. One publishes its entire portfolio voluntarily yet manages the timing of its mandatory filings; another discloses neither assets under management nor periodic returns yet files in full; a third does not appear in the filer universe at all. That range is what makes the institutional section of this paper possible, but the empirical results rest on the stratified sample of 26 managers, because single-filer designs cannot support the inference they appear to offer.

2.6 Inference

The econometric issues are standard but consequential here. MacKinlay (1997) sets out the event-study framework. Securities disclosed in the same filing share an event date and a filer, so their abnormal returns are cross-sectionally dependent; Petersen (2009) shows that ignoring such dependence understates standard errors severely, and I cluster by filing throughout and report a portfolio-level statistic alongside. Cameron, Gelbach and

Miller (2008) show that cluster-robust tests over-reject substantially when the number of clusters is between five and thirty. That result governs the reporting convention adopted below: every cell estimated from fewer than thirty filings is flagged, and no significance is claimed for coefficients identified off small cluster counts, however favourable the nominal statistic.

I also report the minimum detectable effect for each specification. A null result is interpretable only against the size of effect the design could have found, and in a literature where sample sizes range from dozens of filings to tens of thousands of security-events, that quantity varies by two orders of magnitude across studies that report similar-looking coefficients.

3 Institutional setting

3.1 The filing requirement

Form 13F is filed within 45 days of each calendar quarter end. Filers report each Section 13(f) security over which they exercise investment discretion, giving issuer name, class, CUSIP, market value, share or principal amount, discretion type and voting authority. A manager may file a combination report covering affiliated entities, in which case the cover page lists the other included managers. A manager whose holdings are reported entirely by another filer submits a notice report instead.

Two features of the mechanics matter for the empirical design. Values were reported in thousands of dollars until filings made before 3 January 2023 and in whole dollars thereafter; mixing the two inflates post-2023 figures by a factor of one thousand, which is invisible in summary statistics and fatal in any size-based test. The change is visible in the record: Norges Bank reports a portfolio total of 376,129,364 for the third quarter of 2022 and 414,862,674,207 for the fourth, a change of units rather than of portfolio value. Care is also required with the acceptance timestamp that determines when a filing becomes public. This study takes it from the acceptanceDateTime field of the submissions endpoint at data.sec.gov, which carries an explicit Z suffix and is Coordinated Universal Time; converting it to Eastern Time reproduces the value shown on the corresponding EDGAR filing index page. The index pages display Eastern Time without a timezone label, as do the Commission's stated operating hours and filing-day cut-offs, so a value read from a page must not be converted a second time. The basis of the stored field was verified against the whole sample rather than assumed: EDGAR accepts filings only between 06:00 and 22:00 Eastern (Securities and Exchange Commission 2026b), and of 2,112 stored timestamps none falls between 06:00 and 10:00, which is impossible were the values Eastern and expected if they are UTC, while 101 fall between 22:00 and midnight, which is impossible were they Eastern.

3.2 Confidential treatment and the empty filing

Rule 24b-2 permits a manager to request confidential treatment for some or all holdings. The request does not remove the public filing obligation: the manager must still submit a public Form 13F whose summary page indicates that information has been omitted, and must amend the public filing within six business days of any denial or expiry of the grant (Securities and Exchange Commission 2026a).

Norges Bank has used this provision throughout the sample. Its filings from 2009 onward state that the Form 13F is submitted pursuant to a request for confidential treatment, coordinated with Commission staff in light of the extensive public disclosure requirements applicable to the fund under Norwegian law. The resulting pattern is systematic. Original quarterly reports carry a single information-table entry while reporting the full portfolio value on the summary page: the second quarter of 2013 reports one entry against a stated total of USD 154.9 billion, and the fourth quarter of 2017 one entry against USD 279.0 billion. The fund therefore discloses on schedule what its US equity book is worth while withholding what it contains, which is the precise sense in which these filings are empty for the purposes of this study. Complete information tables, of between roughly 1,500 and 2,400 positions, arrive by amendment about twelve months later for first, second and third quarter report dates, and in early March for fourth quarter dates, shortly after the fund publishes its own annual holdings list. From the fourth quarter of 2022 the fourth-quarter original carries a full table at the statutory deadline, and from the second quarter of 2024 the second-quarter original does likewise, matching the fund's move to semi-annual voluntary reporting.

Appendix A lists 53 original filings, one per report period from 2013 to 2026, which are listed with accession numbers, acceptance timestamps, entry counts and delays. Of these, 46 are single-entry filings carrying the confidential-treatment flag, six are complete originals, and the original for the first quarter of 2013 is a zero-entry, unflagged exception. Within the machine-readable analysis sample, all thirteen first-quarter originals from

2014Q1 through 2026Q1 and all thirteen third-quarter originals from 2013Q3 through 2025Q3 are single-entry confidential-treatment filings. The regime changes only in the second- and fourth-quarter series: eleven of thirteen second-quarter and nine of thirteen fourth-quarter originals are single-entry, and the six complete originals first appear for the fourth quarter of 2022 and the second quarter of 2024. No complete original carries the confidential-treatment flag.

The delay from original to first complete filing has a median of 366 days. For first, second and third quarter reports it ranges from 331 to 385 days; for fourth-quarter reports it ranges from 20 to 36 days, falling shortly after the fund's own annual publication. The two most recent single-entry quarters have no amendment yet and are right-censored.

This is worth stating plainly. A sovereign investor that publishes extensive information on its holdings, returns and voting records (Norges Bank Investment Management 2026) nonetheless uses a confidentiality provision to control when its mandatory filings become public. Comprehensive disclosure of portfolio levels does not imply acceptance of externally timed disclosure of portfolio flows. This distinction is what makes the negative control available.

3.3 A filer absent from the record

A scan of the sampling frame returns no public Form 13F under the GIC names and identifiers searched. The confidential-treatment route can be excluded on structural grounds, since a confidentiality request does not remove the public filing obligation and no public Form 13F of any type appears. Beyond that, the public filing record does not establish whether relevant positions are reported through a separate legal entity or an external manager, fall outside the Section 13(f) universe, or are structured in some other way that generates no filing under the entities searched. The absence is a fact about the filings retrieved, not about the fund; I draw no inference about which explanation obtains. The point retained for the argument is narrow: the regime observes the managers it observes, and absence from the filer universe is not itself evidence about a manager's conduct.

4 Data

4.1 Sampling frame and filer selection

The frame is the population of Form 13F-HR filers on EDGAR in the fourth quarter of 2025 and the first quarter of 2026: 17,408 filings from 8,992 distinct managers. I draw a random subsample of 400 filers under a fixed seed and profile each one, recording position count, reported portfolio value, share of value in the ten largest positions and the Herfindahl index of position weights. Profiling uses the cover page and information table of the manager's most recent filing.

I stratify the profiled filers on portfolio breadth into four bands and draw six from each: the three largest by reported value plus three at random. The floor is USD 250 million in reported portfolio value. The rule is stated in advance and applied mechanically, so the sample is drawn from a defined frame rather than assembled from recognisable names. To this I add three sovereign wealth funds — Temasek Holdings, the Public Investment Fund and Norges Bank — which motivate the design and supply the negative control, and GIC, which returns the null result described above.

Table 1 Sampling frame and stratified selection

Stratum (positions)	Profiled	Median value	Median HHI	Selected
Concentrated (≤ 40)	109	$0.26bn	0.139	6
Mid (41–200)	181	$0.29bn	0.038	6
Broad (201–1,000)	88	$0.93bn	0.018	6
Index-like (> 1,000)	22	$9.37bn	0.010	6
Sovereign funds (added)	—	—	—	4

Notes. Frame is 8,992 distinct 13F-HR filers across 2025Q4 and 2026Q1. 400 filers profiled under a fixed random seed. Median HHI is the Herfindahl index of position weights computed from the information table. Selection floor is USD 250 million reported portfolio value. Selected filers span 19 positions (Akre Capital Management) to 6,077 (SEI Investments), a thirteen-fold range in median concentration across strata.

Values are normalised to whole dollars before any use. This matters: reported values are in thousands of dollars on filings made before 3 January 2023 and in whole dollars thereafter, so a panel that mixes the two inflates every post-2023 figure by a factor of one thousand. The convention is applied on filing date rather than report period, since the change attaches to when a filing was made. Section 3.1 shows the switch in Norges Bank's own record.

Collection returns 2,112 filing-level events across the 27 filer keys. Position-level records are parsed from the XML information tables, which become mandatory from the second quarter of 2013; that date sets the start of the sample. Values are normalised to dollars using the filing date and the January 2023 rule change. Positions reported by several included managers under one parent are summed before classification, which matters for combination filers such as Temasek, whose reports cover up to 41 subsidiary vehicles, and SEI, whose cover pages name 260 distinct managers across the sample.

4.2 Prices and factors

Security identifiers are mapped to tickers through the OpenFIGI mapping API (version 3, accessed without an API key), queried in July 2026. Identifiers are submitted as CUSIPs or, where the leading character is a letter, as CINS. The two are distinct identifier types at OpenFIGI and a CINS submitted as a CUSIP returns nothing, which silently discards foreign-domiciled issuers that trade normally on US exchanges: 321 identifiers in this sample take that form and none resolved until the distinction was applied. The rule is a heuristic rather than a definition, since the CUSIP issuer field is alphanumeric and a leading letter does not by itself establish that an identifier is a CINS; identifiers that fail under one type are retried under the other before being recorded as unresolved.

Where a query returns several records, they are ranked and the first is taken: records carrying a ticker are preferred to those without; records on a primary US venue are preferred to others; equity market sector is preferred; and common stock, ordinary shares, depositary receipts and similar types are preferred to other instrument types. Where no US record exists the first record with a ticker is accepted, which for securities that have ceased trading in the United States can return a foreign venue line. That fallback is the reason the unpriced set divides into 627 identifiers resolving only to a foreign line and 125 resolving to a US venue: the first group is a limit of identifier resolution, the second is delisting. Two further caveats apply. The mapping returns current identifiers, so a security that changed ticker during the sample may be matched to its present symbol rather than the one it carried on the event date, and reused symbols are a possibility. Appendix B reports an audit of every mapping in both arms against the checks that can be made from the data held: whether a ticker was returned at all, whether the issuer name on the filing resembles the name behind the ticker, whether the listing venue is in the United States, whether several identifiers resolve to one symbol, and whether the price series spans the events it is used to price.

Daily prices are taken from the Yahoo Finance chart endpoint (query1.finance.yahoo.com/v8/finance/chart), retrieved in July 2026. Returns are computed from the split- and dividend-adjusted closing series supplied in the adjclose field; the unadjusted close is retained alongside it, and the ratio of the two is used to identify corporate actions when classifying position changes. Non-trading days are absent from the series rather than carried forward, and returns are formed between consecutive available observations. Tickers are normalised by upper-casing and converting separators to hyphens, so that share classes reported as one form resolve to the vendor's convention. Securities that have ceased trading are not carried by the source at all, which is the survivorship limitation quantified below. An alternative public source was tested and found to serve an automated-access challenge rather than data, so no second series is available for cross-validation; this is a weakness of working from free sources and is stated as such.

Factor returns are the Fama–French five factors and the momentum factor from the Kenneth French data library, retrieved in July 2026, giving 15,833 daily observations.

The price panel covers 3,389 securities and 9.74 million security-days. Coverage is incomplete and the incompleteness is not random, so I decompose it rather than reporting a single rate. Two mechanisms operate. Securities that stopped trading — through acquisition, renaming or bankruptcy — are absent from public price sources; in the sovereign-fund subsample where I measured this directly, exited positions were priced at 29.1 per cent against 67.3 per cent for newly disclosed positions, which is the expected pattern given that exits are disproportionately caused by acquisition. Separately, a substantial share of identifiers no longer resolve to any US listing: of 752 unpriced tickers in that subsample, 627 resolved only to a foreign venue line and 125 to a US venue. The second mechanism is identifier resolution, not survivorship, and conflating them overstates delisting.

4.3 Sample flow

Table 2 traces the panel from candidate security-events to estimated abnormal returns. Every candidate event is taken through the estimation pipeline; none is sampled: 88.2 per cent resolve to a ticker, 81.5 per cent of those have a usable price series, and 98.7 per cent of those have sufficient estimation history, for an overall retention of 71.0 per cent. An earlier version sampled the large control classes at 250 securities per filing to save computation; Section 7.2 reports why that was abandoned. Attrition is not random with respect to position size: included positions average USD 242 million against USD 133 million for excluded ones, because large positions sit disproportionately in large, liquid, still-listed securities. The sample therefore over-represents the cases in which a disclosure should be easiest to absorb. This works towards the null and is treated as the study's principal limitation.

Table 2 Sample flow from candidate events to estimated abnormal returns

Stage	Events	Lost	Retained
Candidate security-events	155,168	—	—
With a mapped ticker	136,853	18,315	88.2%
With a usable price series	111,586	25,267	81.5%
With sufficient estimation history	110,085	1,501	98.7%
Overall	110,085	45,083	71.0%

Notes. Candidate events are new, increased, decreased, exited and negative-control positions above the USD 10 million floor. Mapping loss reflects identifiers that do not resolve to a US-listed ticker; price loss reflects mapped tickers without a usable US price series, principally foreign-venue fallback and, secondarily, delisting; estimation loss reflects insufficient history in the window [−140, −21] or absent prices through day twenty. Mean position value is USD 242 million for retained events and USD 133 million for discarded ones.

The estimation stage discards 18,315 security-events for want of a mapped ticker, 25,267 for want of a price series and 1,501 for insufficient estimation history, out of 155,168 candidates. The direction of the resulting bias is towards surviving securities. Because the negative control is drawn from the same price universe and subject to the same attrition, the difference-based estimates are less exposed to it than the levels are.

5 Method

5.1 Event definition

The event is the public dissemination of a filing on EDGAR. I take the acceptance timestamp, convert it from Coordinated Universal Time to Eastern Time with daylight saving handled explicitly, and assign day zero as the first NYSE trading session on or after that moment. Filings accepted after the 16:00 close trade on the next session. The trading calendar is computed rather than tabulated, including Good Friday from the Gregorian Easter algorithm, Juneteenth from 2022, Saturday holidays observed on the preceding Friday, and six special closures. Event dates are never taken from the period of report: filings are sometimes submitted years after the reporting quarter.

Where several filings cover the same filer-period, the earliest populated dissemination is the primary event and later filings are classified as restatements and excluded. Filings disclosed more than 365 days after the period end are classified as archival backfills and excluded: a manager submitting a decade of past quarters on a single day is not making dozens of disclosures. This rule removes one filer's bulk submission of its 2008–2019 history in 2026 and several late refilings by others. Norges Bank is exempt, since its delayed amendments are the object of study and carry their own classes.

5.2 Position classification

Each primary filing is compared with the same filer's primary filing for the previous quarter. Positions are classified as new, increased where shares rise by at least 20 per cent, decreased where they fall by at least 20 per cent, exited, or continuing. Positions in a filer's first primary filing carry no prior quarter and are flagged as sample-start artefacts rather than counted as new; without this the first filing of a large index-like manager

contributes thousands of spurious new positions. The tests use new and increased positions with a floor of USD 10 million.

Share counts are examined for corporate actions before classification. A holding reported as 1,000 shares one quarter and 2,000 the next has not necessarily been increased: a two-for-one split produces the same arithmetic. The price panel stores both the raw and the split- and dividend-adjusted close, and their ratio is the cumulative adjustment factor, so comparing that factor at two quarter ends recovers the corporate action between them. The comparison is made only where a session lies within ten days of each quarter end; without that restriction a security with no nearby price history is matched to the start of its series and the ratio reflects years of dividend reinvestment rather than an action. That restriction leaves 203,000 of 251,816 security-quarters with a usable factor.

The factor alone is not sufficient. A split, a spinoff and a special dividend all shift the adjusted price series, and only the split changes the share count. Of 46 security-quarters where the factor moved by more than fifteen per cent, most show share counts that did not follow, and several distinct securities share an identical factor ratio with unchanged holdings, which is the signature of a distribution rather than a split. Reclassification is therefore applied only where the share count moves with the price factor, within a tolerance of twenty per cent. Under that rule the number of reclassified positions is negligible and the results are unchanged. The limitation is stated plainly: adjusted-close ratios cannot separate splits from distributions without an events file, and this study has none, so corporate-action adjustment here is a check rather than a correction.

5.3 Abnormal returns

Abnormal returns are computed under two models: a market model, and the Fama and French (2015) five factors together with the momentum factor of Carhart (1997). Parameters are estimated over trading days −140 to −21 relative to day zero, requiring at least 120 observations. Cumulative abnormal returns are formed over [0, +1], [0, +3], [0, +5] and a drift window of [0, +20]. Abnormal volume is log volume less its estimation-window mean, a within-security measure rather than turnover, since I have no shares-outstanding series.

Standard errors are clustered by filing, since securities disclosed in the same filing share a date and a filer, following Petersen (2009); cross-sectional dependence within event windows is the concern identified by Kolari and Pynnönen (2010). Estimates weighting each filing equally are reported in Section 7.6; manager weighting is unavailable for the treatment-minus-control contrast, because every negative-control filing comes from one manager. Ordinary least squares and the cluster-robust covariance are implemented directly rather than through a statistical package and verified against a reference implementation to fifteen significant figures. The analysis code uses only the Python standard library. During preparation, the author used a generative AI assistant to help draft and debug data-collection and estimation code and to improve exposition. All outputs were reviewed, tested and edited by the author, who takes full responsibility for the analysis.

5.4 The negative control and its identifying assumption

For each Norges Bank placeholder filing I carry forward the securities in the fund's most recently disclosed information table and attach them to the placeholder's day zero. These events pass through the same estimation procedure, factor models and statutory calendar as informative filings, and differ in containing no positions. The expected abnormal return is zero.

The case for this benchmark rests on what the placeholder filings are, not merely on what they lack. They are not arbitrary non-events. They are mandatory filings by a large institutional manager, submitted on the same statutory calendar as informative filings, disseminated through the same EDGAR mechanism, and concentrated on the same quarter-end deadlines. Whatever a filing date does to returns mechanically — attention drawn by the dissemination feed, clustering of institutional activity around deadlines, quarter-boundary effects in the factor models, any error in assigning the event to a trading session — operates on both arms. A calendar-time portfolio of matched securities on random dates would share the securities but not the institutional setting; this control shares the setting. The systematic difference between the arms is the one of interest: whether the filing conveys position-level information.

Identifying assumption. Conditional on security characteristics, calendar timing and market conditions, residual abnormal returns around Norges Bank placeholder filings provide the counterfactual benchmark for residual abnormal returns around informative filings. The assumption is not innocuous. The control comprises 48 filings from one manager whose portfolio is broad and index-like, while the treatment comprises 894 filings from 26 managers holding between 19 and 6,077 positions. Section 7.3 reports estimates restricted to securities held in

both arms and under security and calendar fixed effects, and Section 7.7 splits the sample by filing session, so that the assumption is examined rather than asserted.

The observed control return may reflect a bundle of factor-model error, calendar effects, event-date error and characteristics specific to the securities carried forward from Norges Bank. The comparison does not separately identify these components, and composition differences between the two arms remain an identifying concern: the control is one manager with a broad, index-like book, the treatment 26 managers with books ranging from nineteen to 6,077 positions. A common event-dating error may partly cancel in the difference, although the different filing-session distributions prevent exact cancellation.

The estimand is therefore the difference between treatment and control, obtained by regressing the cumulative abnormal return on a treatment indicator across the pooled treatment and negative-control samples, with standard errors clustered by filing. I report the negative control level, the difference, its standard error and the minimum effect the design could detect at five per cent size and eighty per cent power. Inference uses exact Student-t tails; a normal approximation understates p-values materially when the number of clusters is small.

6 Results

6.1 Sample

The estimation produces 220,170 abnormal-return observations, one for each security-event under each of the two factor models. The treatment arm comprises 29,421 security-events from 894 filings, of which 5,676 events from 721 filings are newly disclosed positions and 23,745 events from 842 filings are material increases. The negative control comprises 53,542 security-events from 48 placeholder filings. Two further Norges Bank classes appear for context: 481 new positions from 45 stale amendments, disclosed about a year after the reporting quarter, and 54 from six at-deadline full filings.

Filing counts differ across stages of the pipeline and the differences are stated here rather than left implicit. The appendix lists one original per report period: 53 in total, of which 46 report a single entry, one (the first quarter of 2013) reports none at all, and six carry complete tables. The negative-control class is defined on filings rather than on periods, so it also admits single-entry amendments; it contains 48 filings, which after carrying forward the most recently disclosed holdings and applying the price and estimation-window requirements contribute 53,542 security-events. Complete tables carry between 1,577 and 2,351 positions. One reported total is anomalous: the third quarter of 2015 states 189,418 against a median of roughly 300 million in the same units, which is an error in the filing rather than in the data handling, and the summary total is not used in estimation. The stale amendments follow the same path: 55 are classified, eight are later filings for a period already covered and are reclassified as restatements, leaving 47 primary amendments, of which 46 contribute decreased positions and 45 contribute newly disclosed ones. Counts therefore vary by position class because not every filing contains every class. Appendix A records one original per report period; the filing-level path, including amendments, can be reproduced from the public filings using the analysis code.

6.2 The main result

Table 3 reports the central finding. The first row of each panel gives the negative control, the second the treatment net of it. Under the six-factor model the two-day difference is −7.0 basis points with a t-statistic of −1.88, a p-value of 0.06 and a 95 per cent confidence interval from −14.4 to 0.3; the design has 80 per cent power against a difference of 10.5 basis points. The control itself returns 0.7 basis points, close to zero in this specification. The three- and five-day windows give −7.2 and −12.2 basis points, with intervals of −18.0 to 3.6 and −26.1 to 1.7. Under the market model the two-day difference is −6.9 basis points but the standard error is two and a half times larger, and nothing in that panel is distinguishable from zero.

Table 3 Treatment net of the empty-filing negative control, basis points

Window	Control	Diff.	Std. err.	t	95% CI	MDE
Panel A: market model						
[0, +1]	3.7	−6.9	9.5	−0.73	−25.5 to 11.7	26.5
[0, +3]	−12.5	8.3	11.7	0.71	−14.6 to 31.2	32.8

Window	Control	Diff.	Std. err.	t	95% CI	MDE
[0, +5]	−10.0	4.2	13.6	0.31	−22.5 to 30.9	38.1
[0, +20]	16.3	−56.9	35.9	−1.59	−127.3 to 13.5	100.6
Panel B: five factors + momentum						
[0, +1]	0.7	−7.0	3.7	−1.88*	−14.4 to 0.3	10.5
[0, +3]	−2.9	−7.2	5.5	−1.31	−18.0 to 3.6	15.3
[0, +5]	−2.9	−12.2	7.1	−1.72*	−26.1 to 1.7	20.0
[0, +20]	−18.9	−42.3	16.9	−2.50**	−75.4 to −9.2	47.4

Notes. Each row regresses the cumulative abnormal return on a treatment indicator across pooled treatment and negative-control observations, with standard errors clustered by filing. Treatment is newly disclosed and materially increased positions; the negative control is securities carried forward to Norges Bank placeholder filings, which disclose no positions. Every eligible control security is estimated; none is sampled. n = 82,963 security-events from 942 filings at the three shorter windows and 80,505 from 920 at [0, +20]. MDE is the effect size the design would detect with 80 per cent power at 5 per cent size; it is a prospective property of the design, not a limit on the effects the data exclude, for which the confidence interval is the relevant quantity. A small MDE indicates a well-powered comparison: at 10.5 basis points the two-day six-factor specification would detect an effect of that size in four cases out of five, so a finding near zero there is more informative than the same finding from a design able to detect only large effects. * denotes 10 per cent and ** 5 per cent under an exact Student-t tail.

The comparison with the raw abnormal returns is instructive. Pooled treatment abnormal returns under the six-factor model are −6.4 basis points at [0, +1] (t = −3.33), −10.8 at [0, +3] and −16.1 at [0, +5]. Measured against zero these are significant at one per cent; measured against the control the two-day figure is slightly larger in magnitude but far less precisely estimated, because the control contributes its own sampling error to the difference. The t-statistic falls from −3.33 to −1.88 while the point estimate moves from −6.4 to −7.0. That is the cost of the benchmark and it is worth paying: a raw estimate carries whatever offset the estimation introduces, and without a control there is no way to tell an offset from a finding.

The cross-regime regression in Table 4 gives the same answer with controls. Taking the placeholder filings as the omitted category, the coefficient on treatment filings is −4.30 basis points with a t-statistic of −0.89. That is smaller in magnitude than the −7.0 of Table 3 and further from significance. Two things differ: the regression conditions on position value, past return, illiquidity and dollar volume, and it admits the stale-amendment and at-deadline classes to the sample, raising it to 90,426 observations. The coefficient is a further point in the range the estimate occupies across defensible specifications rather than a refutation of Table 3, and is reported as such. The continuous controls are reported as conditional associations and are not interpreted as validation tests of the specification. Past six-month return enters at −21.6 basis points per standard deviation and the Amihud (2002) measure at 1.3. These controls are applied symmetrically to both arms, but conditioning on them reduces the treatment coefficient from −7.0 to −4.30 basis points. The adjusted estimate remains negative and statistically indistinguishable from zero.

Table 4 Cross-regime regression, CAR[0, +1], five factors + momentum

Term	Coef. (bp)	Std. error	t	
Treatment filings	−4.30	4.85	−0.89	
Stale amendments	−0.73	6.94	−0.10	
At-deadline full filings	3.18	7.20	0.44	† 6 filings
New position indicator	1.12	4.18	0.27	
Log position value (sd)	0.63	2.11	0.30	
Past six-month return (sd)	−21.60	3.53	−6.12	***
Amihud illiquidity (sd)	1.33	0.69	1.93	*
Log dollar volume (sd)	−1.05	2.14	−0.49	

Notes. Omitted category is the empty-filing negative control. Standard errors clustered by filing. n = 90,426 security-events from 993 filings. The dagger marks a coefficient identified off six filings; cluster-robust standard errors are downward biased at that cluster count and no significance is reported for it. Continuous controls are standardised, so their coefficients are basis points per standard deviation; t-statistics are unaffected by the rescaling. * denotes 10 per cent; * denotes 1 per cent.**

6.3 Where the effect should be, and is not

If disclosure carries an execution cost, it should appear where a position is large relative to the stock's trading volume, and where the stock is thinly traded. Table 5 sorts the treatment arm on both. Neither sort produces a gradient. The difference between the top and bottom quartiles of position size measured in days of average dollar volume is 0.7 basis points (t = 0.14), with an approximate 95 per cent confidence interval from −9.1 to 10.5; for liquidity it is 7.6 basis points (t = 1.54), interval −2.1 to 17.3.

These nulls are informative rather than merely inconclusive. The intervals are narrow enough to rule out economically meaningful execution-cost gradients: a difference between the most and least demanding quartiles larger than about twenty basis points in either direction lies outside both intervals. If disclosure imposed an execution cost scaling with the difficulty of replicating a position, a gradient of that order is the smallest that would be economically interesting, and the data exclude it. What the data cannot exclude is a gradient of a few basis points, which would be of little practical consequence for a manager and none for a regulator.

Table 5 Mechanism sorts, CAR[0, +1], five factors + momentum

Quartile	Median	CAR (bp)	t	MDE	Filings
Panel A: position size in days of average dollar volume					
Q1 (smallest)	0.06	−4.7	−1.41	9.2	693
Q2	0.18	−8.2	−2.71	8.6	615
Q3	0.48	−7.0	−2.24	8.7	619
Q4 (largest)	2.16	−5.3	−1.58	9.4	625
Q4 − Q1		0.7	0.14		
Panel B: stock liquidity, mean daily dollar volume					
Q1 (least liquid)	$24m	−7.7	−2.32	9.3	603
Q2	$85m	−11.7	−3.80	8.7	648
Q3	$205m	−6.2	−2.10	8.3	687
Q4 (most liquid)	$658m	0.5	0.14	9.9	763
Q4 − Q1		7.6	1.54		

Notes. 7,353 to 7,356 security-events per quartile. Quartile CARs are measured against zero and therefore do not net out the control, which returns 0.7 basis points at this window; the quartile differences, which do not, are the interpretable quantities. The univariate continuous slope on days-of-volume is 0.26 basis points (t = 0.25) and is robust across winsorisation, trimming, tail-dropping, piecewise and quadratic specifications. Quartile contrasts are computed from unrounded estimates and may differ slightly from the difference of the rounded quartile means shown.

An earlier specification regressing the abnormal return jointly on log days-of-volume and log dollar volume returned slopes of 3.76 (t = 2.07) and 4.94 (t = 2.67) basis points per log unit, which appeared to contradict the flat quartile sorts. The contradiction is only apparent. Days-of-volume is defined as log position value less log dollar volume, so including both regressors is an algebraic reparametrisation of a regression on position value and dollar volume. The coefficient labelled days-of-volume in that specification is the position-value slope holding dollar volume fixed, and is not a measure of execution difficulty. The quantity the execution-cost hypothesis predicts on is the univariate slope on days-of-volume, and that slope is 0.26 basis points with a t-statistic of 0.25.

A broader set of sensitivity specifications confirms it. Winsorisation at zero, 0.5, 1, 2.5 and 5 per cent gives slopes between 0.13 and 1.80 basis points; one per cent trimming gives 0.77; dropping the extreme one and five per cent of days-of-volume gives 0.27 and −0.82; the standardised regressor gives 0.44; a piecewise-linear specification broken at the median gives 0.48 below and a further 0.10 above; and a quadratic gives −0.47 on both the linear and squared terms. No t-statistic in this set exceeds one. The null does not depend on outlier treatment or on functional form.

Binned means across deciles of days-of-volume tell the same story directly. The first decile, with a median position of 0.025 days of average dollar volume, averages −8.3 basis points; the tenth, at 5.4 days, averages −7.1. The interior deciles range between −13.0 and 1.5 basis points without pattern. Every decile sits near the pooled

treatment level, and the spread across them is small relative to their standard errors. Positions requiring more than five days of average volume to replicate move prices no differently from those requiring a fortieth of a day.

Abnormal volume around treatment filings is 0.017 log points, equivalent to trading about 1.7 per cent above the estimation-window average, with a cluster-robust t-statistic of 1.60. Disclosure is not followed by detectable trading activity either.

6.4 Manager weighting, and a caution about small portfolios

The estimates above weight security-events equally, which is standard but consequential when filers differ by more than two orders of magnitude in portfolio breadth. Six index-like managers contribute 16,285 of the 23,745 increased-position observations, so the pooled figure is close to a statement about the largest portfolios rather than about managers in general. This section reports what happens when each manager, rather than each security-event, carries equal weight.

Splitting the stratified sample by portfolio breadth produces an ordering consistent with the proprietary-cost hypothesis. For increased positions, treatment minus control is 26.9 basis points for concentrated managers, 18.1 for mid, −0.7 for broad and −3.4 for index-like: monotone in concentration, with the two upper strata individually significant at five per cent under filing-level clustering. The difference between the extreme strata, estimated within the treatment arm and therefore requiring no negative control at all, is 29.6 basis points with a cluster-robust t-statistic of 2.68 across 433 filings.

That result does not survive manager-level inference. Clustering by manager rather than by filing raises the standard error from 11.1 to 18.2 basis points and reduces the t-statistic to 1.62. A wild cluster bootstrap at manager level returns p = 0.17. Collapsing to one observation per manager — the mean across that manager's filings — and enumerating all 924 possible splits of twelve managers into two groups of six gives a difference of roughly 4 basis points with p near 0.9. The same exercise on new positions gives −4.4 basis points (p = 0.99) and on the pooled sample 2.8 basis points (p = 0.84). Inference becomes weaker when dependence is recognised at the manager level, while changing to equal-manager weighting also reduces the point estimate. The apparent concentration gradient is therefore sensitive to both dependence assumptions and observation weighting.

Table 6 The concentration contrast under alternative weighting, CAR[0, +1], increases

Procedure	Difference	Std. error	t or p	Units
Clustered by filing	29.6 bp	11.1	t = 2.68	433 filings
Clustered by manager	29.6 bp	18.2	t = 1.62	12 managers
Wild cluster bootstrap, manager	29.6 bp	—	p = 0.17	12 clusters
Exact permutation, manager means	4.3 bp	—	p ≈ 0.9	924 splits

Notes. The first three procedures retain the security-event-weighted point estimate and vary the treatment of within-manager dependence. The final procedure collapses the data to one mean per manager and therefore estimates an equal-manager-weighted contrast; its 4.3-basis-point estimate is not the same estimand as the 29.6-basis-point security-event-weighted contrast. The exact permutation enumerates every division of twelve managers into groups of six and requires no distributional or clustering assumption.

The manager means explain the divergence. Ranking the twelve managers by mean abnormal return on increased positions, the six concentrated managers occupy the three highest and the three lowest positions, while the six index-like managers occupy the entire middle. The rank sums are identical. Concentrated managers range from +99.7 to −112.6 basis points, a spread of 212 basis points; index-like managers range from +16.1 to −10.4, a spread of 26. The two groups have the same location and radically different dispersion.

The pattern can be explained mechanically by heterogeneous portfolio-level precision and observation weighting. A manager reporting eight securities in a filing produces a noisy average; one reporting two thousand produces a precise one. Portfolio concentration therefore predicts the variance of a manager's estimated abnormal return, not its level. Because observation weighting gives the managers with more filings and more securities greater influence, and because the concentrated managers with positive means happen to have more filings than those with negative means, an apparent gradient emerges from the weighting alone. I report the observation-weighted figure alongside the others precisely so that the divergence is visible, and I do not treat the gradient as a finding.

7 Robustness and falsification

7.1 The negative control at the drift window

The one window where the result is not clean is the twenty-day drift. Net of the negative control the treatment effect is −42.3 basis points with a t-statistic of −2.50 under the six-factor model, with a 95 per cent confidence interval from −75.4 to −9.2. Under the market model it is −56.9 basis points and not significant.

The control behaves differently at this horizon. Empty filings generate an abnormal return of −18.9 basis points over twenty days under the six-factor model and +16.3 under the market model, a spread of 35 basis points between two specifications that agree closely at the shorter windows. A well-behaved negative control would ideally remain close to zero and reasonably stable across plausible models. The divergence raises concern about model sensitivity and accumulated specification error at the longer horizon, though sampling variation and imperfect benchmark comparability remain possible explanations. The twenty-day results are therefore reported as robustness rather than as findings, and no claim of post-disclosure drift is made on this evidence. A design intended to test drift would need a different benchmark, most plausibly a characteristic-matched control portfolio rather than a factor model.

7.2 Sampling the negative control

The control classes are large. A placeholder filing carries forward up to 2,400 securities, and estimating all of them across 48 filings is costly for a falsification test whose expected value is zero. An earlier version therefore sampled them at 250 securities per filing under a fixed seed. That choice materially changed the result, and the episode is reported here because the same shortcut is available to any study using a large control portfolio.

Table 7 The estimate under complete and sampled control populations, CAR[0, +1]

Quantity	Estimate	Note
Complete control population	−7.0 bp	53,542 control security-events
Median across 100 sampled draws	−6.8 bp	250 securities per filing
5th to 95th percentile of draws	−9.9 to −3.4 bp	sign never differs
Minimum and maximum of draws	−12.4 to −0.5 bp	
Original sampled estimate	−0.9 bp	above the 95th percentile

Notes. Abnormal returns are computed per security and do not depend on which other securities are in the sample, so drawing repeated subsamples from the complete estimated population is equivalent to repeating the estimation under different seeds. Each draw takes 250 control securities per filing, the size used in the earlier version, and retains the treatment arm in full. Significance at 5 per cent is attained in 30 per cent of draws. The original sampled estimate of −0.9 basis points lies outside the fifth-to-ninety-fifth percentile interval but within the full range of the 100 draws: it is unusually close to zero relative to most draws, not outside the simulated distribution.

The sampled estimate used in the earlier version, −0.9 basis points, lies above the ninety-fifth percentile of the hundred-draw distribution and is among the estimates closest to zero, though it falls within the full range, whose maximum is −0.5. That draw was therefore unusual rather than impossible. The consequence was not a change of sign, since every draw is negative, but a change of substantive conclusion: at −0.9 the treatment estimate was indistinguishable from the control and the paper concluded that disclosure conveyed no price-relevant information, while at −7.0 it reports a marginal negative response. All results in this paper use the complete control population.

7.3 Specification grid

The identifying assumption of Section 5.4 is examined directly. Estimating the treatment-minus-control difference at [0, +1] under the six-factor model gives −7.0 basis points in the baseline specification, −5.4 with calendar-quarter fixed effects (t = −2.25), −6.6 when the sample is restricted to the 1,707 securities held in both arms (t = −1.68), and −7.7 when security fixed effects are added to that restricted sample (t = −1.98). The estimate is insensitive to time-invariant security composition among securities observed in both arms, and to calendar effects; neither rules out an event-dating error common to both arms, which is why the design relies on the control absorbing such errors rather than on their absence.

The pooled treatment result is reported across both factor models and all four windows in the main output, without selection among them. The models give similar negative point estimates at the two-day window but diverge at the three- and five-day windows, where the market model becomes positive and imprecise; at the drift window the control shows the specification to be unreliable in both. Sub-period estimates for 2013–2019, 2020–2022 and 2023–2026 are −3.7, −4.0 and −9.3 basis points under the six-factor model. These are treatment-arm returns measured against zero, not differences against the control, and are reported to bound period heterogeneity rather than to measure it; no net figure is computed from them.

The sub-periods illustrate the weighting problem of Section 6.4 in a second dimension. For 2013–2019 the filing-clustered statistic is −1.63 while the portfolio-level statistic, which gives each filing equal weight rather than each security, is +1.78; under the market model the pair is −1.86 and +1.30. The two disagree in sign on the same data, because in that period filings carrying many securities have negative average abnormal returns while filings carrying few have positive ones. Both statistics are treatment-arm returns measured against zero and neither is a difference against the control; the divergence is reported because it shows the unit of observation to be a substantive choice in this literature rather than a technicality.

7.4 Filer-level heterogeneity

The sovereign funds that motivated the design are individually uninformative and I do not lean on them. Temasek contributes 43 to 49 filings depending on position class, and its estimates carry minimum detectable effects of 47 to 127 basis points. The Public Investment Fund contributes 15 filings. Some individual-filer estimates reach conventional significance, but estimates based on fewer than thirty filings are not interpreted as significant because cluster-robust inference is unreliable at such small cluster counts. This is the reason the paper is about the regime rather than about any individual fund: single-filer designs of this kind cannot support the inference they appear to offer.

7.5 Mapping quality

The identifier audit of Appendix B quantifies observable warning signs and tests the sensitivity of the results to excluding flagged mappings. The question that bears on the results is narrower: does the estimate move when doubtful mappings are removed? Table 8 re-estimates the headline difference on progressively cleaner samples.

Table 8 Treatment minus control under mapping-quality exclusions, CAR[0, +1]

Sample	Diff (bp)	Std. err.	t	95% CI	Treat kept	Ctrl kept
All priced	−7.0	3.7	−1.88	−14.4 to 0.3	100%	100%
Drop similarity < 0.40	−6.9	3.7	−1.84	−14.2 to 0.5	98%	98%
Drop any audit flag	−6.9	3.8	−1.81	−14.3 to 0.6	94%	96%
High confidence only	−7.2	3.9	−1.84	−14.9 to 0.5	87%	92%

Notes. High confidence requires issuer name similarity of at least 0.85, a US listing venue, no other identifier resolving to the same ticker under a different issuer name, and a price series spanning the events. The final two columns give the share of treatment and control observations retained, since an exclusion falling unevenly on the arms would change what is compared.

The estimate stays within a third of a basis point of the baseline throughout, the interval barely moves, and the exclusions fall almost evenly on the two arms: the strictest sample retains 87 per cent of treatment and 92 per cent of control observations. Removing every mapping the audit flags, and then removing every mapping short of the highest confidence, leaves the estimate between −6.9 and −7.2 basis points. The headline estimate is insensitive to the mapping-quality exclusions implemented here.

7.6 Weighting and the unit of inference

Section 6.4 reports the concentration contrast under four procedures. The first three differ only in the unit treated as independent; the fourth, the exact permutation on manager means, also changes the estimand, since it compares managers rather than security-events. The estimate falls from 29.6 basis points to roughly 4 as the unit moves from the security-event to the manager. The same pattern appears in the main result, though with less consequence: the treatment-minus-control estimate at [0, +1] is −7.0 basis points under every clustering scheme, and only its standard error changes. Event-date clustering gives t = −1.86, close to the filing-clustered result of t = −1.88. Because each filing belongs to exactly one event date, filing clusters are nested within event-date clusters, so two-

way filing-and-date clustering adds no independent clustering dimension; it returns a standard error of 3.8 basis points against 3.7 for filing clustering alone. Two-way clustering on filing and security, which are not nested, gives t = −1.87. Clustering never alters the point estimate, which remains −7.0 basis points throughout. Manager-clustered inference is not credible because every control observation belongs to one manager, leaving only a single control-side manager cluster. Clustering on security alone gives a larger t-statistic but ignores the dependence among securities disclosed in the same filing; it is reported as a diagnostic and not preferred. The filing-clustered figure is used throughout.

The same question arises for the headline estimate itself, and there it matters more. The control arm holds 53,542 security-events across 48 filings, about 1,115 each; the treatment arm holds 29,421 across 894, about 33 each. Weighting security-events equally therefore gives one control filing roughly thirty-four times the influence of a typical treatment filing on the point estimate. Clustering by filing corrects the standard error for this and leaves the point estimate as it is.

Table 9 The headline estimate under alternative weighting, CAR[0, +1]

Estimand and weighting	Diff. (bp)	Std. err.	t	95% CI	Units
Average disclosed position					
security-event weighted	−7.0	3.7	−1.88	−14.4 to 0.3	82,963
Average filing					
difference of filing means	−5.7	5.3	−1.08	−16.0 to 4.6	942
filing means, 1% winsorised	−2.2	4.4	−0.51	−10.9 to 6.4	942
filing means, ≥ 5 securities	−3.6	4.0	−0.91	−11.5 to 4.2	708

Notes. Filing weighting takes the mean abnormal return of the securities in each filing and treats that as one observation. The difference of filing means applies no adjustment and is the primary filing-level figure; the two rows below it are sensitivity checks. Winsorising caps the extreme one per cent of filing means at the first and ninety-ninth percentiles rather than deleting them, and is a stronger intervention at filing level than at security level because a filing holding few securities has a volatile mean. Treatment filings average −4.89 basis points and control filings 0.78. Manager weighting is unavailable because every control filing comes from one manager.

Neither weighting is universally correct, because they answer different questions. Security-event weighting estimates the response of the average disclosed position; filing weighting estimates the response associated with the average filing. A manager who discloses two hundred positions contributes two hundred observations to the first and one to the second, and both are defensible depending on whether the object of interest is a position or a disclosure event. I report both and do not translate one into the other.

No estimate in the table is significant at five per cent and all four carry the same sign. For the average disclosed position the estimate is −7.0 basis points with an interval of −14.4 to 0.3; for the average filing it is −5.7 with an interval of −16.0 to 4.6. The two agree in magnitude more closely than the winsorised variants suggest, and the divergence among the filing-level rows reflects how much a small number of filings with few securities influences their mean.

One structural limit cannot be resolved by any procedure. The negative control comprises filings from a single manager, so the control arm contains only one manager cluster and manager-clustered inference is not credible. Manager fixed effects would also absorb the treatment contrast because treatment status is determined at the manager level in this design. No resampling scheme repairs this. It is a property of the design and is stated rather than worked around.

7.7 Filing session

A filing accepted while the market is open produces a day-zero close-to-close return that contains mostly pre-disclosure trading, so the three sessions are examined separately. Treatment minus control is −5.5 basis points for filings accepted before the open (t = −0.86; 148 treatment and 34 control filings) and −6.5 basis points for those accepted during the session (t = −1.37; 514 and 13 filings). Neither is individually significant, both carry the sign of the pooled estimate, and their magnitudes bracket it.

The after-close cell is not reported. It returns 7.0 basis points with a cluster-robust t-statistic of 1.88, carrying the opposite sign to every other cell in the study. It rests on a single control filing. Of the 48 placeholder filings, 34 were accepted before the open, 13 during the session and one after the close, so within that session the 232

treatment filings are being compared with the 1,264 securities of one placeholder disclosure on one date. The quantity is a difference between a large sample and a single cross-section, not a treatment effect, and no amount of clustering repairs it. I record it here because it is exactly the sort of result that survives a conventional significance filter: the t-statistic is large, the sample is not obviously small, and only the arm counts reveal the problem.

7.8 What the falsification tests establish

Three falsification checks are built into the design. The empty-filing negative control is the principal one and is discussed above. Stale amendments, which disclose positions approximately a year old and substantially stale relative to the current portfolio, are reported in Table 4 against the same baseline as the treatment filings, where the coefficient is −0.73 basis points with a t-statistic of −0.10.

That near-zero is a stronger result than it appears. A stale amendment is a filing in every mechanical respect: it is disseminated through EDGAR, it draws whatever attention the feed generates, and it contains a full information table of one to two thousand positions. What it does not contain is current information, since the holdings it reports are roughly a year old. If the reaction to a disclosure were driven by the act of filing — by mechanical price pressure, by attention, or by the arrival of a large document — stale amendments should produce it too. They do not. Whatever the treatment arm measures is tied to the information a filing carries rather than to the event of its filing.

At-deadline full filings, the third Norges Bank class, comprise six filings and are reported for completeness only.

8 Limitations

Six limitations bear on interpretation.

- **Partial observation.** The regime observes long positions in securities on the Official Section 13(f) List. Short positions and written options are excluded, as are holdings outside the list. A disclosed position may therefore be one leg of a structure whose other legs are unobserved, in which case its information content is lower than the reported size implies. This attenuates any true effect and works towards the null.
- **Selective attrition.** Of the candidate security-events, 71.0 per cent survive mapping, pricing and the estimation-window requirement. The remaining attrition is not random. Included positions average USD 242 million against USD 133 million for excluded ones, because large positions sit in large, liquid, still-listed securities. The sample therefore over-represents exactly the cases in which a disclosure should be easiest to absorb and least likely to move a price. This works towards the null and is the most serious limitation of the study.
- **Survivorship.** Prices come from a public source that drops securities which stopped trading. The bias runs towards survivors. Because the negative control is drawn from the same universe and suffers the same attrition, the difference estimates are less exposed than the levels, but the sample of exited positions is too depleted to support the secondary tests on decreases and exits, which I therefore do not report.
- **Unverified mappings.** Identifier-to-ticker mappings were screened using issuer-name similarity, listing venue, duplicate-symbol checks and price-series coverage. Excluding all flagged mappings leaves the headline estimate essentially unchanged (Table 8). Historical security-master data would nevertheless be required to determine the true mapping-error rate.
- **Identifier resolution.** Roughly a quarter of CUSIPs in a thirteen-year panel no longer resolve to a US-listed ticker in a public identifier database. This is a distinct mechanism from delisting and, in the subsample where I measured both, the larger of the two.
- **Cross-sectional breadth.** The treatment sample contains 26 filers: 24 selected through the stratified procedure and two purposively added sovereign funds. It spans the concentration distribution by construction, but filer-level inference rests on that number, and the estimates are best read as descriptive of a wide cross-section rather than as population parameters.

9 Conclusion

Mandatory disclosure under Section 13(f) is followed by abnormal returns that are negative, small, and not reliably distinguishable from zero at magnitudes this design can detect. Across 29,421 security-events from 894 filings by 26 managers spanning nineteen to 6,077 positions, newly disclosed and materially increased holdings earn two-day abnormal returns whose difference from filings disclosing nothing is negative but not conventionally significant in the principal specifications. In the preferred specification, the six-factor model at the two-day window, the estimate is −7.0 basis points with a 95 per cent confidence interval from −14.4 to 0.3 and a p-value of 0.06, and the design has 80 per cent power against a difference of 10.5 basis points. The market model at the same window gives −6.9 basis points with an interval two and a half times wider, and the twenty-day window is unreliable in both arms. Precision is specification-specific and each interval is reported with its specification. The average disclosed position and the average filing are distinct estimands and both are reported: a regulator asking about the average disclosure event may prefer the filing-weighted figure of −5.7 basis points, while an investor asking whether the average disclosed stock is traded against wants the security-weighted −7.0. Neither translates into the other. There is no gradient in position size relative to trading volume, nor in stock liquidity, which is where an execution cost should appear if one exists: across the reported sensitivity specifications the slope on days of average dollar volume lies between −0.8 and 1.8 basis points and no t-statistic exceeds one, and the first and tenth deciles of that variable differ by 1.2 basis points.

The policy reading is narrow and I state it narrowly. The analysis finds no robust positive price pressure of the kind expected if disclosure enabled front-running of continuing accumulation. This does not show that disclosure is costless: costs may arise through later execution, altered trading strategies or channels outside the event window, none of which this design observes. The preferred six-factor estimates suggest that the immediate market-visible component is small, although the market model is substantially less precise. For a statute defended partly on the ground that it burdens filers, that is a relevant fact.

The methodological reading is broader. The raw six-factor abnormal returns are negative and significant at one per cent across all three short windows. Against the complete control population they are not explained away: an empty filing returns 0.7 basis points at the two-day window, and the difference is −7.0. The methodological lesson is not the one this study set out to draw. The negative control does not account for the short-window treatment estimates once the whole control population is used. What the exercise does show is that estimating a large control portfolio from a single random subsample can introduce enough error to change the substantive interpretation. The 13F regime happens to supply such a benchmark through the interaction of Rule 24b-2 with the public filing obligation, though its comparability must be argued rather than assumed. Where such a benchmark is not available, a confidence interval and the design's power should at least be reported alongside the estimate, and the two should not be conflated: the interval states what the data exclude, the power calculation what the design could have found.

One institutional observation sits alongside the main result. A sovereign fund that publishes its complete holdings voluntarily nonetheless uses a confidentiality provision to control when its mandatory filings become public, which indicates that disclosure of portfolio levels and disclosure of portfolio flows are separate choices. A further fund does not appear in the filer universe under the identifiers searched, and the record does not establish why. Both illustrate that a regime applying one rule to filers with widely differing disclosure costs will meet varied filing behaviour, and that the behaviour is observable in the filing record even where its cause is not.

Statements and Declarations

Funding. The author received no external funding.

Competing interests. The author declares no competing interests. The author is employed outside academia; the research was conducted independently using only public data, and no employer resources, systems or data were used at any stage.

Author contributions. Boon Chuan Lim is the sole author and was responsible for conceptualization, methodology, software, data curation, formal analysis, investigation, writing—original draft, and writing—review and editing.

Data availability. The underlying regulatory filings, factor returns and market prices are publicly available from SEC EDGAR, the Kenneth French Data Library and the public price source described in Section 4.2. Derived datasets supporting the reported results are available from the author upon reasonable request.

Code availability. The complete collection, event-construction and estimation code is written using the Python standard library and is available from the author upon reasonable request.

Use of generative AI tools. During preparation, the author used a generative AI assistant to help draft and debug data-collection and estimation code and to improve the clarity of the exposition. All outputs were reviewed, tested and edited by the author, who takes full responsibility for the analysis and the final manuscript.

Appendix A Norges Bank Form 13F filings, 2013–2026

Table 10 lists one original filing per report period, 53 in total. For each it gives the accession number of the original, the EDGAR acceptance timestamp, the number of information-table entries, whether the confidential-treatment flag is set, and the accession number of the first filing for that period to carry a complete table, with the delay in days.

Table 10 Norges Bank Form 13F originals by report period

Quarter	Original accession	Accepted (UTC)	Entries	CT	First complete filing	Delay
2013Q1	0001374170-13-000014	2013-05-08 12:46	0		0001374170-14-000011	331
2013Q2	0001374170-13-000023	2013-08-07 12:46	1	Y	0001374170-14-000032	372
2013Q3	0001374170-13-000032	2013-11-12 12:25	1	Y	0001374170-14-000044	367
2013Q4	0001374170-14-000002	2014-02-07 14:05	1	Y	0001374170-14-000008	28
2014Q1	0001374170-14-000013	2014-05-09 11:42	1	Y	0001374170-15-000013	368
2014Q2	0001374170-14-000031	2014-08-08 10:59	1	Y	0001374170-15-000014	367
2014Q3	0001374170-14-000043	2014-11-07 12:29	1	Y	0001374170-15-000025	370
2014Q4	0001374170-15-000002	2015-02-11 12:11	1	Y	0001374170-15-000011	36
2015Q1	0001374170-15-000012	2015-05-11 11:09	1	Y	0001374170-16-000038	368
2015Q2	0001374170-15-000015	2015-08-10 11:34	1	Y	0001374170-16-000043	368
2015Q3	0001374170-15-000021	2015-11-06 15:18	1	Y	0001374170-16-000049	374
2015Q4	0001374170-16-000030	2016-02-10 11:28	1	Y	0001374170-16-000034	35
2016Q1	0001374170-16-000037	2016-05-10 11:44	1	Y	0001374170-17-000021	365
2016Q2	0001374170-16-000044	2016-08-12 12:36	1	Y	0001374170-17-000026	362
2016Q3	0001374170-16-000048	2016-11-14 11:49	1	Y	0001374170-17-000036	364
2016Q4	0001374170-17-000005	2017-02-09 17:03	1	Y	0001374170-17-000010	21
2017Q1	0001374170-17-000020	2017-05-10 12:21	1	Y	0001374170-18-000019	366
2017Q2	0001374170-17-000025	2017-08-02 12:11	1	Y	0001374170-18-000021	356
2017Q3	0001374170-17-000034	2017-11-07 13:46	1	Y	0001374170-18-000026	371
2017Q4	0001374170-18-000012	2018-02-09 13:24	1	Y	0001374170-18-000013	24
2018Q1	0001374170-18-000016	2018-05-09 13:21	1	Y	0001374170-19-000012	371
2018Q2	0001374170-18-000020	2018-07-23 11:36	1	Y	0001374170-19-000014	385
2018Q3	0001374170-18-000025	2018-11-09 12:27	1	Y	0001374170-19-000020	368
2018Q4	0001374170-19-000008	2019-02-08 13:06	1	Y	0001374170-19-000009	24
2019Q1	0001374170-19-000010	2019-05-09 14:53	1	Y	0001374170-20-000031	368
2019Q2	0001374170-19-000013	2019-08-09 10:50	1	Y	0001374170-20-000033	370
2019Q3	0001374170-19-000018	2019-11-08 12:17	1	Y	0001374170-20-000036	371
2019Q4	0001374170-20-000021	2020-02-12 14:52	1	Y	0001374170-20-000022	20
2020Q1	0001374170-20-000030	2020-05-08 10:43	1	Y	0001374170-21-000031	364
2020Q2	0001374170-20-000032	2020-08-07 20:20	1	Y	0001374170-21-000042	371
2020Q3	0001374170-20-000035	2020-11-10 13:48	1	Y	0001374170-21-000048	367
2020Q4	0001374170-21-000020	2021-02-10 11:33	1	Y	0001374170-21-000028	20
2021Q1	0001374170-21-000030	2021-05-07 12:57	1	Y	0001374170-22-000025	367
2021Q2	0001374170-21-000038	2021-08-10 14:05	1	Y	0001374170-22-000029	370

Quarter	Original accession	Accepted (UTC)	Entries	CT	First complete filing	Delay
2021Q3	0001374170-21-000047	2021-11-05 11:58	1	Y	0001374170-22-000035	374
2021Q4	0001374170-22-000019	2022-02-10 15:03	1	Y	0001374170-22-000020	27
2022Q1	0001374170-22-000024	2022-05-09 12:59	1	Y	0001374170-23-000034	368
2022Q2	0001374170-22-000028	2022-08-11 14:41	1	Y	0001374170-23-000040	368
2022Q3	0001374170-22-000034	2022-11-14 13:38	1	Y	0001374170-23-000043	364
2022Q4	0001374170-23-000006	2023-02-14 11:47	2130		0001374170-23-000006	0
2023Q1	0001374170-23-000036	2023-05-15 13:25	1	Y	0001374170-24-000035	366
2023Q2	0001374170-23-000039	2023-08-14 11:05	1	Y	0001374170-24-000039	366
2023Q3	0001374170-23-000042	2023-11-13 14:36	1	Y	0001374170-24-000072	366
2023Q4	0001374170-24-000018	2024-02-09 13:26	1942		0001374170-24-000018	0
2024Q1	0001374170-24-000034	2024-05-14 14:53	1	Y	0001374170-25-000050	365
2024Q2	0001374170-24-000038	2024-08-14 13:57	1930		0001374170-24-000038	0
2024Q3	0001374170-24-000071	2024-11-13 13:46	1	Y	0001374170-25-000102	364
2024Q4	0001374170-25-000027	2025-02-12 13:58	1845		0001374170-25-000027	0
2025Q1	0001374170-25-000049	2025-05-13 11:34	1	Y	0001374170-26-000021	363
2025Q2	0001374170-25-000079	2025-08-13 13:35	1707		0001374170-25-000079	0
2025Q3	0001374170-25-000104	2025-11-13 15:29	1	Y	-	—
2025Q4	0001374170-26-000012	2026-02-10 14:22	1577		0001374170-26-000012	0
2026Q1	0001374170-26-000023	2026-05-11 13:13	1	Y	-	—

Notes. Entries is the number of rows in the original filing's information table. CT marks originals carrying the confidential-treatment flag. Delay is days from the original to the first filing for that period with a complete table; a delay of zero means the original itself was complete. Two periods have no complete filing yet and are right-censored.

The table is the evidence for the characterisation in Section 3.2 and can be checked against EDGAR filing by filing. It shows 46 single-entry originals, one original with no entries (the first quarter of 2013), and six complete originals, all falling in 2022 or later; the confidential-treatment flag on all 46 single-entry originals and none of the six complete ones; complete tables of between 1,577 and 2,351 positions; and delays with a median of 366 days, ranging from 331 to 385 for first, second and third quarter reports and from 20 to 36 for fourth-quarter reports.

Appendix B Audit of identifier mappings

Mapping security identifiers to tradable tickers is a key limitation in a study built on public data. This appendix quantifies observable warning signs across both arms and tests whether the results change when flagged mappings are excluded.

Table 11 Automated audit of identifier mappings

Check	All	Share	Priced only
Identifiers examined	5,953	—	3,314
No ticker returned	1,451	24.4%	—
No usable price series	1,188	20.0%	—
Issuer name similarity below 0.60	420	7.1%	11.2%
Issuer name similarity below 0.40	—	—	4.8%
Resolved to a non-US venue	1,005	16.9%	0.4%
Ticker shared across different issuers	40	0.7%	1.0%
Price series does not span the events	18	0.3%	0.5%
Any of the last four	—	—	12.4%

Notes. The first column covers all 5,953 distinct identifiers appearing in newly disclosed or materially increased positions or in the negative control. The final column covers the 3,314 that resolved to a ticker with a usable price series and therefore contributed to the estimates; for those, the first two rows are empty by construction. Issuer name similarity compares the

name as filed with the name returned by the identifier service after stripping corporate suffixes, using a sequence-matching ratio. Identifiers whose issuer contributes several identifiers, 22.3 per cent of the total, are not counted as a fault: multiple share classes and parallel ordinary and depositary lines are ordinary features of the data.

Three readings follow. First, the losses in the first two rows are the attrition already reported in Table 2 and are not mapping errors; an identifier that returns no ticker is absent from the results rather than wrong within them. Second, among the identifiers that did contribute, the resolution problems that dominate the full population largely disappear: non-US venue falls from 16.9 to 0.4 per cent, because a mapping resolved to a foreign line generally has no usable US price series and is dropped before estimation. Third, the one check that speaks to a mapping being wrong rather than missing is issuer name similarity, and it flags 11.2 per cent of contributing identifiers at a threshold of 0.60 and 4.8 per cent at 0.40.

These percentages should not be read as error rates. The similarity measure is crude: a company renamed after a merger, an abbreviation, a truncated filing name or a holding-company style scores poorly while the mapping is correct. Over a thirteen-year panel such cases are common, since the filing records the issuer name in force at the time and the identifier service returns the name in force now. Automated checks may also fail to flag some incorrect mappings.

A worklist of 297 identifiers was assembled — the 209 severe name discrepancies, 26 tickers reached from identifiers with differing issuer names, 17 non-US resolutions, 15 coverage failures and 30 randomly drawn unflagged mappings as a control on the screen — each with links to the relevant filings and identifier records. Adjudicating the flagged cases requires the issuer's corporate history at the filing date, including renamings, reincorporations, acquisitions and splits. Because that history is not reliably recoverable from public filings alone at the scale required, the flagged cases are reported as undetermined rather than classified as errors or correct mappings.

The audit therefore quantifies observable warning signs rather than the true mapping-error rate. Among contributing identifiers, 11.2 per cent carry an issuer-name similarity flag at the 0.60 threshold and 4.8 per cent at the 0.40 threshold; cross-issuer ticker collisions and coverage failures each affect around one per cent. Rates are reported separately for the treatment and negative-control arms, since differences across arms would bias the contrast rather than add symmetric noise. These diagnostics are summarised in Table 11; the underlying identifier-level audit can be provided by the author upon reasonable request.

References

Admati AR, Pfleiderer P (2000) Forcing firms to talk: financial disclosure regulation and externalities. Rev Financ Stud 13(3):479–519. https://doi.org/10.1093/rfs/13.3.479

Agarwal V, Jiang W, Tang Y, Yang B (2013) Uncovering hedge fund skill from the portfolio holdings they hide. J Finance 68(2):739–783. https://doi.org/10.1111/jofi.12012

Amihud Y (2002) Illiquidity and stock returns: cross-section and time-series effects. J Financ Mark 5(1):31–56. https://doi.org/10.1016/S1386-4181(01)00024-6

Aragon GO, Hertzel M, Shi Z (2013) Why do hedge funds avoid disclosure? Evidence from confidential 13F filings. J Financ Quant Anal 48(5):1499–1518. https://doi.org/10.1017/S0022109013000483

Bortolotti B, Fotak V, Megginson WL (2015) The sovereign wealth fund discount: evidence from public equity investments. Rev Financ Stud 28(11):2993–3035. https://doi.org/10.1093/rfs/hhv036

Brown SJ, Schwarz C (2020) Do market participants care about portfolio disclosure? Evidence from hedge funds' 13F filings. SSRN Working Paper 1683628. https://doi.org/10.2139/ssrn.1683628

Cameron AC, Gelbach JB, Miller DL (2008) Bootstrap-based improvements for inference with clustered errors. Rev Econ Stat 90(3):414–427. https://doi.org/10.1162/rest.90.3.414

Carhart MM (1997) On persistence in mutual fund performance. J Finance 52(1):57–82. https://doi.org/10.1111/j.1540-6261.1997.tb03808.x

Christoffersen SEK, Danesh E, Musto DK (2015) Why do institutions delay reporting their shareholdings? Evidence from Form 13F. Rotman School of Management Working Paper 2661535. https://doi.org/10.2139/ssrn.2661535

Dewenter KL, Han X, Malatesta PH (2010) Firm values and sovereign wealth fund investments. J Financ Econ 98(2):256–278. https://doi.org/10.1016/j.jfineco.2010.05.006

Fama EF, French KR (2015) A five-factor asset pricing model. J Financ Econ 116(1):1–22. https://doi.org/10.1016/j.jfineco.2014.10.010

Frank MM, Poterba JM, Shackelford DA, Shoven JB (2004) Copycat funds: information disclosure regulation and the returns to active management in the mutual fund industry. J Law Econ 47(2):515–541. https://doi.org/10.1086/422982

Griffin JM, Xu J (2009) How smart are the smart guys? A unique view from hedge fund stock holdings. Rev Financ Stud 22(7):2531–2570. https://doi.org/10.1093/rfs/hhp026

Kolari JW, Pynnönen S (2010) Event study testing with cross-sectional correlation of abnormal returns. Rev Financ Stud 23(11):3996–4025. https://doi.org/10.1093/rfs/hhq072

Kotter J, Lel U (2011) Friends or foes? Target selection decisions of sovereign wealth funds and their consequences. J Financ Econ 101(2):360–381. https://doi.org/10.1016/j.jfineco.2011.03.007

Kwon D (2022) The differential effects of the 13F disclosure rule on institutional investors. SSRN Working Paper 4095482. https://doi.org/10.2139/ssrn.4095482

MacKinlay AC (1997) Event studies in economics and finance. J Econ Lit 35(1):13–39. https://www.jstor.org/stable/2729691

Musto DK (1999) Investment decisions depend on portfolio disclosures. J Finance 54(3):935–952. https://doi.org/10.1111/0022-1082.00132

Norges Bank Investment Management (2026) Our voting records. Oslo. Voting instructions are published five days before each shareholder meeting; the fund's holdings are published with its periodic reports. https://www.nbim.no/en/responsible-investment/voting/our-voting-records/. Accessed July 2026

Petersen MA (2009) Estimating standard errors in finance panel data sets: comparing approaches. Rev Financ Stud 22(1):435–480. https://doi.org/10.1093/rfs/hhn053

Securities and Exchange Commission (2020) SEC proposes to update Form 13F reporting threshold for institutional investment managers. Press Release 2020-152. https://www.sec.gov/newsroom/press-releases/2020-152

Securities and Exchange Commission (2026a) Frequently asked questions about Form 13F. Division of Investment Management staff guidance. https://www.sec.gov/rules-regulations/staff-guidance/division-investment-management-frequently-asked-questions/frequently-asked-questions-about-form-13f

Securities and Exchange Commission (2026b) Determine the status of my filing. Filer support resources, including EDGAR operating hours and filing cut-offs. https://www.sec.gov/submit-filings/filer-support-resources/how-do-i-guides/determine-status-my-filing

Shi Z (2017) The impact of portfolio disclosure on hedge fund performance. J Financ Econ 126(1):36–53. https://doi.org/10.1016/j.jfineco.2017.06.001

Sias RW (2004) Institutional herding. Rev Financ Stud 17(1):165–206. https://doi.org/10.1093/rfs/hhg035

Verbeek M, Wang Y (2013) Better than the original? The relative success of copycat funds. J Bank Finance 37(9):3454–3471. https://doi.org/10.1016/j.jbankfin.2013.05.003

Verrecchia RE (1983) Discretionary disclosure. J Account Econ 5:179–194. https://doi.org/10.1016/0165-4101(83)90011-3