The Price of Passive Ownership *

Pouya Behmaram†

May 3, 2026

Abstract

Highly indexed U.S. stocks look like winners because indexed demand has raised their valuations, leaving them with lower expected returns. I measure this exposure with the Indexing Inclusion Ratio (IXI), a monthly holdings-based measure of realized indexed ownership built from best-fit benchmark assignment, fund holdings, and institutional 13F holdings. Stocks with high IXI have implied costs of capital about two percentage points per year below stocks with low IXI, yet earned higher realized returns during the active-to-passive transition as passive ownership grew faster than prices had incorporated. Once that growth is anticipated or slows, the realized premium fades and the lower expected-return wedge is what remains.

Keywords: asset pricing, institutional investors, index effect, passive investing, mutual funds, ETFs
JEL Codes: G11, G12, G23

*An earlier version of this paper circulated under the title "From Realized to Expected: The Passive Investing Impact." I am grateful to Laurent Barras, David Schumacher, Daniel Andrei, Sebastien Betermier, Evan Jo, Jiacui Li, Charlotte Haendler, Vincent Grégoire, Thomas Rivera, Gregory Weitzner, Begum Ipek Yavuz, and seminar participants at McGill University, UQAM ESG, and the FMA for guidance on earlier versions of this work. I acknowledge research support from the Social Sciences and Humanities Research Council of Canada. All errors are my own.

†Université du Québec à Montréal, École des sciences de la gestion (behmaram.pouya@uqam.ca)

1 Introduction

Over the past two decades, the most heavily indexed U.S. stocks earned higher realized returns than less heavily indexed stocks, yet they are priced today to earn lower expected returns. Both facts arise from the same active-to-passive transition. Indexed demand has been capitalized into prices, raising valuations and compressing forward expected returns. Realized returns were positive while passive ownership was still growing faster than prices had incorporated. Once that growth is anticipated or slows, the realized premium fades and the lower expected-return wedge is what remains. Testing this account requires a stock-level measure of indexed ownership.

The empirical challenge is that index-linked ownership is not the same thing as index membership, fund labels, or aggregate passive AUM. A stock can sit in many index-like portfolios with different weights, and index-linked capital can arrive through index funds, ETFs, active funds whose portfolios closely follow benchmarks, pensions, insurers, and other institutions. Existing measures capture either investor incentives, narrow passive-fund ownership, or membership in a particular index. They do not provide a monthly stock-level measure of realized indexed ownership that can be taken directly to expected-return and realized-return tests.

I measure this exposure with the Indexing Inclusion Ratio (IXI), a monthly holdings-based measure of how much index-linked capital is held in each stock. IXI uses best-fit index assignment recovered from observed portfolios and extends beyond mutual funds and ETFs through an institutional 13F branch. Because IXI is built from observed positions, it separates realized indexed ownership from fund labels and avoids assigning index-linked capital to stocks an institution does not actually hold. The resulting measure rises from roughly 4 percent in the early 2000s to above 14 percent by May 2024, with the steepest increase after 2013.

Index-linked capital absorbs part of a stock's residual supply, raising prices and lowering expected returns. Unexpected increases in that ownership produce realized revaluation

shocks. IXI therefore treats indexing depth as a continuous stock-level characteristic rather than as index membership.

Stocks with high IXI carry an implied cost of capital about two percentage points per year lower than stocks with low IXI. The portfolio spread is mirrored in a firm-level panel with firm and month fixed effects, size, value, and momentum controls. The wedge survives size conditioning, alternative ICC constructions, industry adjustment, and the exclusion of technology stocks. This ex-ante wedge gives the most direct evidence that indexed ownership is priced. Deeper indexed ownership is associated with lower forward returns.

Realized returns line up with the same ownership characteristic. A value-weighted high-minus-low IXI portfolio earns 0.43 percentage points per month during the passive era and positive alphas against standard factor models. The premium is concentrated after 2013, when the passive share of U.S. equity funds accelerated most sharply, and it remains visible in size-conditioned tests, so it is not only a mega-cap pattern. This creates the paper's main tension. Stocks with high IXI are priced to earn lower future returns, yet they delivered high realized returns while passive capital was still moving into the market.

The reconciliation lies in the distinction between expected and unexpected returns. Following Pástor et al. (2022), I decompose the realized high-minus-low IXI spread into an expected-return component built from the ICC wedge and an unexpected residual. The expected component averages -0.17 percentage points per month, mirroring the ex-ante wedge. The realized $+0.43$ pp spread is positive because the unexpected residual runs $+0.60$ pp in the opposite direction. In the aggregate data, that residual co-moves with unanticipated growth in the passive-ownership share and with earnings news concentrated on high-IXI stocks. This is the return signature of a transition: high-IXI stocks are priced to earn less going forward, yet earned more during a window in which passive ownership repeatedly exceeded what was already priced in.

At the stock-month level, passive-fund purchases are associated with positive same-month returns that substantially reverse over two to three months, consistent with short-run price

pressure. At the aggregate level, the first-difference decomposition is the paper's quantitative synthesis: subtracting the estimated co-movement between the cumulative high-minus-low IXI return path and growth in the passive-ownership share turns a 171-percent gain into a 35-percent loss. Because the exercise is reduced-form, the synthesis should be read as co-movement with passive-share growth rather than as a structural no-passive benchmark. The realized premium is a transition premium, not a stand-alone anomaly: it was earned while passive ownership was still growing faster than prices had already incorporated.

S&P 500 inclusions give a discrete-shock test of the same mechanism. In a modern passive market, most large stocks already sit in benchmark-linked portfolios through Russell, CRSP total-market, style, and sector products before they enter the S&P 500, so inclusion changes the composition of indexed ownership rather than turning it on. IXI lets me measure the gap between predicted S&P-family demand at inclusion and pre-existing passive ownership through other benchmark families. This S&P demand gap predicts the realized fund-branch IXI jump at the announcement and the short-window cumulative abnormal return; gross S&P-family demand does not consistently predict both margins. The long-horizon evidence disciplines the interpretation: treated firms' implied cost of capital does not fall after inclusion, so the inclusion design identifies short-window demand pressure rather than an average permanent decline in expected returns.

The evidence points to a single underlying ownership characteristic. Indexed ownership is not just a fund classification or an event indicator; it is the stock-level share of supply absorbed by benchmark-linked capital. In levels, that share compresses discount rates. In changes, unexpected increases generate revaluation. IXI measures it at the stock-month level, so the same object organizes the ex-ante return wedge, the transition-era realized premium, and the cross-sectional response to inclusion shocks. This is not another index-effect estimate but a return-side account of how the active-to-passive transition was capitalized into U.S. equity prices.

The closest precedents measure benchmarked or passive ownership at the stock level:

Pavlova and Sikorskaya (2023)'s incentive-based benchmarking-intensity measure built from index weights, and Chinco and Sammon (2024)'s aggregate accounting of the post-2010 expansion of the U.S. passive-ownership share. The contribution here is to show that realized indexed ownership, measured directly from holdings at the stock-month, organizes the expected-return wedge, the transition-era realized premium, and the cross-sectional inclusion-event evidence through the active-to-passive transition. The mechanism rests on demand-based asset pricing, in which inelastic benchmarked demand is priced in the cross section (Koijen and Yogo, 2019; Haddad et al., 2025; Gabaix and Koijen, 2021; Koijen et al., 2024), and complements the elasticity evidence in Behmaram (2026) from the return side. The inclusion evidence connects to Greenwood and Sammon (2025)'s account of the disappearing average index effect by showing that inclusion effects remain visible cross-sectionally when measured through the demand gap: predicted S&P-family demand relative to prior passive saturation predicts both the realized IXI jump and the short-window inclusion return. The cross-sectional design is the return-side counterpart of Sammon and Shim (2026)'s firm-side evidence on share supply to passive-fund buying.

Section 2 sets out the framework, and Section 3 defines IXI and compares it with existing measures. Sections 4 through 6 present the realized-return, expected-return, and aggregate-flow evidence. Sections 7 and 8 report the event-study and demand-gap evidence from S&P 500 inclusions.

2 A benchmarked-demand framework for effective benchmarked ownership

Effective benchmarked ownership is the share of a stock's supply held by institutions whose capital tracks a benchmark. In a standard delegated-asset-management economy, that share is the stock-level ownership characteristic linking the paper's expected-return, realized-return, and event-study evidence. It lowers the residual supply borne by price-sensitive

investors and therefore lowers expected returns. When it rises unexpectedly, it generates realized revaluation. The framework below formalizes those two margins and fixes notation for the IXI measure used in the rest of the paper. The equilibrium is the standard CARA-normal benchmarking environment of Brennan (1993) and Basak and Pavlova (2013), developed in the asset-pricing applications of Kashyap et al. (2021) and Pavlova and Sikorskaya (2023).

2.1 Setup and equilibrium

Consider a two-period economy with $t \in \{0, 1\}$. The investment menu contains a risk-free bond paying zero and N risky assets indexed by $n = 1, \ldots, N$, with period-1 cash flows $D \sim \mathcal{N}(\mu, \Sigma)$ and period-0 prices S. Risky assets have fixed share supply $\bar{x}$. Period-1 prices equal cash flows.

Benchmarks are indexed by $i = 1, \ldots, I$. Benchmark i is a portfolio with share weights $\omega_i \in \mathbb{R}^N_+$ per unit of benchmark capital. The economy is populated by three investor types, all with CARA utility and common risk aversion γ. Direct investors, share λ_D, manage their own capital. Active managers benchmarked to i, share λ_i^A, receive compensation $w_i = aR_i + b(R_i - R_i^B) + c$, where $R_i = x_i'(D - S)$ is the portfolio return and $R_i^B = \omega_i'(D - S)$ is the benchmark return. Passive managers on benchmark i, share λ_i^P, hold ω_i exactly.

Mean-variance optimization yields direct-investor demand $x_D = \gamma^{-1}\Sigma^{-1}(\mu - S)$, active-manager demand

$$x_i^A = \frac{1}{\gamma(a+b)}\Sigma^{-1}(\mu - S) + \frac{b}{a+b}\,\omega_i, \tag{1}$$

and passive demand ω_i. Aggregating across types, market clearing $\lambda_D x_D + \sum_i(\lambda_i^A x_i^A + \lambda_i^P \omega_i) = \bar{x}$ delivers the equilibrium price vector

$$S = \mu - \gamma A\,\Sigma\,(\bar{x} - B), \tag{2}$$

with aggregation constant $A = \left[\lambda_D + \sum_i \lambda_i^A/(a+b)\right]^{-1}$ and stock-level benchmarked share

demand

$$B_n = \frac{b}{a+b} \sum_i \lambda_i^A \, \omega_i(n) \; + \; \sum_i \lambda_i^P \, \omega_i(n). \tag{3}$$

Appendix E gives the full derivation.

2.2 Effective benchmarked ownership as the central object

Let $\bar{x}_n$ denote shares outstanding of stock n. Define the stock's benchmarked ownership share as

$$\mathcal{I}_n \equiv \frac{B_n}{\bar{x}_n}. \tag{4}$$

$\mathcal{I}_n$ is the share of the stock's supply absorbed by benchmarked capital. The complement $(1 - \mathcal{I}_n)$ is the residual share borne by price-sensitive investors, which equation (2) identifies as the driver of the stock's equilibrium risk premium.

The aggregation in (3) can equivalently be expressed at the institution level. Let e index institutions, each with capital A_e, benchmark exposure ω_e in shares per unit of capital, and tracking intensity $\alpha_e \in [0, 1]$. Then

$$B_n = \sum_{e \in \mathcal{E}} A_e \, \alpha_e \, \omega_e(n). \tag{5}$$

The compensation-contract derivation above corresponds to the special case where α_e takes only two values, $b/(a+b)$ for active managers and 1 for passive managers, uniformly within type. Equation (5) accommodates the more general case in which α_e varies at the institution level, consistent with heterogeneous effective tracking across institutions that nominally share a benchmark.

In the model, what matters for cross-sectional pricing is $\mathcal{I}_n$, the share of the stock's supply absorbed by benchmarked capital. Empirically, the paper's IXI measure is constructed to recover that same object from realized holdings, using best-fit benchmark assignment and institution-level tracking intensity rather than declared-benchmark labels and uniform

scaling. Because benchmark portfolios are implemented in value weights, the stock price cancels in the ratio of dollar benchmarked holdings to market capitalization, leaving the share ratio $\mathcal{I}_n$ up to measurement noise. Pavlova and Sikorskaya (2023) work with the special case in which $\alpha_e = 1$ on declared-benchmark AUM and the share is aggregated across a smaller menu of indices; Behmaram (2026) reports the detailed comparison between the two constructions on the elasticity side.

2.3 Cross-sectional expected-return wedge

Equation (2) implies a simple cross-sectional relation between expected returns and effective benchmarked ownership.

Proposition 1 (Cross-sectional expected-return wedge) *In the equilibrium of Section 2.1 and under diagonal Σ, the expected per-share excess payoff on stock n satisfies*

$$E[D_n - S_n] = \gamma A \, \sigma_n^2 \, \bar{x}_n \left(1 - \mathcal{I}_n\right). \tag{6}$$

Dividing by the positive price S_n gives the corresponding percentage expected return, which inherits the negative sign in $\mathcal{I}_n$. The expected excess payoff decomposes into an unconstrained baseline $\gamma A \sigma_n^2 \bar{x}_n$ and the residual share fraction $(1 - \mathcal{I}_n)$ borne by price-sensitive investors. Stocks more deeply embedded in the benchmarked-ownership system leave a smaller residual share, trade at a higher equilibrium price, and carry a lower ex-ante expected return. With a general covariance matrix, increasing a stock's own benchmarked-ownership share still lowers its expected return, and cross-stock covariance terms generate factor-structure components for which the empirical specifications include size, value, momentum, and industry controls as observable proxies. Appendix E carries both cases.

2.4 Unanticipated benchmark-demand shocks

Proposition 1 is a statement about expected returns given realized benchmarked ownership. To connect to realized returns in a setting where benchmarked ownership is itself evolving, extend the economy by adding an earlier period $t = -1$ at which $\mathcal{I}_n$ is uncertain and the market's pricing expectation of its $t = 0$ realization is $E^*_{-1}[\mathcal{I}_n]$. Holding A, Σ, expected cash flows, and share supply fixed, $\mathcal{I}_n$ realizes at $t = 0$ and prices clear at each date under equation (2).

Proposition 2 (Unanticipated benchmark-demand shocks) *In the extended equilibrium and under diagonal Σ, the realized per-share revaluation on stock n from $t = -1$ to $t = 0$ satisfies*

$$S_{n,0} - S_{n,-1} = \gamma A \, \sigma_n^2 \, \bar{x}_n \left(\mathcal{I}_n - E^*_{-1}[\mathcal{I}_n] \right). \tag{7}$$

Realized revaluation depends on the unanticipated component of benchmarked ownership; the anticipated component is already priced into $S_{n,-1}$. A closely related single-step result for the aggregate B_n appears in Appendix A of Pavlova and Sikorskaya (2023). Appendix E derives (7) by differencing the $t = -1$ and $t = 0$ equilibria.

2.5 Transition interpretation and empirical roadmap

Proposition 1 speaks to expected returns given realized $\mathcal{I}_n$. Proposition 2 speaks to realized returns at a single news resolution. The paper's empirical window is a sustained transition from active to passive management, over which benchmarked ownership has grown for roughly two decades. Read together, the two propositions organize that setting.

Each period of the transition is a realization of the mechanism in Proposition 2. When realized growth in $\mathcal{I}_n$ exceeds what was priced in at $t-1$, high-$\mathcal{I}$ stocks register positive revaluation returns; when it matches expectations, the benchmark-demand component is already in price and no additional demand-news revaluation occurs. Over a window of repeated positive surprises, cumulative realized returns on high-$\mathcal{I}$ stocks exceed their static ex-ante

expected returns even as Proposition 1's expected-return wedge remains negative. In the modern asset-pricing idiom of Campbell and Shiller (1988), the unanticipated compression of future expected returns is discount-rate news, and the realized return absorbs that news on the date it resolves. The Pástor et al. (2022) decomposition in Section 6.5 is empirically analogous to this split, separating realized HILI into an ex-ante expected component built from the Section 5 ICC wedge and an unanticipated residual that co-moves with passive-share growth. When the transition slows and realized $\mathcal{I}_n$ shocks no longer exceed expectations, the revaluation premium dissipates: high-$\mathcal{I}$ stocks continue to carry compressed expected returns but no longer earn an additional transition premium from benchmark-demand news. Aggregate growth in benchmark-linked capital is allocated through benchmark weights, so stocks already deep in benchmark-linked portfolios are the natural cross-sectional carriers of aggregate passive-demand news. The level of $\mathcal{I}_n$ is therefore the empirical proxy for transition-news exposure in the realized-return tests, while changes in $\mathcal{I}_n$ and inclusion shocks test the mechanism directly.

Discrete benchmark shocks. Index inclusions are a discrete-jump realization of Proposition 2. An inclusion-induced jump in $\mathcal{I}_n$ generates a short-window revaluation return when the jump was not fully priced before the announcement. After the jump is priced in, Proposition 1 predicts a lower ex-ante expected return for the included stock, all else equal, while indexed ownership remains elevated. Section 7 tests the short-window prediction. Section 8 measures the cross-sectional S&P demand gap, defined as predicted S&P-family demand at inclusion net of pre-existing passive ownership through other benchmark families, and tests whether it predicts the realized inclusion shock and the short-window return.

Section 3 describes the construction of IXI as the empirical proxy for $\mathcal{I}_n$. Section 4 documents the realized-return premium on high-$\mathcal{I}$ stocks that Proposition 2's transition reading predicts. Section 5 estimates the cross-sectional wedge of Proposition 1 through implied cost of capital. Section 6 presents the stock-level and aggregate passive-flow evidence

and the Pástor-Stambaugh-Taylor decomposition. Sections 7 and 8 test the discrete-shock prediction around S&P 500 inclusions.

3 Data and measurement

IXI is the paper's monthly stock-level measure of indexed ownership, constructed from observed fund and institutional holdings. It is the empirical proxy for the share $\mathcal{I}_n$ in Section 2's framework, and the object on which the realized-return, expected-return, and inclusion-event tests of the rest of the paper rely.

Notation and standard-error conventions. From Section 3 onward, i indexes stocks and t indexes calendar months; e indexes institutions here and in Section 6 but is reused as an event index locally in Section 7. The γ appearing inside empirical regression equations is a control-coefficient vector, distinct from the CARA risk aversion of Section 2. Throughout the paper, time-series regressions use Newey-West (1987) standard errors with six lags, and panel regressions use standard errors double-clustered by firm and year-month unless otherwise noted.

3.1 Data sources

The empirical analysis combines holdings, returns, fundamentals, and event data. CRSP monthly stock files supply returns, prices, shares outstanding, and market capitalization for U.S. common stocks. FactSet security-level monthly holdings data supply the fund-level position data used to construct the fund branch of IXI; FactSet's coverage spans mutual funds and ETFs with AUM data from Morningstar. SEC Form 13F quarterly institutional holdings supply the input for the 13F institutional branch of IXI. Morningstar Direct supplies the benchmark identifier used for initial fund assignment and the monthly aggregate active and passive AUM series used in the aggregate flow regressions. Compustat quarterly fundamentals supply book equity, total assets, and the earnings inputs for the implied-cost-of-capital

construction. IBES supplies earnings per share forecasts and analyst estimate counts used for the ICC construction. Siblis Research supplies the true S&P 500 announcement dates used in the event-study design of Section 7. Factor returns are from the Kenneth French Data Library. The return universe is U.S. common stocks on NYSE, AMEX, and NASDAQ; standard CRSP delisting adjustments apply throughout.

3.2 IXI construction

IXI is a stock-month measure of benchmarked ownership. It captures how strongly a stock is embedded in the system of benchmark-driven capital, combining mutual fund and ETF holdings with institutional holdings outside the fund sector. The current construction is a revised version of the earlier fund-based IXI studied in Behmaram (2026), which I refer to as the *legacy IXI* where the comparison matters. Three improvements distinguish the current construction from the legacy version.

The leading innovation is a best-fit benchmark assignment. The legacy IXI attributed each fund's capital through its *declared* prospectus benchmark, typically the Morningstar primary prospectus benchmark identifier, and computed the fund's Active Share (Cremers and Petajisto, 2009) relative to that declared benchmark. That convention is vulnerable to benchmark hacking. A fund that actually tracks the S&P 500 Growth index, for example, can declare the broader S&P 500 as its prospectus benchmark; the fund then shows a higher Active Share relative to its declared benchmark and its capital is still attributed through the broader benchmark's weights. The current construction replaces declared benchmarks with a non-negative least squares (NNLS) projection of each entity's observed holdings onto the full menu of active benchmarks in the quarter, and uses the resulting best-fit benchmarks to attribute benchmarked capital across stocks. NNLS imposes non-negativity on the benchmark loadings but does not impose a unit-sum constraint, so an entity whose holdings only partially match any benchmark receives a loading vector that sums below one, and an entity with clean multi-benchmark alignment receives a mixture. The unconstrained sum

is rescaled to the entity's Active-Share-adjusted attributable mass downstream rather than enforced in the NNLS step. The same best-fit assignment applies to 13F institutions in the non-fund branch below, where no declared prospectus benchmark exists in the first place.

The second improvement is broader ownership coverage. The legacy measure was fund-sector-only and did not cover passively configured capital outside mutual funds and ETFs. The current construction adds a non-fund branch built from quarterly 13F institutional filings that applies the same best-fit assignment to institutional portfolios. The non-fund branch picks up benchmark-oriented capital in pensions, insurance companies, sovereign wealth pools, and similar large institutional accounts that the legacy fund-sector measure omitted.

The third improvement is held-stock normalization in the non-fund branch. Each 13F institution's implied benchmarked dollars are distributed only across the stocks the institution actually holds in the quarter, using the best-fit benchmark's weights renormalized within that held-stock set. No benchmarked dollars are attributed to stocks the institution does not hold. This preserves per-institution benchmarked mass inside the stock universe and eliminates phantom attribution to benchmark members that sit outside the institution's actual portfolio.

The *fund branch* carries the revised logic in the fund sector. For each stock n and month t, the fund branch aggregates, across fund-benchmark pairs identified by the best-fit assignment, each fund's AUM share at the benchmark multiplied by one minus the fund's Active Share relative to that benchmark, further multiplied by the stock's weight in the benchmark at month t. A fund that tracks a single benchmark closely contributes close to its full AUM to that benchmark's pool. A fund with a clean NNLS fit on a style subset contributes to that subset's weights rather than to a broader declared benchmark.

The *non-fund branch* extends the same logic to institutional ownership outside the mutual fund sector. It uses quarterly 13F filings for institutions that pass a broad eligibility screen and a persistence filter. For each institution e in each quarter q, NNLS projects the

institution's holdings onto the benchmark menu to recover the institution's best-fit benchmark loadings, and the institution's implied benchmarked dollars are distributed over the stocks in its actual held-stock set using the benchmark-implied portfolio shape.

Formally, for stock i at month t,

$$\mathrm{IXI}_{i,t} = \min\left(\frac{\Phi^F_{i,t} + \Phi^N_{i,t}}{ME_{i,t}},\ 1\right), \tag{8}$$

where $\Phi^F_{i,t}$ is the fund component, $\Phi^N_{i,t}$ is the 13F institutional component, and $ME_{i,t}$ is the monthly market capitalization denominator. The fund component $\Phi^F_{i,t}$ aggregates, for each fund-benchmark pair identified by the NNLS assignment, the fund's AUM scaled by one minus the fund's Active Share relative to its best-fit benchmark, weighted by the benchmark's weight on stock i at month t. The 13F institutional component is

$$\Phi^N_{i,t} = \sum_{e \in \mathcal{E}} A_{e,q^*(t)} \cdot \beta_{e,q^*(t),t,i}, \tag{9}$$

where $\mathcal{E}$ is the set of qualifying 13F institutional entities, $A_{e,q}$ is entity e's AUM at quarter q, and $q^*(t)$ is the most recently observable quarter at month t under the information-date rule. The benchmark-allocation weight is

$$\beta_{e,q,t,i} = \frac{w^{\mathrm{NNLS}}_{e,q,t,i} \cdot \mathbf{1}\{i \in H_{e,q}\}}{\sum_{j \in H_{e,q}} w^{\mathrm{NNLS}}_{e,q,t,j}} \cdot \bar{\alpha}_{e,q}, \tag{10}$$

where $H_{e,q}$ is the set of stocks actually held by entity e at quarter-end q, $w^{\mathrm{NNLS}}_{e,q,t,i}$ is the weight of stock i under entity e's NNLS best-fit benchmark mixture at month t, and $\bar{\alpha}_{e,q}$ is the entity's passive-intensity scalar, equal to the sum of its NNLS benchmark loadings across all benchmarks in the quarter. The scalar $\bar{\alpha}_{e,q}$ measures how much of entity e's portfolio is explained by benchmark exposures rather than active stock selection; an entity whose holdings mirror one benchmark exactly contributes $\bar{\alpha}_{e,q} = 1$. By construction, $\sum_{i \in H_{e,q}} \beta_{e,q,t,i} = \bar{\alpha}_{e,q}$, so entity e's total 13F-based passive-like capital attribution sums to $A_{e,q} \cdot \bar{\alpha}_{e,q}$ and is dis-

tributed only across stocks the entity actually holds. The final cap at one is a definitional backstop and binds in less than 0.1 percent of stock-months over the sample.

IXI is saved under two timing conventions. Rule-I uses only 13F filings publicly available by month t (due forty-five days after quarter-end); Rule-P maps month t to the most recent quarter-end at or before t. All forward-looking return tests in the paper use Rule-I by default under a conservative two-month lag, so the main return sample begins in August 2002. The canonical panel truncates at May 2024 under Rule-I and at March 2024 for any Rule-I versus Rule-P comparison. The two conventions produce mean monthly cross-sectional Spearman correlations of about 0.998 across the sample (Table 1), so the choice is conventional.

The pipeline inputs are reproducible from publicly described primitives. The benchmark menu covers the major U.S. equity index families (S&P, Russell, CRSP total-market, MSCI, and large style and sector indices) at quarterly frequency. The 13F universe includes institutions with at least four consecutive filings and a portfolio value above the de minimis Form 13F filing threshold, screening out vehicles whose 13F coverage is intermittent. Each entity's AUM is the sum of reported 13F-eligible long positions in U.S. equities at quarter-end. Rule-I times each month t to the most recent quarter publicly available under the forty-five-day filing deadline, with an additional one-month buffer that pushes return tests to a two-month lag. Appendix A documents each step.

Two sources of measurement error are intrinsic to the construction. The NNLS benchmark assignment replaces declared benchmarks with the best-fitting mixture at the entity-quarter level, which recovers the true benchmark allocation only up to sampling noise in holdings. The non-fund 13F branch imputes passive-equivalent exposure from portfolio shape under the held-stock normalization, which is a Sharpe-style style-analysis assumption rather than a direct observation of tracking intent. Both sources introduce error in ownership levels, but the empirical tests rely mainly on stable rankings and within-firm variation. The construction checks in Behmaram (2026) show high rank stability across NNLS-assignment variants (Spearman 0.979), across the non-fund enrichment (Spearman 0.993), and across

the Rule-I and Rule-P timing conventions used here (Table 1), suggesting that the paper identifies a stable indexed-ownership relationship rather than exact ownership-level effects.

IXI is not a measure of explicit index-fund ownership alone. The fund branch captures the passive portion of all benchmarked mutual funds and ETFs, including actively managed funds that hold close to benchmark weights. The non-fund branch captures institutional capital whose portfolio shape is consistent with a benchmarked allocation, even when the institution is not publicly labeled as an index fund. Together, the two branches characterize how strongly a stock sits inside the system of benchmark-driven ownership. Even when IXI is correlated with market capitalization, as it is for any stock-level benchmark-ownership measure in the modern U.S. market, the measure still quantifies an economically distinct concept: the fraction of a stock's market cap held by investors whose capital allocation is structurally tied to benchmarks rather than to independent active decisions.

3.3 Summary statistics

Figure 1 summarizes the cross-sectional monthly mean of IXI under the information-date rule from January 2000 through the validated Rule-I safe end of May 2024. The figure plots the 12-month trailing average of the full revised IXI alongside the fund-branch-only series, which is the closest analog within the current paper to the fund-only legacy IXI studied in Behmaram (2026). The full measure rises from roughly 4 percent in the early 2000s to over 14 percent by the validated sample end, with the trajectory accelerating sharply after 2013. The fund branch carries most of the level in every year, reflecting the continued role of mutual fund and ETF holdings as the dominant visible passive component. The revised measure extends the legacy fund-only coverage by adding benchmarked capital held outside the mutual fund and ETF sector; Appendix Figure 7 shows the raw monthly branch decomposition, including the 13F institutional branch drawn as a step function to reflect its quarterly filing cadence. At the aggregate level the Rule-P line is visually very close to Rule-I over the shared validated sample, consistent with the near-unit rank correlation reported in

Table 1.

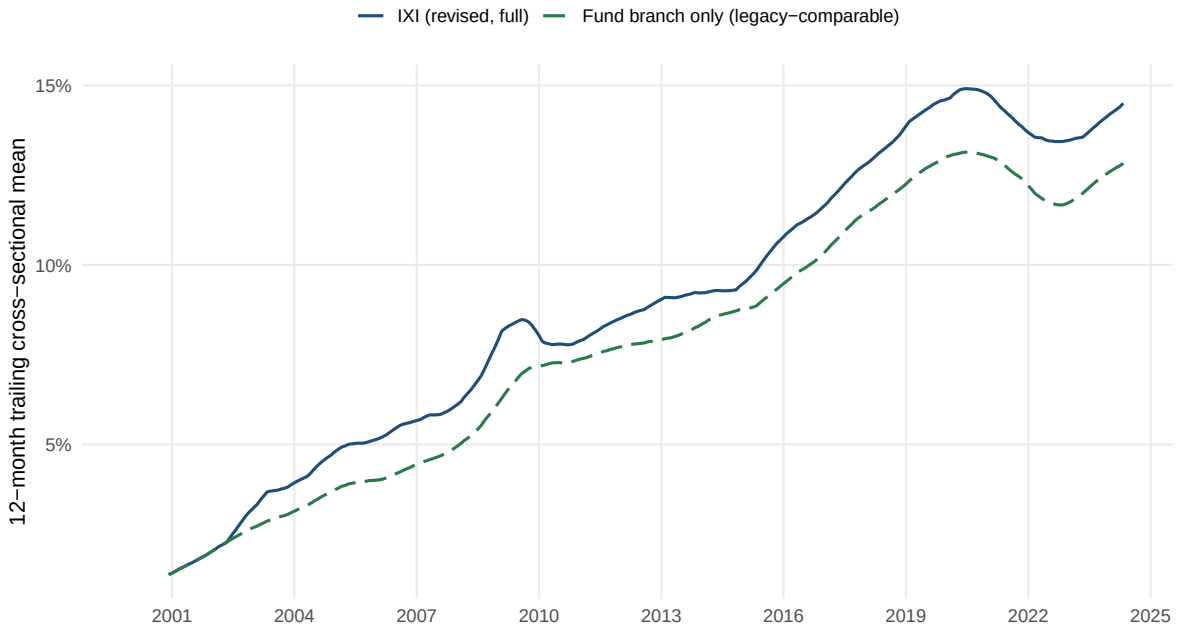

Figure 1: 12-month trailing cross-sectional mean of IXI under the information-date rule, January 2000 through May 2024. The y-axis is IXI as a share of market capitalization. The solid blue line is the full revised IXI; the dashed green line is the fund-branch-only series, which is the closest analog within the current paper to the fund-only legacy IXI studied in Behmaram (2026). The trailing-average presentation emphasizes the secular rise and smooths out the quarter-to-quarter stepping that the 13F branch introduces in the raw monthly series. Appendix Figure 7 plots the raw monthly decomposition, where the 13F institutional branch is drawn as a step function to reflect the quarterly 13F filing cadence.

Table 1 summarizes the cross-sectional distribution of IXI by subperiod, and decomposes the pre-cap numerator into fund and non-fund components. Panel A shows that the mean and median of IXI roughly triple across the sample, with the standard deviation rising from 0.057 in 2002-2006 to 0.156 in 2019-2024-05 and the 99th percentile rising from 0.268 to 0.646. The growth is therefore not confined to a level shift; it reflects an expansion of the right tail and increasing concentration of benchmarked ownership in a smaller set of widely held names. Cap-hit stock-months remain rare throughout, reaching 0.112 percent in the late subperiod and staying at 0.047 percent over the full validated sample, so the IXI ≤ 1 backstop is an occasional definitional constraint rather than a routine operating bound.

Panel B reports the decomposition of the pre-cap numerator into fund and non-fund components under two weighting schemes. Equal-weighted across stock-months, the fund

Table 1: Properties of IXI. Panel A reports descriptive statistics for the Rule-I full measure across subperiods through the Rule-I safe end (2024-05). Panel B reports the fund-branch and non-fund-branch shares of the pre-cap numerator and the mean monthly cross-sectional Spearman correlation between the Rule-I and Rule-P full measures through the shared safe end (2024-03).

Panel A: Descriptive statistics on IXI (Rule-I)

Period	Stock-months	Permnos	Mean	Median	SD	p90	p99	Max	Cap-hit (%)
2002-2006	300,327	7,257	0.0450	0.0271	0.0568	0.1068	0.2678	1.0000	0.0107
2007-2012	352,872	7,529	0.0786	0.0591	0.0831	0.1773	0.3500	1.0000	0.0215
2013-2018	370,525	7,579	0.1125	0.0829	0.1172	0.2671	0.4840	1.0000	0.0343
2019-2024-05	375,023	8,729	0.1415	0.0819	0.1556	0.3482	0.6463	1.0000	0.1123
Full sample to 2024-05	1,398,747	15,192	0.0972	0.0533	0.1176	0.2594	0.5064	1.0000	0.0469

Panel B: Fund and non-fund decomposition; Rule-I vs Rule-P correlation

Period	Equal-weighted share (%)		Value-weighted share (%)		Spearman
	Fund	Non-fund	Fund	Non-fund	Rule-I vs Rule-P
2002-2006	94.33	5.67	75.63	24.37	0.9972
2007-2012	95.60	4.40	83.47	16.53	0.9975
2013-2018	96.47	3.53	82.82	17.18	0.9986
2019-2024-03	96.51	3.49	77.22	22.78	0.9989
Full sample to 2024-03	95.80	4.20	79.36	20.64	0.9981

Notes: Panel A: IXI is the capped passive-ownership ratio at the permno-month under Rule-I (information-date timing). Cap-hit (%) is the share of stock-months at which the final IXI ≤ 1 backstop binds. Panel B: The equal-weighted share is the within-subperiod average of the stock-month ratio of the fund (respectively, non-fund) branch to the pre-cap numerator, restricted to stock-months with a positive pre-cap numerator. The value-weighted share is the aggregate-dollar analog across the same stock-months. The Spearman column is the mean monthly cross-sectional Spearman rank correlation between the Rule-I and Rule-P full IXI within the subperiod.

branch accounts for 94 to 97 percent of the numerator in every subperiod. In aggregate passive-dollar terms, however, the non-fund branch is materially larger, contributing 17 to 24 percent of the total numerator in every subperiod. The gap between the equal-weighted and value-weighted shares is itself informative: the non-fund branch is economically concentrated in a comparatively small set of widely held large-cap names rather than spread uniformly across the cross-section. This pattern is what one would expect if the non-fund branch primarily captures benchmark-driven institutional allocations that concentrate on benchmark-dominant stocks. Panel B also reports that the Rule-I and Rule-P timing conventions produce near-identical rank orderings, with the mean monthly cross-sectional Spearman correlation between the two versions averaging about 0.998 across the sample. Rule-I is therefore used as the default timing convention for forward-looking tests on conceptual grounds; the choice between Rule-I and Rule-P is not an empirical distinction.

3.4 Comparison with Benchmarking Intensity

Benchmarking Intensity (BMI) in Pavlova and Sikorskaya (2023) is the closest incentive-based predecessor to IXI. BMI is built from index-composition weights and fund AUM benchmarked to each index, so it captures benchmark-contract exposure without requiring holdings data. IXI instead measures realized indexed ownership from observed fund and 13F holdings. The two measures are related but capture different dimensions of benchmarked ownership, with Pearson and Spearman correlations of 0.53 and 0.62 on the matched ICC panel.

The conceptual distinction translates into a large empirical gap, because two features of IXI's construction screen out exposure that is nominal rather than realized. First, active managers often declare benchmarks that do not match their holdings: in 65 percent of active fund-date observations the benchmark that best fits observed holdings differs from the declared prospectus benchmark, and the divergence holds at 47 percent of active fund AUM with a 9.1 percentage-point AUM-weighted Active Share gap. Second, not every dollar in a benchmarked fund behaves like a passive dollar. An active manager's discretionary positions do not move with index weights, so counting the fund's full AUM at benchmark weights, as BMI does, attributes passive behavior to capital that does not exhibit it. The Active Share adjustment in IXI retains only the index-tracking portion. Table 2 reports the resulting gap. On the matched ICC/BMI sample, BMI averages 18.5 percent of market capitalization against 12.9 percent for IXI, so BMI is about 43 percent higher than IXI on the matched panel. The divergence is most pronounced early on: the BMI/IXI ratio is 1.69 in 2002-2012 and narrows to 1.17 by 2013-2018, consistent with the gap between benchmark-contract exposure and realized indexed ownership narrowing as the passive transition matures.

This gap is not a defect in BMI. It reflects that the two measures answer different questions: BMI is closer to benchmark-contract exposure, whereas IXI is constructed for realized ownership, flow, and return tests.

The ICC horse race confirms that the measures are complementary. In separate regres-

Table 2: Declared benchmarks, best-fit benchmarks, and realized indexed ownership.

Panel A: Fund-side construction diagnostics

	Fund-dates	% diff. canonical	Median stated−best-fit AS (pp)	% diff. canonical, AUM-wt.	Stated−best-fit AS (pp), AUM-wt.
1999–2024	741,590	65.2	11.6	47.4	9.1
1999–2012	215,874	53.3	3.2	36.8	6.6
2013–2024	525,716	70.1	13.1	50.7	9.9

Panel B: Stock-side matched-sample diagnostics

	Obs	Mean IXI (%)	Mean BMI (%)	BMI/IXI	Pearson ρ	Spearman ρ
2002–2018	31,523	12.9	18.5	1.43	0.53	0.62
2002–2012	19,705	10.5	17.7	1.69	0.47	0.56
2013–2018	11,818	17.0	19.9	1.17	0.58	0.64

Notes: Panel A uses the IXI fund construction panel of active fund-dates, 1999-03 through 2024-12. Stated−best-fit Active Share is the Active Share measured against the fund's declared prospectus benchmark minus the Active Share measured against the benchmark that best fits observed holdings, in percentage points. AUM-weighted columns weight each fund-date by the fund's total AUM at that date. Panel B uses the matched ICC/BMI June stock-year panel, with BMI from the public release in Pavlova and Sikorskaya (2023). Mean IXI and mean BMI are expressed as a percentage of market capitalization. BMI/IXI is the ratio of sample means.

sions, both IXI and BMI predict lower implied costs of capital. In the joint specification with firm and year fixed effects and the paper's standard controls, IXI remains negative at -0.0293 and BMI remains negative at -0.0354, both retaining individual significance. IXI is the organizing measure for this paper because the monthly return, flow, event-time, and decomposition designs require observed holdings, Active-Share-adjusted attribution, and the fund/non-fund branch structure.

4 Realized returns

The realized-return evidence is the rewarding side of the same mechanism that lowers expected returns. The most heavily indexed U.S. equities outperformed the least heavily indexed throughout the passive era, the empirical counterpart of Proposition 2's transition reading. Over a window of positive unanticipated shocks to $\mathcal{I}_n$, high-$\mathcal{I}$ stocks can earn positive realized returns even as Proposition 1's static expected-return wedge remains negative.

4.1 Value-weighted IXI quintile spread

Sorting stocks each month on lagged IXI and forming value-weighted quintile portfolios produces a sizeable spread between the highest-IXI and lowest-IXI portfolios. I label this top-minus-bottom IXI quintile spread HILI (high-indexing minus low-indexing) and use this label throughout the paper. Figure 2 plots the cumulative performance of Q5, Q1, and the market from August 2002 through May 2024. The headline configuration sorts on full IXI under Rule-I with a two-month lag, uses within-month empirical quintile breakpoints, and value-weights by lagged market capitalization.

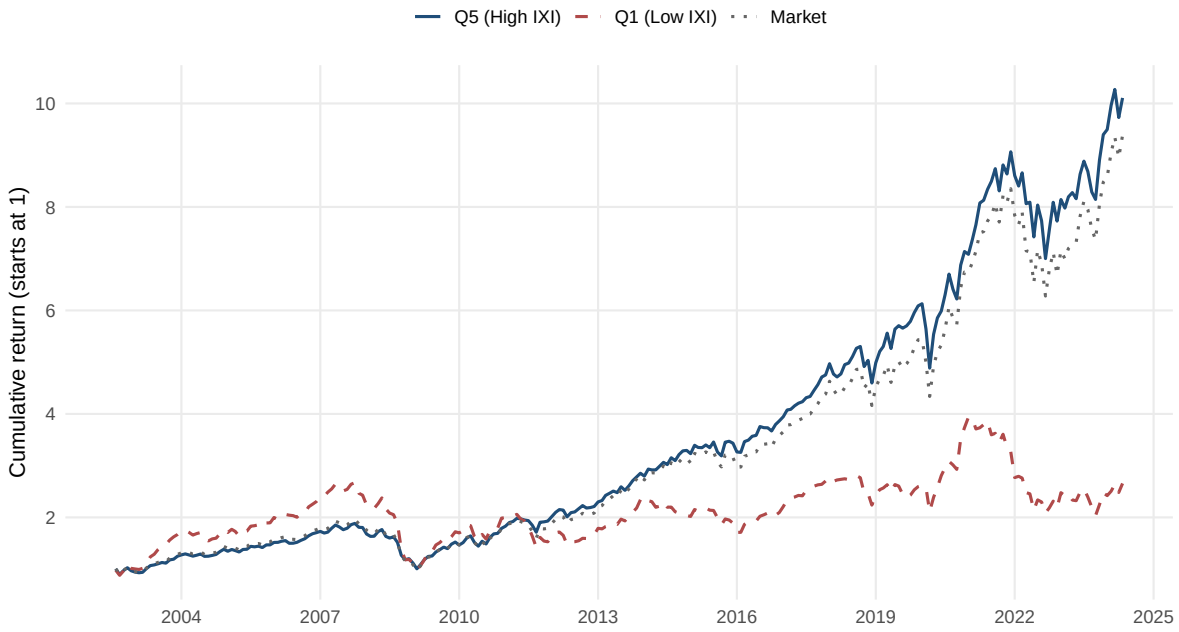

Figure 2: Cumulative value-weighted returns of Q5 (high-IXI), Q1 (low-IXI), and the market portfolio, August 2002 through May 2024. Each series starts at 100. Portfolios are formed monthly under the main specification (full IXI, Rule-I, two-month lag, within-month empirical quintile breakpoints).

In the post-2013 passive era, the HILI spread delivers a FF5-plus-momentum alpha of 0.53 percentage points per month. The full-sample alpha of 0.37 averages the post-2013 period with a quieter pre-2013 period (alpha 0.10). Aggregate passive AUM grew from roughly 17 to 61 percent of equity-fund assets over 2002-2024, with the sharpest acceleration after 2013 (Section 6); Appendix IA.N reports the period breakdown.

HILI has a mean return of 0.43 percentage points per month (annualized Sharpe ratio

Table 3: HILI monthly return on factor models. Main specification: IXI under Rule-I with a two-month lag, within-month empirical quintile breakpoints, value-weighted, August 2002 through May 2024 (262 months). Standard errors in parentheses are Newey-West standard errors with six lags.

	Dependent variable: HILI (pp per month)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Constant	0.433*	0.627***	0.600***	0.405**	0.382**	0.585**	0.365**
	(0.247)	(0.224)	(0.199)	(0.165)	(0.164)	(0.230)	(0.163)
MKT−Rf		−0.230***	−0.145***	−0.106***	−0.097***	−0.182***	−0.068*
		(0.042)	(0.045)	(0.035)	(0.034)	(0.047)	(0.036)
SMB			−0.492***	−0.332***	−0.320***		−0.318***
			(0.097)	(0.080)	(0.078)		(0.078)
HML			0.172***	0.115**	0.060		0.094
			(0.065)	(0.057)	(0.063)		(0.062)
RMW				0.535***	0.557***		0.533***
				(0.069)	(0.068)		(0.076)
CMA					0.146		0.131
					(0.104)		(0.106)
MOM						0.115***	0.085**
						(0.044)	(0.033)
Observations	262	262	262	262	262	262	262
R^2	0.000	0.126	0.289	0.410	0.416	0.154	0.430

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. HILI is the value-weighted top-minus-bottom IXI quintile spread. MKT-Rf, SMB, HML, RMW, CMA from Fama-French (2015); MOM from Carhart (1997). All from the Kenneth French Data Library.

0.52) and earns positive alphas against CAPM, FF3, FF5, and FF5-plus-momentum (Table 3). The factor loadings are large-cap and profitability-tilted, consistent with HILI being a value-weighted long-short on widely held benchmark-dominant stocks. The low-IXI quintile earns negative CAPM and FF3 alphas of -0.44 pp/month, so both tails contribute to the spread.

4.2 Cross-sectional evidence: Fama-MacBeth regressions

A Fama-MacBeth cross-sectional regression of individual stock returns on lagged IXI provides complementary evidence that does not rely on portfolio construction or value-weighting. The unconditional IXI slope is negative because IXI and size are highly correlated; conditioning on log(me) flips the slope positive (0.013 with momentum and book-to-market controls), so heavily indexed stocks outperform once size is held fixed. Appendix A.3 (Table 10) reports the full grid.

4.3 Size-tertile robustness

The value-weighted HILI sort places most of its weight on the very largest stocks: the Big NYSE tertile, which contains the Magnificent Seven, accounts for 91 percent of total market capitalization in 2023. Splitting the universe into three NYSE size tertiles and re-running the IXI sort within each, the mid-tertile, where IXI varies cleanly and no single firm dominates the cap weighting, delivers an FF5-plus-momentum alpha of 0.28 percentage points per month. Mid-tertile membership is determined by the NYSE size distribution each month, so the result is not selected on any specific firm's performance.

The Big and Small tertiles are less informative on their own. Standard factor models absorb the Big-tertile alpha, and the Small tertile is too economically noisy for a clean within-tertile read.[1] Appendix IA.O reports the full within-tertile HILI for all three groups.

[1] A weight-capping exercise (1 to 10 percent maximum weight per stock within each quintile) leaves the CAPM alpha at 0.47-0.55 percentage points per month but attenuates the FF5-plus-momentum alpha. The attenuation reflects weight redistribution among large-cap stocks with similar factor exposures rather than

A coarser binary big/small sort at the NYSE median yields weaker alphas within each bucket (Big 0.16, Small 0.31) because it pools mid-cap and mega-cap stocks together; the mid-cap IXI signal is diluted by the mega-cap segment that the factor models span. Appendix IA.P reports the binary sort.

4.4 Between-firm versus within-firm decomposition

The HILI sort blends between-firm and within-firm variation in IXI, and Appendix A.4 (Table 11) shows that essentially all of the realized HILI premium lives on the between-firm dimension; the within-only HILI is a null. Within-firm return predictability is recoverable once size is conditioned on, with positive slopes in the big and middle NYSE tertiles and a noisy small-stock segment. The two-margin nature of the ICC wedge is taken up in Section 5.

Realized HILI is a joint size-and-indexing dimension rather than a size-independent IXI anomaly: equal-weighted HILI is near zero, the binary NYSE-median double sort is weak within buckets, and the premium is concentrated after 2013. Appendix IA.D reports the full sensitivity grid.

5 Indexing premium: expected returns and the implied cost of capital

The implied cost of capital measures the expected-return cost of heavy indexed ownership. Section 2's framework predicts that if index-linked capital absorbs part of a stock's residual supply, high-IXI stocks carry lower ex-ante expected returns than low-IXI stocks. The question here is whether high-IXI stocks carry lower expected returns after conditioning on size, value, and momentum, separate from the realized-return pattern of Section 4.

I answer this question with an implied cost of capital (ICC) measure constructed at the firm-month level using the Hou et al. (2012) methodology (HVZ), aligned to the IXI panel;

a change in portfolio composition.

Appendix B.1 gives the construction details. The main-sample ICC panel runs from August 2002 through May 2024, with mean ICC of 9.3 percent per year.[2] The evidence comes from three pieces: a portfolio-level ICC series under the HILI sort, a size-conditional ICC check that mirrors the double sort on the realized-return side, and a firm-level panel regression with firm and month fixed effects and size, value, and momentum controls.

The ICC evidence is a discount-rate test of Proposition 1, distinct from the announcement-informativeness mechanism studied in a parallel literature.[3]

5.1 Portfolio ICC under the HILI sort

Each month stocks are sorted into value-weighted quintiles on lagged IXI using the same HILI sort parameters as Section 4. Within each quintile and month I compute the value-weighted ICC using only stock-months with non-missing ICC, and then take the time-series mean over the 262-month sample.

Coverage of the HILI universe varies by quintile, rising from roughly 42 percent in Q1 to 72 percent in Q5. The skew reflects high-IXI stocks being widely held large-caps with complete Compustat inputs, whereas low-IXI stocks are disproportionately small-caps with missing accruals or balance-sheet items. A Heckman two-stage selection correction (Heckman, 1979) attenuates the pooled-OLS IXI coefficient but leaves the ICC-IXI relationship negative; the wedge is also robust to two alternative ICC constructions (the residual-income model of Gebhardt et al. 2001 and the PEG approach of Easton 2004). Appendix IA.H reports the Heckman correction and Appendix IA.I reports the alternative ICC constructions.

Average value-weighted ICC declines monotonically from 7.93 percent per year in Q1 to

[2]The ICC panel has 704,565 valid firm-months, 8,072 unique permnos, and 262 monthly cross-sections, with a solver-failure rate of 3.9 percent. Table 4 regressions require non-missing lagged log(ME), lagged log(B/M), and twelve-to-two-month momentum, reducing the estimation sample to approximately 657,000 firm-months in the firm-plus-month fixed-effects specifications and 575,617 firm-months with the full control set.

[3]Sammon (2025), Israeli et al. (2017), and Bennett et al. (2020) document an information-production decline under rising passive ownership; Bond and Garcia (2022) provide an equilibrium-theory companion. A direct Sammon-style test of the announcement-informativeness channel on IXI is inconclusive (Appendix IA.A), so the ICC evidence below is read as a discount-rate test rather than as identification of the informativeness channel.

5.86 percent in Q5. The HILI-ICC spread (Q5 minus Q1) averages −2.07 percentage points per year, with $t = -8.82$. The widely held top-IXI quintile carries expected returns roughly two percentage points per year below the narrowly held bottom-IXI quintile, consistent with the indexing-premium hypothesis that benchmark-driven demand compresses expected returns on widely held stocks (Basak and Pavlova, 2013; Kashyap et al., 2023). Figure 3 plots the spread.

Figure 3: HILI-ICC spread (Q5 minus Q1) of value-weighted ICC under the HILI sort, August 2002 through May 2024. The y-axis is the spread in annualized percentage points; a negative value indicates that high-IXI stocks carry a lower implied cost of capital than low-IXI stocks. The time-series mean is −2.07 percentage points per year.

5.2 Firm-level panel evidence

Table 4 reports a firm-level panel regression of ICC on lagged IXI with escalating specifications. Standard errors are double-clustered by firm and year-month.

The lagged-IXI coefficient attenuates when size and value controls are added but remains negative (−0.026 in the fully controlled specification, −0.111 with firm-plus-month fixed effects). Most of the raw ICC-IXI relation is shared with size, but the residual effect after

Table 4: Panel regression of firm-month ICC on lagged IXI. ICC is constructed using the HVZ methodology. Sample: August 2002 through May 2024 under the Rule-I two-month lag. Column (1) is pooled OLS; columns (2)–(4) add firm and month fixed effects and controls. Standard errors in parentheses, double-clustered by firm and year-month.

	Dependent variable: ICC (decimal)			
	(1)	(2)	(3)	(4)
Lagged IXI (2m)	-0.1835^{***}	-0.1112^{***}	-0.0266^{***}	-0.0259^{***}
	(0.0095)	(0.0072)	(0.0050)	(0.0048)
Lag $\log(me)$			-0.0292^{***}	-0.0283^{***}
			(0.0011)	(0.0011)
Lag $\log(B/M)$			0.0066^{***}	0.0043^{***}
			(0.0008)	(0.0009)
Momentum (t-12 to t-2)				-0.0098^{***}
				(0.0006)
Firm FE		Y	Y	Y
Month FE		Y	Y	Y
Observations	657,615	657,557	657,555	575,617
Within R^2		0.0239	0.2592	0.2666

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Column (1) is pooled OLS with no fixed effects. Column (2) adds firm and month fixed effects. Column (3) adds lagged log market cap and log book-to-market. Column (4) adds the 12-to-2 month cumulative excess return. ICC is in decimal units (e.g., 0.07 denotes 7 percent per year).

size, value, and momentum controls is economically meaningful: a within-firm increase of 0.20 in lagged IXI, roughly the 50th-to-90th-percentile gap, maps to a 52 basis-point reduction in ICC. A pooled fully interacted specification confirms that the implied post-2013 slope is also negative, with no era-specific change.

A Mundlak decomposition (Table 11, Panel B) separates the between-firm and within-firm IXI components in a single regression. Both are negative, so the ICC wedge is a two-margin phenomenon: stocks that sit persistently deep in the benchmarked-ownership system are priced for lower expected returns, and stocks whose IXI rises over time see a parallel compression. The within-firm margin links to the Section 6 flow mechanism and the Section 8 demand-gap design; the between-firm margin links to the Section 4 HILI sort.

The ICC wedge is present in both Big and Small size buckets under a NYSE-median double sort: -1.04 percentage points within Big and -2.87 within Small, with $|t| > 9$ in both buckets. This stands in direct contrast to the realized-return side, where the NYSE-

median double sort produces HILI alphas below conventional significance within each bucket; the realized premium is a joint size-and-indexing stylized fact, while the expected-return wedge is an IXI effect that survives size conditioning cleanly. Appendix tables show that the wedge is stable across size conditioning, FF12-industry adjustment, technology exclusion, the BMI matched-sample horse race, alternative ICC constructions, and Heckman selection correction; it is negative and significant in every cut.

Realized returns capture a size-and-indexing composite that standard factor models only partially span and that strengthens after 2013; the expected-return wedge is a cleaner IXI effect that survives size, value, momentum, industry adjustment, and tech exclusion.

6 Stock-level passive flow and the price-pressure mechanism

Passive flows matter for the transition story only if they show up in prices and in the aggregate HILI path. I test both margins. At the stock-month level, the question is whether holdings-based passive-fund purchases generate short-run price pressure and reversal, as in the mutual-fund-flow evidence of Coval and Stafford (2007), Lou (2012), and Frazzini and Lamont (2008) and the ETF mechanics studied by Ben-David et al. (2018). At the aggregate level, the question is whether growth in the passive-ownership share accounts for the unexpected component of HILI, in the spirit of Koijen and Yogo (2019), Pavlova and Sikorskaya (2023), and Chinco and Sammon (2024). The first-difference aggregate decomposition is the paper's quantitative synthesis. Subtracting the estimated passive-share-growth contribution from the realized cumulative HILI path turns a +171-percent gain over 2002-09 through 2024-05 into a −35-percent loss.

6.1 Holdings-based passive flow measure

The main explanatory variable in this section is a stock-month passive-flow proxy built directly from changes in passive-fund holdings. For permno i at month t,

$$\text{Flow}_{i,t}^{\text{passive}} = \frac{\sum_{j \in J^{\text{passive}}} \left(\text{shares}_{i,j,t} - \text{shares}_{i,j,t-1}\right) \cdot p_{i,t-1}}{ME_{i,t-1}}, \tag{11}$$

where $\text{shares}_{i,j,t}$ is the number of shares of stock i held by passive fund j at the end of month t, $p_{i,t-1}$ is stock i's adjusted price at the end of the prior month, and $ME_{i,t-1}$ is stock i's lagged market capitalization. The passive-fund set J^{passive} is defined as the 20,292 FactSet funds whose FactSet style is tagged as Index, and covers both index ETFs and index mutual funds. Share counts come from the monthly-filled FactSet holdings panel. Fund entries into a stock are handled by treating the prior-period share count as zero; fund exits are reconstructed in a symmetric robustness variant reported alongside the main measure.

This is a holdings-based flow measure, not a change in the indexing inclusion ratio. The numerator is the dollar value of the share change across passive funds valued at the prior-month price. Lagged-price valuation isolates the share-change component from the contemporaneous price movement, so a month with large market moves does not mechanically produce a large aggregate flow. The ratio is therefore a dimensionless share of lagged market capitalization that represents the dollar change in passive-fund holdings of stock i during month t; the construction parallels the flow-induced-trading measure in Lou (2012) but applies it to the passive subset of the fund universe and uses monthly holdings filled from Ben-David et al. (2018)'s evidence base. Aggregated across stocks and months, the value-weighted cross-sectional mean of $\text{Flow}^{\text{passive}}$ has a Pearson correlation of -0.002 with the contemporaneous value-weighted market return over the paper sample, confirming that the proxy is not mechanically tied to the market direction.

6.2 Contemporaneous panel regression

For each permno i and month t in the paper sample (August 2002 through May 2024), I estimate

$$\begin{aligned}\text{Ret}_{i,t} &= \beta_0 \text{Flow}^s_{i,t} + \beta_1 \text{Flow}^s_{i,t-1} + \beta_2 \text{IXI}^s_{i,t-1} + \beta_3 \big(\text{IXI}^s_{i,t-1} \cdot \text{Flow}^s_{i,t}\big) \\ &\quad + \gamma' X^s_{i,t} + \alpha_i + \lambda_t + \varepsilon_{i,t}, \end{aligned} \tag{12}$$

where superscript s denotes cross-sectional standardization within each calendar month (subtract month-t mean, divide by month-t standard deviation), $\text{IXI}^s_{i,t-1}$ uses the two-month-lag paper convention, and the control vector $X^s_{i,t}$ contains the earnings-release return (ER, the sum of three-day market-adjusted returns around all IBES earnings announcements falling in month t, in the spirit of the post-earnings-announcement-drift literature summarized in Bernard and Thomas 1989 and DellaVigna and Pollet 2009), the twelve-month change in the IBES consensus long-term-growth forecast (ΔEF), log lagged market equity, log lagged book-to-market, and the cumulative return from month $t-12$ to $t-2$. Every right-hand-side regressor except the IXI level is winsorized within month at one and ninety-nine percent before standardization. Fixed effects are firm (α_i) and year-month (λ_t). Standard errors are double-clustered by firm and year-month. Table 5 reports the contemporaneous panel regression (column t) and the horizon decomposition characterized in Section 6.3.

In the main specification, a one-standard-deviation within-month movement in passive flow, roughly 0.5 percent of a stock's lagged market capitalization, is associated with a contemporaneous return of approximately 47 basis points and a next-month reversal of approximately -19 basis points.

The result is robust to three alternative passive-fund definitions. The exit-reconstructed variant, which adds a symmetric negative-flow row when a fund-stock pair stops reporting after having held a positive position, leaves the contemporaneous coefficient essentially unchanged. Restricting the passive-fund set to index ETFs, the structurally tightest passive

Table 5: Decay profile of the holdings-based passive flow effect. Panel A reports single-month-horizon regressions of $\text{ret}_{i,t+k}$ on $\text{flow}_{i,t}$ and the D5 control set; Panel B reports cumulative-horizon regressions with dependent variable the sum of log returns from month $t+1$ through $t+K$. All right-hand-side regressors are winsorized within-month at 1/99 and then cross-sectionally standardized at month t, except the IXI level which is standardized only (D5 convention). Fixed effects are permno and the year-month of the dependent-variable observation. Standard errors double-clustered on (permno, date); t-statistics in parentheses (this table reports t-statistics rather than standard errors because several cells have coefficients that round to zero at four decimal places, for which SEs are not reliably back-computable from the reported coefficient).

Panel A: Single-month-horizon return

	t	t+1	t+2	t+3	t+4	t+5	t+6	t+12
$\text{Flow}_{i,t}$	0.0047^{***}	-0.0021^{***}	-0.0011^{***}	-0.0004	-0.0004	-0.0005	-0.0006^{*}	-0.0007^{**}
	(7.41)	(−5.28)	(−2.60)	(−0.94)	(−1.01)	(−1.23)	(−1.67)	(−2.17)
$\text{Flow}_{i,t-1}$	-0.0019^{***}	-0.0009^{**}	-0.0002	-0.0003	-0.0005	-0.0006^{*}	-0.0006^{*}	-0.0006^{*}
	(−5.72)	(−2.53)	(−0.53)	(−0.74)	(−1.29)	(−1.89)	(−1.72)	(−1.91)
$\text{IXI} \times \text{Flow}_{i,t}$	-0.0011^{***}	0.0004	-0.0003	-0.0000	-0.0000	-0.0002	-0.0003	0.0002
	(−3.06)	(1.62)	(−1.04)	(−0.05)	(−0.02)	(−0.74)	(−1.01)	(0.96)
Observations	798,565	791,533	784,631	777,752	770,909	764,095	757,301	717,524
Permnos	7,104	7,074	7,033	6,989	6,943	6,904	6,869	6,574

Panel B: Cumulative-horizon log return

	t+1 : t+3	t+1 : t+6	t+1 : t+12
$\text{Flow}_{i,t}$	-0.0042^{***}	-0.0058^{***}	-0.0086^{***}
	(−5.73)	(−6.34)	(−6.72)
$\text{Flow}_{i,t-1}$	-0.0019^{***}	-0.0039^{***}	-0.0054^{***}
	(−3.00)	(−5.16)	(−4.70)
$\text{IXI} \times \text{Flow}_{i,t}$	0.0001	-0.0003	0.0001
	(0.16)	(−0.31)	(0.11)
Observations	777,472	756,633	716,193
Permnos	6,987	6,867	6,571

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$ under standard errors double-clustered on (permno, date). All specifications include the full D5 right-hand-side: flow, lagged flow, IXI level (2-month lag), ER, ΔEF, log(me_lag), log(BM), momentum, and IXI $\times$ flow interaction. Only flow, lagged flow, and IXI $\times$ flow are displayed for conciseness. Cumulative-window observations with any missing month in the window are dropped.

trackers because they must settle in-kind to authorized participants daily and therefore admit very limited closet-indexing, attenuates the coefficient by only about 11 percent. The index-ETF subset and the broader 20,292-fund index-labeled set therefore deliver nearly identical contemporaneous coefficients. Restricting further to index mutual funds delivers a smaller but still positive coefficient. All four variants are positive with t-statistics above 5.

The flow response is concentrated in small stocks. The small-stock contemporaneous coefficient is roughly twice the large-stock value, so a dollar of passive flow moves small-stock returns about twice as much as large-stock returns. The IXI $\times$ Flow interaction is a small-stock phenomenon: sharply negative among small stocks and near zero among large stocks. Within small stocks, the price-pressure response declines as a stock's existing IXI rises.

The panel coefficient on the IXI level itself is essentially zero ($\hat{\beta}_2 = -0.00013$). This is consistent with the realized HILI premium living on the between-firm dimension (absorbed by firm fixed effects in this regression) and with the within-firm IXI-return relation being recoverable once the sample is stratified by NYSE size tertile (Table 11). The cross-sectional IXI-level evidence of Sections 4 and 5 remains the paper's headline; this regression identifies a different object, namely the within-stock flow-driven response. The negative lagged-flow coefficient is a single point on the decay profile characterized in the next subsection, where the partial reversal pattern becomes explicit.

Within-month passive-fund holdings changes may be partly driven by same-month returns through mechanical rebalancing, so the contemporaneous coefficient is not a clean identification of a demand shock. The S&P 500 event study (Section 7) and the demand-gap design (Section 8) provide the paper's event-level evidence on benchmarked-ownership shocks.

6.3 Decay of the flow effect

The horizon-specific regressions in Table 5 characterize the time profile of the flow response. Each column regresses the return at $t+k$ on flow and controls measured at t, with year-month fixed effects indexed to the return observation; Panel A reports horizons $k \in \{0, 1, 2, 3, 6, 12\}$ for single-month returns, and Panel B reports cumulative log returns over $t+1$ through $t+K$ for $K \in \{3, 6, 12\}$.

The contemporaneous response substantially reverses within three months. The cumulative $t+1$ through $t+3$ coefficient roughly offsets the $k = 0$ positive response, so the net three-month effect of a flow shock is close to zero. Single-month coefficients at $k = 6$ and $k = 12$ are not significant. The reversal is concentrated in small stocks, mirroring the small-stock concentration of the contemporaneous response: the small-stock $t+1$ reversal is roughly six times the big-stock reversal. The IXI $\times$ Flow interaction is a same-month phenomenon and does not persist at any future horizon.

The combined evidence is consistent with a short-run passive-demand price-pressure mechanism in the tradition of Scholes (1972), Shleifer (1986), Coval and Stafford (2007), and Lou (2012). Passive flow into a stock co-moves with a positive same-month return that substantially reverses over the following one to three months, with the small-stock segment carrying both the largest same-month response and the fastest reversal.

The cumulative post-flow coefficient remains marginally negative at $t + 12$, implying a residual drift of roughly -40 basis points per standard deviation of flow that is small relative to the same-month impact. The panel estimates a within-stock association rather than a between-stock level relationship, so the decay profile complements the cross-sectional evidence of Sections 4 and 5. Short-run price pressure and slow-moving ex-ante compression are the two margins at which passive demand shows up in this paper, and the decay result is the signature of the first.

6.4 Aggregate passive flow and the HILI counterfactual

The stock-level panel documents the within-stock association between returns and passive-flow changes. A natural companion question is whether the time series of HILI returns is also tied to the time series of aggregate passive flow and the aggregate level of passive ownership. I address this using monthly aggregate data from Morningstar Direct on the US-equity universe of open-end and ETF funds (ex money-market, ex fund-of-funds, ex feeder), split into actively managed and passively managed aggregates.

The Morningstar series is internally consistent (ETF and open-end splits sum exactly to the combined total each month, Organic Growth Rate matches implied flow at correlation 1.000, TNA decomposition into flow and market appreciation leaves a median residual under 0.1 percent). The passive share of US-equity fund TNA rises from roughly 17 to 61 percent over the sample, with the post-2013 acceleration documented in Chinco and Sammon (2024) and Sammon (2025) visible.

From the Morningstar series I construct two aggregate regressors at monthly frequency on the paper sample,

$$\text{AggFlow}_t \;=\; 100 \times \frac{\text{passive net flow}_t}{\text{CRSP total market cap}_{t-1}}, \tag{13}$$

$$\text{AggAUM}_t \;=\; 100 \times \frac{\text{passive TNA}_t}{\text{CRSP total market cap}_t}. \tag{14}$$

AggAUM_t is the share of CRSP market capitalization held in passive U.S.-equity mutual funds and ETFs, a fund-side aggregate measure rather than the holdings-level passive ownership share of the stock market. From the same series I construct two HILI-specific aggregates of the stock-level earnings variables from the previous subsections: $\text{ER}_t^{\text{HILI}}$ is the value-weighted month-t Q5 minus Q1 of stock-level earnings-release returns, and $\Delta\text{EF}_t^{\text{HILI}}$ is the corresponding value-weighted Q5 minus Q1 of long-term-forecast changes. The preferred

aggregate specification is the time-series regression

$$\text{HILI}_t = \alpha + \beta_{\text{ER}}\text{ER}_t^{\text{HILI}} + \beta_{\Delta\text{EF}}\Delta\text{EF}_t^{\text{HILI}} + \beta_{\text{flow}}\text{AggFlow}_{t-1} + \beta_{\text{AUM}}\text{AggAUM}_t + \varepsilon_t, \quad (15)$$

where passive flow enters at one-month lag. The lag avoids a within-month reverse-causality concern, because Morningstar's Estimated Net Flow in month t is affected by same-month returns through the mechanical TNA identity and HILI itself is a return series, so the contemporaneous coefficient on AggFlow_t is biased toward zero, the classic concern raised by Warther (1995) at the aggregate level and refined by Lou (2012) at the stock level. Table 6 reports the preferred specification in column (1) alongside five robustness columns.

The primary aggregate result comes from the first-difference specification, which replaces AggAUM_t in Equation (15) with $\Delta\text{AggAUM}_t = \text{AggAUM}_t - \text{AggAUM}_{t-1}$, the month-to-month change in the passive-ownership share. That coefficient is 11.39 with $t = 4.51$. Because AggAUM is measured in percentage points, a 0.01-percentage-point faster monthly increase in the passive-ownership share is associated with an 11.39-basis-point increase in HILI that month. The first-difference specification removes any trend from the passive-ownership series by construction, so the result does not depend on the particular sample having monotone passive-share growth.

Figure 4 translates this coefficient into a statistical cumulative decomposition. The FD-decomposition HILI each month removes only the estimated passive-share-growth contribution:

$$\text{HILI}_t^{\text{cf}} \;=\; \text{HILI}_t - \hat{\beta}_{\Delta\text{AggAUM}}\,\Delta\text{AggAUM}_t, \quad (16)$$

leaving the intercept, earnings news, long-term-forecast changes, and lagged flow contributions intact. Starting both paths at 100 in September 2002 (the first month for which ΔAggAUM is defined) and compounding monthly through May 2024, realized HILI reaches 271 while the FD-decomposition path reaches 65. The estimated passive-share-growth component averages 0.55 percentage points per month; subtracting it from realized HILI reduces

Table 6: HILI on aggregate passive flow and earnings components (Morningstar Direct). Dependent variable is monthly realized HILI return under the main specification (percent per month). ER_t and ΔEF_t are HILI-specific aggregates (value-weighted Q5 minus Q1 of stock-level ER and ΔEF). Passive flow and TNA are Morningstar US-equity active-vs-passive aggregates, scaled by lagged CRSP total market capitalization for flow and contemporaneous CRSP total market capitalization for TNA. Column (1) is the preferred specification with one-month-lagged flow to avoid within-month reverse causality. Column (2) is the contemporaneous reduced form reported for comparison. Column (3) is a 2SLS estimate that instruments contemporaneous flow with its one-month lag; the first-stage partial F on the instrument is 1.28, well below the weak-instrument threshold, so column (3) is reported for transparency only. Columns (4)-(6) substitute ETF-only, OEF-only, and the Morningstar organic growth rate for the combined passive flow, all with the one-month-lagged flow timing. Newey-West standard errors with six lags in parentheses. Sample: August 2002 through May 2024.

	(1) Preferred (RF, flow_{t-1})	(2) RF, flow_t	(3) IV 2SLS	(4) ETF only	(5) OEF only	(6) OGR
ER_t	131.341***	132.133***	105.188**	131.980***	132.565***	131.323***
	(21.737)	(22.180)	(46.178)	(21.837)	(22.111)	(21.983)
ΔEF_t	−0.034	−0.033	−0.074	−0.034	−0.030	−0.033
	(0.021)	(0.022)	(0.055)	(0.022)	(0.021)	(0.021)
Passive flow_{t-1} / CRSP MV (%)	7.269**					
	(3.363)					
Passive flow_t / CRSP MV (%)		2.284				
		(3.358)				
Passive flow_t (IV)			103.410			
			(140.696)			
ETF flow_{t-1} / CRSP MV (%)				7.673**		
				(3.290)		
OEF flow_{t-1} / CRSP MV (%)					−0.553	
					(14.221)	
Passive OGR_{t-1} (%/mo)						0.185
						(0.220)
Passive TNA_t / CRSP MV (%)	0.100*	0.112*	−0.113			0.122**
	(0.054)	(0.057)	(0.321)			(0.056)
ETF TNA_t / CRSP MV (%)				0.169*		
				(0.099)		
OEF TNA_t / CRSP MV (%)					0.246**	
					(0.110)	
Constant	−0.722	−0.632	−2.530	−0.405	−0.811	−0.726
	(0.471)	(0.476)	(2.657)	(0.369)	(0.558)	(0.494)
Observations	262	262	262	262	262	262
R^2	0.189	0.173	–	0.190	0.170	0.174

Notes: $^*p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$ under Newey-West standard errors with six lags. The preferred headline specification is column (1), with one-month-lagged passive flow. Column (2) reports the contemporaneous reduced form; the contemporaneous flow coefficient is not statistically different from zero and is likely biased toward zero by within-month reverse causality of flow on same-month returns. Column (3) reports a 2SLS estimate; the instrument (one-month-lagged flow) is weak (first-stage F = 1.28), so column (3) is not used for paper interpretation. Columns (4), (5), and (6) replace the combined passive-flow variable with its ETF-only, OEF-only, and organic-growth-rate analogs respectively, all timed one month ahead of the return observation.

the cumulative path from a 171 percent gain to a 35 percent loss. The figure shows that the gap between the two paths widens steadily post-2013 as the passive-share growth rate entered its steepest phase. The decomposition is reduced-form rather than causal: because ΔAggAUM_t has a same-month mechanical link to returns through the TNA identity, the gap should be read as statistical co-movement between HILI and passive-share growth rather than as a structural no-passive counterfactual.

The level specification in column (1) of Table 6 delivers a coefficient of 0.100 on AggAUM_t with $t = 1.85$, and at sample means the passive-TNA-share channel dominates the channel-by-channel contribution to HILI. However, the passive-TNA-share series is nearly monotone across the sample, rising from roughly 2.5 to 17 percent of CRSP market capitalization. Adding a linear time trend to the level specification absorbs the AggAUM coefficient: the coefficient flips from $+0.100$ to -0.520 ($t = -1.89$) and the time-trend coefficient enters at 0.033 ($t = 2.23$). The level coefficient and a monotone time trend are not separately identified over this sample. The level-specification results should therefore be read as descriptive rather than structural. The first-difference coefficient of 11.39 is the paper's preferred trend-free aggregate specification, and Figure 4 is built from that coefficient. Columns (4) through (6) of Table 6 report ETF-only passive flow, OEF-only passive flow, and the organic-growth-rate variant, all as robustness. The ETF-only lagged-flow coefficient is 7.67; OEF-only lagged flow is essentially zero; the organic-growth-rate variant is not significant.

Adding contemporaneous market, SMB, HML, RMW, CMA, and momentum returns as direct controls leaves the ΔAggAUM coefficient at 9.52, compared with 11.39 in the baseline (Internet Appendix IA.Q). Factor returns do not absorb the relation.

The stock-level decay profile and the aggregate time-series regression locate the paper's demand-based mechanism in two adjacent pieces of recent work. Ben-David et al. (2022) use the 2002 Morningstar rating-methodology reform as a natural experiment and show that correlated investor demand at the investment-style level (driven by rating chasing) generated style-level price pressure that reverted over multi-year horizons; once the reform

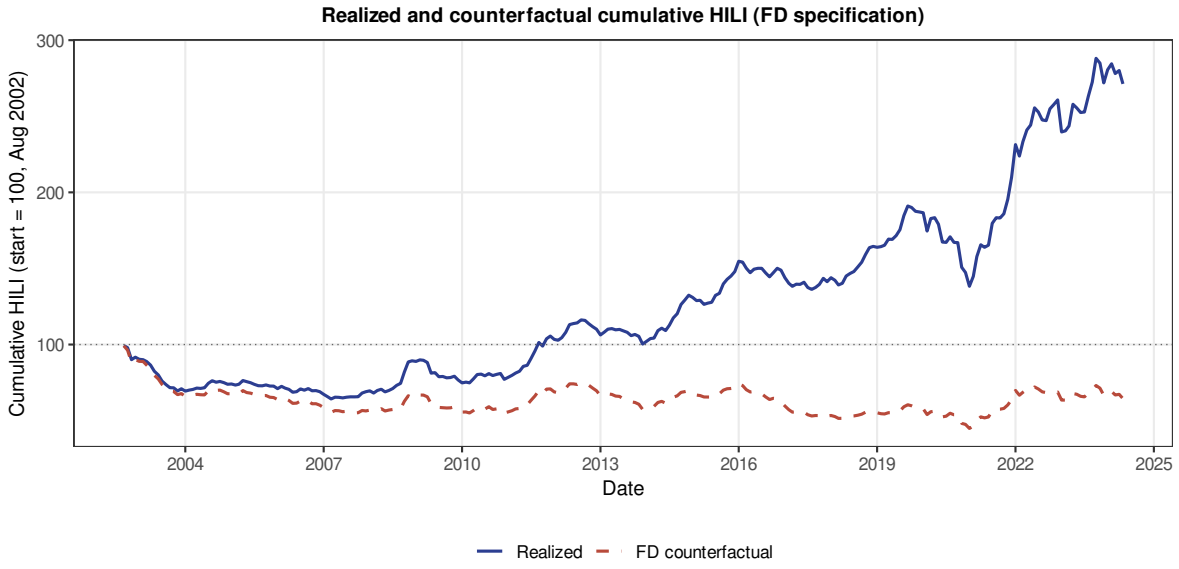

Figure 4: Realized HILI and the reduced-form FD-decomposition path under the first-difference specification, 2002-09 through 2024-05. Both paths start at 100. The dashed line subtracts $\hat{\beta}_{\Delta\text{AggAUM}}\Delta\text{AggAUM}_t$ each month from realized HILI, leaving earnings-release, long-term-forecast, lagged-flow, and intercept contributions intact. Realized HILI reaches 271; the FD-decomposition path reaches 65. The decomposition is reduced-form, not a structural no-passive counterfactual. See Table 6 for the underlying coefficient estimates.

broke the cross-style rating correlation, the style-level reversal pattern disappeared. Their mechanism is demand-based and behavioral, and their reversal horizon is measured in years. The mechanism studied here is narrower in object but shorter in horizon: a monthly holdings-based flow at the stock level is associated with a same-month price-pressure coefficient that substantially reverses within two to three months. The two results are not in conflict. The 2- to 3-month reversal here is a stock-level analog of the broader demand-pressure logic that Ben-David et al. document at the style-return frequency. Related work by Ben-David et al. (2024) and Duarte et al. (2025) places the mechanism in a broader account of how the rise of passive investing has reshaped price-discovery dynamics and aggregate elasticity. What is distinctive in this paper's setting is that the central *aggregate* correlate of HILI is the passive-ownership share rather than the passive-flow rate, a result closer in spirit to the equilibrium demand-system predictions of Koijen and Yogo (2019) and Haddad et al. (2025) than to a flow-chasing story.

The stock-level mechanism and the aggregate evidence operate on different margins and

at different frequencies. The stock-level flow coefficient captures a within-stock association between returns and monthly passive-demand movements; it is large at month zero and substantially reverses by month three. The first-difference aggregate result captures a between-stock, slow-moving co-movement between the HILI factor and the month-to-month change in the passive-ownership share. Passive demand therefore operates at two margins: short-run price pressure that fades at the stock level, and a slow-moving repricing of the benchmarked ownership base that compresses high-IXI ex-ante expected returns in the cross section. The first is a demand-pressure story with reversal in the tradition of Coval and Stafford (2007) and Lou (2012); the second is the inelastic-demand-equilibrium implication of Koijen and Yogo (2019) and Haddad et al. (2025) at the passive-ownership margin. The paper's cross-sectional ICC wedge of Section 5 is the direct expected-return consequence of the second margin; the stock-level flow reversal is the price-impact signature of the first. Both are consistent with the paper's indexing-premium reading of the passive era, and neither, taken alone, is sufficient.

6.5 Decomposing realized HILI into expected and unexpected components

A natural question is whether the realized HILI spread of Section 4 (positive, +0.43 percentage points per month) and the ex-ante ICC wedge of Section 5 (negative, −2.07 percentage points per year) can be reconciled in a single analytical frame. Pástor et al. (2022) show for ESG returns that a realized positive factor return is consistent with a negative ex-ante expected factor return when the sample contains persistent unexpected news about the relevant-factor exposure. The analog here decomposes realized monthly HILI into an expected-return component built from the Section 5 ICC wedge, annualized-to-monthly, and an unexpected-return residual.

Table 7 reports the decomposition over the 262-month sample. The expected-return component averages −0.170 percentage points per month, significant at the 1% level, in line

with the Section 5 ICC wedge.[4] The unexpected-return residual averages +0.603 percentage points per month, significant at the 5% level. The decomposition identity +0.433 = −0.170+ 0.603 holds by construction. The ratio of the unexpected to the realized component is 1.39, so the unexpected component dominates the realized spread and the ex-ante negative expected component is more than offset.

Table 7: Pastor-Stambaugh-Taylor-style decomposition of realized HILI. Panel A reports the decomposition of realized monthly HILI into an expected-return component (built from the Section 5 ICC wedge, annualized-to-monthly) and an unexpected-return residual over August 2002 through May 2024 (262 months). Panel B reports the regression of the unexpected-return residual on the aggregate covariates from Section 6. Newey-West standard errors in parentheses.

Panel A: Decomposition means (pp per month)	
Realized HILI	0.433^{*}
	(0.247)
Expected (ICC-based)	-0.170^{***}
	(0.020)
Unexpected (residual)	0.603^{**}
	(0.255)
Unexpected-to-realized ratio = 1.39	
Panel B: Unexpected residual regression on aggregate covariates	
ΔAggAUM_t	10.91^{***}
	(2.60)
$\text{ER}_t^{\text{HILI}}$	124.43^{***}
	(20.95)
AggFlow_{t-1}	6.81^{*}
	(3.51)
Intercept	−0.27
	(0.24)
n	261
R^2	0.220

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$ under Newey-West standard errors. Expected HILI is $(\text{ICC}_{Q5,t} - \text{ICC}_{Q1,t})/12$, where the ICC series is from the Section 5 portfolio-level construction (annualized). Unexpected HILI is Realized minus Expected. Panel B regressors: ΔAggAUM_t is the month-over-month change in the aggregate passive-ownership share, $\text{ER}_t^{\text{HILI}}$ is the HILI-specific aggregate earnings-release return, AggFlow_{t-1} is the lagged aggregate passive flow as a share of lagged CRSP market capitalization.

Panel B of Table 7 regresses the unexpected-return residual on the aggregate covariates

[4]The monthly PST expected-return component is constructed from calendar-month cross-sectional averages of the portfolio-level ICC spread rather than from 1/12 of the annual headline. The small discrepancy between $-0.170 \times 12 = -2.04$ and the −2.07 percentage-points-per-year figure reported in Section 5 reflects timing and weighting differences between the two constructions; the qualitative reading is identical.

already used in the HILI time-series regression: the monthly change in the passive-ownership share ΔAggAUM_t, the HILI-specific aggregate earnings-release return $\text{ER}_t^{\text{HILI}}$, and lagged aggregate passive flow AggFlow_{t-1}. All three carry the predicted positive signs; ΔAggAUM and ER^{HILI} deliver t-statistics above 4. The three covariates together explain 22 percent of the monthly variation in the unexpected-return residual. Positive news about passive-ownership accumulation and HILI-portfolio earnings news are therefore closely related to the residual that turns a negative expected-return component into a positive realized spread.

The decomposition is the bridge between the paper's realized and expected sides. HILI's realized return is positive, but the expected-return component built from the ICC spread is negative. The difference is the unexpected residual. Passive-share growth and HILI earnings news explain a meaningful, not exhaustive, share of that residual. The S&P 500 event study in Section 7 supplies the stock-level counterpart: an inclusion shock raises indexed ownership and produces a revaluation response. The apparent conflict between positive realized returns and lower ex-ante expected returns is therefore not a contradiction. It is the return signature of a transition in which passive ownership repeatedly rises by more than prices had already incorporated.

The decomposition reads the same way under an alternative ICC construction. Re-running the decomposition with the Gebhardt-Lee-Swaminathan (GLS) residual-income ICC in place of HVZ, which uses analyst EPS forecasts for years 1-3 and mean-reverts to the industry-median ROE at year 12, delivers an expected-return component averaging -0.173 percentage points per month, nearly identical to the HVZ value (the GLS and HVZ portfolio ICC spreads differ by only 4 basis points per year on the saved sample). The unexpected-return residual averages $+0.597$, versus $+0.603$ under HVZ. Panel B under GLS delivers a ΔAggAUM coefficient of 10.99 and an adjusted R^2 of 0.216, in line with the HVZ baseline. The GLS exercise is an alternative-ICC stability check rather than an independent expected-return source. It leaves the two core conclusions unchanged: high-IXI stocks carry a negative ex-ante expected-return component, and the positive realized spread comes from a larger

unexpected component related to passive-ownership growth and HILI-specific earnings news.

6.6 The transition premium and forward expected returns

The transition-era realized premium and the ICC wedge are two sides of the same mechanism. The realized cumulative high-minus-low IXI path reaches +171 percent over the 262-month sample. Compounded over the same window, the −2.07-percentage-point per year ICC wedge implies a cumulative path of approximately −36 percent. The difference gives an economic scale for the transition premium: it compares what investors received during the active-to-passive transition with what the ex-ante expected-return wedge would imply absent continued positive demand news. That gap should not be extrapolated as permanent alpha.

The realized premium can persist only while passive ownership grows faster than prices anticipate; once passive growth is fully anticipated or slows, the lower expected-return wedge is what remains. The paper's evidence ties HILI to the flow at conventional horizons: in flow-reversal episodes, mean HILI falls from +0.54 percentage points per month in high-flow months to −0.03 in the bottom-quintile months of aggregate passive flow, and a non-overlapping six-month deceleration-window test delivers a 9.75-percentage-point cumulative HILI differential (Appendix D). Behmaram (2026) estimates that high-IXI stocks are about 40 percent less elastic than low-IXI stocks, identifying the inelasticity channel through which flows propagate to prices in the framework of Section 2. The flow itself is anchored by defined-contribution defaults into low-fee balanced and target-date vehicles, the persistent fee differential between active and passive funds against a near-zero average active alpha (French, 2008; Fama and French, 2010; Berk and van Binsbergen, 2015), and the gradual re-allocation of pension, endowment, and insurance assets toward factor and passive strategies; the Grossman and Stiglitz (1980) active-floor argument supplies the natural termination condition. A reduced-form projection in Internet Appendix IA.L translates the Section 6 coefficient into 60-month horizons under three flow scenarios; bootstrap bands bracket zero under deceleration and stagnation, so the projection should be read as a conditional what-if

rather than a forecast.

The historical realized premium on heavily indexed stocks during the active-to-passive transition should not be extrapolated as permanent alpha.

7 Evidence from S&P 500 inclusions

S&P 500 additions provide the paper's most direct event-level evidence on the short-run price-pressure margin of the IXI channel. Inclusion events add a discrete ownership shock that mechanically raises benchmarked passive ownership. In the framework of Section 2, this is the short-window test of the discrete-shock prediction following Proposition 2. The design extends the tradition of Harris and Gurel (1986), Shleifer (1986), Chang et al. (2015), and Greenwood and Sammon (2025), but instead of the binary inclusion indicator it uses the cross-sectional magnitude of the IXI jump at the event as a continuous treatment. The evidence should be read as supportive of the paper's short-run price-pressure interpretation, not as a clean elasticity estimate.

7.1 Event design and treatment

The event sample consists of 386 S&P 500 additions between August 2002 and May 2024 with full coverage of announcement-window daily returns and IXI lookup months, matched against Siblis Research's roster of true announcement dates (Greenwood and Sammon, 2025). The match rate against the CRSP effective date is 96.3 percent; the 18 unmatched events fall back to the standard five-trading-day proxy. Match-rate accounting is detailed in Appendix C.1. Figure 5 plots the average market-adjusted cumulative abnormal return path around inclusion on two event-time anchors. The announcement-aligned path shows a jump of 2.07 percentage points on the single trading day after announcement (from 1.43 at day zero to 3.50 at day +1), consistent with the market pricing in inclusion on the first full trading day after the after-market announcement. The effective-aligned path peaks around

effective day zero at 3.11 percentage points and reverts partially over the following month.

The main continuous treatment is the fund-branch IXI jump over a tight event window:

$$T1_e \;=\; \mathrm{IXI}^{\mathrm{fund}}_{e,\,m+1} \;-\; \mathrm{IXI}^{\mathrm{fund}}_{e,\,m-1}, \tag{17}$$

where m is the calendar month containing the Siblis announcement date. Three considerations motivate $T1$. First, the fund branch captures mutual funds and ETFs explicitly tracking the S&P 500, which rebalance within one to two months of inclusion because they are mandated to hold constituents at benchmark weight. Second, the two-month bracket $[m-1,\, m+1]$ is tight enough to exclude the slower non-fund 13F branch, which adjusts over several quarters. Third, the measure is IXI-centric, so its variation isolates the part of the inclusion shock that operates through the paper's measure. The sample-mean $T1$ is 3.7 percentage points with a standard deviation of 3.4 percentage points. Slower treatments using the full IXI jump over a wider window are retained as robustness and are not the main pilot estimand.

The short-window outcome is the market-adjusted CAR from announcement day -1 through effective day $+1$ (call this CAR_short). The reversal window is effective day $+2$ through effective day $+21$, and cumulative horizons through effective day $+252$ are used for the decay profile below.

7.2 Baseline and augmented-control results

Table 8 reports the main regression results. The pooled Siblis baseline with log me as the only control yields $\hat{\beta}_{T1} = 22.04$ with cluster-robust $t = 1.93$, just short of the five percent bar. Adding log BM and 12-1 momentum to the same 372-event Siblis sample tightens the coefficient to $\hat{\beta}_{T1} = 31.94$ with cluster-robust $t = 2.44$, column (1). I report both specifications and treat the augmented-control version as the primary pooled-sample number. The Q1 best-timing subsample described below reports a sharper coefficient on the events

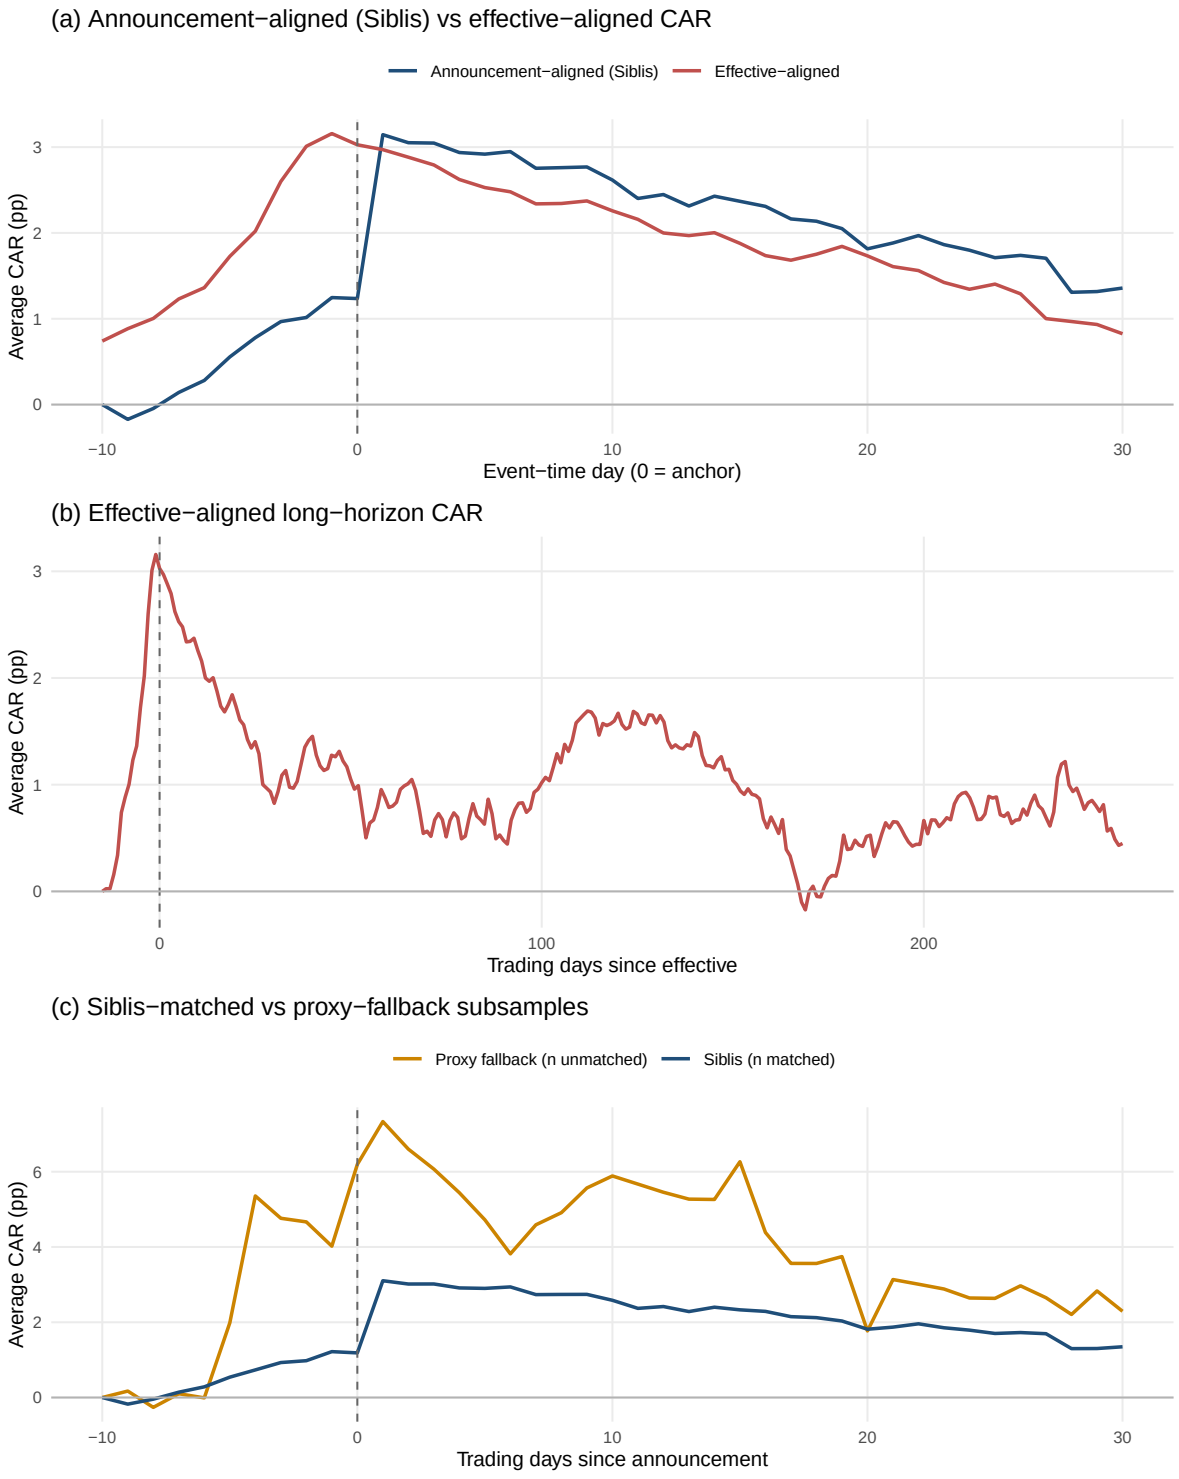

Figure 5: Average cumulative abnormal return around S&P 500 inclusion for the 386-event Siblis-matched sample. The y-axis is cumulative abnormal return in percentage points. The left panel uses the announcement-date anchor, the middle panel uses the effective-date anchor, and the right panel compares the Siblis-matched and proxy-fallback samples. Dotted lines mark the announcement day (left) and the effective day (middle).

where Siblis-proxy timing alignment is cleanest; I treat Q1 as a timing-accuracy sharpening check rather than as a primary estimate, so that the paper's headline short-window claim rests on the pooled Siblis result.

Table 8: S&P 500 inclusion CAR regressions with controls and size heterogeneity. The dependent variable is the market-adjusted CAR from announcement day -1 through effective day $+1$ (percentage points). Columns (1) and (2) add log BM and 12-1 momentum to the baseline controls. Columns (3) and (4) add a continuous $T1 \times \log$ me interaction. Columns (5) and (6) add a discrete $T1\times$ top-size-decile interaction.

	(1) Sib.	(2) Q1	(3) Sib.	(4) Q1	(5) Sib.	(6) Q1
$T1$ (fund, short)	31.944^{**}	56.979^{***}	-443.839^{**}	-537.504	27.427^{**}	55.555^{***}
	(13.087)	(16.948)	(203.641)	(351.505)	(13.580)	(19.028)
$T1 \times \log$ me			51.230^{**}	63.547^{*}		
			(21.775)	(38.090)		
$T1 \times \mathbf{1}_{\text{top10}}$					-3.298	-13.314
					(32.327)	(38.814)
$\mathbf{1}_{\text{top10}}$					6.602^{***}	8.351^{***}
					(1.981)	(3.029)
log me	1.145^{*}	0.156	-1.438	-2.704^{*}	-0.699	-1.873^{**}
	(0.681)	(0.785)	(0.929)	(1.607)	(0.803)	(0.766)
log BM	-0.055	-0.134	-0.030	-0.036	-0.193	-0.232
	(0.278)	(0.403)	(0.283)	(0.402)	(0.283)	(0.421)
Momentum (12-1)	1.548	2.454^{**}	1.741^{*}	2.623^{**}	1.760	2.796^{**}
	(1.114)	(1.174)	(1.057)	(1.257)	(1.132)	(1.183)
Observations	372	157	372	157	372	157
R^2	0.091	0.130	0.122	0.170	0.135	0.193

Notes: The sample covers S&P 500 additions between August 2002 and May 2024 (Siblis Research roster). Sib. columns use the pooled Siblis-matched sample; Q1 columns use the best-timing quartile. These are cross-sectional event regressions with no panel fixed effects. $T1$ enters in decimal units, so coefficients are CAR percentage points per unit increase in $T1$. Standard errors are in parentheses and are clustered by announcement year. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

The Siblis announcement date carries some measurement error. Splitting the Siblis sample into four quartiles based on the signed proxy-to-Siblis gap, the Q1 best-timing subsample (157 events) delivers $\hat{\beta}_{T1} = +56.98$, cluster-robust $t = 3.36$ (column (2)). The Q1 coefficient is 1.78 times the pooled coefficient, consistent with classical attenuation from event-time measurement error in the broader sample. I read Q1 as the sharpest cut of the short-window evidence and the pooled Siblis estimate as consistent but attenuated evidence on the broader sample.

7.3 Matched-control robustness

The matched-control robustness, reported as Table 16 in Appendix C.2, matches each treated event to up to five non-S&P 500 control firms on Fama-French 12 industry, size decile within ±1, and book-to-market quintile within ±1, excluding firms that were S&P 500 members within ±12 months of the event. The size-weighted mean CAR of the control pool is subtracted from the treated CAR. Of 382 Siblis events with full controls, 326 have at least one valid match.

The matched-control Q1 coefficient preserves 86 percent of the augmented Q1 headline. The pooled Siblis matched-control coefficient preserves only 57 percent of the headline and is not significant; an FF3-adjusted variant raises the Siblis preservation to 67 percent and recovers significance at the 5% level. The FF3 variant absorbs sample-wide size-value-market drift in the control pool. Read together, the matched-control evidence supports short-horizon identification on Q1 under both raw and FF3 variants, and on the pooled Siblis sample under FF3 once style-tilt drift is removed. A 1,000-iteration placebo permutation of $T1$ rejects the specification-search null on all three cells at the five percent level (Internet Appendix IA.E).

7.4 Decay profile

Figure 6 plots the horizon-by-horizon coefficient on $T1$ at $h \in \{1, 5, 21, 63, 126, 252\}$ trading days past the effective date, under both the market-adjusted and matched-control outcomes, on the Siblis-all and Q1 subsamples. Under market-adjusted CAR on Siblis-all, the coefficient decays from 32 at one day to 18 at one month to 10 at three months, consistent with short-run price pressure followed by partial reversal in the tradition of Harris and Gurel (1986), Coval and Stafford (2007), and Lou (2012). Under market-adjusted CAR on the Q1 subsample, the pattern is different: the coefficient grows from 57 at one day to 176 at twelve months (cluster-robust $t = 2.38$), implying a persistent rather than reverted inclusion effect for events where event-time alignment is cleanest. Once matched-control adjustments remove industry-size-book-to-market drift, the Q1 long-horizon persistence shrinks sharply: the matched-control

twelve-month Q1 coefficient is 42 with cluster-robust $t = 0.39$.

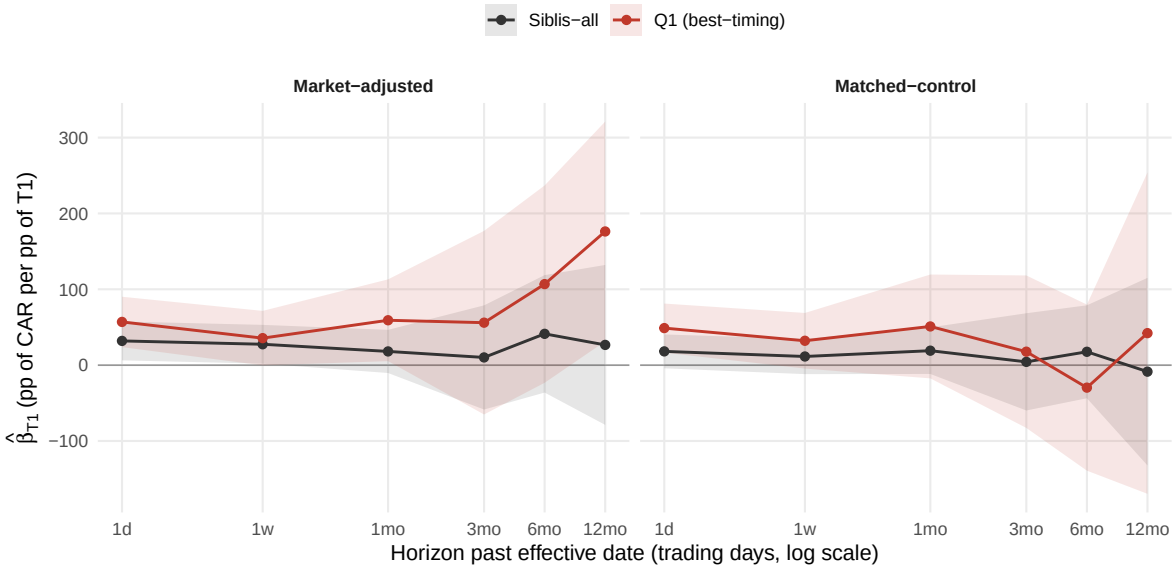

Figure 6: Decay of the $T1$ coefficient at horizons past the effective date, under the market-adjusted outcome (left) and the matched-control outcome (right). Each line plots $\hat{\beta}_{T1}$ from a regression of the CAR through effective day h on $T1$ plus log me, log BM, and 12-1 momentum. Ribbons show 95 percent cluster-robust confidence intervals. Horizon axis is on a log scale.

The two subsamples give different pictures. Siblis-all shows the textbook short-run price pressure with reversal pattern. Market-adjusted Q1 looks like persistence rather than reversal, but the matched-control adjustment shows that much of that long-horizon Q1 persistence reflects contemporaneous industry-style drift rather than a pure stock-level persistence effect. The intersection of the two readings is that the short-horizon price-pressure channel is present on both samples, and that the long-horizon story is more consistent with pooled reversal than with Q1-only persistence once appropriate controls for industry drift are applied.

7.5 Mega-firm and family-concentration heterogeneity

Columns (3) through (6) of Table 8 report two mega-firm heterogeneity cuts. The continuous $T1 \times \log \text{me}$ interaction on Siblis is 51.23, significant at the 5% level; the discrete $T1 \times \mathbf{1}_{\text{top10}}$ interaction is a null at -3.30 (top decile contains only 41 events). The two cuts therefore do

not yield a clean mega-firm-elasticity ordering. A parallel cut on pre-event passive-family concentration, reported in Internet Appendix IA.G, shows that top-five-family share acts as a level shifter on inclusion CARs (0.8 pp per ten-percentage-point increase, marginally significant) rather than as a $T1$ elasticity amplifier.

The event-study evidence supports the short-run price-pressure mechanism and anchors the aggregate results in the cross-section. The design does not on its own identify a causal demand elasticity. The cross-sectional demand-gap design in Section 8 provides the complementary first-stage and reduced-form evidence on which inclusion events generate large realized IXI shocks and large short-window price responses.

8 S&P 500 inclusions and the demand gap

In a modern passive market, S&P 500 inclusion is not a binary demand event. Most large stocks already sit in benchmark-linked portfolios through Russell, CRSP total-market, style, and sector products before they ever enter the S&P 500. Inclusion changes the composition of this benchmark-linked ownership rather than turning it on. The cross-sectional object that matters for the inclusion shock is therefore the *S&P demand gap*: predicted S&P-family demand at inclusion, net of the passive ownership the stock had already absorbed through other benchmark families. IXI makes the gap observable. The gap predicts the realized fund-branch IXI jump at the announcement and the short-window cumulative abnormal return; gross S&P-family demand does not consistently predict both margins.

8.1 Constructing the S&P demand gap

For each addition event e, define gross predicted S&P-family demand as

$$\text{PredDemand}_e^{SP} \;=\; w_e^{SP}\,\text{AUM}_{m(e)}^{SP}, \tag{18}$$

where w_e^{SP} is the first non-zero post-effective S&P 500 weight from Siblis Research, restricted to a 45-calendar-day window after the effective date, and $\text{AUM}_{m(e)}^{SP}$ is the dollar value of S&P 500-family passive holdings in the announcement month. The S&P-family AUM series is built by aggregating the family-contribution decomposition that underlies IXI over benchmarks classified to the S&P 500 family. It grows from \$0.4 trillion in early 2010 to \$3.7 trillion at the end of the sample, roughly half of total U.S. passive equity AUM in recent years.

The S&P demand gap nets out pre-existing passive ownership through non-S&P benchmarks:

$$\text{Gap}_e^{SP} \;=\; \text{PredDemand}_e^{SP} \;-\; \text{PassiveOwnership}_{e,\,m-1}^{ex\text{-}SP}, \tag{19}$$

where the second term sums passive ownership at month $m-1$ across Russell, CRSP, style, sector, and other benchmark families using the same family-contribution decomposition. The gap is an exposure measure, not a literal net order-flow object: existing Russell or CRSP passive holders do not mechanically sell when a stock joins the S&P 500. The gap describes incremental S&P-specific demand relative to prior passive saturation.

Three descriptive facts about Gap_e^{SP} frame the cross-sectional analysis. The median entrant carries about \$0.5 billion in gross S&P-family predicted demand, but the median demand gap is roughly negative \$150 million because most modern entrants already hold substantial passive ownership through other families. Cross-sectional rank correlation between ME-scaled gross demand and the ME-scaled demand gap is 0.16, so the gap captures genuinely different cross-sectional variation rather than rescaled inclusion size. I report ADV-scaled, ME-scaled, and dollar-scaled versions of both gross demand and the demand gap; the three scalings give three different views of demand pressure relative to a stock's trading capacity, market capitalization, or absolute exposure size.

8.2 First-stage validation

The first stage tests whether ex-ante predicted demand predicts the realized fund-branch IXI jump $T1_e$ at inclusion. The regression runs on the cross section of events:

$$T1_e \;=\; \alpha + \beta\, \text{Demand}_e^{(z)} + \gamma' X_e + \tau_{y(e)} + \varepsilon_e, \tag{20}$$

where $\text{Demand}_e^{(z)}$ is one of the z-scored predicted-demand variables, X_e collects log(ME), log(B/M), twelve-month-minus-one-month momentum, and the size decile, and $\tau_{y(e)}$ are announcement-year fixed effects. Standard errors cluster by announcement year. Panel A of Table 9 reports the result.

Table 9: S&P demand gaps, realized IXI shocks, and inclusion returns. Each cell reports a separate cross-sectional regression at the event level. Panel A's outcome is the realized fund-branch IXI jump in the announcement window, in percentage points. Panel B's outcome is the short-window cumulative abnormal return in percentage points. Regressors are z-scored, so coefficients are read as the change in the outcome per one-standard-deviation increase in the predicted-demand variable.

	Gross S&P-family demand		S&P demand g	
Regressor	ADV-scaled	ME-scaled	ADV-scaled	M
Panel A: First stage. Outcome = T1 fund-branch IXI shock (pp)				
Coefficient	0.002	0.253	1.041***	
	(0.295)	(0.239)	(0.359)	
Observations	360	360	360	
Panel B: Short-window CAR. Outcome = $\text{CAR}_{\text{short}}$ *(pp)*				
Coefficient	0.276	1.747***	1.186**	
	(0.655)	(0.545)	(0.439)	
Observations	370	370	370	

Notes: Cross section of S&P 500 inclusion events, 2002-2024. The outcome in Panel A is the realized fund-branch IXI jump from the month before to the month after the announcement, scaled to percentage points. The outcome in Panel B is the announcement-window cumulative abnormal return $\text{CAR}_{\text{short}}$ in percentage points. All five right-hand-side variables are z-scored. *Gross S&P-family demand* is the first post-effective Siblis S&P 500 weight times monthly S&P 500-family passive dollars from the IXI family-contribution decomposition. *S&P demand gap* subtracts pre-existing ex-S&P passive ownership at $m-1$ (the month before announcement) from gross S&P-family demand and is therefore an exposure-gap measure rather than a literal net-flow object. ADV-scaled divides by pre-event average daily dollar volume, ME-scaled divides by lagged market capitalization, and dollar-scaled is the raw demand variable in dollars. Controls are log(ME), log(B/M), 12-1 momentum, and the size decile, with announcement-year fixed effects throughout. Standard errors clustered by announcement year are in parentheses. Siblis weights are restricted to the first non-zero post-effective observation within 45 calendar days. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Gross S&P-family predicted demand is essentially uninformative for the realized inclusion

shock once these controls are absorbed. The ADV-scaled and ME-scaled gross variables both sit near zero with $|t| < 1.1$. The S&P demand gap predicts the realized shock on every scaling. A one-standard-deviation higher ADV-scaled demand gap raises the announcement-window fund-branch IXI jump by 1.04 percentage points ($t = 2.90$); the ME-scaled and dollar-scaled gaps both clear five-percent significance under year-clustered inference. Pre-existing passive ownership sharpens which inclusion events generate large realized IXI shocks: stocks already deeply held through Russell or CRSP funds experience smaller realized IXI jumps than gross S&P demand alone would predict. The first stage is therefore the cleanest direct event-level evidence in the paper that IXI captures realized indexed ownership. IXI in levels organizes the cross-sectional ICC wedge of Section 5; the demand-gap first stage validates that IXI also registers the discrete cross-sectional shocks the model implies it should at inclusion.

8.3 Price response

Panel B of Table 9 repeats the cross section with the announcement-window CAR (defined in Section 7) as the outcome. The ME-scaled demand gap predicts CAR positively and significantly: a one-standard-deviation higher gap raises the short-window CAR by 1.71 percentage points ($t = 4.25$). The ADV-scaled and dollar-scaled gaps deliver 1.19 and 4.07 percentage points respectively, both significant at the five-percent level under year-clustered inference. Gross ME-scaled demand is also significant in this panel, but the gross ADV-scaled variable is not, so the gross demand result is fragile to scaling while the demand-gap result is not.

The price response loads on the same cross-sectional object as the first stage. Inclusion events with larger predicted demand relative to prior passive saturation generate larger realized IXI shocks and larger short-window price responses. Gross index demand does not consistently predict both margins, which is what makes the gap the right shock object. This is the stock-level event-study counterpart of the paper's aggregate evidence: unanticipated growth in benchmark-linked demand moves prices when it creates new ownership pressure,

not when it simply rebalances which benchmark family holds a given stock.

8.4 What the inclusion design does not show

The inclusion event identifies short-window demand pressure cleanly. It does not, on its own, deliver an average permanent revaluation through lower expected returns. Three pieces of long-horizon evidence discipline the interpretation.

The event-time log market-equity path rises sharply through the announcement window and partly reverts over the following three years. Relative to a far-pre baseline of months $[-36, -6]$ before the effective date, the treated-minus-matched-control gap reaches 33 percent at $[+2, +5]$ months, fades to roughly 4 percent at $[+25, +36]$ months, and is no longer statistically distinguishable from zero at the longest horizon. The price level partially persists; the long-horizon estimate is too noisy to support a clean permanent-revaluation claim.

The implied cost of capital does not fall after inclusion. In a parallel event-time DiD on the same matched-control panel, with permno-by-event and year-quarter fixed effects, treated firms' ICC rises by 1.17 percentage points ($t = 5.89$) at $[+25, +36]$ months relative to the immediately-pre-event reference and by 0.62 percentage points ($t = 3.08$) relative to the far-pre reference. The pattern holds in both the pre-2013 and the post-2013 subsamples and is unchanged by the demand-gap interaction. Inclusion does not produce an average decline in expected returns, so it cannot stand on its own as a clean permanent discount-rate shock.

Split-adjusted shares outstanding do not respond at long horizons. Once log(SHROUT) is corrected for stock splits using the CRSP share-adjustment factor, the event-time path is statistically flat at $[+25, +36]$, in contrast to the unadjusted measure used in earlier drafts of the paper. I therefore do not interpret S&P inclusion as evidence on equity issuance.

The inclusion design isolates demand pressure as a mechanism. The expected-return wedge in Section 5 carries the paper's discount-rate claim; inclusion events validate the demand mechanism without standing as an independent permanent-revaluation result.

8.5 Linking back to the cross-sectional and aggregate evidence

The paper's main pricing results live on different margins. The ICC wedge in Section 5 prices stocks in levels: deeper indexed ownership is associated with lower expected returns whether the stock recently changed benchmark or not. The aggregate first-difference and PST decomposition in Section 6 prices the time series of unanticipated growth in indexed ownership; months in which passive ownership grew faster than prices had incorporated produced positive realized HILI returns. Inclusion events are the discrete cross-sectional version of the second margin. The demand gap predicts which events generate large realized IXI shocks and large short-window price responses, while the long-horizon ICC evidence shows that the broader expected-return wedge does not reduce mechanically to the response from a single index event. The two margins are complementary: stocks priced for low expected returns in the cross section are not the same stocks that experience large realized revaluation around any specific inclusion event. The demand-gap evidence validates the event-level demand-pressure margin that links indexed ownership to the transition-era return evidence.

9 Conclusion

The active-to-passive transition has been capitalized into U.S. equity prices through a stock-level ownership characteristic. IXI measures realized indexed ownership from observed holdings; stocks with high IXI carry implied costs of capital about two percentage points per year below otherwise similar low-IXI stocks. The same characteristic organizes the transition-era realized premium: high-IXI stocks earned positive realized returns when passive ownership rose by more than prices had already incorporated, and the unexpected component of the realized spread co-moves with growth in aggregate passive ownership. The first-difference aggregate decomposition is the paper's quantitative synthesis: subtracting the estimated passive-share-growth contribution changes the cumulative high-minus-low IXI path from a

171-percent gain to a 35-percent loss. S&P 500 inclusion shocks corroborate the demand mechanism cross-sectionally; the inclusion design does not carry a permanent-revaluation claim on its own. For forward expected returns, the transition-era realized premium reflects repeated unexpected demand growth, not permanent alpha; the roughly two-percentage-point ICC wedge is what remains once that demand growth is anticipated or slows. The paper does not forecast when the transition slows; how active-floor and concentration-limit mechanisms (Grossman and Stiglitz, 1980; Pástor et al., 2026) bind at the top of the cross-section is the boundary condition for the next stage.

References

S. Basak and A. Pavlova. Asset prices and institutional investors. *American Economic Review*, 103(5):1728–1758, 2013.

P. Behmaram. Indexing and the elasticity of stock demand. Working paper, 2026.

I. Ben-David, J. Li, A. Rossi, and Y. Song. Ratings-driven demand and systematic price fluctuations. *Review of Financial Studies*, 35(6):2790–2838, 2022.

I. Ben-David, J. Li, A. Rossi, and Y. Song. Discontinued positive feedback trading and the decline of return predictability. *Journal of Financial and Quantitative Analysis*, 59(7): 3062–3100, 2024.

B. Bennett, R. M. Stulz, and Z. Wang. Does joining the s&p 500 index hurt firms? Technical Report 27593, National Bureau of Economic Research, 2020.

I. Ben-David, F. Franzoni, and R. Moussawi. Do ETFs Increase Volatility? *The Journal of Finance*, 73(6):2471–2535, 2018. ISSN 1540-6261. doi: https://doi.org/10.1111/jofi.12727. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12727. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12727.

J. B. Berk and J. H. van Binsbergen. Measuring skill in the mutual fund industry. *Journal of Financial Economics*, 118(1):1–20, 2015.

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

P. Bond and D. Garcia. The equilibrium consequences of indexing. *The Review of Financial Studies*, 35(7):3175–3230, 2022.

M. J. Brennan. Agency and asset pricing. Working paper, University of California, Los Angeles, 1993.

J. Y. Campbell and R. J. Shiller. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1(3):195–228, 1988.

Y.-C. Chang, H. Hong, and I. Liskovich. Regression Discontinuity and the Price Effects of Stock Market Indexing. *The Review of Financial Studies*, 28(1):212–246, Jan. 2015. ISSN 0893-9454. doi: 10.1093/rfs/hhu041. URL https://doi.org/10.1093/rfs/hhu041.

A. Chinco and M. Sammon. The passive ownership share is double what you think it is. *Journal of Financial Economics*, 157:103860, 2024.

J. L. Coles, D. Heath, and M. C. Ringgenberg. On index investing. *Journal of Financial Economics*, 145(3):665–683, 2022.

J. Coval and E. Stafford. Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 86(2):479–512, Nov. 2007. ISSN 0304-405X. doi: 10.1016/j.jfineco.2006.09.007. URL http://www.sciencedirect.com/science/article/pii/S0304405X07001158.

K. M. Cremers and A. Petajisto. How active is your fund manager? a new measure that predicts performance. *The review of financial studies*, 22(9):3329–3365, 2009.

S. DellaVigna and J. M. Pollet. Investor inattention and friday earnings announcements. *Journal of Finance*, 64(2):709–749, 2009.

V. Duarte, M. Kargar, J. Li, and D. Silva. Dissecting the aggregate market elasticity. Technical report, Working Paper, SSRN, 2025. Available at SSRN 6025014.

P. D. Easton. Pe ratios, peg ratios, and estimating the implied expected rate of return on equity capital. *The Accounting Review*, 79(1):73–95, 2004.

E. F. Fama and K. R. French. Luck versus skill in the cross-section of mutual fund returns. *Journal of Finance*, 65(5):1915–1947, 2010.

A. Frazzini and O. A. Lamont. Dumb money: Mutual fund flows and the cross-section of stock returns. *Journal of Financial Economics*, 88(2):299–322, 2008.

K. R. French. Presidential address: The cost of active investing. *Journal of Finance*, 63(4): 1537–1573, 2008.

X. Gabaix and R. S. J. Koijen. In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis. Working Paper 28967, National Bureau of Economic Research, 2021.

W. R. Gebhardt, C. M. Lee, and B. Swaminathan. Toward an implied cost of capital. *Journal of accounting research*, 39(1):135–176, 2001.

R. Greenwood and M. Sammon. The disappearing index effect. *Journal of Finance*, 80(2): 657–698, 2025.

S. J. Grossman and J. E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

V. Haddad, P. Huebner, and E. Loualiche. How competitive is the stock market? theory, evidence from portfolios, and implications for the rise of passive investing. *American Economic Review*, 115(3):975–1018, 2025.

L. Harris and E. Gurel. Price and Volume Effects Associated with Changes in the S&P 500 List: New Evidence for the Existence of Price Pressures. *The Journal of Finance*, 41(4):815–829, 1986. ISSN 1540-6261. doi: https://doi.org/10.1111/j.1540-6261.1986.tb04550.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.1986.tb04550.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.1986.tb04550.x.

J. J. Heckman. Sample selection bias as a specification error. *Econometrica*, 47(1):153–161, 1979.

K. Hou, M. A. Van Dijk, and Y. Zhang. The implied cost of capital: A new approach. *Journal of Accounting and Economics*, 53(3):504–526, 2012.

D. Israeli, C. M. C. Lee, and S. A. Sridharan. Is there a dark side to exchange traded funds? an information perspective. *Review of Accounting Studies*, 22(3):1048–1083, 2017.

A. K. Kashyap, N. Kovrijnykh, J. Li, and A. Pavlova. The benchmark inclusion subsidy. *Journal of Financial Economics*, 142(2):756–774, 2021.

A. K. Kashyap, N. Kovrijnykh, J. Li, and A. Pavlova. Is there too much benchmarking in asset management? *American Economic Review*, 113(4):1112–1141, 2023.

R. S. J. Koijen and M. Yogo. A Demand System Approach to Asset Pricing. *Journal of Political Economy*, 127(4):1475–1515, 2019. ISSN 0022-3808. doi: 10.1086/701683. URL https://www.journals.uchicago.edu/doi/full/10.1086/701683. Publisher: The University of Chicago Press.

R. S. J. Koijen, R. J. Richmond, and M. Yogo. Which investors matter for equity valuations and expected returns? *Review of Economic Studies*, 91(4):2387–2424, 2024.

J. Li. What drives the size and value factors? *Review of Asset Pricing Studies*, 12(4): 845–885, 2022.

D. Lou. A flow-based explanation for return predictability. *Review of Financial Studies*, 25 (12):3457–3489, 2012.

L. Pástor, R. F. Stambaugh, and L. A. Taylor. Dissecting green returns. *Journal of Financial*

Economics, 146(2):403–424, 2022.

Ľ. Pástor, T. Sikorskaya, and L. A. Wang. The hidden cost of stock market concentration: When funds hit regulatory limits. NBER Working Paper 35007, National Bureau of Economic Research, 2026.

A. Pavlova and T. Sikorskaya. Benchmarking intensity. *The Review of Financial Studies*, 36 (3):859–903, 2023.

M. Sammon. Passive ownership and price informativeness. *Management Science*, 71(6): 4582–4598, 2025.

M. Sammon and J. J. Shim. Who clears the market when passive investors trade? *Review of Financial Studies*, 2026. Advance article.

M. S. Scholes. The Market for Securities: Substitution Versus Price Pressure and the Effects of Information on Share Prices. *The Journal of Business*, 45(2):179–211, 1972. ISSN 0021-9398. URL https://www.jstor.org/stable/2352030. Publisher: University of Chicago Press.

A. Shleifer. Do Demand Curves for Stocks Slope Down? *The Journal of Finance*, 41(3):579–590, 1986. ISSN 1540-6261. doi: https://doi.org/10.1111/j.1540-6261.1986.tb04518.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.1986.tb04518.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.1986.tb04518.x.

V. A. Warther. Aggregate mutual fund flows and security returns. *Journal of Financial Economics*, 39(2-3):209–235, 1995.

A IXI construction details and realized-return robustness

This appendix first works through the held-stock-normalized non-fund branch of IXI with a stylized two-benchmark, three-stock example, and then collects supporting robustness on the realized-return section.

A.1 Worked example: held-stock normalization with two benchmarks and three stocks

Consider a single 13F institution e at quarter-end q with \$100 in AUM, holding two of three stocks in the universe. Stocks A, B, C have month-t market weights of 0.5, 0.3, 0.2 in a broad index h_1, and 0.6, 0.4, 0.0 in a style-subset index h_2. The institution holds \$50 in stock A and \$50 in stock B and does not hold stock C.

The NNLS projection of the institution's \$50/\$50 portfolio onto (h_1, h_2) recovers loadings $(\lambda_{h_1}, \lambda_{h_2}) = (0.0, 1.0)$, because the style subset fits the observed A-B concentration perfectly and assigns zero weight to C. Without held-stock normalization, the institution's benchmarked-dollar allocation would be \$0 to A, \$0 to B, and \$0 to C from h_1 (since $\lambda_{h_1} = 0$) and \$60 to A, \$40 to B, and \$0 to C from h_2, for a total of \$100. The attribution already sums to the institution's AUM, and no held-stock renormalization is needed in this special case.

Now change the NNLS result to loadings $(\lambda_{h_1}, \lambda_{h_2}) = (0.5, 0.5)$, which would arise if the institution's holdings were partially consistent with both benchmarks. Without held-stock normalization, the implied benchmarked dollars would be \$25 to A, \$15 to B, and \$10 to C from h_1 (using h_1's 0.5/0.3/0.2 weights on a \$50 share), plus \$30 to A, \$20 to B, and \$0 to C from h_2. The total benchmarked mass is \$100 correctly, but \$10 of it is attributed to stock C, which the institution does not hold. Held-stock normalization corrects this phantom attribution. The benchmarked dollars over the held-stock set $\{A, B\}$ are renormalized to

sum to the institution's AUM within that set: from h_1, the A/B share within held stocks is $0.5/(0.5 + 0.3) = 0.625$ and $0.3/0.8 = 0.375$, so the \$50 assigned through h_1 becomes \$31.25 to A and \$18.75 to B; from h_2, the A/B share is 0.6 and 0.4 on the held set, so the \$50 assigned through h_2 becomes \$30 to A and \$20 to B. Total benchmarked dollars now attribute \$61.25 to A, \$38.75 to B, and \$0 to C, with the total preserved at \$100.

The example makes two design choices visible. First, the NNLS loadings do not sum to one in general; the \$100 total is preserved through the institution's Active-Share-adjusted attributable mass, not through a sum-to-one constraint on the loadings themselves. Second, held-stock normalization eliminates phantom attribution to stocks that the institution does not hold while preserving the institution's total benchmarked mass within its actual portfolio. The two design choices together are what distinguishes the revised IXI construction's 13F institutional branch from a naive attribution that would spread benchmarked dollars across the entire benchmark universe.

A.2 Raw monthly branch decomposition of IXI

Figure 7 plots the raw monthly cross-sectional mean of IXI under the information-date rule, decomposed into the fund branch and the 13F institutional branch. The fund branch and the full revised IXI update monthly as FactSet fund holdings refresh and are drawn as connected lines. The 13F institutional branch updates on the quarterly 13F filing schedule. Within each quarter the entity-level AUM scalar carries forward the most recently observable 13F data under the information-date rule, and the series is drawn as a step function to reflect that discrete update cadence. No spline, Kalman, or interpolation smoother is applied to the 13F series; the step function is the honest rendering of a quarterly-updating input. The visible quarter-to-quarter stepping in the full revised IXI is primarily driven by the 13F branch, but is not exclusively driven by it. The revised measure is also somewhat more active month to month than the legacy fund-only IXI studied in Behmaram (2026), because the NNLS best-fit benchmark assignment is recomputed each quarter and coverage is broader than the

declared-benchmark fund universe. The main-text trend presentation in Figure 1 uses a 12-month trailing mean to emphasize the secular rise and to smooth over the quarterly stepping shown here.

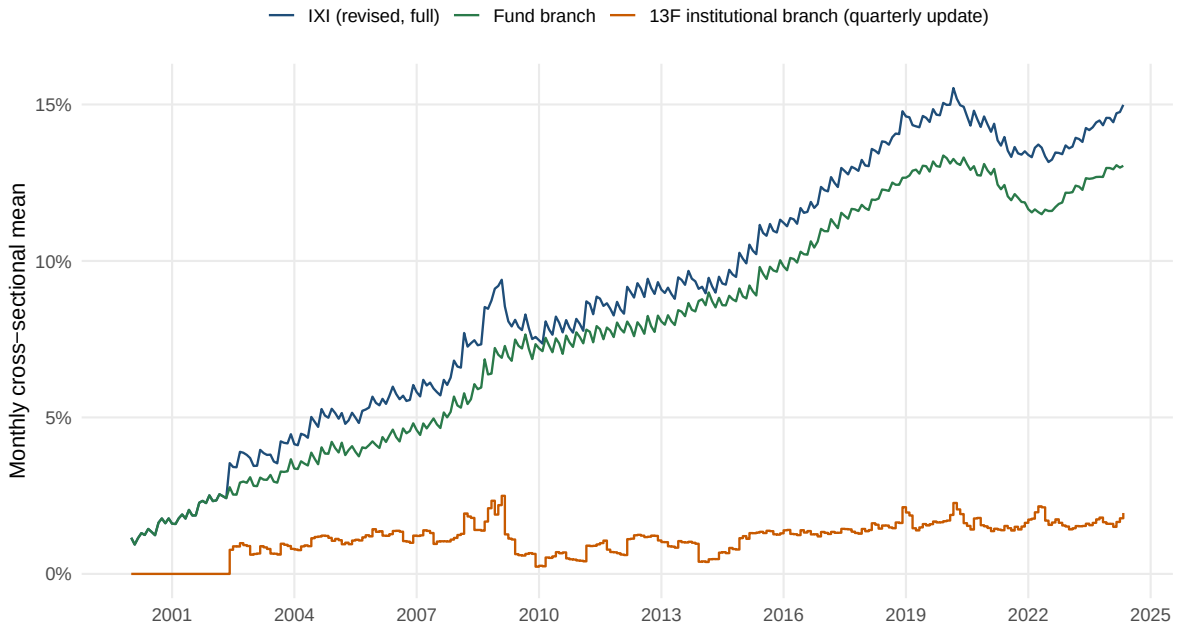

Figure 7: Raw monthly cross-sectional mean of IXI under the information-date rule, January 2000 through May 2024. The y-axis is IXI as a share of market capitalization. The full IXI (blue) and the fund branch (green) update monthly as FactSet fund holdings refresh and are drawn as connected lines. The 13F institutional branch (orange) updates on the quarterly 13F filing schedule and is drawn as a step function to reflect that discrete update cadence. No interpolation is applied to the 13F series.

A.3 Fama-MacBeth cross-sectional regressions

Table 10 reports Fama-MacBeth cross-sectional regressions of monthly CRSP returns on the two-month-lag Rule-I IXI signal with size, value, and momentum controls. The unconditional cross-sectional slope on IXI is essentially zero ($\bar{\lambda}_{\text{IXI}} = -0.002$, $t = -0.19$, Column 1), reflecting the strong unconditional correlation between IXI and size. Conditioning on log me alone (Column 2) flips the slope to 0.020 with $t = 4.00$. Adding log BM (Column 3) and momentum (Column 4) leaves the IXI slope at 0.013 with $t = 3.11$. A winsorized variant (Column 5) at the 1 and 99 percent cross-sectional quantiles of the right-hand-side characteristics delivers 0.012 with $t = 2.68$. The positive conditional slope on IXI is con-

sistent with the positive realized-return HILI spread documented in Section 4 on a classical Fama-MacBeth cross-sectional lens. The Fama-MacBeth specification does not test the ex-ante ICC wedge of Section 5, which uses firm-level implied cost of capital as the dependent variable and a different panel-regression design.

Table 10: Fama-MacBeth cross-sectional regressions of monthly CRSP stock returns on lagged IXI with standard controls. Each column reports the time-series average of the monthly cross-sectional OLS slope on each characteristic, with Newey-West standard errors on the time series of slopes in parentheses. The IXI signal uses the Rule-I two-month lag. Log market equity and log book-to-market are lagged one month; momentum is the 12-to-2 cumulative log return. Column (5) is Column (4) with within-month winsorization of the non-return regressors at the 1st and 99th percentiles. Sample: August 2002 through May 2024.

	(1)	(2)	(3)	(4)	(5) wins.
IXI (lag 2)	−0.00204	0.02019***	0.01400***	0.01268***	0.01214***
	(0.01074)	(0.00505)	(0.00414)	(0.00408)	(0.00453)
log me		−0.00140*	−0.00068	−0.00084	−0.00082
		(0.00073)	(0.00066)	(0.00055)	(0.00055)
log BM			0.00252***	0.00217***	0.00228***
			(0.00091)	(0.00078)	(0.00083)
Momentum (12-2)				0.00252	0.00325
				(0.00283)	(0.00301)
Constant	0.00992**	0.01661**	−0.00283	−0.00088	−0.00173
	(0.00491)	(0.00810)	(0.00858)	(0.00880)	(0.00911)
Avg. N	4,277	4,277	4,093	3,970	3,970
Avg. R^2	0.004	0.009	0.016	0.027	0.027
Months T	262	262	262	262	262

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$ under Newey-West standard errors on the time series of slopes. The dependent variable is the monthly CRSP return in decimal.

A.4 Between-firm versus within-firm decomposition

Table 11 reports the three-panel decomposition of the HILI and ICC panel referenced in Section 4. Panel A decomposes the realized HILI spread into between-firm and within-firm components; Panel B reports a parallel Mundlak decomposition of the ICC panel; Panel C interacts within-firm IXI with NYSE-size tertiles in both the realized-return and ICC regressions. Read together, the three panels show that the realized HILI premium is essentially between-firm, the ICC wedge is two-margin, and the within-firm return-IXI relation is recoverable once size is conditioned on.

Table 11: Between-firm versus within-firm decomposition of the HILI premium. Panel A decomposes the value-weighted HILI sort into three variants: the headline sort on current lagged IXI (which blends between-firm and within-firm variation), a between-only sort in which each firm is assigned once to a quintile based on its time-average IXI (static across months), and a within-only sort in which each firm-month is sorted on the within-firm IXI deviation $\text{IXI}_{i,t} - \overline{\text{IXI}}_i$. Firms with at least 24 months of IXI observations enter the between-only and within-only sorts. Panel A alphas use Newey-West standard errors; Panels B and C use standard errors double-clustered by firm and year-month. Panel B reports a Mundlak decomposition of the ICC-IXI panel regression: ICC is regressed on the firm time-average IXI and on the within-firm IXI deviation, with the same size, value, and momentum controls as Table 4 and month fixed effects. Panel C reports the return-panel IXI coefficient interacted with NYSE size tertiles, identified from within-firm variation with firm and month fixed effects. Standard errors in parentheses. Sample: August 2002 through May 2024 (262 months).

Panel A: HILI spread by variation source (pp per month)

	Mean	CAPM α	FF5+MOM α
Headline HILI (sort on current lagged IXI)	0.430	0.642***	0.385**
		(0.232)	(0.174)
Between-only HILI (sort on firm-mean IXI)	0.414	0.600***	0.450**
		(0.227)	(0.223)
Within-only HILI (sort on firm-deviation IXI)	−0.042	0.021	−0.044
		(0.140)	(0.112)

Panel B: ICC Mundlak decomposition

	Between (firm-mean IXI)	Within (firm deviation)
Coefficient	−0.1054***	−0.0952***
	(0.00878)	(0.00663)
Controls: log me, log BM, MOM	Y	Y
Month FE	Y	Y
Observations	652,045	

Panel C: Within-firm return-panel IXI by NYSE size tertile

	Small tertile	Middle tertile	Big tertile
Implied within-firm IXI slope on return	−0.016	0.026**	0.062***
	(0.011)	(0.010)	(0.014)
Firm + Month FE		Y	
Observations		1,227,690	

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Panel A factor-model alphas are from regressions of each HILI series on the six right-hand-side factors. Panel B: the Mundlak decomposition adds the firm time-average IXI and the within-firm IXI deviation to the ICC panel regression with month FE and no firm FE; both components are identified in the same regression. Panel C: small, middle, big tertiles are defined on the NYSE-only distribution of lagged market equity each month. The small-tertile slope is the main-effect coefficient on IXI; the middle- and big-tertile slopes are the main effect plus the corresponding interaction coefficient, with standard errors computed by the delta method from the stored clustered variance-covariance matrix.

B Additional ICC robustness

Supporting robustness on the ex-ante implied-cost-of-capital analysis of Section 5.

B.1 Implied cost of capital construction

For each stock-month I solve for the internal rate of return r_e that equates the current market cap with the present value of projected future residual income, following Hou et al. (2012). Earnings forecasts for one, two, and three years ahead come from pooled cross-sectional regressions of net income on total assets, dividends, a dividend-payer dummy, current earnings, a negative-earnings dummy, and accruals. The regressions are estimated on a rolling ten-year window of Compustat annual fundamentals; regressors are winsorized at one and 99 percent each year. Return on equity is linearly interpolated from the year-three forecast to the 48-industry median in year twelve, where the industry median is computed on the same rolling window excluding loss-making firms. Book equity is projected forward through the clean-surplus relation under a constant firm-specific payout ratio. Stock-month ICC is winsorized at one and 99 percent each month.

B.2 Size-conditional ICC and the realized-vs-expected contrast

The ICC wedge survives size conditioning cleanly. Splitting each month at the NYSE median of lagged market equity and computing value-weighted ICC within each (size, IXI quintile) cell, the HILI-ICC spread is negative in both size buckets: within Big, the Q5-minus-Q1 spread averages -1.04 percentage points per year with $t = -10.29$; within Small, -2.87 percentage points with $t = -9.07$. Table 12 reports the underlying quintile-level value-weighted ICC means. Panel A gives the (size, IXI quintile) ICC cells; the within-bucket pattern is not strictly monotonic in IXI within either size group, but Q5 carries a lower ICC than Q1 in both Big and Small. Panel B gives the Q5-minus-Q1 spreads.

The ex-ante ICC wedge therefore sits differently from the realized-return premium. The

binary NYSE-median double sort on realized returns (Table IA.14 in Appendix IA.P) shows HILI alphas below conventional significance within both Big and Small under FF5-plus-momentum, and the three-tertile sort in Section 4 (Table IA.13) delivers significance mainly in the mid-cap tertile (FF5-plus-momentum alpha 0.28, $t = 2.19$). The ex-ante ICC wedge, by contrast, clears $|t| > 9$ in both size buckets. The realized-return premium is a joint size-and-indexing stylized fact, while the ex-ante expected-return compression is an IXI effect that is not absorbed by size conditioning. Both patterns are consistent with one channel discussed in Sammon (2025) and Bennett et al. (2020), and the ICC evidence is the clean within-size IXI result.

Table 12: Size-conditional ICC. Each month, stocks are split into Big (above NYSE median me_{lag}) and Small (at or below). Within each size bucket, stocks are sorted into 5 within-bucket quintiles on lagged IXI using monthly empirical 1/5 breakpoints. Panel A reports the time-series mean of the within-cell value-weighted ICC. Panel B reports the HILI-ICC spread (Q5 minus Q1) with Newey-West standard errors in parentheses. Sample: August 2002 through May 2024 under the Rule-I 2-month lag.

Panel A: Time-series mean of within-cell value-weighted ICC

Size bucket	Q1 (Low IXI)	Q2	Q3	Q4	Q5 (High IXI)	Q5 minus Q1
Big	0.0668	0.0603	0.0615	0.0610	0.0564	−0.0104
Small	0.0880	0.0809	0.0641	0.0562	0.0593	−0.0287

Panel B: HILI-ICC spread (Q5 minus Q1) with inference

Size bucket	Mean spread	Standard error	Months
Big	−0.0104***	(0.00101)	262
Small	−0.0287***	(0.00316)	262

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Within-cell VW ICC weights by me_{lag}. The spread is in decimal ICC units; a value of −0.020 is −200 basis points per year.

B.3 Industry-neutral and tech-excluded ICC wedge

Table 13 reports the Q5-minus-Q1 portfolio ICC spread and the firm-level panel coefficient under three cuts. Baseline is the headline construction from Section 5 (Table 4); the Baseline column reproduces the canonical −2.07 percentage-point spread ($t = -8.80$) and the R4 firm-panel coefficient of −0.0259 ($t = -5.37$) on the 2002-08 through 2024-05 sample. Industry-neutral forms the portfolio quintiles within FF12 industry and then value-weights across

industries; it adds FF12 × date fixed effects to the firm-panel specification. Tech-excluded drops FF12 BusEq (computer hardware, software, and electronic equipment, roughly 14 percent of HILI-universe firm-months and 23 percent of lagged market capitalization) from both the portfolio sort and the firm-panel regression. All three variants share the canonical HILI-universe quintile breakpoint construction, the ICC-valid sample for outcomes, and the Rule-I two-month lag.

The wedge attenuates modestly and retains sign and significance in every cut. The portfolio spread moves from −2.07 (baseline) to −1.91 (industry-neutral, 8 percent attenuation) to −1.75 (tech-excluded, 16 percent attenuation), each well outside a conventional significance band. The firm-panel coefficient moves from −0.0259 to −0.0238 (industry-neutral, 8 percent attenuation) to −0.0220 (tech-excluded, 15 percent attenuation), all with double-clustered t-statistics above 4. The ICC wedge is not a cross-industry compositional pattern and is not a tech-growth-optimism artifact; it is present in non-tech stocks alone with a −1.75 percentage-point value-weighted spread that clears $t = -7.23$, and it survives FF12 × date fixed effects that absorb all industry-month common variation.

B.4 IXI and Benchmarking Intensity in the ICC panel

Table 14 reports the direct ICC-panel horse race between IXI and Benchmarking Intensity (BMI) from Pavlova and Sikorskaya (2023) described in Section 3.4. The sample is the matched June stock-year panel from 2002 through 2018, the years for which the public BMI release is available and the paper's lagged Rule-I IXI signal can be aligned to the ICC panel. After merging, winsorizing the standard controls within year, and dropping fixed-effect singletons, the estimation sample contains 30,930 firm-year observations across 4,516 unique PERMNOs. On this sample, IXI and BMI are related but not interchangeable: the stock-year Pearson correlation is 0.53 and the Spearman correlation is 0.62. Both measures predict lower ICC in separate regressions. In the jointly-included specification, the IXI coefficient is −0.0293 with double-clustered $t = -5.67$ (significant at the one percent level), and the

Table 13: Industry-neutral and tech-excluded robustness of the ICC wedge. Panel A reports the value-weighted Q5-minus-Q1 ICC spread (percentage points per year) under the Rule-I two-month lag, 2002-08 through 2024-05. Newey-West t-statistics with six lags are in parentheses. Panel B reports the firm-level panel R4 coefficient from regressing ICC on lagged IXI plus log(ME), log(B/M), and 12-1 momentum with firm and month fixed effects; standard errors are double-clustered by firm and year-month. All three columns use the canonical HILI-universe quintile breakpoint construction from Section 5. Industry neutralization uses FF12 industries; tech exclusion drops FF12 BusEq.

	Baseline	Industry-neutral	Tech-excluded
Panel A: Portfolio Q5-Q1 value-weighted ICC spread (pp/year)			
Spread	−2.07	−1.91	−1.75
t	(−8.80)	(−9.93)	(−7.23)
Months	262	262	262
Panel B: Firm-panel R4 coefficient on lagged IXI (firm + month FE)			
$\hat{\beta}_{\text{IXI}}$	−0.0259	−0.0238	−0.0220
Double-clustered t	(−5.37)	(−5.07)	(−4.20)
Observations	575,617	575,617	468,957
Added fixed effects	none	FF12 × date	none

Notes: Sample runs from August 2002 through May 2024 (262 monthly cross-sections). Panel A value-weights ICC within each IXI quintile using lagged market capitalization. Newey-West t-statistics with six lags are in parentheses in Panel A. Panel B regresses firm-month ICC on lagged IXI plus standard controls; t-statistics are double-clustered by firm and year-month and shown in parentheses. Significance levels are the paper default: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$; all tabled coefficients clear the one percent level.

BMI coefficient is −0.0354 with $t = -2.64$ (significant at the five percent level). Adding BMI to the IXI specification attenuates IXI by about twenty percent without eliminating its significance; adding IXI to the BMI specification attenuates BMI by about sixteen percent with the same pattern. The result does not support a claim that IXI subsumes BMI; it shows that realized benchmarked ownership and benchmarking incentives carry overlapping but non-identical return-side information, with IXI more precisely estimated in the joint regression on this sample.

C Additional event-study robustness

Supporting robustness on the S&P 500 inclusion event study of Section 7.

Table 14: IXI and Benchmarking Intensity in the ICC panel.

	IXI only	BMI only	Both
IXI	-0.0367^{***}		-0.0293^{***}
	(0.0059)		(0.0052)
BMI		-0.0421^{***}	-0.0354^{**}
		(0.0138)	(0.0134)
Lagged log(ME), log(B/M), 12-2 momentum	Y	Y	Y
PERMNO fixed effects	Y	Y	Y
Year fixed effects	Y	Y	Y
Observations	30,930	30,930	30,930
Within R^2	0.2123	0.2128	0.2144

Notes: The sample is the matched June stock-year panel from 2002 through 2018, where BMI from Pavlova and Sikorskaya (2023) is available and the paper's lagged Rule-I IXI signal aligns to the ICC panel; 4,516 unique PERMNOs. The dependent variable is winsorized implied cost of capital (decimal, Hou-van-Dijk-Zhang construction). Controls are lagged log(ME), lagged log(B/M), and twelve-to-two-month momentum, each winsorized within year at the first and ninety-ninth percentiles. Standard errors in parentheses are double-clustered by PERMNO and year. On the estimation sample the Pearson correlation between IXI and BMI is 0.53 and the Spearman correlation is 0.62. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

C.1 Event-count accounting

The event-study and demand-gap sections reference several event counts as different filters are applied. Table 15 reconciles the sequence of sample sizes used across Section 7 and Section 8.

C.2 Matched-control robustness

Table 16 reports the short-window matched-control event-study regressions referenced in Section 7. Columns report the one-day, one-week, one-month, three-month, six-month, and twelve-month $\hat{\beta}_{T1}$ under both the market-adjusted matched-control outcome and the FF3-adjusted matched-control variant, on the Siblis and Q1 subsamples. The sample-construction details (up to five controls per event, FF12 industry plus size-decile plus book-to-market-quintile matching, ±12-month non-member exclusion, size-weighted control-pool CAR subtraction) are described in Section 7.

Table 15: Event-count accounting across the S&P 500 event-study and demand-gap samples

Count	Label	Filter
490	Raw Siblis roster	All S&P 500 additions between August 2002 and May 2024 with Siblis announcement dates.
472	CRSP-matched	Siblis events successfully matched to a CRSP permno by ticker/name; 96.3 percent match rate.
386	Full-coverage Siblis-matched	Matched events with complete announcement-window daily returns and IXI lookup months; working sample for the short-window event design.
384	Log(me)-only control sample	Matched events entering the short-window baseline with log me as the only control (Column 1 of Table IA.5).
382	Full-controls sample	Matched events with non-missing log me, log BM, and 12-1 momentum.
372	Siblis augmented-controls regression sample	Events entering the augmented-controls regression (Column 1 of Table 8); pooled Siblis headline is $\hat{\beta}_{T1} = 31.94$, $t = 2.44$.
326	Matched-control pool & demand-gap cohort	Subset of 382 full-controls events with at least one valid industry-size-B/M-matched control firm; also the cohort used for the Section 8 demand-gap design (326 treated events, up to five matched controls each).
318	Matched-control regression sample at $h = 1$	Subset of the matched-control pool with complete controls at the one-day horizon (Table 16).
157	Q1 best-timing subsample	Smallest quartile of the Siblis-proxy timing gap; reported in Column 2 of Table 8 as a timing-accuracy sharpening check.

Notes: Counts are unique event-level. Siblis announcement dates are from Siblis Research; fallback proxy dates are the standard five-trading-day offset from the CRSP effective date (used for 18 of 490 events, reported separately as a robustness subsample). Each row inherits the filters of the row above. The 326-event matched-control cohort is the working sample for the Section 8 demand-gap design.

Table 16: Matched-control robustness for the event-study decay profile. Each cell reports $\hat{\beta}_{T1}$ from a regression of the CAR outcome on $T1$ plus log me, log BM, and momentum (12-1). Columns 1-2 use market-adjusted CAR. Columns 3-4 subtract the size-weighted mean market-adjusted CAR of up to five industry-size-BM matched control firms per event. Columns 5-6 report the FF3-adjusted matched-control variant: treated and control returns are first residualized on MKT, SMB, and HML using daily FF3 loadings estimated over effective days −252 through −21, then the same matched-control differencing is applied.

	Market-adjusted		Matched-control		FF3-matched	
h (TD)	Siblis	Q1	Siblis	Q1	Siblis	Q1
1 (1d)	31.94**	56.98***	18.14	48.84***	21.55**	44.25**
5 (1w)	27.54**	35.60*	11.41	32.14*	18.65*	32.87*
21 (1mo)	18.09	59.18**	19.06	51.01	16.53	41.57
63 (3mo)	10.16	55.99	4.32	17.74	−5.19	−10.57
126 (6mo)	41.28	106.81	17.60	−29.66	−25.28	−127.08
252 (12mo)	26.67	176.24**	−8.57	42.20	−79.61	−188.17
Observations ($h = 1$)	372	157	318	132	318	132

Notes: The sample covers S&P 500 additions between August 2002 and May 2024 (Siblis Research roster). All samples are Siblis-matched; Q1 is the best-timing quartile. $T1$ enters in decimal units, so coefficients are CAR percentage points per unit increase in $T1$. The dependent variable is the market-adjusted CAR through effective day $+h$. The table uses a coefficient-only display for compactness across six specifications and six horizons; standard errors are clustered by announcement year. Preservation ratios at $h = 1$ relative to the augmented-controls headline (31.94 Siblis, 56.98 Q1): market-adjusted matched-control is 56.8 percent (Siblis) and 85.7 percent (Q1); FF3-matched is 67.5 percent (Siblis) and 77.7 percent (Q1). The FF3-matched Siblis result clears the pre-specified 60 percent preservation threshold. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

D Flow-sustainability tests

Table 17 reports the four conditional tests referenced in Section 6.6. Panel A compares the mean value-weighted HILI return in months of weak aggregate passive flow to the remaining months. The flow series is the Morningstar estimated-net-flow scaled by lagged passive TNA (rather than by CRSP market capitalization as in the main aggregate regressions of Section 6); the two scaling choices give similar rankings of low-flow months but are not the identical object. The bottom quintile of contemporaneous aggregate passive flow contains 52 months concentrated in the 2008-2009 financial crisis, spring 2020 during the COVID shock, and the 2022 bear market. In these 52 months, mean HILI is −0.03 percentage points per month; in the remaining 209 months, mean HILI is +0.54. The unconditional difference is −0.57 percentage points with a Welch two-sample t-statistic of 1.33. The statistical significance is marginal given only 52 low-flow months, but the economic magnitude is large,

because essentially the entirety of the HILI premium disappears in weak-flow months.

Panel B regresses HILI on the bottom-quintile indicator with the full set of aggregate controls used in Section 6; the residual low-flow coefficient is -0.26 with $t = -0.62$, pointing in the predicted direction but absorbed by the AggAUM level control that already captures most of the flow-related signal. AggAUM and the low-flow indicator are substantially redundant given that low-flow months are concentrated in episodes of broader aggregate passive-share stagnation, and a single-month test has limited statistical power at the 52-observation window.

Panel C sharpens the reading through a non-overlapping six-month-window test. The 262-month sample is grouped into 43 non-overlapping six-month windows classified by within-window sum of ΔAggAUM. The nine bottom-quintile deceleration windows deliver a cumulative HILI of -5.03 percentage points; the remaining 34 windows deliver 4.73. The 9.75-percentage-point differential translates to approximately 1.6 percentage points per month, with Welch $t = 3.44$ ($p = 0.003$). The multi-month aggregation lifts statistical power above the single-month Panel A test, and non-overlapping windows avoid the inflated-inference problem of rolling-window tests.

Panel D reports a symmetry test on the first-difference specification of Section 6. Augmenting that regression with $\mathbf{1}\{\Delta\text{AggAUM} < 0\}$ and its interaction with ΔAggAUM, the ΔAggAUM slope in positive-innovation months is 9.56 ($t = 2.72$) and the implied slope in negative-innovation months is 15.85. The joint Wald test gives $\chi^2(2) = 0.71$ with $p = 0.70$, so the linear response to ΔAggAUM is the same in positive- and negative-innovation months up to sampling variability. The co-movement is therefore bi-directional and linear, consistent with the reduced-form specification used throughout Section 6 and with the scenario projection reported below.

Table 17: HILI performance in weak aggregate passive-flow episodes. Panel A reports the mean value-weighted HILI return in the bottom quintile of contemporaneous aggregate passive-flow months versus the remaining months and in outright negative-flow months versus positive-flow months. Panel B reports a monthly time-series regression of HILI on an indicator for bottom-quintile flow months and on the Section 6 controls (ER^{HILI}, ΔEF^{HILI}, AggAUM). Panel C reports the multi-month-episode test on non-overlapping six-month windows. Panel D reports the symmetry-of-response test obtained by augmenting the Section 6 first-difference specification with an indicator for ΔAggAUM < 0 and its interaction with ΔAggAUM. Newey-West standard errors in parentheses. Sample: August 2002 through May 2024 (262 months).

Panel A: HILI mean by flow-episode status (monthly)	
Full sample mean (all months, $n = 262$)	0.433
Bottom-quintile (low-flow) months ($n = 52$)	−0.032
Remaining (high-flow) months ($n = 209$)	0.538
Difference (low − high)	−0.571
t statistic (Welch two-sample t)	1.33
Negative-flow months ($n = 52$)	−0.032
Positive-flow months ($n = 209$)	0.538
Difference (neg − pos)	−0.571
t statistic	1.33
Panel B: HILI on low-flow indicator with controls (monthly)	
$\mathbf{1}\{$low-flow month$\}$ coefficient	−0.259
	(0.421)
Observations	261
Panel C: Six-month window deceleration episodes	
Non-overlapping 6-month windows ($n = 43$)	
Bottom-quintile deceleration windows, cumulative HILI ($n = 9$)	−5.03 pp
Remaining windows, cumulative HILI ($n = 34$)	4.73 pp
Difference (decel. − other)	−9.75 pp
t statistic (Welch two-sample t)	3.44^{***}
p-value	0.003
Panel D: Asymmetric-flow augmented first-difference spec	
ΔAggAUM coefficient (positive-innovation months)	9.562^{**}
	(3.520)
ΔAggAUM $\times$ $\mathbf{1}\{\Delta$AggAUM $< 0\}$ interaction	6.286
	(10.488)
Implied ΔAggAUM slope in negative-innovation months	15.848
Joint Wald test of asymmetry ($\chi^2(2)$, Newey-West)	0.71
Wald p-value	0.70

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Low-flow months are defined as the bottom quintile of contemporaneous aggregate passive flow (Morningstar monthly aggregate). Negative-flow months are months with aggregate passive net outflow. Panel C windows are non-overlapping six-month blocks; deceleration windows are the bottom quintile of the sum of ΔAggAUM across the window. Panel D augments the Section 6 first-difference specification with an indicator for ΔAggAUM < 0 and its interaction with ΔAggAUM; the joint Wald test tests whether both augmentation coefficients are zero.

D.1 Stationarity and bandwidth robustness of the FD regression

The first-difference aggregate regression of Section 6 survives three time-series diagnostics on input stationarity, Newey-West bandwidth choice, and the sampling distribution of the test statistic. The specification is $\text{HILI}_t = \alpha + \beta_\Delta \Delta\text{AggAUM}_t + \gamma_{\text{ER}}\text{ER}_t^{\text{HILI}} + \gamma_{\Delta\text{EF}}\Delta\text{EF}_t^{\text{HILI}} + \gamma_{\text{flow}}\text{AggFlow}_t^{\text{lag}} + \varepsilon_t$ over the 261-month sample September 2002 through May 2024 (one observation is lost to the first difference).

Unit-root and stationarity tests on the four input series address the spurious-regression concern that the levels specification of Section 6 acknowledges. The ADF test with drift (AIC-selected lag) rejects the unit-root null at the one percent level on HILI (ADF statistic -9.0), AggFlow (ADF statistic -10.7), and ΔAggAUM (ADF statistic -8.3). The KPSS test with intercept fails to reject the stationarity null on HILI at the ten percent level and rejects only at the five percent level on AggFlow and ΔAggAUM. The level series AggAUM fails the ADF test and is rejected by KPSS at the one percent level, consistent with its near-monotone drift over the sample and the Section 6 caveat that the level specification cannot separate the AggAUM coefficient from a deterministic time trend. The first-difference specification operates on HILI and ΔAggAUM, both of which pass the ADF stationarity test.

Table 18 reports Newey-West bandwidth sensitivity on β_Δ across lag choices of 0 (OLS), 3, 6 (paper default), 9, and 12 months. The coefficient is 11.39 at every bandwidth, matching the Section 6 headline. The Newey-West t-statistic ranges narrowly from 4.36 to 4.54 across the five bandwidths, clearing the one percent significance bar on every choice. The headline result is insensitive to the bandwidth selection.

A stationary block bootstrap under the null of no relation between HILI and ΔAggAUM, conditional on the three controls, addresses a weaker but complementary concern: that small-sample distortion in the NW(6) sampling distribution of the observed t-statistic could be driving the significance. I draw 2,000 block-bootstrap replications of the restricted-model residuals with a six-month block length, re-estimate the FD specification on each replication, and tabulate the sampling distribution of the NW(6) t-statistic on β_Δ. The observed NW(6)

Table 18: Newey-West bandwidth sensitivity on $\beta_{\Delta\text{AggAUM}}$ in the FD regression of Section 6

	OLS	NW(3)	NW(6)	NW(9)	NW(12)
$\hat{\beta}_{\Delta}$	11.39	11.39	11.39	11.39	11.39
Standard error	2.56	2.54	2.53	2.51	2.61
t-statistic	4.46	4.48	4.51	4.54	4.36
Two-sided p-value	<0.001	<0.001	<0.001	<0.001	<0.001

Notes: Dependent variable is monthly realized HILI return (percent per month). Regressor of interest is ΔAggAUM_t; controls held fixed across columns are $\text{ER}_t^{\text{HILI}}$, $\Delta\text{EF}_t^{\text{HILI}}$, and lagged aggregate passive flow, following the first-difference specification in Section 6. Sample runs from September 2002 through May 2024 (261 monthly observations; one month lost to the first difference). OLS column uses classical standard errors; NW(k) columns use Newey-West standard errors with the stated lag. The coefficient is identical to four decimal places across bandwidths because only the variance estimator changes.

$t = 4.51$ sits well outside the bootstrap 95 percent band of $[-2.02, +2.01]$, and the two-sided bootstrap p-value is 0.0005. The bootstrap null is rejected at the one-in-two-thousand level.

The three diagnostics suggest that the FD coefficient and its t-statistic are not artifacts of a non-stationary input, a bandwidth choice, or a small-sample distortion in the sampling distribution. What the FD regression cannot do, and does not try to do, is identify a structural elasticity of HILI to ΔAggAUM. The coefficient is a reduced-form time-series co-movement, and the counterfactual in Figure 4 should be read as a statistical decomposition of the realized monthly HILI path into a component attributable to ΔAggAUM innovations and a residual, not as a causal counterfactual.

E Derivation of the equilibrium price, expected-return, and shock equations

Derivation of Equations (2), (6), and (7) from the market-clearing condition, filling in the steps skipped in Section 2.

E.1 Agent demands

Direct investors. A direct investor with CARA utility $U(W) = -e^{-\gamma W}$ and initial wealth W_0 holds portfolio x_D of the N risky assets. Terminal wealth is $W_D = W_0 + x_D'(D - S)$.

Under joint normality of D, the optimal portfolio solves

$$\max_{x_D} \; E[W_D] - \tfrac{\gamma}{2}\mathrm{Var}(W_D) = x_D'(\mu - S) - \tfrac{\gamma}{2}x_D'\Sigma x_D.$$

The first-order condition $(\mu - S) - \gamma\Sigma x_D = 0$ yields

$$x_D = \frac{1}{\gamma}\Sigma^{-1}(\mu - S). \tag{21}$$

Active managers. An active manager benchmarked to index i receives compensation $w_i = aR_i + b(R_i - R_i^B) + c$, where $R_i = x_i'(D - S)$ is fund return and $R_i^B = \omega_i'(D - S)$ is benchmark return. Substituting, terminal wealth is

$$W_i = W_0 + c + (a + b)\, x_i'(D - S) - b\,\omega_i'(D - S).$$

The mean-variance objective is

$$\max_{x_i} \; (a + b)\, x_i'(\mu - S) - b\,\omega_i'(\mu - S) - \tfrac{\gamma}{2}\big[(a + b)x_i - b\omega_i\big]'\Sigma\big[(a + b)x_i - b\omega_i\big].$$

Taking the first-order condition with respect to x_i:

$$(a + b)(\mu - S) - \gamma(a + b)\Sigma\big[(a + b)x_i - b\omega_i\big] = 0 \implies (a + b)x_i - b\omega_i = \tfrac{1}{\gamma}\Sigma^{-1}(\mu - S),$$

which gives

$$x_i^A = \frac{1}{\gamma(a + b)}\Sigma^{-1}(\mu - S) + \frac{b}{a + b}\omega_i. \tag{22}$$

As $b \to \infty$, the manager's penalty for benchmark deviation dominates and $x_i^A \to \omega_i$, the benchmark portfolio. This is the passive-manager limit $x_i^P = \omega_i$.

E.2 Market clearing and equilibrium prices

Aggregating across agent types with population shares λ_D, λ_i^A, and λ_i^P, market clearing requires

$$\lambda_D x_D + \sum_{i=1}^{I} \lambda_i^A x_i^A + \sum_{i=1}^{I} \lambda_i^P x_i^P = \bar{x}.$$

Substituting Equation (21), Equation (22), and $x_i^P = \omega_i$:

$$\lambda_D \cdot \frac{1}{\gamma}\Sigma^{-1}(\mu - S) + \sum_i \lambda_i^A \left[\frac{1}{\gamma(a+b)}\Sigma^{-1}(\mu - S) + \frac{b}{a+b}\omega_i\right] + \sum_i \lambda_i^P \omega_i = \bar{x}.$$

Collecting the $\Sigma^{-1}(\mu - S)$ terms:

$$\underbrace{\left[\lambda_D + \sum_i \frac{\lambda_i^A}{a+b}\right]}_{A^{-1}} \frac{1}{\gamma}\Sigma^{-1}(\mu - S) = \bar{x} - \frac{b}{a+b}\sum_i \lambda_i^A \omega_i - \sum_i \lambda_i^P \omega_i.$$

Let $A \equiv \left[\lambda_D + \sum_i \lambda_i^A/(a+b)\right]^{-1}$ and $B_n \equiv \frac{b}{a+b}\sum_i \lambda_i^A \omega_i(n) + \sum_i \lambda_i^P \omega_i(n)$. Multiplying both sides by $\gamma A \Sigma$ gives $\mu - S = \gamma A \Sigma(\bar{x} - B)$, which is Equation (2).

E.3 Expected returns and the cross-sectional wedge

From Equation (2), the per-share expected excess payoff $E[D_n - S_n] = \mu_n - S_n$ satisfies

$$E[D_n - S_n] = \gamma A \sum_{m=1}^{N} \sigma_{nm}\,(\bar{x}_m - B_m), \tag{23}$$

where σ_{nm} is the (n, m) element of Σ. Under the diagonal leading case $\Sigma = \mathrm{diag}(\sigma_n^2)$, equation (23) reduces to $E[D_n - S_n] = \gamma A \sigma_n^2(\bar{x}_n - B_n)$. Dividing the parenthesized term by $\bar{x}_n$ and introducing the benchmarked ownership share $\mathcal{I}_n \equiv B_n/\bar{x}_n$ yields

$$E[D_n - S_n] = \gamma A \sigma_n^2 \bar{x}_n \big(1 - \mathcal{I}_n\big),$$

which is Equation (6). The expected excess payoff is the unconstrained baseline $\gamma A\sigma_n^2\bar{x}_n$ scaled by the residual share fraction $(1-\mathcal{I}_n)$. Stocks with a larger benchmarked ownership share have a smaller residual share, command a higher equilibrium price, and carry a lower ex-ante expected return.

In the general-Σ case, equation (23) introduces cross-stock terms $\sigma_{nm}(\bar{x}_m - B_m)$ for $m \neq n$. The empirical specifications include size, value, momentum, and industry controls to address the leading factor components of this cross-stock channel, and the diagonal case carries the cross-sectional interpretation.

E.4 Unanticipated benchmark-demand shocks

To derive Equation (7), extend the economy by adding a period $t = -1$ at which B_n is stochastic and the market's pricing expectation of its $t = 0$ realization is $E^*_{-1}[B_n]$. Holding A, Σ, μ, and $\bar{x}$ fixed across the two dates, B_n realizes at $t = 0$ and the equilibrium described in Section 2.1 obtains. Because equation (2) is linear in B, equilibrium prices at each date satisfy

$$S_{n,-1} = \mu_n - \gamma A\sigma_n^2\left(\bar{x}_n - E^*_{-1}[B_n]\right), \qquad S_{n,0} = \mu_n - \gamma A\sigma_n^2\left(\bar{x}_n - B_n\right),$$

in the diagonal leading case. Differencing,

$$S_{n,0} - S_{n,-1} = \gamma A\sigma_n^2\left(B_n - E^*_{-1}[B_n]\right) = \gamma A\sigma_n^2\bar{x}_n\left(\mathcal{I}_n - E^*_{-1}[\mathcal{I}_n]\right),$$

where the second equality uses $\mathcal{I}_n = B_n/\bar{x}_n$ and the fact that shares outstanding $\bar{x}_n$ is not stochastic. This is Equation (7). The realized revaluation is proportional to the unanticipated component of benchmarked ownership; the anticipated component is priced into $S_{n,-1}$ and does not contribute to realized returns at $t = 0$. A closely related single-step result for the aggregate B_n appears in Appendix A of Pavlova and Sikorskaya (2023).

The empirical Pástor-Stambaugh-Taylor decomposition in Section 6.5 separates realized HILI into an ex-ante expected component, proxied by the ICC wedge of Section 5, and

an unanticipated residual that co-moves with aggregate passive-share growth. This split is empirically analogous to the decomposition in Equation (7) under a physical-expectation proxy for the pricing expectation $E^*_{-1}[\cdot]$; the two differ by a second-order risk-premium shift in the CARA-normal setting that does not affect the qualitative comparative statics.

F Long-horizon IV diagnostic appendix

This appendix retains the long-horizon firm-by-cohort IV decomposition as a diagnostic reference. The headline Section 8 evidence uses the cross-sectional demand-gap design rather than the IV decomposition reported here; readers who want the underlying IV machinery, event-study coefficients, and horizon and period breakdowns will find them below.

The IV log(SHROUT) row in particular should be read with caution: as discussed in Section 8, the unadjusted CRSP log(SHROUT) measure is partly mechanical because of stock splits, and the split-adjusted version is statistically flat at $[+25, +36]$. I therefore do not interpret S&P inclusion as evidence on equity issuance.

F.1 Event-study coefficients

Table 19 tabulates the event-study coefficients underlying the long-horizon IV diagnostic. Pre-event IXI coefficients are statistically zero across all three pre-event bins, and the post-event jump is significant at every horizon through $[+25, +36]$. Pre-event log(ME) coefficients rise from 0.059 to 0.268, reflecting the well-documented S&P committee-selection pattern (Chang et al., 2015; Greenwood and Sammon, 2025); the post-event peak is 0.368 at $[+2, +6]$, and the level at $[+25, +36]$ relative to the omitted bin is 0.140.

F.2 Horizon decomposition of the IV effect

Table 20 decomposes the IV estimate by post-event horizon bucket ($[0, 2]$, $[3, 6]$, $[7, 12]$, $[13, 24]$, $[25, 36]$) on the four main outcomes log(ME), log(BE), log(SHROUT), and log(M/B).

Table 19: Event-study coefficients of treated × event-month-bin indicators on log(ME) (reduced form, left column) and IXI (first stage, right column) in the S&P 500 DiD panel. Baseline omitted bin is $[-24, -19]$ months relative to inclusion. The key identification check is that the IXI pre-event coefficients are statistically indistinguishable from zero while the post-event coefficients are significantly positive and persistent.

Event-month bin relative to inclusion	log(ME)	IXI
[-18, -13]	0.059***	−0.00130
	(0.015)	(0.00196)
[-12, -7]	0.148***	−0.00154
	(0.022)	(0.00309)
[-6, -2]	0.268***	−0.00009
	(0.027)	(0.00395)
[+2, +6]	0.368***	0.03208***
	(0.032)	(0.00405)
[+7, +12]	0.318***	0.03512***
	(0.034)	(0.00444)
[+13, +24]	0.242***	0.02970***
	(0.040)	(0.00488)
[+25, +36]	0.140***	0.02083***
	(0.048)	(0.00517)
Observations	102,352	102,352

Notes: The S&P 500 inclusion panel covers 2002-2024. Regressions include firm-by-cohort and year-quarter fixed effects. Standard errors are in parentheses and are double-clustered by firm and year-quarter. Event-month bins are measured relative to the S&P 500 inclusion month; the baseline omitted bin is $[-24, -19]$, normalized to zero. The table reports the IXI first-stage path and the log(ME) reduced-form path; the log(SHROUT) reduced-form event study is not tabulated separately. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

The IV log(ME) effect peaks at 3-6 months and fades through 36 months. The unadjusted log(SHROUT) coefficients are positive across horizons, but as noted above, this measure is partly mechanical because of stock splits.

Table 20: Horizon decomposition of S&P 500 inclusion IV on valuation outcomes, using IXI as the endogenous passive-ownership measure. Each column restricts post-event observations to the stated horizon bucket and runs the IV specification with firm-by-cohort and year-quarter fixed effects, double-clustered by firm and year-quarter.

Horizon (months)	3-6	6-12	12-24	24-36
log(ME)	4.759^{***}	4.236^{***}	4.275^{***}	3.852^{*}
	(0.961)	(0.908)	(1.228)	(2.250)
log(BE)	2.734^{**}	5.199^{***}	6.812^{***}	9.357^{***}
	(1.084)	(1.347)	(1.930)	(3.355)
log(SHROUT)	1.551^{**}	1.537^{**}	1.414^{*}	2.339^{*}
	(0.685)	(0.635)	(0.787)	(1.355)
log(M/B)	2.407^{**}	-0.606	-2.172	-5.744^{*}
	(1.009)	(1.041)	(1.512)	(2.932)
First-stage F	71.1	73.4	43.9	17.9

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Standard errors in parentheses, double-clustered by firm and year-quarter.

F.3 Period heterogeneity by event-year subsample

Table 21 reports the log(ME) and log(SHROUT) IV estimates by event-year subsample, with the 2019-2024 window split into 2019-2021 and 2022-2024 to align with Greenwood and Sammon (2025). Pre-2013 coefficients are positive but insignificant. The 2013-2018 coefficient is 5.94 ($t = 2.03$). The 2019-2021 subsample covers only 43 events and delivers a first-stage F of 2.2; the IV point estimate of 5.70 is therefore uninformative rather than null. The 2022-2024 subsample covers 22 events with a much stronger first stage ($F = 47.5$) and yields a significant IV log(ME) coefficient of 3.33 ($t = 2.92$). The unadjusted log(SHROUT) IV varies across subsamples but is again subject to the split-adjustment caveat above.

Table 21: Period subsamples of S&P 500 inclusion IV on log market capitalization and log shares outstanding, using IXI. The 2019-2024 window is split into 2019-2021 and 2022-2024 to align with Greenwood and Sammon (2025). The log(ME) effect is significant in 2013-2018 and in 2022-2024 and underpowered in the 43-event 2019-2021 window (first-stage $F = 2.2$); magnitudes moderate with passive-capital saturation while the direction persists.

Period	2002-2007	2008-2012	2013-2018	2019-2021	2022-2024
log(ME)	3.167	3.969	5.936**	5.703	3.326***
	(2.025)	(2.549)	(2.919)	(4.851)	(1.141)
log(SHROUT)	3.165*	1.332	2.415	−0.516	−0.501
	(1.735)	(1.795)	(1.886)	(2.365)	(0.306)
First-stage F	25.7	17.0	9.9	2.2	47.5
Events	167	167	94	43	22
Observations	32,437	31,756	36,010	16,419	5,982

Notes: Each column estimates the same 2SLS specification on the event-year subsample named in the header, with firm-by-cohort and year-quarter fixed effects. The endogenous regressor is IXI and the excluded instrument is Treated × Post. Standard errors are in parentheses and are double-clustered by firm and year-quarter. First-stage F is reported by subsample. The 2019-2021 first-stage F of 2.2 is below the weak-instrument threshold; the log(ME) point estimate of +5.70 in that window is therefore best read as uninformative rather than as null evidence. The 2022-2024 subsample has a stronger first stage ($F = 47.5$) and a significant log(ME) IV despite a smaller event count (22), reflecting that each post-2021 inclusion event moves a larger share of a stock's ownership as passive-capital scale has grown. First-stage F values are Kleibergen-Paap statistics; the Stock-Yogo 10 percent critical value for a single endogenous regressor is $F = 10$. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

F.4 IV decomposition table (diagnostic)

Table 22 reports the original long-horizon IV decomposition on log(ME), log(BE), log(SHROUT), log(PRC), log(AT), and log(M/B) across IXI, IXI_{pass}, and the fund-branch sub-measure. The decomposition is retained here as a diagnostic for the long-horizon firm-by-cohort design; the headline Section 8 evidence is the cross-sectional demand-gap design. The unadjusted log(SHROUT) row in particular should not be read as evidence on equity issuance, for the split-adjustment reason discussed above.

Table 22: S&P 500 inclusion IV decomposition on valuation components. Each cell reports a 2SLS coefficient from a separate outcome regression. The endogenous regressor is the IXI measure named in the column header; the excluded instrument is Treated × Post.

	IXI (main)	IXI_{pass}	Fund branch
log(ME)	4.298***	7.977***	3.562***
	(1.170)	(2.167)	(0.896)
log(BE)	5.970***	11.239***	4.955***
	(1.651)	(2.959)	(1.216)
log(SHROUT)	1.654**	3.069**	1.370**
	(0.803)	(1.510)	(0.654)
log(PRC)	2.944***	5.464***	2.440***
	(1.073)	(1.947)	(0.849)
log(AT)	2.750**	5.114**	2.282***
	(1.079)	(2.039)	(0.869)
log(M/B)	−1.539	−2.899	−1.278
	(1.340)	(2.484)	(1.092)
First-stage F	52.2	108.7	305.3
Observations	122,604	122,604	122,604

Notes: The monthly S&P 500 inclusion panel covers 2002-2024 and includes firm-by-cohort and year-quarter fixed effects. The endogenous regressor is the IXI measure named in the column header; the excluded instrument is Treated × Post. Each row is estimated on its own non-missing sample. Standard errors are in parentheses and are double-clustered by firm and year-quarter. First-stage F is the Kleibergen-Paap first-stage statistic for the excluded instrument; all three values exceed the Stock-Yogo 10 percent maximal IV size critical value of 10. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Internet Appendix for "The Price of Passive Ownership"

IA.A Announcement-informativeness evidence: a null direct test

A parallel strand of the passive-ownership literature studies firm-level informativeness around earnings announcements. Sammon (2025) shows, using Russell 1000 / 2000 reconstitution, that higher passive ownership reduces the fundamental information incorporated into prices ahead of announcements. Section 5 cites this work as part of the literature background for the ICC wedge. Because the direction of the Sammon mechanism is the same as the paper's expected-return result, a natural follow-on is whether a Sammon-style announcement test is informative on IXI directly. This appendix reports that test and the reason the paper does not treat it as a mechanism pillar.

I build a firm-quarter panel of quarterly earnings announcements over 2001Q3-2023Q4, using IBES `actu_epsus` for announcement dates and times (after-close announcements are rolled to the next trading day), CRSP daily returns for event windows, and the IXI panel for the covariate. The primary outcome is $\text{UNINF} = |\text{CAR}^{[0,+1]}|/\sigma^{\text{pre}}$, the two-day absolute market-adjusted announcement return scaled by pre-event volatility measured over the $[-60, -6]$ trading-day window, the Sammon-analog reading of low pre-announcement informativeness. The analysis sample restricts to `shrcd` $\in \{10, 11\}$, `exchcd` $\in \{1, 2, 3\}$, and announcements with IXI and the standard controls non-missing, yielding 285,919 firm-quarter observations on 9,246 unique permnos. The panel specification is

$$\text{UNINF}_{i,q} = \alpha_i + \lambda_q + \beta\, \text{IXI}_{i,q-1} + \Gamma' X_{i,q} + \varepsilon_{i,q},$$

with firm and calendar-quarter fixed effects, controls for $\log(me)$, twelve-to-two-month momentum, and pre-event turnover, and standard errors two-way clustered by firm and calendar quarter.

The panel coefficient is positive as predicted but does not clear conventional significance on the primary two-day outcome. The full-sample coefficient is 0.231 with clustered $t = 1.50$; it strengthens to 0.338 ($t = 2.12$) on the three-day robustness window, to 0.791 ($t = 4.85$) on the pre-2020 subsample, and to 1.60 ($t = 2.98$) and 0.64 ($t = 2.19$) on the Small and Mid NYSE size terciles respectively, while the Big tercile is statistically zero. The cross-section of results is consistent with the Sammon prediction but does not deliver a clean full-sample rejection of the null on the primary two-day specification.

The S&P 500 inclusion difference-in-differences on the same outcome runs on the same matched-control panel used in Section 8, with 326 treated events, 1,592 matched-control pairs, the $[-36, +36]$-month window, and the $[-1, +1]$-month anticipation window excluded. The primary DiD coefficient on Treated $\times$ Post is -0.044 ($t = -0.43$); with the standard controls it is 0.018 ($t = 0.18$); the three-day-window DiD is -0.074 ($t = -0.71$); the pre-2020 subsample DiD is 0.065 ($t = 0.60$). The event-study decomposition across six month-offset bins also returns null pre-event and post-event coefficients on UNINF $\times$ Treated. The matched-control panel delivers a tight IXI first stage in Section 8, so the DiD null on this outcome is not attributable to weak identification.

The panel-FE correlation and the null inclusion DiD do not jointly identify a within-firm causal informativeness channel; the panel correlation is most consistent with cross-sectional sorting that the firm and quarter fixed effects only partially absorb. The paper therefore does not treat announcement informativeness as a mechanism pillar.

IA.B IXI and the size and value factors

This appendix places the IXI-sorted HILI factor in the broader size- and value-factor discussion and reports the two-factor decomposition of HML and SMB with HILI. I regress HML and SMB on the excess market return with and without HILI, where HILI is the value-weighted top-minus-bottom IXI quintile spread from Section 4. Table IA.1 reports the decomposition over the full IXI-aligned sample (Panel A), the legacy 2011-2021 window from

an earlier working-paper version of this analysis (Panel B), and a 2013-through-May-2024 modern-passive-era window (Panel C).

Table IA.1: Two-factor decomposition of HML and SMB using the market and HILI. HML and SMB are the Fama-French factors. HILI is the monthly high-minus-low IXI value-weighted quintile spread under the main specification from Section 4. All series are in percentage points per month. Panel A uses the full-aligned sample (August 2002 through May 2024). Panel B uses the 2011-2021 window for direct comparison with the legacy pre-revision draft. Panel C uses the 2013-modern window (January 2013 through May 2024), which starts after the post-crisis recovery period and extends to the Rule-I safe end. Newey-West standard errors in parentheses.

	HML		*SMB*	
	(1)	(2)	(3)	(4)
Panel A: Full sample, August 2002 through May 2024				
Constant	−0.106	−0.159	−0.090	0.128
	(0.250)	(0.239)	(0.157)	(0.152)
MKT-Rf	0.092	0.111	0.205***	0.125***
	(0.068)	(0.068)	(0.036)	(0.040)
HILI		0.085		−0.348***
		(0.074)		(0.054)
Observations	262	262	262	262
R^2	0.018	0.023	0.124	0.257
Panel B: 2011-2021 (legacy window)				
Constant	−0.394	−0.445	−0.364*	−0.049
	(0.368)	(0.379)	(0.213)	(0.239)
MKT-Rf	0.116	0.124	0.235***	0.184***
	(0.085)	(0.089)	(0.049)	(0.052)
HILI		0.051		−0.316***
		(0.073)		(0.070)
Observations	132	132	132	132
R^2	0.025	0.028	0.135	0.253
Panel C: 2013-modern window (January 2013 through May 2024)				
Constant	−0.167	−0.310	−0.368	−0.032
	(0.439)	(0.398)	(0.236)	(0.241)
MKT-Rf	0.032	0.057	0.202***	0.142***
	(0.099)	(0.093)	(0.051)	(0.054)
HILI		0.162		−0.381***
		(0.099)		(0.061)
Observations	137	137	137	137
R^2	0.002	0.020	0.102	0.271

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Columns (1) and (3) regress HML and SMB on the excess market return. Columns (2) and (4) add HILI. HILI is constructed under the main HILI specification (full IXI, Rule-I + 2 month lag, N=5 within-month quintile breakpoints, value-weighted, raw returns).

HILI co-varies strongly and negatively with SMB in every window. Adding HILI to the CAPM regression of SMB on the market loads HILI at -0.348 with $t = -6.43$ in the full sample, moves the SMB CAPM alpha from -0.364 ($t = -1.71$) to -0.049 ($t = -0.20$) in the legacy window, and loads HILI at -0.381 ($t = -6.29$) in the modern window. The right reading of the SMB pattern is mechanical. IXI-sorted HILI is a large-cap-tilted long-short by construction, so it must share variation with a small-minus-big factor in any sample where small caps underperform large caps. A formal size-neutral test confirms this: averaging the within-NYSE-median-big HILI and the within-small HILI into a 0.5/0.5 size-neutral factor drops the HILI-on-SMB loading to -0.06 ($t = -1.04$), close to zero. The HILI-SMB co-movement is a composition effect of the large-cap tilt, not an independent absorption of SMB by IXI.

The value side of the legacy decomposition does not survive under the revised IXI. In the full sample, the HML CAPM alpha of -0.106 ($t = -0.42$) is essentially unchanged when HILI is added (0.085, $t = 1.15$). The null is consistent across all three windows. The legacy finding that HILI explained HML's underperformance was estimated on an older IXI construction that loaded negatively on HML; the revised IXI loads positively (0.172, $t = 2.66$) because best-fit benchmark assignment reallocates growth-declared funds to the Growth subset and because the held-stock-normalized 13F non-fund branch adds benchmarked capital from pensions and insurers whose portfolios tilt toward large-cap value names. The revised measure spreads high-IXI attribution more evenly across the value-growth dimension than the legacy declared-benchmark-only fund-sector measure.

The reading across both factors is that HILI shares a large-cap composition dimension with SMB but captures a phenomenon distinct from HML. This is consistent with Section 4's joint size-and-indexing interpretation of the HILI spread and with Li (2022)'s account of how factor returns blend risk premia with demand-driven components in the passive era.

IA.C Magnificent Seven exclusion

Table IA.2 reports the HILI spread with the seven largest mega-cap stocks excluded (AAPL, MSFT, Alphabet, AMZN, META, NVDA, TSLA). Excluding these seven names reduces the mean HILI from 0.43 to 0.33 percentage points per month and the FF5-plus-momentum alpha from 0.37 to 0.27 with $t = 1.65$. The three-tertile mid-cap sort reported in Section 4 (Table IA.13) is the more principled size-robustness test: it removes mega-cap concentration by construction through the NYSE size-percentile breakpoints rather than by excluding specific firms by name.

Table IA.2: HILI realized-return robustness to exclusion of the Magnificent Seven. Each row reports the mean HILI spread and the Fama-French five-factor plus momentum alpha on the headline sort (Rule-I, 2-month lag, within-month empirical quintiles, value-weighted, August 2002 through May 2024). The second row re-runs the same sort after dropping AAPL, MSFT, Alphabet (GOOGL and GOOG Class C), AMZN, META, NVDA, and TSLA from the universe. Newey-West standard errors in parentheses.

Sample	Mean HILI (pp/mo)	FF5+M alpha (pp/mo)
Headline	0.433*	0.365**
	(0.247)	(0.163)
Excluding Mag7	0.334	0.270*
	(0.246)	(0.164)

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

IA.D HILI sensitivity grid

Table IA.3 reports the full sensitivity grid across timing (Rule-I vs Rule-P, lag structure), breakpoint structure (empirical vs NYSE, within-month vs recalibrated), signal construction (full IXI vs fund-branch vs IXI_{pass}), and weighting (value-weighted vs equal-weighted). The headline configuration (Rule-I, two-month lag, within-month empirical quintile breakpoints, value-weighted full IXI) is carried forward from the body. The grid documents that the headline FF5-plus-momentum alpha is not an artifact of any one construction choice.

Table IA.3: HILI sensitivity grid. Mean return and Fama-French five-factor plus momentum alpha of the top-minus-bottom IXI quintile spread under alternative specification choices. Returns and alphas are in percentage points per month. Newey-West standard errors in small parentheses. Annualized Sharpe ratios and the number of months enter the last two columns. The main specification (Panel A, first row) uses full IXI under Rule-I, two-month lag, within-month empirical quintile breakpoints, value-weighted, raw monthly returns. The grid is estimated on the Rule-I/Rule-P shared validated window (August 2002 through March 2024) for direct cross-comparison and therefore covers 260 months rather than the 262-month Rule-I safe window used elsewhere.

Specification	Mean return	FF5m alpha	Sharpe (ann.)	Months
Panel A: Timing choice at the main breakpoint, weighting, and quintile structure				
Rule-I, 2m lag (main spec)	0.446* (0.248)	0.381** (0.163)	0.535	260
Rule-I, 1m lag	0.354 (0.257)	0.291* (0.172)	0.421	260
Rule-P, 1m lag	0.344 (0.253)	0.280 (0.171)	0.410	260
Panel B: Breakpoint, quintile-structure, and size-trim variants (Rule-I + 1m baseline)				
Rule-I + 1m, N=5, within-month bp, VW, raw	0.367 (0.253)	0.286* (0.171)	0.436	263
Rule-I + 1m, N=3 quintile → tercile	0.212 (0.206)	0.147 (0.144)	0.284	263
Rule-I + 1m, N=5, NYSE breakpoints	0.289 (0.204)	0.239* (0.137)	0.391	263
Rule-I + 1m, N=5, drop top 2.5% me	0.159 (0.252)	0.164 (0.173)	0.188	263
Panel C: Fund-only branch of IXI (drops the non-fund branch)				
Fund-only IXI, N=5	0.238 (0.251)	0.165 (0.168)	0.291	263
Fund-only IXI, N=3 tercile	0.183 (0.203)	0.112 (0.140)	0.247	263
Panel D: Portfolio weighting and return winsorization				
Rule-I + 1m, N=5, equal-weighted	−0.107 (0.282)	−0.429* (0.244)	−0.097	263
Rule-I + 1m, N=5, returns winsorized at 2.5%	0.457* (0.244)	0.373** (0.164)	0.567	263

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$ under Newey-West standard errors. Rows in Panel A report the main specification (Rule-I, two-month lag) alongside the Rule-I one-month-lag and Rule-P one-month-lag variants. Rows in Panel B vary the breakpoint convention (empirical within-month quintiles vs. NYSE breakpoints), the number of portfolios (quintile vs. tercile), and the size trim (drop the top 2.5% of lagged market equity each month). Panel C replaces the full signal with the fund-only branch (mutual fund and ETF holdings, no 13F non-fund branch). Panel D reports equal-weighted and return-winsorized variants at the Rule-I + 1m baseline. Panels B, C, and D use Rule-I one-month lag as the baseline to preserve a comparable sample to the grid's original construction.

IA.E Placebo permutation of $T1$

Table IA.4 reports a placebo permutation test on the event-study headline specification. For each of three samples (Siblis augmented, Q1 best-timing, and the FF3-adjusted matched-control variant), the $T1$ vector is randomly shuffled across events, the augmented-controls regression is re-estimated, and the cluster-robust t-statistic on $T1$ is recorded. The procedure is repeated 1,000 times with a fixed seed to generate a null distribution of t-statistics under the hypothesis that $T1$ carries no information about short-window CAR.

The actual Siblis augmented t-statistic of 2.44 sits in the 97th percentile of the null $|t|$ distribution, delivering a two-sided placebo p-value of 0.026. The Q1 best-timing t of 3.36 is in the 99th percentile ($p = 0.006$). The FF3-matched Siblis t of 2.41 is in the 96th percentile ($p = 0.035$). All three cells reject the placebo null at the five percent level. The test verifies that the headline t-statistics are not artifacts of correlated noise between $T1$ and the control variables that would show up equivalently under a random reassignment of $T1$.

Table IA.4: Placebo permutation test on the event-study headline. For each sample, $T1$ is randomly shuffled across events 1,000 times, the augmented-controls regression is re-estimated each iteration, and the cluster-robust t-statistic on $T1$ is recorded. The reported null-distribution percentiles are from the null t-statistic distribution. The two-sided placebo p-value is the share of permutations with $|t^{\text{null}}| \geq |t^{\text{actual}}|$. Seed: 20260419.

Sample	n	t^{actual}	Null p_5	Null median	Null p_{95}	Placebo p
Siblis (augmented)	372	2.44	−1.75	−0.02	1.74	0.026**
Q1 (best-timing)	157	3.36	−1.95	−0.05	1.89	0.006***
Siblis (FF3-matched)	318	2.41	−1.87	0.02	1.75	0.035**

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$ on the two-sided placebo p-value.

IA.F Treatment-variable robustness

The paper's short-window event study in Section 7 uses $T1^{\text{short}}_{\text{fund}} = \text{IXI}_{\text{fund}}[m+1] - \text{IXI}_{\text{fund}}[m-1]$, the fund-branch IXI jump over a two-month window around the announcement month. Table IA.5 reports three alternative treatment constructions on the Siblis, pooled, and Q1 best-timing samples.

Replacing the fund branch with the full IXI over the same short $[m-1, m+1]$ window delivers a treatment coefficient of -0.71 ($t = -0.15$) on Siblis, -1.31 ($t = -0.29$) pooled, and $+6.80$ ($t = 0.99$) on Q1 best-timing, none statistically distinguishable from zero. The pattern reflects measurement timing. The fund branch of IXI rebalances within days of an inclusion announcement because ETFs and index mutual funds trade mechanically into the new benchmark, so the full price-relevant passive-demand shock is captured within the $[m-1, m+1]$ window. The non-fund branch, by contrast, is computed from quarterly 13F institutional filings due forty-five days after quarter-end, so the institutional component of passive-demand adjustment typically shows up only one quarter after the event at the earliest. Including the non-fund component in a two-month treatment window dilutes the treatment signal.

The wider $[m-2, m+3]$ window partially closes this gap. The full-IXI wide-window treatment delivers 4.97 ($t = 0.80$) on Siblis, 4.46 ($t = 0.71$) pooled, and 14.85 ($t = 2.02$) on Q1 best-timing, the last of which clears the five-percent significance bar. The fund-branch wide-window treatment delivers 37.91 ($t = 2.56$) on Siblis, 34.90 ($t = 2.36$) pooled, and 39.92 ($t = 2.44$) on Q1 best-timing, statistically significant throughout. The fund branch is the appropriate short-window treatment for the Section 7 identification design; the full measure is informative at longer windows.

IA.G Pre-event passive-family concentration

Table IA.6 reports the continuous $T1 \times \text{conc}_{10\text{pp},\,c}$ specification referenced in Section 7, on both the market-adjusted and the matched-control outcomes across the Siblis and Q1 samples. The concentration measure is the top-five fund-family share of passive ownership at month $m-1$, centered at the Siblis-sample mean of 57.2 percent and rescaled to ten-percentage-point units. The interaction coefficient on $T1 \times \text{conc}_{10\text{pp},\,c}$ is null in every cell, with cluster-robust t-statistics between -0.26 and $+0.67$. The main effect of concentration is positive: a ten-percentage-point higher pre-event concentration is associated with between 0.74 and 0.90

Table IA.5: Event-study treatment robustness. The headline treatment $T1^{\text{short}}_{\text{fund}} = \text{IXI}_{\text{fund}}[m+1] - \text{IXI}_{\text{fund}}[m-1]$ uses the fund branch of IXI over the short $[m-1, m+1]$ window, where m is the announcement month. Three alternative constructions replace the fund-branch-over-short-window object with the full IXI over the short window, the fund branch over a wider $[m-2, m+3]$ window, and the full IXI over the wider window. The dependent variable is the market-adjusted short-window CAR. Regressions include $\log(me)$ as the control. Standard errors in parentheses, clustered by calendar year of announcement. Sample: S&P 500 additions from August 2002 through May 2024.

Treatment	Siblis	Pooled	Q1 best-timing
	(1)	(2)	(3)
$T1^{\text{short}}_{\text{fund}}$ (paper headline)	22.04*	20.24*	42.18***
	(11.40)	(11.32)	(15.41)
$T1^{\text{short}}_{\text{full}}$ (full IXI, short window)	−0.71	−1.31	6.80
	(4.69)	(4.51)	(6.89)
$T1^{\text{wide}}_{\text{fund}}$ (fund branch, wide window)	37.91**	34.90**	39.92**
	(14.81)	(14.77)	(16.39)
$T1^{\text{wide}}_{\text{full}}$ (full IXI, wide window)	4.97	4.46	14.85**
	(6.21)	(6.23)	(7.35)
Observations	384	388	160

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$ under standard errors clustered by announcement year. The short window is the event-month pair $[m-1, m+1]$. The wider window is $[m-2, m+3]$. The paper's headline specification uses the fund-branch short-window object because ETFs and mutual funds rebalance holdings within days of index-inclusion announcements, whereas 13F institutional holdings that populate the non-fund branch adjust on quarterly reporting cycles with a forty-five-day filing lag. The fund branch therefore captures the full price-relevant passive-demand shock inside the short window, while the full measure mixes in a structurally delayed non-fund component that has not yet responded by month $m+1$.

percentage points of additional short-window CAR. A discrete high-tercile concentration indicator predicts an additional 4.83 percentage points of CAR on Siblis (cluster-robust $t = 3.34$) and 4.44 on Q1 (cluster-robust $t = 4.65$). Pre-event family concentration operates as a level shifter on inclusion CARs rather than as an elasticity amplifier, consistent with Pavlova and Sikorskaya (2023)'s reading of benchmarked ownership as a persistent equity-market feature that acts on price levels.

Table IA.6: Phase 2C family-concentration heterogeneity. The concentration measure is top5pctPass_{m-1}, the percentage of pre-event passive ownership held by the top-5 fund families at month $m-1$. $\text{conc}_{10\text{pp,c}}$ is this measure, centered at the Siblis-sample mean (57.2 percent) and rescaled to 10-percentage-point units. Columns (1) and (2) use the market-adjusted CAR from announcement -1 through effective $+1$. Columns (3) and (4) use matched-control CAR at the same horizon. Columns (5) and (6) report the discrete T1 x high-tercile interaction as robustness. Sib. columns are Siblis-matched; Q1 are the best-timing quartile. Standard errors in parentheses, clustered by announcement year.

	(1) Sib. CAR_short	(2) Q1 CAR_short	(3) Sib._{MC} CAR_matched	(4) Q1_{MC} CAR_matched	(5) Sib. CAR_short	(6) Q1 CAR_short
$T1$	25.759**	49.135***	7.209	36.321***	33.472*	49.295***
	(12.671)	(15.619)	(11.291)	(12.515)	(18.087)	(18.721)
$T1 \times \text{conc}_{10\text{pp,c}}$	−1.790	2.205	6.097	9.385		
	(6.768)	(9.028)	(10.228)	(13.948)		
$T1 \times \mathbf{1}_{\text{highConc}}$					−32.629	−20.276
					(23.898)	(30.397)
$\text{conc}_{10\text{pp,c}}$	0.772**	0.839**	0.897*	0.744		
	(0.369)	(0.345)	(0.487)	(0.454)		
$\mathbf{1}_{\text{highConc}}$					4.830***	4.443***
					(1.447)	(0.955)
log me	1.083*	0.060	0.994	0.032	0.998	0.150
	(0.657)	(0.798)	(0.638)	(0.982)	(0.667)	(0.841)
log BM	−0.077	−0.236	−0.084	−0.248	−0.137	−0.269
	(0.262)	(0.369)	(0.302)	(0.343)	(0.240)	(0.333)
Momentum (12-1)	1.511	1.829*	1.368	1.598*	1.660	1.996
	(1.101)	(1.082)	(1.119)	(0.947)	(1.052)	(1.228)
Observations	372	157	318	132	372	157
R^2	0.115	0.176	0.097	0.167	0.157	0.203

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. MC = matched-control.

IA.H Heckman selection correction

Section 5 reports that the HILI-universe ICC coverage rises monotonically from 41.5 percent in Q1 to 72.0 percent in Q5, a selection asymmetry driven by Compustat and IBES input availability for the Hou-van-Dijk-Zhang ICC construction. Table IA.7 applies a Heck-

man two-stage selection correction to the firm-level panel regression. The first-stage probit regresses an ICC-availability indicator on lagged log me, log price, NYSE and NASDAQ exchange dummies, and the rolling 11-month standard deviation of monthly returns on the wide CRSP permno-month base. The second stage regresses ICC on lagged IXI and standard controls, optionally augmented with the inverse Mills ratio from the first stage and optionally with month fixed effects.

The Stage 1 probit delivers strongly significant coefficients in the expected directions: larger firms, higher-priced stocks, NYSE and NASDAQ firms, and higher-volatility firms are all more likely to have ICC available. Column (H0) of Table IA.7 reproduces the Section 5 pooled OLS baseline and delivers $\hat{\beta}_{\text{IXI}} = -0.02602$ ($t = -4.45$), comparable to the baseline column (4) of Table 4. Column (H1) adds the inverse Mills ratio. The IXI coefficient attenuates to -0.01976 ($t = -3.49$), a 24 percent reduction in magnitude, and the inverse Mills ratio itself loads strongly negatively (-0.0798, $t = -11.71$), confirming that selection is systematic. Column (H2) adds month fixed effects alongside the inverse Mills ratio; the IXI coefficient strengthens to -0.0904 ($t = -15.16$), reflecting that month fixed effects absorb substantial time-varying selection in addition to what the inverse Mills ratio absorbs cross-sectionally.

Selection attenuates the pooled OLS IXI coefficient by roughly 24 percent but leaves the ex-ante ICC-IXI relationship negative and statistically significant at the one percent level. The Section 5 headline conclusion that high-IXI stocks carry lower ex-ante expected returns is preserved.

IA.I Alternative ICC constructions

The Section 5 headline ICC uses the Hou-van-Dijk-Zhang methodology, which projects earnings from a pooled cross-sectional regression on Compustat inputs rather than using analyst forecasts. Table IA.8 reports the Q5-minus-Q1 portfolio ICC spread and the firm-level panel regression of ICC on lagged IXI plus standard controls under three ICC constructions: HVZ

Table IA.7: Heckman selection correction on the firm-level ICC panel. Stage 1 is a probit of ICC-availability on lagged log me, log price, exchange dummies (NYSE, NASDAQ relative to AMEX baseline), and rolling 11-month return standard deviation, estimated on the wide CRSP permno-month base (August 2002 through May 2024) with lagged IXI non-missing. Stage 2 regresses ICC on lagged IXI plus standard controls and the inverse Mills ratio from Stage 1. Column (H0) is the pooled OLS baseline without Heckman correction. Column (H1) adds the inverse Mills ratio. Column (H2) adds month fixed effects. Standard errors in parentheses, double-clustered by firm and year-month.

	(H0) Pooled OLS	(H1) + IMR	(H2) + IMR + month FE
IXI (lag 2)	-0.02602^{***}	-0.01976^{***}	-0.09040^{***}
	(0.00585)	(0.00566)	(0.00596)
log me	-0.01568^{***}	-0.02162^{***}	-0.01891^{***}
	(0.00076)	(0.00077)	(0.00071)
log BM	0.00557^{***}	0.00543^{***}	0.00703^{***}
	(0.00079)	(0.00077)	(0.00075)
Momentum (12-2)	-0.02347^{***}	-0.01806^{***}	-0.01783^{***}
	(0.00151)	(0.00151)	(0.00105)
Inverse Mills ratio		-0.07977^{***}	-0.05895^{***}
		(0.00681)	(0.00657)
Observations	649,495	649,495	649,495
R^2	0.3193	0.3310	0.3927
Month FE	no	no	yes

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. The second-stage sample conditions on ICC being non-missing. Stage 1 probit coefficients are reported in the text of Appendix IA.H.

(the headline), Easton PEG (Easton 2004, a closed-form ICC using one- and two-year-ahead IBES EPS forecasts and current price), and Gebhardt-Lee-Swaminathan (GLS 2001, a residual-income model with a 12-year horizon using IBES forecasts for years one through three and industry-median ROE as the terminal anchor). All three ICCs are winsorized cross-sectionally at the 1st and 99th percentiles each month.

The GLS spread of -2.03 percentage points per year with $t = -6.58$ is nearly identical to the HVZ baseline of -2.07 percentage points with $t = -8.81$, reinforcing that the portfolio-level ICC wedge is not a HVZ-model artifact. The Easton PEG spread is substantially larger in magnitude at -3.92 percentage points with $t = -7.16$; PEG is known to pick up near-term forecast optimism cycles and therefore amplifies level differences across high- and low-IXI stocks. At the firm-level panel, the IXI coefficient is negative and statistically significant at the one percent level under all three ICCs (-0.026 HVZ, -0.062 PEG, -0.010 GLS), with GLS attenuating the magnitude relative to HVZ but retaining the sign and significance. Across the three ICC constructions, the ex-ante ICC-IXI relationship is robust in sign and statistical strength. Magnitudes differ across constructions as expected, with PEG larger and GLS slightly smaller than HVZ in the firm-panel specification.

IA.J OLS companion to the 2SLS decomposition

Table IA.9 reports the OLS within-firm coefficients alongside the 2SLS results of Table 22. OLS uses the same firm-by-cohort and year-quarter fixed effects as the 2SLS, without the Treated $\times$ Post instrument. For IXI (the paper's main measure), OLS delivers a positive, significant within-firm coefficient on all three main outcomes. For IXI_{pass} (the strict passive sub-measure), OLS is near zero on every outcome while 2SLS remains large and significant; the role of the instrument is to isolate the inclusion-induced component of IXI variation from noisy non-inclusion-related within-firm IXI movements.

Table IA.8: ICC wedge under alternative ICC constructions. Panel A reports the time-series mean of value-weighted ICC by IXI quintile and the Q5-minus-Q1 spread under three ICC models: Hou-van-Dijk-Zhang (HVZ, the Section 5 baseline), Easton PEG (Easton 2004), and Gebhardt-Lee-Swaminathan (GLS 2001). All ICCs are winsorized at the 1st and 99th percentiles cross-sectionally each month. Panel B reports pooled-OLS firm-level panel regressions of each winsorized ICC on lagged IXI plus log me, log BM, and 12-2 momentum. Standard errors in parentheses, double-clustered by firm and year-month. Sample: August 2002 through May 2024.

Panel A: Value-weighted mean ICC by IXI quintile (decimal)

	Q1	Q2	Q3	Q4	Q5	Q5-Q1
HVZ	0.0793	0.0663	0.0621	0.0597	0.0586	-0.0207^{***}
						(0.00235)
Easton PEG	0.1236	0.1111	0.0985	0.0886	0.0844	-0.0392^{***}
						(0.00547)
GLS	0.0827	0.0702	0.0641	0.0632	0.0624	-0.0203^{***}
						(0.00308)

Panel B: Firm-level pooled OLS, ICC on lagged IXI + controls

	HVZ	Easton PEG	GLS
IXI (lag 2)	-0.02613^{***}	-0.06243^{***}	-0.00968^{***}
	(0.00582)	(0.00673)	(0.00312)
log me	-0.01555^{***}	-0.01254^{***}	-0.00082^{***}
log BM	0.00542^{***}	0.00773^{***}	0.01976^{***}
Momentum (12-2)	-0.02312^{***}	-0.02931^{***}	-0.00343^{***}
N	649,495	590,208	401,281
R^2	0.319	0.087	0.220

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Easton PEG: $r = \sqrt{(\text{EPS}_{t+2} - \text{EPS}_{t+1})/P_t}$, requires $\text{EPS}_{t+2} \geq \text{EPS}_{t+1} > 0$. GLS: residual-income model with 12-year horizon, IBES EPS forecasts years 1 through 3, linear interpolation to industry-median ROE at year 12. Sample coverage in Panel A uses the full joint support of each ICC method; Panel B conditions on each method's non-missing sample plus non-missing control variables. For brevity, Panel B reports standard errors for the IXI coefficient only; controls are shown with significance stars.

Table IA.9: OLS versus 2SLS on the S&P 500 inclusion panel. Each cell reports the within-firm panel coefficient on the endogenous regressor (the IXI measure named in the column header) in a regression of the row outcome with firm-by-cohort plus year-quarter fixed effects. OLS (pooled within-firm) is reported directly; 2SLS uses Treated × Post as the instrument. The 2SLS numbers reproduce the 4.30 headline log(ME) response and its decomposition into log(SHROUT) and log(PRC) in Table 22. Standard errors in parentheses, double-clustered by firm and year-quarter.

	IXI (paper's main)		IXI_{pass} (strict)		Fund branch	
Outcome	OLS	2SLS	OLS	2SLS	OLS	2SLS
log(ME)	0.538***	4.298***	−0.639	7.977***	1.442***	3.562***
	(0.144)	(1.171)	(0.652)	(2.168)	(0.378)	(0.895)
log(SHROUT)	0.163*	1.654**	−0.150	3.069**	0.444**	1.370**
	(0.086)	(0.803)	(0.357)	(1.511)	(0.214)	(0.652)
log(PRC)	0.423***	2.944***	−0.217	5.464***	1.148***	2.440***
	(0.122)	(1.074)	(0.586)	(1.945)	(0.346)	(0.850)
Observations	122,604		122,604		122,604	

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. OLS is the pooled within-firm coefficient on the endogenous regressor with the same firm-by-cohort and year-quarter fixed effects as the 2SLS; no instrument is used. 2SLS instruments the endogenous regressor with Treated × Post. For IXI (paper's main measure), OLS is positive and significant for all three outcomes, and 2SLS scales the coefficient up by a factor of roughly eight as the instrument isolates the inclusion-induced component of IXI variation. For IXI_{pass} (strict passive only), OLS is statistically indistinguishable from zero while 2SLS is large and significant, the classical signature of IV correcting attenuation from measurement noise in the endogenous within-firm variation.

IA.K Russell 1000-to-2000 switcher cohort

Table IA.10 reports the same IV specifications on the Russell 1000-to-2000 switcher cohort. Russell switchers are firms whose May 31 market capitalization fell below the 1000th rank, so the treatment selection is mechanically correlated with declining firm size. The Russell log(SHROUT) IV is statistically zero across all three passive measures, indicating that the unadjusted-share-count IV response on the S&P sample reflects the S&P committee-selection margin rather than index membership in general.

IA.L Conditional flow projection

Figure IA.1 projects cumulative HILI under three flow scenarios using the first-difference specification from Section 6, with controls held at their post-2013 means. Bootstrap 95 percent bands come from a 6-month block bootstrap with 1,000 replications. Bootstrap uncertainty reflects the regression's coefficient uncertainty rather than uncertainty about

Table IA.10: Russell 1000-to-2000 switcher IV on valuation components, parallel to Table 22. Russell switchers are mechanically selected on declining market capitalization because they fall out of the R1000. The log(ME) and log(PRC) IV coefficients are therefore strongly negative, while the log(SHROUT) coefficients are not statistically distinguishable from zero.

	IXI	IXI_{pass}	Fund branch
log(ME)	-26.02^{**}	-38.93^{***}	-30.51^{***}
	(11.83)	(11.15)	(10.42)
log(SHROUT)	-1.63	-2.44	-1.91
	(2.23)	(3.16)	(2.54)
log(BE)	-12.76^{**}	-21.06^{***}	-16.39^{***}
	(5.41)	(7.02)	(6.13)
First-stage F	5.9	15.3	11.7

Notes: The Russell panel uses 544 R1000-to-R2000 switcher events between 2002 and 2019, identified using the banding rule of Coles et al. (2022); control firms are R1000 stayers within ±100 ranks of the cutoff. Each cell reports a 2SLS coefficient from a separate outcome regression. The endogenous regressor is the IXI measure named in the column header and the excluded instrument is the Russell switch indicator (Treated × Post). Specifications include firm-by-cohort and year-quarter fixed effects. Standard errors are in parentheses and are double-clustered by firm and year-quarter. First-stage F is the Kleibergen-Paap statistic reported by column; the Stock-Yogo 10 percent critical value is $F = 10$. Firm-month observation counts vary by outcome due to non-missing requirements and are not tabulated in this display. The table is read as an absence of a share-issuance response in the Russell setting, not as evidence against the S&P share-issuance response. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

the realized flow path itself, and the projections should be read as conditional what-ifs rather than forecasts.

IA.M Industry-neutral HILI

Table IA.11 reports an industry-neutral HILI construction. Each month, stocks are sorted into empirical within-Fama-French-12-industry IXI quintiles and value-weighted within-industry Q5-minus-Q1 spreads are aggregated across industries, with industry weights proportional to total industry market capitalization at $t - 1$. Industry-month cells with fewer than 25 stocks are dropped. The headline pooled sort appears in the upper panel for reference. The industry-neutral spread retains positive CAPM and FF3 alphas but attenuates to roughly half of the pooled FF5-plus-momentum alpha, reflecting the partial industry co-movement of IXI with technology-sector concentration.

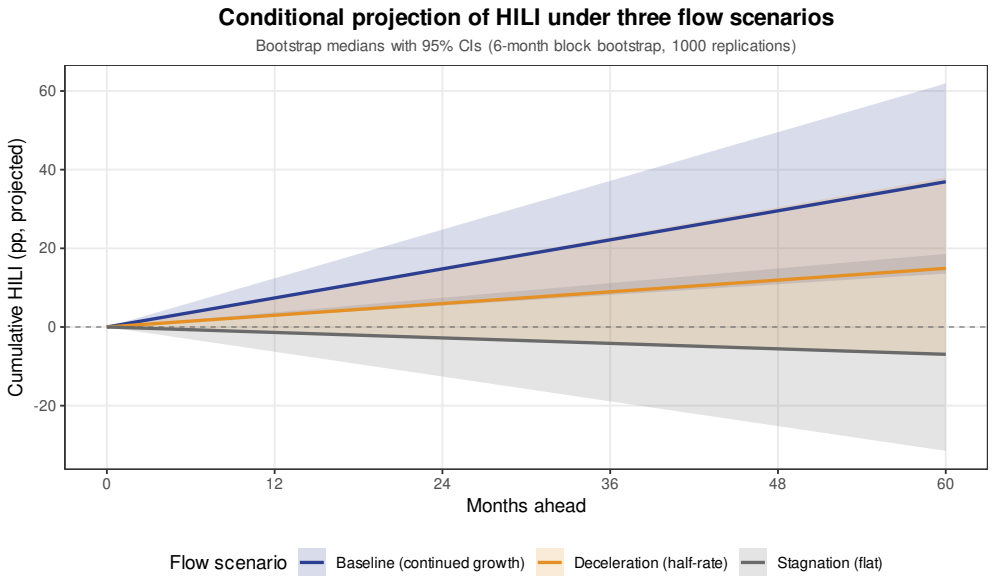

Figure IA.1: Conditional projection of cumulative HILI over a sixty-month horizon under three flow scenarios. The projection uses the first-difference aggregate regression from Section 6 that relates HILI to ΔAggAUM_t with $t = 4.51$. Controls (ER^{HILI}, $\Delta\text{EF}^{\text{HILI}}$, lagged AggFlow) are held at their post-2013 means. Solid lines are bootstrap medians and shaded bands are 95 percent confidence intervals from a 6-month block bootstrap on the first-difference regression with 1,000 replications. Baseline scenario: ΔAggAUM continues at its 2013-2024 average of 0.063 pp per month. Deceleration scenario: monthly ΔAggAUM is halved. Stagnation scenario: ΔAggAUM goes to zero. The projection is a reduced-form application of the existing coefficient and does not estimate or use a demand-system framework.

Table IA.11: Industry-neutral HILI. The industry-neutral HILI is constructed by sorting stocks each month into empirical within-Fama-French-12-industry IXI quintiles, forming value-weighted within-industry quintile portfolios, and aggregating the within-industry Q5 − Q1 spread across industries with industry weights proportional to total industry market capitalization at $t-1$. The table reports the mean return and factor-model alphas of the pooled HILI (Panel A, reference from Table 3) and of the industry-neutral HILI (Panel B) on the same sample (August 2002 through May 2024, 262 months). Newey-West standard errors in parentheses.

	Mean	FF3 alpha	FF5 alpha	FF5+M alpha
Panel A: Pooled HILI (reference)				
Pooled HILI	0.433*	0.393**	0.378**	0.365**
	(0.247)	(0.162)	(0.159)	(0.163)
Panel B: Industry-neutral HILI				
Industry-neutral HILI	0.222	0.359**	0.204	0.193
	(0.216)	(0.162)	(0.146)	(0.144)

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Industry-month cells with fewer than 25 stocks are dropped before sorting.

IA.N HILI period breakdown

Table IA.12 reports the HILI spread split into pre-2013 and post-2013 subperiods across CAPM, FF3, FF5, and FF5-plus-momentum factor models. The post-2013 passive-era FF5-plus-momentum alpha of 0.53 percentage points per month ($t = 2.56$) is the number that appears in the Section 4 body prose. The pre-2013 alpha of 0.10 is essentially zero.

Table IA.12: HILI realized returns by era. The table reports the mean monthly HILI return and factor-model alphas for the full sample and for pre-2013 and post-2013 sub-samples. The post-2013 sample begins in January 2013, coinciding with the acceleration in U.S. passive-equity-fund AUM documented in Section 6. All specifications use the main HILI sort: full IXI under Rule-I, two-month lag, within-month empirical quintile breakpoints, value-weighted. Newey-West standard errors in parentheses.

	n (months)	Mean (pp/m)	CAPM α	FF5+MOM α
Full sample (Aug 2002 – May 2024)	262	0.433	0.627***	0.365**
			(0.225)	(0.163)
Pre-2013 (Aug 2002 – Dec 2012)	125	0.132	0.307	0.102
			(0.284)	(0.217)
Post-2013 (Jan 2013 – May 2024)	137	0.707	0.882***	0.531**
			(0.316)	(0.207)

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Factor models use MKT-Rf, SMB, HML, RMW, CMA (Fama-French 2015) and MOM (Carhart 1997). All factors from the Kenneth French Data Library.

IA.O Three-tertile size sort

Table IA.13 reports the within-tertile HILI spread and factor-model alphas for the three NYSE-size-tertile groups referenced in Section 4. The Big tertile concentrates portfolio weight in mega-caps and carries a small FF5-plus-momentum alpha; the mid-tertile ($t = 2.19$) is the cleanest within-size result; the Small tertile is economically noisy.

IA.P Binary NYSE-median double sort

Table IA.14 reports the binary NYSE-median Big/Small double sort on IXI quintiles. Within Big, the HILI FF5-plus-momentum alpha is 0.16 ($t = 1.29$); within Small, 0.31 ($t = 1.34$); neither individually clears the conventional significance bar. The three-tertile sort in Section 4 (Table IA.13) resolves the tension by separating the mid-cap and mega-cap components

Table IA.13: Within-NYSE-size-tertile HILI. Each month stocks are assigned to Small (below the NYSE 33rd percentile of market equity), Mid (33rd–67th percentile), or Big (above the 67th percentile) size tertiles using NYSE-listed stocks to form the breakpoints. Within each size tertile, stocks are sorted monthly into IXI quintiles using within-tertile empirical breakpoints, and the value-weighted Q5-minus-Q1 spread is computed. The Big tertile captures the top third of the NYSE size distribution and contains the Magnificent Seven in their entirety. The Mid tertile isolates the mid-cap segment where IXI varies most cleanly without mega-cap concentration. Newey-West standard errors in parentheses. Sample: August 2002 through May 2024.

	Small tertile (below p33 NYSE)	Mid tertile (p33–p67 NYSE)	Big tertile (above p67 NYSE)
Mean HILI (pp/month)	0.444	0.286	0.141
CAPM α	0.214	0.361**	0.304**
	(0.261)	(0.175)	(0.143)
FF5+MOM α	0.173	0.284**	0.159
	(0.244)	(0.130)	(0.119)
Permnos (avg/month)	≈6,800	≈2,900	≈1,400
2023 mkt-cap share	2%	7%	91%

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. NYSE size breakpoints computed each month on NYSE-listed stocks only; all CRSP common stocks (share codes 10, 11, 12, 18) are then assigned to size tertiles using those breakpoints. IXI quintile sorts are run within each size-tertile cell using within-cell monthly breakpoints. Factor models use the six Fama-French-plus-momentum factors.

of the binary Big bucket.

Table IA.14: Size-conditional HILI. Each month, stocks are split into Big and Small using the NYSE median of lagged market cap. Within each size bucket, stocks are sorted into 5 within-bucket quintiles on lagged IXI (Rule-I, 2-month lag) using monthly empirical 1/5 breakpoints. HILI is the top-minus-bottom value-weighted spread. Sample: August 2002 through May 2024 (262 months). Newey-West standard errors in parentheses.

Size bucket	Avg. stocks/mo	Months	Mean (pp)	Sharpe (ann)	CAPM α	FF3 α	FF5 α	FF5+MOM α
Big	1,115	262	0.147	0.235	0.303**	0.270*	0.170	0.159
			(0.169)		(0.148)	(0.142)	(0.126)	(0.123)
Small	3,162	262	0.537**	0.534	0.432*	0.478*	0.296	0.307
			(0.233)		(0.260)	(0.250)	(0.228)	(0.229)

Notes: $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. Big is stocks above the NYSE median of me_{lag} each month; Small is at or below.

IA.Q Factor-placebo test on the FD aggregate regression

A complementary placebo test asks whether ΔAggAUM loads specifically on HILI or on any large long-short spread over the same window. Table IA.15 reports the FD specification with the same three controls and five standard Fama-French factor returns (HML, SMB, MOM,

RMW, CMA) as alternative dependent variables, alongside HILI. The HILI coefficient of 11.39 ($t = 4.51$) is the largest across the six regressions and has the strongest statistical significance by a wide margin. SMB, MOM, and RMW load close to zero, consistent with the interpretation that size, momentum, and profitability are not mechanical counterparts of passive-share growth. HML (8.74, $t = 2.55$) and CMA (4.99, $t = 2.46$) do load positively, at 77 percent and 44 percent of HILI's magnitude, indicating that ΔAggAUM also co-moves with value and conservative-investment tilts. The FD coefficient on HILI is therefore the largest of the six candidates rather than uniquely loaded on HILI, and the counterfactual in Figure 4 should be read accordingly: as evidence that passive-share growth is quantitatively strongest on HILI, with a smaller but non-zero co-movement on standard value and conservative-investment factors over the same sample.

Table IA.15: Factor-placebo test on the FD aggregate regression.

LHS	$\hat{\beta}_{\Delta\text{AggAUM}}$	SE	t	p	N
HILI	11.392	2.525	4.51	< 0.001	261
HML	8.738	3.431	2.55	0.012	261
CMA	4.989	2.026	2.46	0.015	261
MOM	2.868	5.042	0.57	0.570	261
RMW	2.741	2.354	1.16	0.245	261
SMB	0.811	2.707	0.30	0.765	261

Notes: Each row is a separate OLS regression of the stated LHS on ΔAggAUM_t plus the three controls in the Section 6 FD specification ($\text{ER}_t^{\text{HILI}}$, $\Delta\text{EF}_t^{\text{HILI}}$, lagged aggregate passive flow). Dependent variables are monthly returns in percent per month. The sample is 2002-09 through 2024-05 (261 months; one month lost to the first difference). Standard errors are Newey-West with six-month bandwidth. The HILI row reproduces the Section 6 headline coefficient. HML, SMB, MOM, RMW, and CMA are the standard Fama-French-plus-momentum factor returns.

The FD coefficient is also stable when the HILI regression directly controls for contemporaneous factor returns: adding Mkt−Rf, SMB, HML, RMW, CMA, and momentum leaves the coefficient on ΔAggAUM at 9.52 with Newey-West $t = 5.37$, compared with 11.39 and $t = 4.56$ in the baseline FD specification.