

ANCHORED OWNERS, DRIFTING DISCOUNTS

INVESTOR SENTIMENT IN SWEDISH CLOSED-END FUNDS

HECTOR FYLKNER

DANIEL ÅKESSON

Bachelor Thesis

Stockholm School of Economics

2026



Anchored Owners, Drifting Discounts: Investor Sentiment in Swedish Closed-End Funds

Abstract:

Using a monthly panel of 18 Swedish closed-end funds from 2017 to 2025, we test whether investor sentiment predicts NAV discount changes in a concentrated-ownership setting. We construct a market-based sentiment index and find that high sentiment is followed by discount widening two to three months later, although this time-series relation weakens under checks for common monthly shocks and sentiment persistence. The pattern is stronger among funds with greater unlisted portfolio exposure, whose portfolios are harder for outside investors to replicate. By contrast, the ownership test does not support a simple retail-investor channel: funds with a larger non-retail shareholder base are not less sensitive to sentiment. The evidence suggests that sentiment-related discount dynamics can arise outside retail-concentrated closed-end fund settings, especially where arbitrage is more limited.

Keywords:

Close-end funds, NAV discounts, Investor sentiment, Arbitrage frictions, Ownership concentration

Authors:

Hector Fylkner (25998)
Daniel Åkesson (25996)

Tutor:

Marieke Bos, Associate Professor, Department of Finance

Examiner:

Ramin Baghai, Professor, Department of Finance

Bachelor Thesis
Bachelor Program in Business and Economics
Stockholm School of Economics
© Hector Fylkner and Daniel Åkesson, 2026

1 Introduction

Closed-end funds trade separately from the portfolios they hold, making the gap between share prices and net asset values directly observable. Lee, Shleifer, and Thaler (1991) argue that this NAV discount partly reflects individual investor sentiment: because U.S. closed-end funds are disproportionately held by retail investors, shifts in retail demand can move fund prices without proportionally moving the prices of the underlying assets. De Long, Shleifer, Summers, and Waldmann (1990) show why such mispricing can persist even when sophisticated arbitrageurs are present: noise-trader demand can push prices further away from fundamentals before correction occurs. Baker and Wurgler (2006) measure aggregate investor sentiment using a composite index of market-based proxies and show that sentiment predicts returns especially for stocks that are hard to value and difficult to arbitrage. Pontiff (1996) similarly emphasizes that closed-end fund discounts are shaped by costly arbitrage. The common implication is that sentiment should matter most for discounts when fund prices are exposed to investor demand and the underlying portfolio is costly to arbitrage. What remains less clear is whether the sentiment-discount link documented in Lee, Shleifer, and Thaler (1991) depends on the retail-investor setting in which it was first established, or whether similar dynamics also appear among closed-end funds with concentrated non-retail ownership.

Swedish investmentbolag provide such a setting. They differ institutionally from the U.S. closed-end funds studied by Lee, Shleifer, and Thaler (1991), but they are conceptually similar for the purpose of this study. Their shares trade separately from the value of their underlying portfolios, which makes NAV discounts and premia observable. We therefore refer to them as closed-end funds throughout the paper. Unlike that U.S. setting, Swedish closed-end funds are typically anchored by stable family, foundation, and institutional owners. These owners often hold board representation and maintain positions over long horizons. Table 3 documents this ownership structure, which is consistent with the broader European pattern of concentrated family control documented by Faccio and Lang (2002). Ma (2024) makes the ownership question concrete by showing that U.S. closed-end funds with higher retail ownership experienced larger discount widening during the COVID-19 shock. However, even his low-retail subsample remains retail-concentrated relative to Swedish closed-end funds. Sweden therefore provides a conservative setting in which to test the sentiment-discount relation. Specifically, we ask whether sentiment-related discount dynamics require retail concentration, or whether they can also arise when the register is held primarily by non-retail owners. Such evidence would extend the classic closed-end fund literature beyond the retail-heavy setting in which it has primarily been studied.

Research question. Does investor sentiment predict NAV discount fluctuations in Swedish closed-end funds, even when the shareholder register is held primarily by non-retail investors?

The ideal test would compare retail-concentrated and non-retail-concentrated closed-end funds within the same market, while also observing who trades at the margin. Our setting is informative, but it does not provide that full comparison. The sample covers Swedish closed-end funds with sufficient NAV history, but the cross-section is small, limiting the power of ownership-based heterogeneity tests. The ownership data do provide meaningful

variation in institutional and family/sphere ownership, from 32% to 88%, but the sample does not include a clean retail-concentrated benchmark comparable to the U.S. closed-end fund setting. Finally, the ownership data describe the shareholder register, not the marginal trader in the float or the household exposure that may sit behind institutional wrappers. We therefore interpret the evidence as showing whether sentiment-related discount dynamics appear in a concentrated-register setting, not as a definitive test of whether retail investors matter at the margin.

We test three related predictions. **H1:** investor sentiment predicts NAV discount changes. **H2:** the sentiment-discount association is larger for funds with greater unlisted exposure. **H3:** funds with higher non-retail ownership concentration are less sensitive to sentiment.

We construct the Composite Sentiment Index (CSI) from five market-based proxies. We test it on a monthly panel of 18 Swedish closed-end funds from January 2017 to December 2025, using first-difference regressions with firm fixed effects. Lagged *CSI* is negatively associated with subsequent monthly discount changes at the two- to three-month horizon. The association is stable in sign but sensitive to stronger common-shock and persistence checks. We therefore interpret H1 as directional evidence consistent with sentiment-related discount dynamics, rather than as a decisive predictive result.

The cross-sectional evidence is strongest for H2. Funds with larger unlisted portfolio shares exhibit larger lagged-CSI slopes in the direction of discount widening, consistent with the Pontiff-style prediction that arbitrage difficulty amplifies sentiment-related mispricing. The absence of a lag-0 interaction reduces concern about the simplest NAV-staleness explanation, although delayed private-valuation adjustments cannot be fully ruled out. H3 is not supported. At lag 2 the ownership interaction is small and insignificant under firm fixed effects. The pooled OLS estimate is significant under firm-clustered inference but does not survive the wild cluster bootstrap. The sign reverses at lag 3. We read the absence of moderation as evidence that the sentiment-discount association in our sample is not driven by residual variation in non-retail ownership concentration. This is not a rejection of the Lee, Shleifer, and Thaler retail-investor mechanism.

The contribution is therefore deliberately bounded. We do not claim that the closed-end structure alone is sufficient for sentiment-related discounts. We do not show that retail investors are irrelevant at the margin. We do not identify a causal sentiment shock. The qualified evidence is that sentiment-related discount dynamics appear in Swedish closed-end funds whose shareholder registers are held primarily by non-retail owners. Sensitivity is stronger among funds with greater unlisted exposure. The remainder of the paper is organized as follows. Section 2 reviews related literature. Section 3 describes the data. Section 4 presents the methodology. Section 5 reports the results. Section 6 discusses interpretation and limitations. Section 7 concludes.

2 Related Literature

The closest papers to ours are Lee, Shleifer, and Thaler (1991), Baker and Wurgler (2006), Pontiff (1996), Gemmill and Thomas (2002), Ma (2024), and Hjelström (2007). Lee, Shleifer, and Thaler establish the sentiment-discount link in a U.S. setting with 6.6% institutional ownership. We ask whether the same relation appears in a contrasting ownership setting. The comparison is not one-for-one. They report institutional ownership, while our Swedish measure combines institutional and family/sphere ownership in the register. Baker and Wurgler develop the composite-sentiment-index methodology for U.S. equity markets. We adapt it to Swedish data and apply it to discount changes rather than to the cross-section of returns. Pontiff and Gemmill and Thomas show that closed-end fund discounts depend on arbitrage costs. We exploit cross-sectional variation in unlisted portfolio share as a Swedish-market proxy for arbitrage difficulty. Ma (2024) finds that U.S. closed-end funds with higher retail ownership experienced larger discount increases during the COVID-19 shock, consistent with a retail-sentiment interpretation. Even his low-retail subsample averages approximately 78% retail, so our Swedish sample falls entirely outside that range. Hjelström (2007) provides a major prior valuation study of Swedish and British closed-end investment companies. We shift the focus from cross-sectional valuation determinants to time-varying aggregate sentiment in a panel framework.

We contribute three things to the literature. First, we construct a Swedish composite sentiment index following Baker and Wurgler (2006), excluding the closed-end fund discount to avoid circularity. Second, we provide a Swedish application of the Lee, Shleifer, and Thaler (1991) sentiment–discount mechanism in a contrasting ownership setting where institutional and family/sphere holders, rather than dispersed retail investors, hold most of the shareholder register. Third, we provide a cross-sectional test of the Pontiff (1996) arbitrage-cost prediction for Swedish closed-end funds, using unlisted portfolio share as a proxy for arbitrage difficulty.

3 Data

3.1 Sample

Our sample consists of ibindex-tracked Swedish closed-end funds with a daily calculated NAV series and at least 24 months of observations. The sample is therefore the ibindex coverage universe with sufficient NAV history, rather than a selection drawn from a broader population. We exclude funds with fewer than 24 months of observations to avoid short panels and ensure that each fund contributes meaningful within-firm variation. This leaves 18 funds over the period January 2017 to December 2025. Of these, 15 are observed throughout the full period, while the remaining funds enter later.

The ibindex calculated NAV (beräknat substansvärde) updates reported NAVs by marking disclosed listed holdings to current market prices, while unlisted holdings are carried forward from the latest available company disclosures. For firms with predominantly listed portfolios, such as Industrivärden, the measured NAV therefore moves closely with the underlying market between reporting dates. Discounts for these firms are less likely

to be driven by stale NAV valuations, and more likely to reflect how the market prices the investment-company wrapper itself.

Figure A.1 in the Appendix plots the average NAV discount over time and illustrates substantial time-series variation in discounts over the sample period.

Mean price-to-NAV deviations vary substantially across firms: First Ventures traded at an average discount of 51.6%, whereas Flat Capital traded at an average premium of 58.0%. Most firms, however, traded below NAV on average.

Table 1: Sample of Swedish closed-end funds. The sample is the population of Swedish closed-end funds with daily NAV data available from ibindex.se for at least 24 months during January 2017–December 2025. First month and Last month indicate the start and end of each firm’s presence in the panel. Mean disc. (%) is the time-series average NAV discount, defined as $(\text{Price} - \text{NAV}) / \text{NAV} \times 100$, where a negative value indicates the share price is below NAV. SD (%) is the standard deviation of monthly discount levels. Obs is the number of firm-month observations. Flat Capital and First Ventures entered the sample in late 2021 with short histories and atypical valuations (a persistent premium and deep discount, respectively); the balanced-panel robustness check in Section 5 excludes these firms.

| Company | First month | Last month | Obs | Mean disc. (%) | SD (%) |
|-------------------|-------------|------------|-----|----------------|--------|
| AJA | 2017-01 | 2025-12 | 108 | -9 | 15.6 |
| Bure | 2017-01 | 2025-12 | 108 | 6.6 | 15.2 |
| Creades | 2017-01 | 2025-12 | 108 | 15 | 21.7 |
| First Ventures | 2021-07 | 2025-12 | 54 | -51.6 | 15.9 |
| Flat Capital | 2021-10 | 2025-12 | 51 | 58 | 58.2 |
| Industrivården | 2017-01 | 2025-12 | 108 | -8.9 | 4.9 |
| Investor | 2017-01 | 2025-12 | 108 | -15.1 | 6.1 |
| Kinnevik | 2017-01 | 2025-12 | 108 | -18.1 | 18.4 |
| Latour | 2017-01 | 2025-12 | 108 | 24.7 | 19.5 |
| Linc | 2021-05 | 2025-12 | 56 | 17.6 | 21.2 |
| Lundbergföretagen | 2017-01 | 2025-12 | 108 | -1.7 | 8 |
| NAXS | 2017-01 | 2025-12 | 108 | -19.4 | 6 |
| Spiltan | 2017-01 | 2025-12 | 108 | -20.8 | 11.3 |
| Svolder | 2017-01 | 2025-12 | 108 | 1.6 | 9.6 |
| Traction | 2017-01 | 2025-12 | 108 | -5.1 | 6 |
| VEF | 2017-01 | 2025-12 | 108 | -27.8 | 21.9 |
| VNV Global | 2017-01 | 2025-12 | 108 | -33.3 | 24.8 |
| Öresund | 2017-01 | 2025-12 | 108 | 5.3 | 6.4 |

3.2 NAV discount construction

We define the NAV discount for firm i in month t as:

$$\text{Discount}_{i,t} = (P_{i,t} - \text{NAV}_{i,t}) / \text{NAV}_{i,t} \times 100.$$

where $P(i,t)$ is the closing share price and $\text{NAV}(i,t)$ is the calculated net asset value, both measured on the last trading day of the month. Negative values indicate discounts, while positive values indicate premia. The dependent variable in all regressions is the within-firm first difference, $\Delta \text{Discount}(i,t) = \text{Discount}(i,t) - \text{Discount}(i,t-1)$. This transformation is motivated by the high persistence of discount levels: an Augmented Dickey-Fuller test on the equally weighted average monthly discount fails to reject the unit-root null ($p = 0.744$; Section 4.2). Before estimation, $\Delta \text{Discount}$ is winsorized at the 1st and 99th percentiles of the pooled firm-month distribution to limit outlier influence.

The first difference is computed within firm. December 2016 is retained only as a lead-in month, allowing January 2017 first differences to be computed for the 15 funds with pre-2017 history. These firms therefore lose no observations to first-differencing. Three funds enter the panel later: Linc in May 2021, First Ventures in July 2021, and Flat Capital in October 2021. Each loses its first firm-specific observation. The result is 1,778 non-missing firm-month observations from 1,781 panel rows.

3.3 Sentiment proxies

We construct the *CSI* from five market-based proxies intended to summarize complementary signals associated with investor sentiment: IPO activity, first-day IPO returns, turnover, equity issuance, and the dividend premium.

***NIPO* (number of IPOs) and *RIPO* (average first-day IPO return).** The two proxies reflect the “windows of opportunity” logic of Baker and Wurgler (2006). Firms are more likely to issue equity when investor demand is strong, and first-day returns capture the intensity of that demand. We source both variables from Capital IQ, filtering for primary-exchange IPOs in Sweden. *RIPO* is computed for months with at least one IPO and an available first-day return. Of 108 months, 32 have missing monthly *RIPO* values before index construction. Of these, 26 are true zero-IPO months and 6 are IPO months with unavailable first-day-return data. In the baseline *CSI*, these missing values are coded as zero, reflecting the absence of an observed first-day-return signal rather than an observed break-even IPO month. Section 5.7 reports robustness checks using alternative codings.

***TURN* (detrended log turnover).** The proxy is intended to capture trading intensity and speculative market activity. Following Baker and Wurgler (2006), we compute $\ln(VA/MV)$ and detrend it using a 60-month trailing moving average. *TURN* is constructed from LSEG Datastream’s TOTMKSD Sweden country-index aggregate using turnover by value (*VA*) and market value (*MV*). TOTMKSD aggregates 70 Swedish-listed equities, capturing the large-cap, actively-traded segment of the market rather than the full universe of listed firms. We interpret *TURN* as a broad-market liquidity proxy.

***S* (equity share of new issues).** The proxy is the ratio of equity issuance to total equity and debt issuance. The proxy follows Baker and Wurgler’s interpretation that equity issuance becomes more attractive when investor demand for equity is high.

***P(D-ND)* (dividend premium).** The proxy is the log difference in value-weighted average market-to-book ratios of dividend-paying versus non-dividend-paying firms, computed from Capital IQ data on Swedish listed companies. Following Baker and Wurgler (2006), it captures time variation in investor demand for dividend-paying stocks relative to non-dividend-paying stocks. Section 4.1 reports how this proxy loads in the Swedish *CSI*.

A Swedish survey-based measure, the Economic Tendency Indicator (*ETI*), produced by Konjunkturinstitutet, is used as an external comparator rather than a *CSI* input. Following Baker and Wurgler (2006), the *CSI* itself is constructed from market-based proxies that capture sentiment through revealed behavior rather than stated confidence.

Table 2: Summary statistics for sentiment proxies, monthly 2017–2025. *NIPO* is the monthly number of Swedish IPOs (Capital IQ). *RIPO* is the average first-day return of IPOs; 32 months have missing monthly *RIPO* before index construction (26 zero-IPO months and 6 IPO months with unavailable first-day-return data) and are coded as zero in the baseline *CSI* (Capital IQ). *TURN* is detrended log market turnover, computed as $\ln(\text{VA}/\text{MV})$ for LSEG Datastream's TOTMKSD Sweden country-index aggregate minus its 60-month trailing moving average. *S* is the equity share of new issues, defined as monthly equity issuance proceeds divided by total equity and debt issuance proceeds (Capital IQ). *P(D-ND)* is the dividend premium, the log difference in value-weighted average market-to-book ratios between dividend payers and nonpayers among Swedish listed firms (Capital IQ). *ETI* is the Economic Tendency Indicator (Konjunkturinstitutet) and is reported as an external survey-based comparator; it is not included in the *CSI*.

| Proxy | Mean | SD | Min | Max | N |
|------------------|---------|--------|---------|--------|-----|
| NIPO | 3.769 | 5.144 | 0 | 27 | 108 |
| RIPO | 5.627 | 18.968 | -45.455 | 77.143 | 108 |
| TURN | -0.119 | 0.179 | -0.49 | 0.744 | 108 |
| S (equity share) | 0.187 | 0.178 | 0.014 | 0.909 | 108 |
| P(D-ND) | -0.233 | 0.402 | -1.023 | 0.575 | 108 |
| ETI | 100.497 | 11.478 | 60.9 | 122.7 | 108 |

3.4 Macro controls

The orthogonalization procedure regresses each sentiment proxy on four macroeconomic variables before extracting the principal component. Adapting Baker and Wurgler (2006) to the available Swedish data, we use four controls. These are growth in industrial production (SCB), growth in total household consumption (SCB consumption indicator), and growth in employment (SCB AKU labor force survey). The fourth is a recession indicator constructed from negative GDP-growth quarters. Section 4.1 details the construction of the orthogonalized index.

3.5 Ownership structure

We document ownership structure using annual shareholder-register data from Modular Finance Holdings, covering 2016-2025. Each shareholder is classified into institutional, family/sphere, or residual other categories using keyword matching on shareholder names, supplemented by manual verification of the largest holders. The institutional category includes pension funds, mutual funds, insurance companies, asset managers, banks, public-sector entities, and matched foreign institutional investors. The family/sphere category includes controlling families, founder- or sphere-controlled blockholders, foundations matched through family/sphere keywords, and named long-horizon blockholders. We define *NonRetailOwnPct* as the sum of the institutional and family/sphere ownership shares. The variable is measured annually, lagged one year in the regressions, and merged into the monthly panel; in descriptive tables we display it in percent, while in regressions it enters as a 0-1 fraction. *NonRetailOwnPct* is a register-level ownership measure: it describes the shareholder base, not the identity of the marginal trader in the float. Robustness to alternative classifier specifications is reported in Appendix Table A.5.

Table 3: Ownership structure of sample closed-end funds, December 2025. Values are computed under the expanded ownership classifier.

| Company | Largest owner | Type | Inst. (%) | Family (%) | NonRetailOwnPct (%) | Within-firm SD (pp) |
|-------------------|-------------------------------------|---------------|-----------|------------|---------------------|---------------------|
| AJA | Mikael Ahlström | Family/Sphere | 10.8 | 39 | 49.8 | 2.42 |
| Bure | Tigerschiöld-familjen | Family/Sphere | 44.3 | 30.1 | 74.3 | 4.49 |
| Creades | Hagströmer-intressen | Family/Sphere | 12.1 | 61.7 | 73.8 | 0.69 |
| First Ventures | Rutger Arnhult | Family/Sphere | 18 | 27 | 45 | 3.34 |
| Flat Capital | Sebastian Siemiatkowski | Family/Sphere | 2.8 | 83.1 | 85.9 | 19.46 |
| Industrivärden | Lundbergföretagen AB | Family/Sphere | 30.9 | 45 | 76 | 1.34 |
| Investor | Wallenberg-sfären | Family/Sphere | 40.6 | 27.9 | 68.4 | 0.58 |
| Kinnevik | Baillie Gifford & Co | Institutional | 54.5 | 13.8 | 68.3 | 3.33 |
| Latour | Investment AB Latour & Fam. Douglas | Family/Sphere | 9.8 | 79 | 88.7 | 0.43 |
| Linc | Bengt Julander | Family/Sphere | 23.7 | 56.4 | 80.1 | 5.35 |
| Lundbergföretagen | Lundberggruppen | Family/Sphere | 15.4 | 72.4 | 87.7 | 0.96 |
| NAXS | Familjen Molse | Family/Sphere | 8.9 | 50.5 | 59.4 | 15.68 |
| Spiltan | Per H Börjesson and family | Family/Sphere | 6.1 | 27.4 | 33.5 | 0.72 |
| Svolder | Lundström-familjen | Family/Sphere | 22.5 | 10.2 | 32.7 | 5.24 |
| Traction | Stillström-sfären | Family/Sphere | 2.1 | 83.3 | 85.3 | 0.31 |
| VEF | Ruane, Cunniff & Goldfarb | Institutional | 51.5 | 0.1 | 51.7 | 10.72 |
| VNV Global | Conifer Management LLC | Institutional | 39.8 | 20.9 | 60.7 | 4.28 |
| Öresund | Qviberg-intressen | Family/Sphere | 10 | 45.1 | 55.1 | 1.69 |

Notes: Inst. and Family/Sphere report the institutional and family/sphere subcomponents of non-retail ownership. *NonRetailOwnPct* is the combined capital share classified as non-retail under the expanded classifier; in regressions, it is scaled to a 0–1 fraction. Within-firm SD is the time-series standard deviation of *NonRetailOwnPct*, in percentage points, across annual Modular Finance Holdings snapshots from 2016 to 2025.

Table 3 documents the central ownership feature of the sample: as of December 2025, institutional and family/sphere holders included in *NonRetailOwnPct* hold a median of 68.4% of equity across the 18 sample funds, with a mean of 65.4% and a range from 32.7% to 88.7%. These primary shareholders often hold board seats in portfolio companies and maintain positions measured in decades, not quarters. This ownership structure differs sharply from the dispersed retail-investor setting emphasized by Lee, Shleifer, and Thaler (1991), who report only 6.6% institutional ownership for U.S. closed-end funds and characterize the remaining investor base as predominantly individual. The comparison is not one-for-one, since our Swedish measure combines institutional and family/sphere blockholders while they report institutional ownership. The institutional contrast is nevertheless substantial: the U.S. closed-end funds they study were characterized by dispersed individual ownership, whereas Swedish closed-end funds are held primarily by identifiable institutional and family/sphere holders. It also falls entirely outside the range of variation in Ma (2024), whose “low retail” subsample still averaged approximately 78% retail ownership. The cross-sectional heterogeneity tests in Section 5 use lagged annual ownership values rather than this single snapshot, but within-firm variation in the combined institutional and family share is limited: the mean within-firm standard deviation across sample funds is 4.5 percentage points (maximum 19.5), which bounds the power of the ownership interaction test.

Ownership concentration describes the shareholder register, not the marginal price-setter. Equity prices are set by trading in the float. The float may include retail investors, mutual funds, and short-horizon institutions even when long-horizon blockholders hold most of the register. Our test therefore speaks to the setting’s ownership structure: concentrated ownership and limited float. It does not speak to who trades at the margin. Evidence consistent with sentiment-discount dynamics in this setting would suggest that retail concentration in the shareholder register is not necessary for such dynamics to appear. It would leave open whether retail investors are important marginal traders.

3.6 Portfolio composition

Portfolio composition is the second main source of cross-sectional variation in the paper. To quantify the listed-unlisted distinction described above, we hand-collect quarterly data from interim and annual reports and define *UnlistedShare* as reported unlisted holdings divided by total NAV. The measure captures unlisted exposure relative to shareholder value and can exceed one when net debt or other net liabilities make unlisted holdings larger than NAV.

Pontiff (1996) links closed-end fund discounts to arbitrage costs, especially the difficulty of replicating the underlying portfolio. We adapt this logic to Swedish closed-end funds by using *UnlistedShare* as a proxy for arbitrage difficulty: portfolios with larger unlisted exposures are harder for outside investors to replicate. To avoid look-ahead bias, we lag *UnlistedShare* by one quarter before merging it into the monthly regression panel.

3.7 Panel summary statistics

Table 4: Firm-month panel summary statistics. The table reports summary statistics for the main variables used in the firm-month regressions.

| Variable | Mean | Median | SD | Min | Max | N |
|-----------------------|--------|--------|--------|---------|--------|------|
| Δ Discount (%) | -0.287 | -0.164 | 7.138 | -25.186 | 22.275 | 1778 |
| CSI | 0.051 | 0 | 1.356 | -4.022 | 3.657 | 1766 |
| CSI-orth | 0.069 | 0.053 | 1.386 | -3.136 | 4.164 | 1766 |
| UnlistedShare (%) | 34.989 | 26.01 | 33.648 | 0 | 170.67 | 1689 |
| NonRetailOwnPct (%) | 62.43 | 67.649 | 21.021 | 7.929 | 90.354 | 1704 |

Notes: Δ Discount is the winsorized within-firm first difference of the NAV discount. *UnlistedShare* and *NonRetailOwnPct* are expressed in percent; *CSI* and *CSI-orth* are standardized index values. *UnlistedShare* is lagged one quarter, and *NonRetailOwnPct* is lagged one year. Values of *UnlistedShare* above 100% are retained because unlisted holdings can exceed NAV when net liabilities are large.

4 Methodology

4.1 Composite Sentiment Index construction

We construct the Composite Sentiment Index (CSI) using principal component analysis (PCA) in three stages, following Baker and Wurgler (2006).

Lag/lead alignment. Following Baker and Wurgler (2006), we consider both the contemporaneous value and the one-month lag of each sentiment proxy, yielding ten candidate variables. We estimate a first-stage PCA on these candidates and, for each proxy, retain the timing with the highest absolute correlation with the first principal component. Appendix Table A.1 reports the selected timings. Four of the five Swedish proxies enter with a one-month lag: *NIPO*, *RIPO*, *S*, and *P(D-ND)*. *TURN* enters contemporaneously. These selections differ from their U.S. annual index, where *TURN*, *RIPO*, and *P(D-ND)* enter lagged, while *NIPO* and *S* enter contemporaneously. Because we apply the same timing-selection procedure, we treat the different Swedish timing pattern as a feature of the sample rather than as a methodological deviation.

Principal component analysis. We extract the first principal component from the five sentiment proxies selected in the lag/lead alignment stage. Because the sign of a principal component is arbitrary, we orient *PC1* so that *NIPO*, *RIPO*, and *S* load positively, following the sentiment-direction convention in Baker and Wurgler (2006). The first component explains 38.8% of the variance in the five standardized proxies, consistent with these market-based measures capturing related but noisy and partly distinct dimensions of sentiment.

Table 5: *PC1* loadings of raw and orthogonalized sentiment proxies. The first column reports loadings from PCA applied to the five raw proxies selected in the lag/lead alignment stage. The second column reports loadings from PCA applied to the orthogonalized proxies, defined as residuals from regressions of each raw proxy on four macroeconomic controls: IP growth, household consumption growth, employment growth, and a recession indicator. All proxies are standardized to zero mean and unit variance before PCA, and both components are oriented so that *NIPO*, *RIPO*, and *S* load positively.

| Proxy | Raw CSI (PC1) | CSI-orth (PC1) |
|---------|---------------|----------------|
| NIPO | 0.195 | 0.179 |
| RIPO | 0.329 | 0.296 |
| TURN | -0.515 | -0.495 |
| S | 0.472 | 0.539 |
| P(D-ND) | 0.604 | 0.587 |

NIPO, *RIPO*, and the equity share *S* anchor the orientation of *PC1*. All three load positively, consistent with the Baker and Wurgler (2006) interpretation that more IPOs, higher first-day IPO returns, and a higher equity share of new issues signal elevated sentiment. *NIPO* has the weakest loading, at 0.195, likely reflecting the lumpiness of Swedish IPO activity; with a mean of 3.8 IPOs per month and frequent zero-IPO months, the series contains substantial idiosyncratic noise relative to the smoother proxies. Two proxies load with signs opposite to their U.S. pattern. *TURN* loads negatively (-0.515 raw, -0.495 orth), whereas they report a positive loading on turnover. *P(D-ND)* loads positively (0.604 raw, 0.587 orth), whereas they report a negative loading on the dividend premium. Because PCA loadings are determined by the covariance structure of the sample rather than imposed by theory, these signs reflect how *TURN* and *P(D-ND)* co-move with the dominant common factor in the Swedish 2017–2025 panel. A full explanation is outside the scope of this paper, but possible factors include differences in U.S. and Swedish market structure, the smaller and shorter Swedish proxy panel, and cross-market differences in how individual proxies relate to latent sentiment. We retain the empirical signs in the baseline rather than imposing the U.S. signs, and Section 5.7 tests the sensitivity of the regression results to the two sign-divergent proxies, reporting specifications excluding *TURN* (R1), excluding *P(D-ND)* (R2), and excluding both (R3); full coefficients are in Appendix Table A.3.

The Swedish *CSI* should therefore be interpreted as a market-based Swedish sentiment factor constructed in the Baker and Wurgler (2006) tradition, not as a literal replication of their U.S. index. Two construction choices are especially important. First, the closed-end fund discount (CEFD) is excluded because the NAV discount is the dependent variable in our regressions; including CEFD would induce a mechanical correlation. Second, *TURN* and *P(D-ND)* have sample-specific signs, so *CSI-orth* captures the dominant common

factor in the Swedish proxy panel but is not directly comparable to the U.S. index proxy by proxy. Finally, *CSI-orth* does not closely track the Swedish survey-based mood measure: its correlation with *ETI* is weak ($r = 0.20$) in our sample (Section 5.1; Appendix Table A.6). We therefore interpret *CSI-orth* as a Swedish market-based proxy for issuance, valuation, and liquidity conditions, rather than as a survey-like consumer confidence index.

Orthogonalization. Following Baker and Wurgler (2006), we remove macroeconomic variation from the sentiment proxies before constructing the orthogonalized sentiment index. Their orthogonalization uses industrial production growth, growth in durable, nondurable, and services consumption, employment growth, and an NBER recession indicator. We adapt this logic to Swedish monthly data using industrial production growth, total household consumption growth, employment growth, and a recession dummy. Since Sweden has no equivalent of the NBER Business Cycle Dating Committee, we adopt the simple sequence rule discussed by Harding and Pagan (2002): the dummy equals one for months in any quarter where both the current and the prior quarter had negative q/q GDP volume growth, and zero otherwise. Total household consumption serves as the Swedish aggregate counterpart to the U.S. consumption-component controls, employment growth captures labor-market fluctuations, and the recession dummy captures broad downturn periods. We then apply PCA to the residualized proxy series, using the same NIPO/RIPO/S-positive orientation convention as for the raw index. *CSI-orth* explains 39.4% of the variation in the orthogonalized proxies and is almost identical to the raw index in this sample ($r = 0.98$). We use *CSI-orth* as the baseline throughout Section 5; raw-*CSI* estimates are reported in Appendix Table A.2 as a comparison.

4.2 Stationarity and the choice of dependent variable

NAV discount levels are persistent. As a specification diagnostic, an Augmented Dickey-Fuller test applied to the equally weighted average monthly discount fails to reject the unit-root null (ADF $p > 0.10$). ADF and KPSS provide complementary evidence because ADF tests the unit-root null, while KPSS tests the stationarity null. First differences of the discount reject the unit-root null at the 1% level under ADF, and KPSS fails to reject stationarity (statistic = 0.104, $p > 0.10$ under both the level- and trend-stationarity nulls). These diagnostics support using $\Delta Discount(i,t)$, rather than discount levels, as the dependent variable in the regressions.

The sentiment index is used in lagged levels in the baseline specification to preserve the predictive structure of Baker and Wurgler (2006). However, *CSI-orth* is persistent in this sample, so the baseline pairs stationary discount changes with a persistent sentiment level. We therefore treat persistence as an inference concern and address it directly in the empirical analysis: Section 5.4 adds contemporaneous $\Delta CSI\text{-orth}(t)$ to the baseline, and Section 5.7 replaces lagged *CSI-orth* levels with lagged $\Delta CSI\text{-orth}$.

4.3 Panel Regression Specification

The baseline specification is:

$$\Delta\text{Discount}_{i,t} = \alpha_i + \beta_k \text{CSI-orth}_{t-k} + \varepsilon_{i,t}. \quad (1)$$

where α_i are firm fixed effects, $\text{CSI-orth}(t-k)$ is the orthogonalized Composite Sentiment Index lagged k months, and $\varepsilon_{i,t}$ is the error term. We estimate equation (1) separately for $k = 1, 2, 3,$ and $6,$ and report all four regressions in a single table. Standard errors are clustered at the firm level to account for serial correlation within firms.

With only 18 firms, however, cluster-robust standard errors may be downward-biased, causing t-statistics to be overstated. Cameron and Miller (2015) report a consensus that roughly 50 clusters are needed for reliable inference in balanced state-year panels, and MacKinnon and Webb (2017) show that the benchmark can fail with even more clusters when cluster sizes are unbalanced. Our setting has both a small cluster count and unbalanced clusters (51 to 108 firm-months), so we report two alternative inference procedures alongside the baseline in Section 5.3: the wild cluster bootstrap of Cameron, Gelbach, and Miller (2008) with Webb weights and 999 replications, and two-way clustered standard errors at the firm and month level. Two-way clustering does not add time fixed effects, which would absorb the *CSI*; it only adjusts standard errors for the cross-sectional dependence induced by the common monthly *CSI* realization.

Firm fixed effects absorb time-invariant firm characteristics, including average discount levels, ownership structure, and portfolio composition. The coefficient β_k is therefore identified from time-series variation in lagged sentiment associated with within-firm changes in discounts. We do not include time fixed effects because *CSI-orth* is identical for all firms within a given month, so time fixed effects would absorb the sentiment variable entirely.

The baseline coefficient is interpreted as a predictive association rather than a causal estimate. For this association to reflect sentiment, *CSI-orth* must not primarily proxy for omitted financial-market conditions that also move discount changes. The macroeconomic orthogonalization removes variation related to industrial production, household consumption, employment, and recession periods, but it does not directly absorb financial channels such as credit conditions, equity risk premia, or market liquidity. If these omitted channels affect both the *CSI* proxies and closed-end fund discounts, the estimated coefficient may combine sentiment effects with residual financial-market effects. We therefore interpret the results as evidence of sentiment-related discount dynamics, not as a causal estimate of a pure sentiment shock.

4.4 Cross-sectional heterogeneity

H2 and H3 are moderation tests implemented in a panel fixed-effects setting. In both tests, lagged *CSI-orth* is interacted with a fund characteristic: *UnlistedShare* for H2 and *NonRetailOwnPct* for H3. The interaction coefficient measures whether the association between sentiment and discount changes differs with that characteristic. Firm fixed effects absorb time-invariant differences across funds, so these specifications should not be read as simple cross-sectional moderation regressions. The moderator main effects

are identified only from within-fund changes over time, while the interactions capture differences in sentiment sensitivity associated with the moderator. Because both moderators are slow-moving, especially ownership concentration, the interaction estimates are descriptive rather than causal and have limited power. We therefore treat the firm fixed-effects specification as the conservative baseline. For H3, we also report pooled OLS with firm-level controls as a descriptive complement that uses more between-fund variation but is more vulnerable to omitted fund-level heterogeneity.

We test two sources of cross-sectional variation in sentiment sensitivity, each motivated by a specific prediction from the literature.

Unlisted holdings and arbitrage difficulty. Following Pontiff (1996) and Gemmill and Thomas (2002), funds that are harder to arbitrage should be more sensitive to sentiment. When portfolios are difficult to replicate or hedge, costly arbitrage limits the correction of sentiment-driven mispricing. We test this prediction using the share of the portfolio held in unlisted assets:

$$\begin{aligned} \Delta \text{Discount}_{i,t} = & \alpha_i + \beta_1 \text{CSI-orth}_{t-k} + \beta_2 \text{UnlistedShare}_{i,t-1Q} \\ & + \beta_3 \text{CSI-orth}_{t-k} \times \text{UnlistedShare}_{i,t-1Q} + \varepsilon_{i,t}. \end{aligned} \quad (2)$$

where $\text{UnlistedShare}(i,t-1Q)$ is the fraction of firm i 's portfolio held in unlisted assets, lagged one quarter to avoid look-ahead bias. Holding UnlistedShare fixed at M , the implied sentiment slope is $\beta_1 + \beta_3 \cdot M$. Firm fixed effects absorb time-invariant differences in portfolio composition, while the lagged UnlistedShare main effect β_2 controls for quarterly within-fund changes in unlisted exposure that fixed effects do not absorb. Because UnlistedShare moves slowly, the interaction is interpreted as descriptive heterogeneity rather than a causal moderator effect. The coefficient of interest is β_3 : a negative β_3 means that the slope on lagged sentiment becomes more negative as UnlistedShare increases, consistent with larger discount widening for harder-to-arbitrage funds.

A potential confound arises because ibindex values unlisted holdings at the last quarterly valuation. Market prices update continuously, while unlisted NAV components remain fixed until the next quarterly report. A market-wide price movement could therefore mechanically widen the measured discount for firms with high unlisted exposure, even without genuine mispricing. To distinguish this mechanical timing effect from a sentiment-related arbitrage-cost effect, we estimate equation (2) separately at lags $k = 0, 1, 2, 3,$ and 6 . Under the staleness interpretation, β_3 should be most visible at short lags, especially lag 0, when market prices have moved but unlisted NAV components have not yet adjusted. The effect should then weaken, and potentially reverse, once quarterly NAV updates incorporate the new valuations. Under an arbitrage-cost interpretation, by contrast, the interaction need not appear immediately; the prediction is that harder-to-arbitrage funds exhibit greater sensitivity to lagged sentiment over subsequent months. We therefore interpret the lag profile as a diagnostic for the staleness story rather than as a positive test of the arbitrage-cost mechanism.

Non-retail ownership and the retail-investor channel. Lee, Shleifer, and Thaler (1991) propose that retail investors drive the sentiment-discount link. If this mechanism operates in our setting, firms with higher non-retail ownership concentration should be less sensitive to sentiment. We test this prediction with:

$$\begin{aligned} \Delta \text{Discount}_{i,t} = & \alpha_i + \beta_1 \text{CSI-orth}_{t-k} + \beta_2 \text{NonRetailOwnPct}_{i,t-1Y} \\ & + \beta_3 \text{CSI-orth}_{t-k} \times \text{NonRetailOwnPct}_{i,t-1Y} + \varepsilon_{i,t}. \end{aligned} \quad (3)$$

where $\text{NonRetailOwnPct}(i,t-1Y)$ is the sum of institutional and family/sphere ownership under the expanded ownership classifier, lagged one year. The moderator enters the regression as a 0-1 fraction; Table 4 reports NonRetailOwnPct in percent for descriptive readability. Holding NonRetailOwnPct fixed at M , the implied sentiment slope is $\beta_1 + \beta_3 \cdot M$. H3 predicts $\beta_3 > 0$. Because the baseline CSI-orth coefficient is negative, a positive β_3 means that higher non-retail ownership makes the sentiment slope less negative, consistent with a dampened discount-widening response after high sentiment.

Firm fixed effects absorb time-invariant differences in ownership concentration. The lagged NonRetailOwnPct main effect β_2 is therefore identified only from within-fund changes in ownership that fixed effects do not absorb. H3, however, is primarily a between-fund prediction, and within-fund variation in ownership is small, with firm-level standard deviations on the order of 5 percentage points. We therefore also report pooled OLS with fund-level controls: mean UnlistedShare , average log assets, and mean discount. This specification uses more of the cross-sectional variation the hypothesis concerns, but is more vulnerable to omitted fund-level differences.

5 Results

5.1 Sentiment Index properties

Figure 1 plots the CSI alongside the average NAV discount across sample funds.

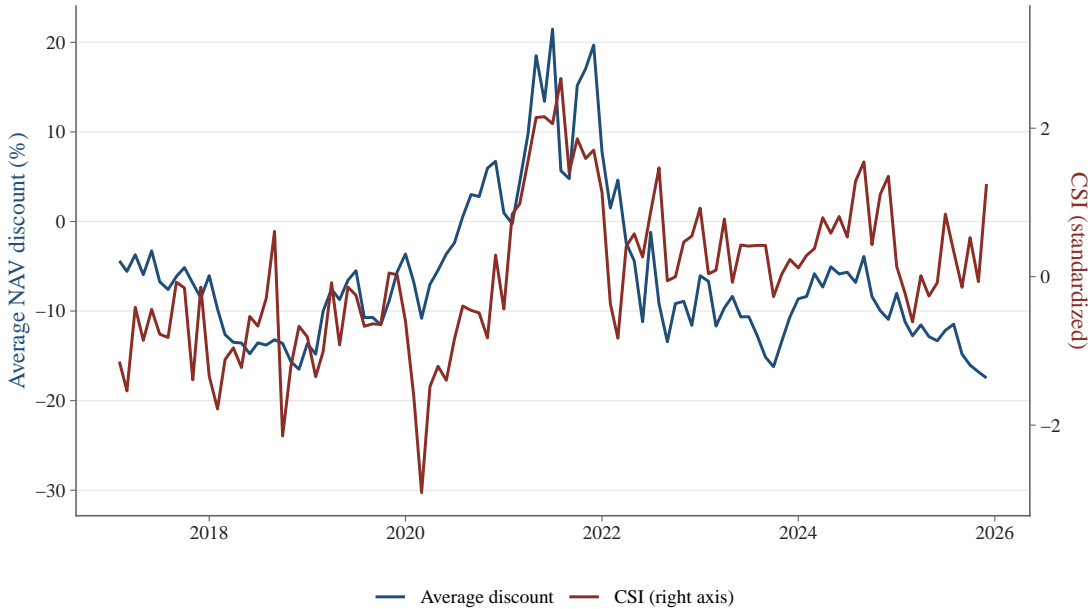


Figure 1. Composite Sentiment Index and average NAV discount, January 2017–December 2025. The CSI is the first principal component of five market-based sentiment proxies. The average NAV discount is the equal-weighted monthly mean across sample firms, defined as $(\text{Price} - \text{NAV}) / \text{NAV} \times 100$. The CSI line is linearly rescaled for display to match the mean and standard deviation of the average NAV discount series; the right axis reports the corresponding standardized CSI units. The rescaling affects only the figure and not the regression analysis.

Because four of the five sentiment proxies enter the index with a one-month lag, the *CSI* primarily reflects lagged information from IPO activity, equity issuance, and the dividend premium. *TURN* is the only proxy that contributes contemporaneous information (Section 4.1; Appendix Table A.1).

Given the sign convention used for *PC1*, *TURN* loads negatively on the *CSI* (-0.515). The index therefore increases when detrended turnover falls. This inverse relationship is discussed in Section 4.1.

The CSI has a modest positive correlation with ETI, at $r = 0.20$ over 107 monthly observations (Appendix Table A.6). The weak correlation suggests that the market-based CSI and the survey-based ETI capture partly distinct dimensions of sentiment and co-move only weakly in this sample.

Several interpretations are possible. First, the two measures reflect different types of information: revealed market behavior on the one hand, and stated expectations about the broader economy on the other. Second, the 2017–2025 period includes macroeconomic shocks, including energy-price movements and interest-rate changes, that may have affected household confidence without producing corresponding movements in IPO activity, equity issuance, or the dividend premium. Third, the relatively short monthly sample limits the precision with which these relationships can be assessed, especially if the underlying series are persistent. We therefore treat the correlation as descriptive rather than diagnostic.

5.2 Discount co-movement

A basic implication of the sentiment hypothesis is that discounts should contain a common component. If aggregate sentiment affects closed-end fund prices, changes in discounts should tend to move together across funds. Figure 2 plots the distribution of the 153 pairwise correlations of monthly discount changes across the sample.

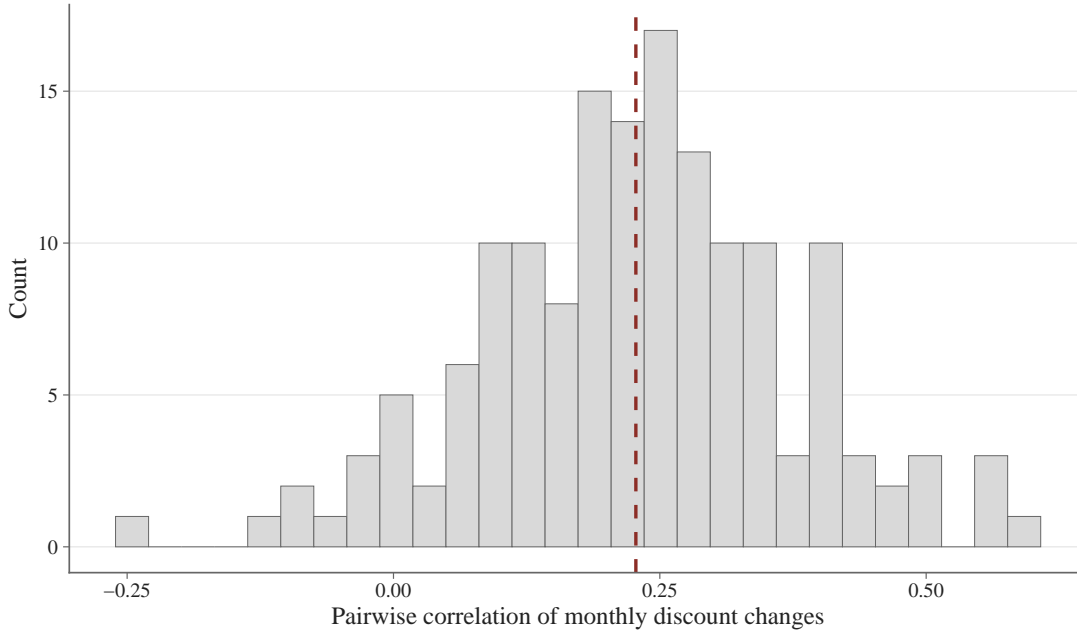


Figure 2. Distribution of pairwise correlations of monthly discount changes across Swedish closed-end funds. The dashed line indicates the sample mean, 0.228. Lee, Shleifer, and Thaler (1991) report a mean pairwise correlation of 0.389 for U.S. closed-end funds.

The mean pairwise correlation is 0.228, and about 93% of pairwise correlations are positive. This pattern indicates that Swedish closed-end fund discount changes contain a common component. The evidence is therefore consistent with the sentiment hypothesis, although it is not unique to that interpretation.

The average co-movement is lower than the 0.389 reported by Lee, Shleifer, and Thaler (1991), suggesting a weaker common component than in their U.S. sample. This difference may reflect variation in sample period, market structure, portfolio composition, ownership structure, or the relative importance of fund-specific shocks. More generally, discount co-movement alone does not identify sentiment as the source of the common factor. It does, however, provide a useful descriptive motivation for the panel regressions below, where the relation between discounts and the constructed sentiment index is examined directly.

5.3 Baseline regressions

The baseline H1 test estimates whether lagged *CSI-orth* levels predict monthly changes in NAV discounts in the firm-month panel. Under the discount convention used throughout the paper, negative values indicate that price is below NAV. A negative coefficient therefore means that higher past sentiment is associated with a subsequent decline in price-to-NAV, or discount widening. This sign convention differs from Lee, Shleifer, and Thaler (1991), who define discounts as positive when price is below NAV.

Table 6, Panel A: Baseline panel regressions of monthly discount changes on lagged *CSI-orth*. The dependent variable is $\Delta Discount$. Each column reports a separate firm-fixed-effects regression with *CSI-orth* lagged k months.

| | Lag 1 | Lag 2 | Lag 3 | Lag 6 |
|-----------------------------|---------------------|----------------------|----------------------|----------------------|
| <i>CSI-orth</i> ($t - k$) | -0.181** (0.066) | -0.369*** (0.098) | -0.255*** (0.088) | -0.335*** (0.104) |
| N | 1,748 | 1,733 | 1,718 | 1,673 |
| Within-R ² | 0.0012 | 0.0049 | 0.0024 | 0.0042 |

Notes: Standard errors in parentheses are clustered by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Within-R² is the within-firm R².

The coefficients on *CSI-orth* are negative at all four lags and statistically significant under firm-clustered standard errors. Given the convention introduced above, the estimates point in the direction predicted by Lee, Shleifer, and Thaler (1991): elevated sentiment is followed by widening NAV discounts.

The largest estimate occurs at lag 2. A one-standard-deviation increase in *CSI-orth* is associated with a two-month-ahead discount change of about -0.51 percentage points, or roughly 7% of the standard deviation of $\Delta Discount$. The lag-1 estimate is about half as large, while the lag-3 and lag-6 estimates remain negative but smaller than the lag-2 coefficient.

We give primary economic interpretation to lags 2 and 3. Lag 1 may be too short for delayed correction to materialize, whereas lag 6 is less clean because *CSI-orth* is persistent in this sample (Section 4.2), and the lag-6 coefficient does not survive the ΔCSI respecification in Section 5.7. The intermediate horizons therefore provide the clearest baseline reading of delayed discount widening after elevated sentiment.

Two qualifications are important. First, Panel B reports alternative inference procedures for the common-month dependence issue. Second, explanatory power is low: within-R² ranges from 0.001 to 0.005 across lags. This is expected in a first-difference panel with substantial firm-specific monthly noise, and motivates interpreting the economic magnitude relative to the standard deviation of $\Delta Discount$.

Table 6, Panel B: Alternative inference for the baseline regression. The table reports alternative p-values for the lagged *CSI-orth* coefficient from the Table 6, Panel A specification.

| | Lag 1 | Lag 2 | Lag 3 | Lag 6 |
|------------------|--------|--------|--------|--------|
| Coefficient | -0.181 | -0.369 | -0.255 | -0.335 |
| p (firm cluster) | 0.015 | 0.002 | 0.010 | 0.005 |
| p (WCB) | 0.013 | <0.001 | 0.005 | <0.001 |
| p (two-way) | 0.580 | 0.320 | 0.452 | 0.299 |

Notes: Entries are p-values for the lagged *CSI-orth* coefficient. WCB denotes wild cluster bootstrap inference with Webb weights; two-way clustering uses firm and month clusters.

Panel B shows that the baseline result is robust to the wild cluster bootstrap but not to two-way clustering. We therefore interpret the firm-clustered estimates as evidence

of a directional association, while treating inference as sensitive to common monthly shocks. This qualification is important because the effective time-series dimension is only 107 months, even though the regression is estimated in a firm-month panel. The block permutation test in Section 5.7 (Appendix Figure A.3) addresses the related persistence concern directly.

5.4 Contemporaneous and reversal patterns

To decompose the H1 pattern, we separate same-month changes in *CSI-orth* from lagged *CSI-orth* levels. The baseline specification in Table 6 only asks whether lagged sentiment levels predict later discount changes; the decomposition adds the current monthly change in the index:

$$\Delta \text{Discount}_{i,t} = \alpha_i + \beta_0 \Delta \text{CSI-orth}_t + \beta_1 \text{CSI-orth}_{t-k} + \varepsilon_{i,t}. \quad (4)$$

A positive β_0 means that when *CSI-orth* rises in month t , price-to-NAV also tends to rise in the same month. This corresponds to discount narrowing, or premium expansion. A negative β_1 means that high *CSI-orth* k months earlier is followed by a decline in price-to-NAV, corresponding to later discount widening.

Table 7: Contemporaneous $\Delta \text{CSI-orth}$ and lagged *CSI-orth* levels. The dependent variable is $\Delta \text{Discount}$. All specifications include firm fixed effects.

| | Lag 1 | Lag 2 | Lag 3 | Lag 6 |
|-----------------------------|---------------------|----------------------|---------------------|----------------------|
| $\Delta \text{CSI-orth}(t)$ | 0.510*** (0.131) | 0.473*** (0.124) | 0.501*** (0.125) | 0.419*** (0.118) |
| $\text{CSI-orth}(t - k)$ | -0.032 (0.064) | -0.325*** (0.097) | -0.244** (0.087) | -0.307*** (0.104) |
| N | 1,748 | 1,733 | 1,718 | 1,673 |
| Within-R ² | 0.0060 | 0.0097 | 0.0078 | 0.0080 |

Notes: Standard errors in parentheses are clustered by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. $\Delta \text{CSI-orth}$ is the contemporaneous first difference of the orthogonalized sentiment index.

$\Delta \text{CSI-orth}(t)$ is positively associated with same-month discount changes: β_0 ranges from approximately 0.42 to 0.51 and is significant at the 1% level across all specifications. The positive contemporaneous coefficient means that increases in *CSI-orth* are associated with contemporaneous increases in price-to-NAV, corresponding to discount narrowing or premium expansion.

The lagged *CSI-orth* level is negative and significant from lag 2 onward. The negative coefficient indicates that elevated past sentiment is followed by later declines in price-to-NAV, consistent with delayed discount widening. At lag 1, the coefficient is small and insignificant (-0.032, $p = 0.616$), suggesting that one month may be too short for this reversal pattern to appear.

Taken together, the estimates are consistent with temporary sentiment-related price pressure followed by partial reversal, in the spirit of Lee, Shleifer, and Thaler (1991). The

pattern remains descriptive rather than causal: both regressors are derived from the same *CSI-orth* series and may also reflect financial-market or macroeconomic conditions not fully removed by the orthogonalization.

5.5 Individual proxy regressions

We estimate equation (1) separately for each standardized sentiment proxy at lag 2. The purpose is diagnostic: to assess which components of the CSI carry predictive content individually, and to compare the market-based proxies with the survey-based ETI, which is excluded from the CSI.

Table 8: Individual proxy regressions at lag 2. The dependent variable is $\Delta Discount$. Each row reports a separate firm-fixed-effects regression using one standardized proxy.

| Proxy | Coefficient | p | N | Within-R2 |
|----------------|----------------------|-------|-------|-----------|
| <i>NIPO</i> | -0.208 (0.194) | 0.298 | 1,748 | 0.0008 |
| <i>RIPO</i> | 0.179* (0.100) | 0.091 | 1,748 | 0.0006 |
| <i>TURN</i> | 0.355** (0.136) | 0.019 | 1,748 | 0.0024 |
| <i>S</i> | -0.209 (0.156) | 0.199 | 1,748 | 0.0007 |
| <i>P(D-ND)</i> | -0.296*** (0.091) | 0.005 | 1,748 | 0.0016 |
| <i>ETI</i> | -0.620*** (0.199) | 0.006 | 1,748 | 0.0076 |

Notes: Standard errors in parentheses are clustered by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. *ETI* is an external survey-based comparator and is not included in the composite sentiment index.

Of the five CSI components, *P(D-ND)* and *TURN* are individually significant at the 5% level. *P(D-ND)* is also the highest-loading proxy in the CSI (Table 5), indicating that the dividend premium carries an important share of the index's predictive content. This should be interpreted cautiously, since its positive loading in the Swedish CSI is opposite to the negative loading reported by Baker and Wurgler (2006) for the U.S. market.

TURN enters with a positive coefficient: higher past turnover is associated with a subsequent rise in price-to-NAV, corresponding to discount narrowing or premium expansion. This is consistent with *TURN*'s negative loading on PC1. In the CSI, higher turnover lowers the index, so the standalone *TURN* result points in the same direction as the baseline CSI reversal pattern.

ETI is the strongest individual predictor in Table 8, despite being excluded from the market-based CSI. *ETI*'s standalone strength suggests that the survey-based measure contains information about discount changes that is not fully captured by the market-based proxies. We nevertheless retain *CSI-orth* as the headline index because the Baker and Wurgler (2006) market-based methodology is the paper's comparability anchor with the international literature.

5.6 Cross-sectional heterogeneity

This section tests whether sentiment sensitivity varies with fund characteristics. H2 predicts stronger sentiment-related discount movements for funds with greater unlisted exposure, used as a proxy for arbitrage difficulty. H3 predicts weaker sentiment sensitivity for funds with higher non-retail ownership concentration, following the retail-investor mechanism in Lee, Shleifer, and Thaler (1991).

Sample sizes are slightly smaller than in the baseline regressions because both moderators enter with lower-frequency lagged data. *UnlistedShare* is collected quarterly and merged into the monthly panel with a one-quarter lag. *NonRetailOwnPct* is observed annually from Modular Finance Holdings and enters with a one-year lag.

5.6.1. Pontiff channel: arbitrage difficulty

We estimate equation (2) at lags $k = 0, 1, 2, 3,$ and 6 . The coefficient of interest is the interaction between lagged *CSI-orth* and *UnlistedShare*. The lag sequence is used as a diagnostic for the two interpretations discussed in Section 4.4. A mechanical NAV-staleness explanation predicts the interaction to be strongest at lag 0, when market prices update before unlisted NAV components. An arbitrage-friction interpretation is instead more consistent with a delayed interaction, as sentiment-related mispricing may persist longer in funds whose portfolios are harder to replicate.

Table 9: *UnlistedShare* \times *CSI-orth* interaction. The dependent variable is $\Delta Discount$. The table reports firm-fixed-effects estimates of the H2 specification across lags.

| | Lag 0 | Lag 1 | Lag 2 | Lag 3 | Lag 6 |
|---|-------------------|--------------------|---------------------|-------------------|---------------------|
| <i>CSI-orth</i> ($t - k$) | 0.199* (0.108) | -0.076 (0.072) | -0.154 (0.103) | -0.095 (0.088) | -0.088 (0.095) |
| <i>UnlistedShare</i> | -0.297 (0.699) | -0.154 (0.660) | -0.173 (0.840) | -0.286 (0.775) | -0.073 (1.051) |
| <i>CSI-orth</i> \times <i>UnlistedShare</i> | -0.136 (0.161) | -0.263* (0.133) | -0.596** (0.228) | -0.397 (0.318) | -0.728** (0.323) |
| N | 1,689 | 1,689 | 1,689 | 1,675 | 1,633 |
| Within-R ² | 0.0009 | 0.0012 | 0.0056 | 0.0024 | 0.0058 |
| p (WCB), interaction | 0.435 | 0.130 | 0.038 | 0.202 | 0.009 |
| p (two-way), interaction | 0.695 | 0.297 | 0.020 | 0.317 | 0.109 |

Notes: Standard errors in parentheses are clustered by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. The WCB and two-way p-values refer to the interaction row. *UnlistedShare* is lagged one quarter.

The interaction is insignificant at lag 0, negative but insignificant at lag 3, and statistically significant at lags 2 and 6 under firm-clustered standard errors. The lag-2 estimate is the central result. It remains significant under two-way clustering by firm and month ($p = 0.020$) and under the wild cluster bootstrap ($p = 0.038$). The leave-one-firm-out results in Appendix Table A.4 also show that the lag-2 coefficient remains negative and of similar magnitude when each firm is excluded in turn.

The lag profile is informative for the NAV-staleness concern discussed in Section 4.4. The interaction is absent at lag 0 but significant at lag 2, which is difficult to reconcile with the simplest stale-NAV story and more consistent with stronger sentiment-related discount widening among harder-to-arbitrage funds, in the spirit of Pontiff (1996). This

diagnostic remains suggestive rather than definitive, since quarterly NAV-update timing varies across firms and delayed private-valuation adjustments cannot be fully ruled out.

The lag profile should not be read as a smooth build-up over time. The effect is concentrated at lag 2 rather than increasing monotonically, and the lag-6 estimate is reported for completeness rather than given economic weight. CSI-orth is persistent (Section 4.2), and the Δ CSI specification in Section 5.7 does not support the long-lag pattern.

5.6.2. Ownership channel (H3): ownership composition

The H3 test estimates equation (3) at lags 2 and 3. The unit of analysis is the firm-month, and the coefficient of interest is the $\text{CSI-orth}(t-k) \times \text{NonRetailOwnPct}(i,t-1Y)$ interaction. Because the baseline lagged CSI-orth coefficient is negative, H3 predicts a positive interaction: higher non-retail ownership should dampen the discount-widening association after elevated sentiment.

Table 10: *NonRetailOwnPct* \times *CSI-orth* interaction. The dependent variable is Δ *Discount*. The table reports firm-fixed-effects estimates and pooled OLS estimates with controls.

| | Firm FE, lag 2 | Firm FE, lag 3 | Pooled OLS, lag 2 | Pooled OLS, lag 3 |
|---|-------------------|--------------------|---------------------|-------------------|
| <i>CSI-orth</i> ($t - k$) | -0.211 (0.267) | -0.619* (0.295) | 0.028 (0.230) | -0.468 (0.284) |
| <i>CSI-orth</i> \times <i>NonRetailOwnPct</i> | -0.347 (0.380) | 0.514 (0.377) | -0.752** (0.330) | 0.195 (0.329) |
| <i>NonRetailOwnPct</i> | -5.076 (5.232) | -5.970 (5.330) | -0.248 (0.466) | -0.445 (0.498) |
| Firm-level controls | – | – | Yes | Yes |
| p (WCB), interaction | 0.517 | 0.152 | 0.148 | 0.536 |
| p (two-way), interaction | 0.189 | 0.097 | 0.006* | 0.359 |
| N | 1,662 | 1,648 | 1,251 | 1,240 |
| Fit statistic | 0.008 | 0.004 | 0.013 | 0.010 |

Notes: Standard errors in parentheses are clustered by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. The WCB and two-way p-values refer to the interaction row. *NonRetailOwnPct* is lagged one year. The asterisk on the pooled-OLS lag-2 two-way p-value indicates that the variance-covariance matrix was not positive semi-definite and was numerically corrected; we do not interpret that value as supportive evidence. The fit statistic is within- R^2 for the firm-FE columns and model R^2 for the pooled-OLS columns.

At lag 2, both specifications produce a negative interaction, opposite to the H3 prediction. Under firm fixed effects, the estimate is statistically insignificant. Under pooled OLS with controls, the interaction is larger and significant under firm-clustered standard errors, but it does not survive the wild cluster bootstrap. At lag 3, the sign reverses in both specifications and neither estimate is statistically significant. The firm-fixed-effects lag-2 interaction also remains insignificant under two-way clustering by firm and month. The pooled-OLS lag-2 two-way p-value is reported for completeness, but the variance-covariance matrix was not positive semi-definite and was numerically corrected, so we do not interpret it as supportive evidence.

Overall, Table 10 does not support the predicted dampening effect of non-retail ownership on sentiment sensitivity. We treat the firm-fixed-effects estimate as the conservative baseline, consistent with Section 4. The pooled OLS estimate is a descriptive complement that uses more between-fund variation, but is more vulnerable to omitted fund-level differences. The result should not be read as a refutation of Lee, Shleifer, and Thaler

(1991), since the sample does not contain a retail-concentrated benchmark and ownership concentration need not identify the marginal price-setter.

Power is limited in both interaction tests. For H2, the lag-2 interaction is large enough to be economically interpretable and is detected despite the small cross-section. For H3, by contrast, the lag-2 firm-fixed-effects estimate is -0.347, below the roughly 0.7 absolute effect size needed for conventional 5% detection given the estimated standard error, and it has the opposite sign of the prediction. We therefore read H3 as unsupported, not as precise evidence that ownership concentration has no moderating role.

5.7 Robustness

This section reports additional robustness and diagnostic tests for the baseline finding. Per-check coefficients are reported in Appendix Table A.3.

The proxy-composition checks no longer weaken the baseline association. Excluding *TURN* (R1), excluding *P(D-ND)* (R2), and excluding both sign-divergent proxies (R3) all preserve negative and statistically significant coefficients at lags 2 and 3. These checks therefore indicate that the baseline result is not mechanically driven by either sign-divergent proxy. Because the reconstructed indices are based on fewer proxies, we interpret the larger absolute coefficients cautiously; the main point is that the sign and significance of the H1 association are not dependent on *TURN* or *P(D-ND)*.

The ΔCSI specification (R6) replaces lagged *CSI* levels with first-differenced *CSI*, addressing the concern that the baseline pairs a stationary dependent variable with a persistent regressor. The coefficient is insignificant at both lags, indicating that the levels association is partly carried by *CSI* persistence. A block permutation test gives a similar caution: under a null where *CSI* autocorrelation is preserved but its temporal alignment with discounts is broken, a coefficient as large as the observed one arises roughly one in eight to ten times (Figure A.3). We keep *CSI* in levels as the baseline regressor to match Baker and Wurgler (2006), but treat persistence as an important qualification in the Discussion.

Other sample and coding checks preserve the baseline pattern. Excluding March–June 2020 (R4), restricting the sample to firms present throughout 2017–2025 (R5), and reconstructing the *CSI* under alternative codings of missing monthly *RIPO* values (R7) all preserve significance at both lags.

6 Discussion

In the baseline firm-fixed-effects specifications, higher lagged *CSI-orth* is associated with subsequent discount widening at the two- to three-month horizon. The lag-2 estimate implies that a one-standard-deviation increase in *CSI-orth* is associated with a 0.51-percentage-point widening of the discount, roughly 7% of the standard deviation of $\Delta Discount$. The coefficient has the sign predicted by H1 and is robust to firm-clustered standard errors and the wild cluster bootstrap, but the evidence weakens when we allow for common monthly shocks, preserve *CSI* persistence in a block-permutation test, or replace *CSI* levels with first differences. We therefore do not claim that *CSI-orth* is a

reliable standalone predictor of discount changes. The evidence instead shows that the sentiment-discount relation has the predicted sign in a setting where institutional and family/sphere holders account for 33% to 89% of equity. The mean pairwise discount-change correlation in our sample is 0.228, below the 0.389 reported by Lee, Shleifer, and Thaler (1991), consistent with a weaker common discount component than in their U.S. sample.

H2's cross-sectional interaction is where the design is most informative. Funds with larger unlisted portfolio shares exhibit significantly larger discount responses to lagged sentiment. The lag-2 interaction survives firm-clustered standard errors, two-way clustering, the wild cluster bootstrap, and a leave-one-firm-out check. The interaction is absent at lag 0, so the most mechanical NAV-staleness explanation is unlikely to drive the result. We argue this is the empirical analogue for Swedish closed-end funds of Pontiff's (1996) finding that arbitrage frictions amplify sentiment-driven mispricing in funds whose portfolios are costliest to replicate. Gemmill and Thomas (2002) report a similar pattern for U.K. investment trusts. H2 is less vulnerable to the persistence concern that weakens H1. To produce the cross-sectional interaction, an omitted persistent factor would have to affect high-unlisted and low-unlisted funds differently when sentiment moves, not simply co-move with sentiment overall. We cannot fully separate the unlisted-share channel from a correlated fund-size channel, since smaller funds in our sample tend to hold proportionally more unlisted assets. Pontiff (1996) treats market capitalization as one of his four arbitrage-cost proxies. A size-driven reading therefore sits within the broader arbitrage-cost interpretation rather than displacing it.

H3 is the direct test of whether the Lee, Shleifer, and Thaler (1991) retail-investor mechanism is the channel through which sentiment moves discounts in our sample. If residual variation in non-retail ownership concentration were doing the work, funds with higher concentration would show weaker sentiment sensitivity. We do not detect such moderation. At lag 2 the firm-fixed-effects interaction is small and insignificant. The pooled OLS interaction is significant under firm-clustered standard errors but does not survive the wild cluster bootstrap. At lag 3 the sign reverses. Two design features bound this null. The cross-sectional range of *NonRetailOwnPct* (33% to 89%) does not span the retail-concentrated values their cross-sectional prediction concerns. Even our least concentrated fund far exceeds the 6.6% institutional ownership they report, and Ma's (2024) "low retail" U.S. subsample averaged 78% retail. And ownership concentration describes the shareholder register, not the marginal price-setter in the float. Within those bounds, the absence of any consistent moderation is informative for the paper's main question. The sentiment-discount association in our sample does not appear to be driven by residual variation in non-retail ownership concentration, and it persists across the full range of concentration we observe.

Taken together, the evidence is consistent with extending the Lee, Shleifer, and Thaler (1991) sentiment-discount association beyond the retail-heavy archetypal setting in which it was first documented. H2 carries the inferential weight. The cross-sectional Pontiff (1996) interaction survives every inference adjustment we apply. H1's time-series association is sensitive to persistence and common-shock corrections, and H3 detects no moderation by ownership concentration. The arbitrage-cost channel does the moderation work that ownership composition does not. The arbitrage-cost pattern suggests the link operates through the closed-end structure plus arbitrage frictions rather than through retail demand specifically. Two limitations bound this reading. First, *CSI* is persistent

in the 2017–2025 sample. The ΔCSI null and the block permutation result both reflect this. Second, the cross-section of 18 funds limits the power of any heterogeneity test, and the unlisted-share channel cannot be cleanly separated from a fund-size channel within the same sample.

7 Conclusion

The most robust cross-sectional finding is that closed-end funds with larger unlisted portfolio shares exhibit larger discount responses to lagged investor sentiment. The interaction between sentiment and unlisted exposure is strongest at lag 2, and the absence of a lag-0 interaction weighs against the simplest mechanical NAV-staleness explanation. We interpret this pattern as consistent with an arbitrage-friction channel in the spirit of Pontiff (1996): funds with harder-to-replicate portfolios appear more sensitive to sentiment-related discount movements. This interpretation remains cautious, since unlisted exposure may also capture opacity, fund size, or delayed private-valuation adjustments. Still, the evidence suggests that sentiment-related discount dynamics can appear outside the retail-concentrated ownership setting emphasized by Lee, Shleifer, and Thaler (1991).

The aggregate time-series evidence is more measured: lagged CSI is associated with discount widening at the two- to three-month horizon, but the association is sensitive to persistence and common-shock adjustments. We therefore read H1 as directional rather than predictive or causal, consistent with slow-moving sentiment regimes accumulating into discount movements. The CSI-orth result suggests that the index carries information not captured by the macro controls used in the orthogonalization, while the survey-based ETI carries additional predictive content; we do not adjudicate which is closer to the marginal trader.

H3 does not support the predicted dampening effect of non-retail ownership concentration on sentiment sensitivity. The lag-2 interaction is small and not robust to wild-cluster inference, and the sign reverses at lag 3. We do not read this as a refutation of Lee, Shleifer, and Thaler (1991): the sample concentration range does not reach the retail-concentrated values their cross-sectional prediction concerns, only 18 funds limit the power of any heterogeneity test, and the shareholder register need not describe the marginal price-setter. The bounded contribution is therefore that sentiment-related discount dynamics persist in closed-end funds whose registers are held primarily by non-retail owners.

On balance, the evidence points most strongly to an arbitrage-friction interpretation of sentiment-related discount dynamics in Swedish closed-end funds.

References

- Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *Journal of Finance* 61, 1645–1680.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2008, Bootstrap-based improvements for inference with clustered errors, *Review of Economics and Statistics* 90, 414–427.
- Cameron, A. Colin, and Douglas L. Miller, 2015, A practitioner’s guide to cluster-robust inference, *Journal of Human Resources* 50, 317–372.
- De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Noise trader risk in financial markets, *Journal of Political Economy* 98, 703–738.
- Faccio, Mara, and Larry H.P. Lang, 2002, The ultimate ownership of Western European corporations, *Journal of Financial Economics* 65, 365–395.
- Gemmill, Gordon, and Dylan C. Thomas, 2002, Noise trading, costly arbitrage, and asset prices: Evidence from closed-end funds, *Journal of Finance* 57, 2571–2594.
- Harding, Don, and Adrian Pagan, 2002, Dissecting the cycle: a methodological investigation, *Journal of Monetary Economics* 49, 365–381.
- Hjelström, Tomas, 2007, The closed-end investment company premium puzzle: Model development and empirical tests on Swedish and British data, Doctoral dissertation, Stockholm School of Economics.
- Ibindex, n.d., *Substansvärden för svenska investmentbolag* (daily price and NAV histories for Swedish closed-end funds), <https://ibindex.se> (accessed 2026-03-10).
- Konjunkturinstitutet, n.d., *Konjunkturbarometern* (Economic Tendency Survey), <https://www.konjunkturinstitutet.se/en/our-publications/economic-tendency-survey/> (accessed 2026-03-06).
- Lee, Charles M.C., Andrei Shleifer, and Richard H. Thaler, 1991, Investor sentiment and the closed-end fund puzzle, *Journal of Finance* 46, 75–109.
- LSEG, n.d., *Datastream* (TOTMKSD Sweden country-index turnover and market capitalization), subscription product (accessed 2026-03-14).
- Ma, Liang, 2024, What drives closed-end fund discounts? Evidence from COVID-19, *Financial Management* 53, 119–143.
- MacKinnon, James G., and Matthew D. Webb, 2017, Wild bootstrap inference for wildly different cluster sizes, *Journal of Applied Econometrics* 32, 233–254.
- Modular Finance, n.d., *Holdings* (annual shareholder registers derived from Euroclear), <https://www.modularfinance.se/en/holdings> (accessed 2026-03-20 to 2026-05-02).
- Pontiff, Jeffrey, 1996, Costly arbitrage: Evidence from closed-end funds, *Quarterly Journal of Economics* 111, 1135–1151.
- S&P Global Market Intelligence, n.d., *Capital IQ Pro* (Swedish IPO and issuance deals; monthly market capitalization, dividend yield, and M/B ratios), <https://www.spglobal.com/marketintelligence/en/capital-iq-pro> (accessed 2026-03-10 to 2026-03-19).

Statistics Sweden (SCB), n.d., *Statistical database* (Industrial Production Index, AKU Labour-Force Survey, Household Consumption Indicator, Quarterly National Accounts), <https://www.scb.se/en/finding-statistics/> (accessed 2026-03-11 to 2026-03-14).

Appendix

This appendix presents six supporting tables and three supporting figures. Table A.1 reports the lag/lead selection diagnostic underlying the Composite Sentiment Index. Table A.2 compares raw-CSI and orthogonalized-CSI baseline regressions. Table A.3 collects the H1 robustness coefficients (alternative CSI compositions, sample restrictions, persistence, and RIPO coding) discussed in Section 5.7. Table A.4 reports a leave-one-firm-out check on the H2 lag-2 interaction. Table A.5 reports the H3 ownership-classifier robustness check used in Section 5.6. Table A.6 reports correlation diagnostics among the raw CSI, CSI-orth, and ETI used in Sections 3.3 and 5.1. Figure A.1 plots the average NAV discount over the sample period. Figure A.2 visualizes the cross-sectional and within-firm distribution of NonRetailOwnPct across sample funds. Figure A.3 reports null distributions from a block permutation test of the H1 lag-2 baseline coefficient.

Table A.1: Lag/lead alignment of sentiment proxies. For each proxy, the correlation with the first principal component of all ten candidate variables (five proxies at t and $t-1$) is reported. The timing with the higher absolute correlation is selected, following Baker and Wurgler (2006). $Cor(t)$ and $Cor(t-1)$ are the correlations of each timing with the first-stage $PC1$.

| Proxy | Cor(t) | Cor(t-1) | Selected |
|---------|---------|----------|----------|
| NIPO | 0.1952 | 0.2083 | t-1 |
| RIPO | 0.2248 | 0.3432 | t-1 |
| TURN | -0.6367 | -0.5665 | t |
| S | 0.6413 | 0.6497 | t-1 |
| P(D-ND) | 0.8852 | 0.8874 | t-1 |

Table A.2: Raw *CSI* vs. four-control *CSI-orth* at lags 2 and 3. The dependent variable is $\Delta Discount$. The table compares the raw and orthogonalized sentiment indices using the baseline firm-fixed-effects specification.

| Specification | Lag 2 | Lag 3 |
|---|----------------------|----------------------|
| Raw <i>CSI</i> | -0.316*** (0.095) | -0.335*** (0.097) |
| <i>CSI-orth</i> | -0.369*** (0.098) | -0.255*** (0.088) |
| N | 1,733 | 1,718 |
| Within-R ² , raw <i>CSI</i> | 0.0036 | 0.0040 |
| Within-R ² , <i>CSI-orth</i> | 0.0049 | 0.0024 |

Notes: Standard errors in parentheses are clustered by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Within-R² is the within-firm R².

Table A.3: Robustness checks for the baseline H1 regression. The dependent variable is $\Delta Discount$. Each row reports one lag-specific robustness specification.

| Specification | Coefficient | p | N |
|-----------------------------|----------------------|--------|-------|
| Baseline (raw CSI), lag 2 | -0.316*** (0.095) | 0.004 | 1,733 |
| Baseline (raw CSI), lag 3 | -0.335*** (0.097) | 0.003 | 1,718 |
| R1: No TURN, lag 2 | -0.416*** (0.111) | 0.002 | 1,733 |
| R1: No TURN, lag 3 | -0.536*** (0.150) | 0.002 | 1,718 |
| R2: No PDND, lag 2 | -0.437*** (0.126) | 0.003 | 1,733 |
| R2: No PDND, lag 3 | -0.368*** (0.122) | 0.008 | 1,718 |
| R3: No TURN, no PDND, lag 2 | -0.619*** (0.138) | <0.001 | 1,733 |
| R3: No TURN, no PDND, lag 3 | -0.512** (0.177) | 0.010 | 1,718 |
| R4: No COVID, lag 2 | -0.283** (0.102) | 0.013 | 1,673 |
| R4: No COVID, lag 3 | -0.308*** (0.106) | 0.010 | 1,658 |
| R5: Balanced panel, lag 2 | -0.266** (0.091) | 0.011 | 1,575 |
| R5: Balanced panel, lag 3 | -0.249*** (0.066) | 0.002 | 1,560 |
| R6: Δ CSI, lag 2 | 0.067 (0.139) | 0.634 | 1,718 |
| R6: Δ CSI, lag 3 | 0.150 (0.169) | 0.388 | 1,703 |
| R7: RIPO NA+fill, lag 2 | -0.307*** (0.094) | 0.005 | 1,733 |
| R7: RIPO NA+fill, lag 3 | -0.359*** (0.094) | 0.001 | 1,718 |
| R7: RIPO cond-mean, lag 2 | -0.279*** (0.092) | 0.008 | 1,733 |
| R7: RIPO cond-mean, lag 3 | -0.319*** (0.090) | 0.003 | 1,718 |

Notes: Standard errors in parentheses are clustered by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table A.4: Leave-one-firm-out robustness for the H2 interaction. The dependent variable is $\Delta Discount$. Each row drops one firm from the lag-2 $CSI-orth \times UnlistedShare$ regression.

| Firm dropped | N | Interaction B | SE | p (clust) |
|-------------------|------|---------------|-------|-----------|
| VNV Global | 1584 | -0.585 | 0.375 | 0.1386 |
| VEF | 1584 | -0.686 | 0.311 | 0.0421 |
| Industrivärden | 1584 | -0.58 | 0.241 | 0.0284 |
| Öresund | 1584 | -0.577 | 0.237 | 0.0272 |
| AJA | 1584 | -0.536 | 0.219 | 0.026 |
| Spiltan | 1623 | -0.567 | 0.23 | 0.0254 |
| Kinnevik | 1584 | -0.564 | 0.228 | 0.0248 |
| Flat Capital | 1641 | -0.47 | 0.189 | 0.0241 |
| First Ventures | 1638 | -0.594 | 0.236 | 0.0228 |
| Investor | 1584 | -0.583 | 0.228 | 0.021 |
| Latour | 1584 | -0.594 | 0.229 | 0.0196 |
| Lundbergföretagen | 1584 | -0.597 | 0.23 | 0.0194 |
| Svolder | 1584 | -0.63 | 0.241 | 0.0187 |
| Creades | 1584 | -0.612 | 0.229 | 0.0166 |
| Linc | 1635 | -0.63 | 0.229 | 0.0143 |
| Bure | 1584 | -0.635 | 0.23 | 0.0138 |
| Traction | 1584 | -0.647 | 0.231 | 0.0129 |
| NAXS | 1584 | -0.721 | 0.228 | 0.006 |

Notes: Standard errors in parentheses are clustered by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

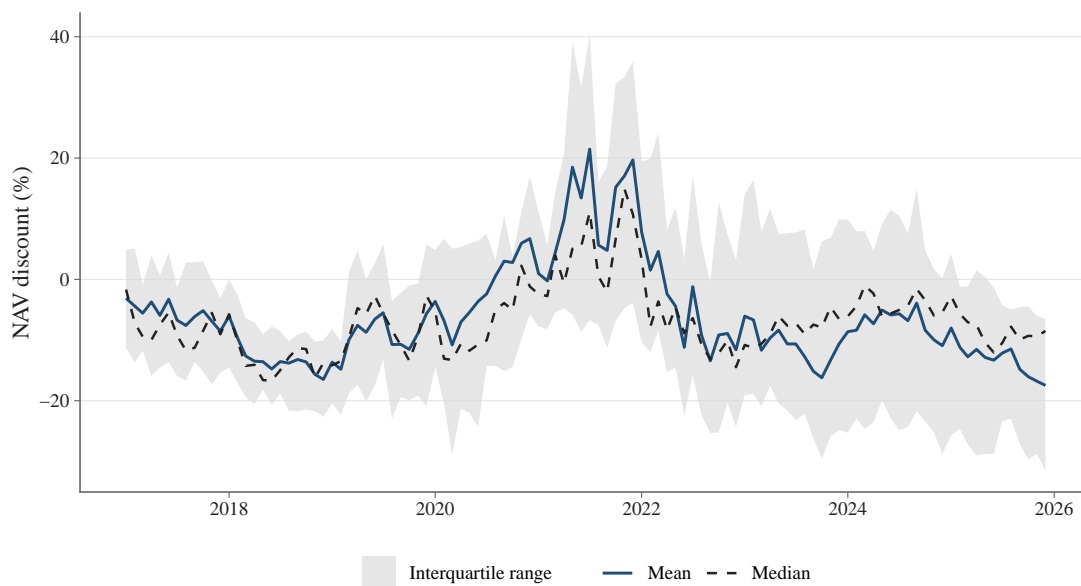


Figure A.1. Average NAV discount across sample Swedish closed-end funds, January 2017 – December 2025. The solid line is the equally-weighted monthly mean discount across all funds in the panel. The dashed line is the median. The shaded region indicates the interquartile range. Discount is defined as $(Price - NAV) / NAV \times 100$.

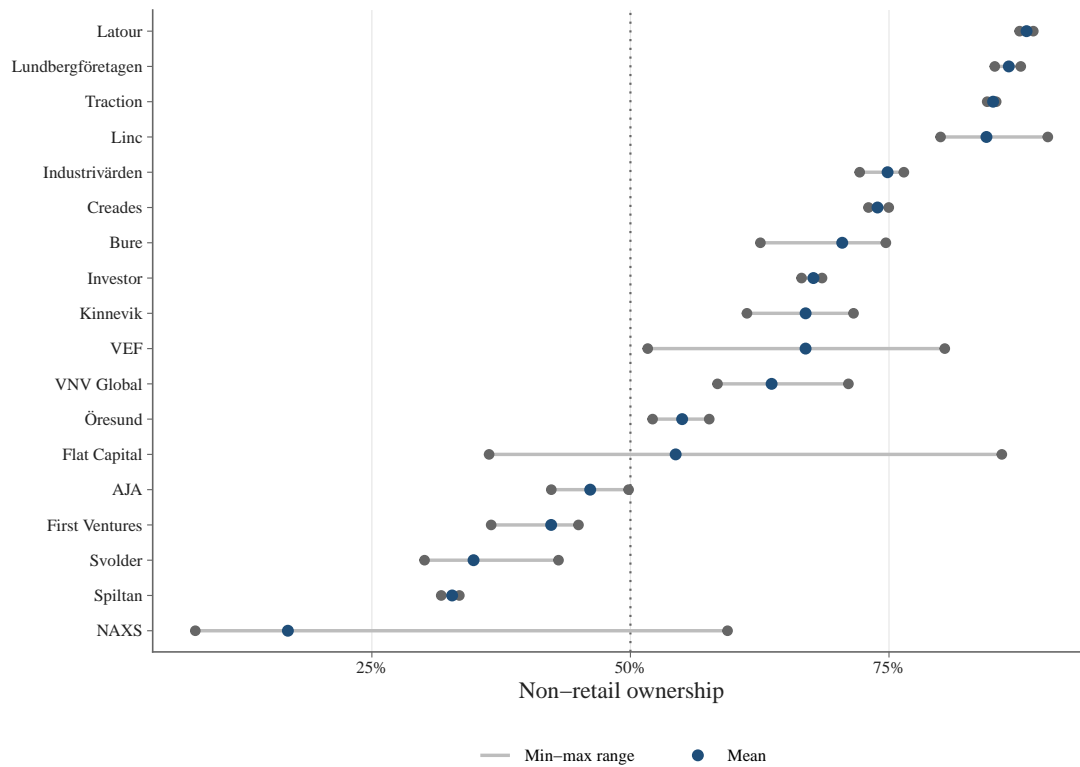


Figure A.2. Cross-sectional and within-firm distribution of non-retail ownership across the 18 sample firms, 2016–2025. Each row represents one firm. The horizontal segment spans the firm’s minimum-to-maximum NonRetailOwnPct across annual Modular Finance Holdings snapshots; the lighter end-points mark the min and max; the dark dot marks the time-series mean. Firms are sorted by mean ownership share. The dotted vertical line at 50% provides a visual reference. The figure makes two points the H3 test depends on: (i) the cross-sectional distribution is concentrated above 30% with no firm approaching retail-concentrated ownership, and (ii) within-firm variation is small for most firms (mean within-firm SD = 4.5 pp; maximum 19.5 pp at Flat Capital), so the within-firm component identified by firm fixed effects in equation (3) is limited.

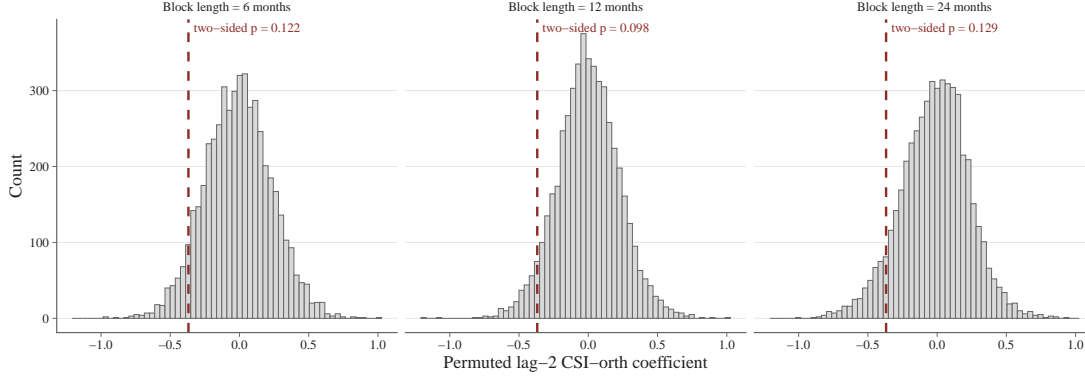


Figure A.3. Block permutation null distributions for the baseline lag-2 CSI-orth coefficient, with block lengths $L = 6$, $L = 12$, and $L = 24$. Each panel shows the distribution of lag-2 CSI-orth coefficients from 5,000 permutations of the monthly CSI-orth series; block permutation preserves local autocorrelation while breaking the temporal alignment with the discount series. The red dashed line marks the observed coefficient (-0.369). Two-sided permutation p-values are 0.122, 0.098, and 0.129 respectively; one-sided (left-tail) p-values are 0.063, 0.048, and 0.078.

Table A.5: Robustness of H3 to ownership-classifier breadth. The table compares the narrow and expanded non-retail ownership classifiers using the lag-2 H3 specification.

| Specification | Classifier | Coefficient | p | p (WCB) |
|-------------------|------------|---------------------|-------|---------|
| Firm FE, lag 2 | Narrow | -0.394 (0.361) | 0.290 | |
| Firm FE, lag 2 | Expanded | -0.347 (0.380) | 0.373 | |
| Pooled OLS, lag 2 | Narrow | -0.770** (0.326) | 0.036 | 0.136 |
| Pooled OLS, lag 2 | Expanded | -0.752** (0.330) | 0.042 | 0.153 |

Notes: Standard errors in parentheses are clustered by firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. WCB p-values are reported for the pooled OLS specifications.

Table A.6: *CSI/ETI correlation diagnostics.* The table reports correlations among the raw *CSI*, *CSI-orth*, and the survey-based *ETI*.

Panel A: Index correlations.

| Index | raw CSI | CSI-orth | ETI (z) |
|----------|---------|----------|---------|
| raw CSI | 1 | 0.978 | 0.202 |
| CSI-orth | 0.978 | 1 | 0.19 |
| ETI (z) | 0.202 | 0.19 | 1 |

Panel B: Correlations with future average discount changes.

| Index | k = 0 | k = 1 | k = 2 | k = 3 |
|----------|-------|--------|--------|--------|
| raw CSI | 0.054 | -0.095 | -0.137 | -0.147 |
| CSI-orth | 0.043 | -0.079 | -0.159 | -0.12 |
| ETI (z) | -0.12 | -0.162 | -0.154 | -0.191 |

Notes: Panel A reports pairwise correlations among sentiment measures. Panel B reports correlations between each sentiment measure and future average discount changes.

AI Disclosure

Claude (Anthropic) and ChatGPT (OpenAI) were used as assistive tools to support R coding, data pipeline development, table formatting, and debugging. All generated code was reviewed, tested, and adapted by the authors. AI tools were also used to assist with improving clarity, grammar, and phrasing. The authors retained responsibility for all research design choices, empirical decisions, interpretation of results, and final text.