

A General-Equilibrium Theory of Target-Date-Fund Rebalancing

Author names withheld for blind review

April 23, 2026

Abstract

A share-target target-date-fund (TDF) sector that interacts with a quadratic-cost arbitrageur sector in general equilibrium produces two testable reduced-form signatures. First, post-news realized variance responds differentially to TDF ownership in dealer-constrained states: the sign prediction is a negative three-way interaction $\text{TDF} \times \text{dealer-tightness} \times \text{post-news}$, matching the sign of a scalar cross-partial proved in the model. We confirm the sign empirically in a matched pre/post triple-difference regression over $N = 118,982$ quarterly EPS events (2005–2018): the $\text{post} \times \text{TDF} \times \text{HKM}$ coefficient is -7.5×10^{-4} with $t = -5.57$ permno- \times -quarter two-way clustered ($t = -10.16$ event-clustered), using the [He et al. \[2017\]](#) intermediary capital factor as the dealer-state proxy. We interpret this as a reduced-form falsifier whose sign is consistent with the theory, not a magnitude match and not a test of any matrix-level identification in the N -asset setting. Second, residual pairwise correlation rises in the count of co-holding TDF sub-funds, a rank- M excess-covariance prediction bridged to a reduced-form co-occupancy count: confirmed at $t = +7.66$ on 6.49M pair-quarters, sitting at the 100th percentile of a 26,685-strategy flex-mining null. A cross-sectional level association between TDF ownership and realized variance is a supporting consistency check with [Parker et al. \[2023\]](#), [Parker and Sun \[2025\]](#), not an independent causal contribution: within-stock and Form 5500 Bartik IV specifications are null. A welfare coda delivers a scaling law: avoidable coordination cost grows as $\alpha^2/(1 - \alpha)$ and becomes first-order as the TDF equity share crosses roughly 0.25; at today's $\alpha \approx 0.08$ the dollar level is small ($\approx \$22,800$ per year).

1 Introduction

Target-date funds (TDFs) hold roughly \$3.5 trillion of U.S. retirement assets and execute share-target rebalancing: when realized equity returns drift the equity share off the glide-path target, the fund mechanically sells (buys) equity to restore the target weight. The rebalancing rule is a primitive of the delegation contract, not a discretionary response to prices. This rule raises two equilibrium questions. First, how do predictable mechanical rebalancing flows propagate to asset prices when the counterparty is a balance-sheet-constrained arbitrageur sector? Second, through what structural cross-partial can the mechanism be falsified, separated from generic passive-demand or benchmarking channels that predict similar reduced-form correlations?

This paper answers both questions empirically, with theory as the interpretive frame. I build a general-equilibrium model in which a mass- α share-target TDF sector interacts with a mass- $(1 - \alpha)$ CARA arbitrageur sector that bears a quadratic inventory cost with coefficient κ , and I derive two reduced-form signatures the model implies. The headline signature is a negative three-way interaction: $\text{post} \times \text{TDF ownership} \times \text{dealer-tightness}$ in realized variance. The model’s scalar cross-partial $\partial^2 \mathbb{E}[\text{RV}]_{\text{post}} / \partial \alpha \partial \kappa < 0$ (Proposition 3.9) pins the sign of this interaction. I confirm it in a matched pre/post triple-difference regression over $N = 118,982$ quarterly earnings events: the $\text{post} \times \text{TDF} \times \text{HKM}$ coefficient is -7.5×10^{-4} with $t = -5.57$ (permno- \times -quarter two-way clustered; $t = -10.16$ event-clustered), using the He et al. [2017] intermediary capital factor as the dealer-state proxy.

The second signature is a reduced-form basket-covariance prediction: residual pairwise correlation rises in the count of co-holding TDF sub-funds. Theorem 3.2 and the Lemma 3.4 bounded-overlap bridge deliver this count-based monotonicity from the Woodbury structure of the share-target-plus-quadratic-arbitrageur equilibrium. I confirm it at $t = +7.66$ on 6.49M pair-quarters, sitting at the 100th percentile of the 26,685-strategy Chen and Zimmermann [2022] flex-mining null.

The mechanism is transparent. Share-target delegation generates a contrarian rebalancing demand proportional to α : a positive price shock pushes the equity share above the target, triggering an equity sale that the arbitrageur sector absorbs. The quadratic inventory cost means that absorption is more expensive when the arbitrageur balance sheet is tight (high κ). Post-news realized variance is therefore a concave function of TDF ownership in the (intermediary-state) dimension: more TDF ownership dampens more when dealers are unconstrained, and dampens less when dealers are constrained. The sign of the interaction pins the mechanism. Generic institutional benchmarking [Basak and Pavlova, 2013, Buffa et al., 2022] or passive ownership [Jiang et al., 2025] is not expected to deliver this specific three-way structure.

My contribution is two-fold. First, on the empirical side, the paper delivers two headline results. The I5 triple-difference is the sharpest falsifier: the sign of the $\text{post} \times \text{TDF} \times \text{HKM}$ coefficient matches the theory and rejects the null of zero asymmetric response at $t = -5.57$ on conservative two-way clustering. The I2 basket-covariance test is the strongest confirmation: $t = +7.66$ on 6.49M pair-quarters sits above all 26,685 data-mined strategies in the flex-mining null. Both tests use a fund-of-fund look-through measurement of TDF exposure that corrects a large direct-only understatement (see Table 3); this look-through measurement is itself a methodological contribution. Second, on the theoretical side, the model provides the interpretive framework that distinguishes these signatures from generic common-ownership or passive-demand alternatives. Theorem 3.2 and Lemma 3.4 are the novel structural content: a rank- M excess-covariance decomposition with a bounded-overlap bridge to a co-occupancy count that generic common-ownership models do not deliver. A welfare coda (§4) develops a scaling law: avoidable coordination cost grows as $\alpha^2 / (1 - \alpha)$ and becomes first-order at $\alpha \gtrsim 0.25$; at today’s $\alpha \approx 0.08$ the level is small ($\approx \$22,800$ per year). A supporting cross-sectional level association (I1) is documented as a consistency check with Parker

et al. [2023], Parker and Sun [2025]; within-stock and Form 5500 Bartik IV specifications are null, and I claim no independent causal identification of the level.

Three qualifications of framing are in order. First, Proposition 3.9 is proved for the scalar-asset equilibrium under the Grossman-Miller identification $\lambda = \kappa$, on the calibrated rectangle $(\alpha, \kappa) \in (0, 0.5] \times (0, 0.1]$. In the N -asset setting the price-impact operator is matrix-valued and the scalar identity is not proved; I treat the I5 test as a *reduced-form* falsifier whose sign is consistent with the theory, and the HKM factor as a literature-standard scalar operationalization adopted for the empirical test. I do not claim the I5 coefficient identifies any structural parameter; the empirical coefficient and the theoretical cross-partial are in different units, and I do not reconcile magnitudes.¹ Second, the cross-sectional level test (I1) is a supporting consistency check, not a headline result: within-stock and IV specifications are null, and I do not claim a within-stock causal identification. Third, the §4 welfare contribution is a *scaling law*, not a current-dollar claim. At today’s $\alpha \approx 0.08$ the avoidable level is small; the contribution is that the level grows as $\alpha^2/(1 - \alpha)$ and becomes policy-relevant at $\alpha \gtrsim 0.25$.

The closest published competitor is Kimball et al. [2020], which develops an OLG general-equilibrium model of rebalancing with heterogeneous risk aversion. That paper does not model TDFs as an institution, does not deliver a cross-section of basket-covariance predictions, and has no arbitrageur balance-sheet channel. The closest *empirical* work is Parker et al. [2023] and the mechanism evidence in Parker and Sun [2025]; my cross-sectional level association (I1) is a consistency check with their dampening finding under a corrected FoF-look-through TDF measure. Generic common-ownership models [Antón and Polk, 2014, Greenwood and Thesmar, 2011] already predict that basket co-holding raises residual covariance; the novel structural content in Theorem 3.2 and Lemma 3.4 is the count-based bridge that emerges from the Woodbury structure of the share-target equilibrium. The cleanest reduced-form falsifier is the intermediary-state three-way interaction: I lead with it because it sharply distinguishes the TDF share-target channel from generic institutional demand. Basak and Pavlova [2013] study institutional benchmarking in GE; the TDF mechanical basket-overlap prediction is structurally distinct from index membership and admits an if-and-only-if characterization under orthogonal sub-fund geometry. Fang and Goldstein [2025] study strategic complementarity in target-allocation funds on the bond side; my predictions are cross-sectional equity-price and variance predictions, not bond-market fragility. Harvey et al. [2025] estimate a \$16B aggregate rebalancing cost empirically; I microfound a scaling law for the TDF-specific coordination-failure component. Pavlova and Sikorskaya [2023], Kojien and Yogo [2019], and Gabaix and Kojien [2021] supply the inelastic-demand language within which TDF rebalancing is a specific primitive. On the dealer side, I extend the Grossman and Miller [1988] / He et al. [2017] framework to a share-target delegation setting, and I adopt their identification of the arbitrageur marginal inventory cost with the date-1 price-impact coefficient. On the optimal-execution side, Almgren and Chriss [2001] establish a single-agent Jensen inequality for calendar

¹The theoretical cross-partial is $O(10^{-12})$ in AUM-fraction units while the regression coefficient is $O(10^{-3})$ in standardized-regressor units; under the reduced-form framing above, I do not claim a magnitude match, only a sign match.

smoothing; Kashyap et al. [2023] derive the planner-internalizes-institutional-externality result in a benchmarking setting; my cross-institutional coordination theorem is absent from both.

The remainder of the paper proceeds as follows. §2 presents the model. §3 establishes the structural results: Theorem 3.2 (rank- M basket covariance), Lemma 3.4 (bounded-overlap bridge), and Proposition 3.9 (scalar cross-partial). Theorem 3.1 (cross-sectional dampening) is retained as a supporting consistency check. §4 develops the welfare coda. §5 describes data and identification; §6 reports the two headline empirical results (I5 mechanism test, I2 basket covariance) and the supporting tests (I1 level, I3 scaling, I4 absorption, I6 flex-mining placebo). §7 presents the Form 5500 IV strategy for the level test. §8 discusses limitations; §9 concludes. Proofs are in Appendix A.

2 Model

2.1 Timing and information

There are three dates, $t \in \{0, \frac{1}{2}, 1\}$. A single risky asset is in unit supply; the N -asset extension is in §2.6. A risk-free asset is in elastic supply with a zero rate. At $t = 0$ agents form portfolios. At $t = \frac{1}{2}$ a public fundamental signal ε is realized and agents retrade. At $t = 1$ the residual dividend u is realized. The terminal payoff is

$$D = \mu + \varepsilon + u, \quad \varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2), \quad u \sim \mathcal{N}(0, \sigma_u^2), \quad \varepsilon \perp u. \quad (2.1)$$

2.2 TDF sector

A mass $\alpha \in (0, 1)$ of target-date funds delegates portfolio management under a share-target contract. The representative TDF at cohort age a is pre-committed to glide-path equity share $s(a) \in (0, 1)$; the wealth-weighted average target is $\bar{s} \in (0, 1)$. At each rebalancing date the manager sets the realized equity share $s_t \equiv P_t x_t^T / W_t^T$ equal to \bar{s} , where x_t^T is the TDF sector's share position and W_t^T is sector wealth.

Lemma 2.1 (Linearized rebalancing). *To first order in $(P_{1/2} - P_0)/P_0$, the TDF sector's rebalancing trade is*

$$\Delta x^T = -\bar{s}(1 - \bar{s}) \frac{W_0^T}{P_0^2} (P_{1/2} - P_0) + O((\Delta P/P_0)^2). \quad (2.2)$$

Defining $\eta \equiv (1 - \bar{s})/P_0$ and aggregating across the TDF mass α (with $W_0^T/P_0 \equiv 1$ as a unit-fixing convention), I obtain the sector-level flow rule

$$\Delta X^T = -\eta \alpha \bar{s} (P_{1/2} - P_0). \quad (\text{T})$$

Equation (T) is the equilibrium primitive of the TDF sector: rebalancing demand is linear in the price change, proportional to α , and opposite-signed to the price innovation.

Normalization and where it is relaxed. The normalization $W_0^T/P_0 \equiv 1$ fixes the unit of TDF-sector inventory as a share-count scaled to sector AUM; equivalently, α is the TDF share of aggregate equity wealth and all inventory quantities are expressed as AUM-fractions. The normalization is maintained throughout §3: every theorem, proposition, and cross-partial in the results section is stated in AUM-fraction units. The welfare calculation in §4 maintains the same convention; the sector-AUM dollar scaling (\$3.5T) enters only at the final dollar-cost step in §4, and the scaling there is an explicit dimensionalization of the AUM-fraction cost to dollars. The empirical coefficient and the theoretical cross-partial are in different units (standardized-regressor units vs. AUM-fraction units); consistent with the reduced-form framing in §6.5, I do not undertake a magnitude reconciliation and I claim a sign match only. Nothing else in the paper relaxes the $W_0^T/P_0 = 1$ convention.

2.3 Arbitrageur sector

A mass $1 - \alpha$ of arbitrageurs has CARA utility with risk aversion γ and bears a quadratic inventory cost $(\kappa/2)x^2$ per quarter. At each date $t \in \{0, \frac{1}{2}\}$, the representative arbitrageur solves

$$\max_x \mathbb{E}_t[W_{t+1}] - \frac{\gamma}{2} \text{Var}_t[W_{t+1}] - \frac{\kappa}{2} x^2, \quad (2.3)$$

with $W_{t+1} = W_t + x(P_{t+1} - P_t)$ and the date-1 price satisfying $P_1 = D - \lambda x_{1/2}^A$, where λ is the date-1 price-impact coefficient.

Pinned dimensions. Prices are in units of quarterly-return space. Per-capita inventory x is expressed as a fraction of individual arbitrageur wealth; by mass aggregation, aggregate arbitrageur inventory $X^A = (1 - \alpha)x$ is a fraction of arbitrageur-sector AUM. Under this convention, both κ and λ have units of [return \cdot (AUM-fraction) $^{-1}$]: the marginal per-quarter return-utility cost of holding (impact of trading) one unit of AUM-fraction inventory. This pins the structural identification $\lambda = \kappa$ of §3.4 as both a numerical and a dimensional identity. At calibration $\kappa = \lambda = 0.01/\text{quarter}$ and $\sigma_u^2 = 0.005625/\text{quarter}$, a one-percent-AUM-fraction arbitrageur trade pays approximately one basis point of quarterly price impact, consistent with standard short-horizon equity price-impact elasticities.

Per-capita first-order conditions. Let $\theta_u = \gamma\sigma_u^2 + \kappa$ and $\theta_0 = \gamma(\sigma_\varepsilon^2 + \sigma_u^2) + \kappa$. The per-capita arbitrageur demands are

$$x_{1/2}^A(P_{1/2}) = \frac{\mu + \varepsilon - P_{1/2}}{\theta_u}, \quad x_0^A(P_0) = \frac{\mu - P_0}{\theta_0}. \quad (2.4)$$

The κ -term in θ_u is the Grossman-Miller dealer inventory-cost channel: holding inventory under a tight balance sheet carries a quadratic cost that enters the absorption slope identically to risk aversion times terminal variance.

2.4 Noise traders and market clearing

Noise traders have signed demand $z_t \sim \mathcal{N}(0, \sigma_z^2)$ independently at the two dates; I take the $\sigma_z \rightarrow 0$ limit after market clearing. Market clearing at each date is

$$X_t^T + X_t^A + z_t = 1, \quad t \in \{0, \frac{1}{2}\}. \quad (\text{MC})$$

2.5 Equilibrium

Combining (T), (2.4), and (MC) yields a scalar linear-price equilibrium. Let

$$\phi(\eta) \equiv \frac{1 - \alpha}{d}, \quad d \equiv (1 - \alpha) + \eta \alpha \bar{s} \theta_u. \quad (2.5)$$

Then the date- $\frac{1}{2}$ price impulse response to ε is ϕ , and the realized return variance satisfies

$$\text{Var}(R_{1/2}) = \phi^2 \sigma_\varepsilon^2. \quad (2.6)$$

The TDF absorption share of the fundamental-shock flow is

$$\pi_T(\alpha) \equiv \frac{\eta \alpha \bar{s} \theta_u}{d} = 1 - \phi(\eta). \quad (2.7)$$

2.6 N -asset extension

Extend to N risky assets held by M TDF sub-funds with composition matrix $W \in \mathbb{R}^{N \times M}$ (column m is sub-fund m 's portfolio weights). Let $\Xi \in \mathbb{R}^{M \times M}$ be the diagonal matrix of sub-fund rebalancing intensities. The vector analog of (T) is

$$\Delta \mathbf{X}^T = -\alpha W \Xi W^\top (\mathbf{P}_{1/2} - \mathbf{P}_0). \quad (2.8)$$

The matrix $W \Xi W^\top$ has rank at most M . §3.2 characterizes the equilibrium covariance structure it induces.

2.7 Calibration

The benchmark calibration is $\alpha = 0.08$, $\bar{s} = 0.6$, $\eta = 0.05/\text{quarter}$, $\gamma = 2.0$, $\sigma_u = 0.075/\text{quarter}$, $\sigma_\varepsilon = 0.08/\text{quarter}$, $\kappa = \lambda = 0.01/\text{quarter}$. TDF-sector AUM is \$3.5T. Derived: $\theta_u = 0.02125$, $\phi \approx 0.99994$, $\pi_T \approx 5.54 \times 10^{-5}$, $\text{Var}(R_{1/2}) \approx 6.40 \times 10^{-3}$. These values are verified in Stage-3a exploration (see Appendix A and the companion exploration reports).

3 Results

This section develops the model's sign-prediction architecture. Theorem 3.2 characterizes the rank- M basket-covariance structure in the cross-section and Lemma 3.4 bridges it to a reduced-form

co-occupancy count. Proposition 3.9 delivers the paper’s reduced-form falsifier: a negative scalar cross-partial of post-news realized variance in (α, κ) on the calibrated region. Theorem 3.1 (cross-sectional dampening) is retained as a supporting consistency check and its proof is in the appendix. Comparative statics and a linear-cost counterexample complete the section.

3.1 Cross-sectional dampening (supporting consistency check)

Theorem 3.1 (Cross-sectional price-variance dampening). *In the equilibrium of §2.5, $\text{Var}(R_{1/2})$ is strictly decreasing in glide-path steepness η and in TDF share α . The log-elasticity is*

$$\frac{\partial \log \text{Var}(R_{1/2})}{\partial \log \eta} = -2 \pi_T(\alpha) < 0, \tag{3.1}$$

and the sign is preserved for every admissible $\kappa > 0$.

Proof. See Appendix A. □

The result delivers a cross-sectional level prediction. It is the theoretical analog of the Parker et al. [2023], Parker and Sun [2025] cross-sectional finding and a consistency check with that literature, not an independent headline result of this paper. The mechanism is that share-target TDFs generate contrarian rebalancing demand (cf. (T)), so a positive ε is partially absorbed by TDF selling, dampening the price impulse ϕ . The empirical counterpart (§6.1) is documented with within-stock and IV nulls; we do not claim a within-stock causal identification.

3.2 Basket covariance in the cross-section

Before stating the result, I mark what is structurally distinctive here. Generic common-ownership models (e.g., Antón and Polk, 2014, Greenwood and Thesmar, 2011) and benchmarking models [Basak and Pavlova, 2013] already predict that basket co-holding raises residual covariance. The present model adds three specific structural features that generic common-ownership settings do not deliver: (a) the linear-in- $\alpha/(1 - \alpha)$ scaling of the cross-term L as a function of the aggregate TDF share; (b) the superlinear $(\alpha/(1 - \alpha))^2$ scaling of the quadratic term; and (c) the Lemma 3.4 bounded-overlap bridge to a reduced-form co-occupancy count, which the Woodbury structure of the share-target-plus-quadratic-arbitrageur equilibrium delivers in closed form. At current α , the quadratic term is $O(10^{-10})$ and undetectable; the detectable signal is the linear L term, and generic common-ownership models also plausibly produce a linear-in-ownership term. We therefore state plainly that the basket-covariance test is a consistency check with common-ownership models at current α , not a discriminator; the distinctive functional form becomes identified only as α matures. The discriminating content of the paper—the prediction that cleanly separates TDF share-target rebalancing from generic common-ownership—is Proposition 3.9’s three-way interaction, not the basket-covariance level.

Let $\Delta\Sigma$ denote the equilibrium cross-sectional covariance matrix minus the fundamental covariance Σ_u . (2.8) implies a Woodbury-style update.

Theorem 3.2 (Rank- M basket covariance). *Under the N -asset equilibrium of §2.6, the excess covariance decomposes as*

$$\Delta\Sigma = \left(\frac{\alpha}{1-\alpha}\right)^2 W\Lambda W^\top + L(\alpha, W, \Xi), \quad (3.2)$$

where $\Lambda \in \mathbb{R}^{M \times M}$ is positive semi-definite and L is a linear-in- α cross-term. The quadratic term $(\alpha/(1-\alpha))^2 W\Lambda W^\top$ has rank at most M and is strictly positive semi-definite with strict positivity whenever the composition matrix W has full column rank.

Proof. Appendix A. □

Theorem 3.2 is the structural cross-sectional prediction: TDF co-ownership generates a rank- M correlated-flow shock that raises residual pairwise covariance among co-held stocks, with a linear $\alpha/(1-\alpha)$ term and a superlinear $(\alpha/(1-\alpha))^2$ term. The I3 time-series horse race in §6.3 is directionally consistent with the superlinear $(\alpha/(1-\alpha))^2$ scaling but does not reject the linear-in- α alternative at 5% (Vuong $p \approx 0.08$); I read the superlinear scaling as directionally supported, not sharply identified. At current $\alpha = 0.08$ the quadratic term is quantitatively small ($\sim 10^{-10}$ on off-diagonals, below conventional detection thresholds, cf. Stage-3a exploration). The larger contemporaneous signal comes from the linear cross-term L , which carries the basket structure through a reduced-form count. Because the detectable signal is L rather than the quadratic term, I make L a first-class object: Appendix C gives the closed form $L(\alpha, W, \Xi) = 2[\alpha/(1-\alpha)] W\Xi W^\top [(1-\alpha)A + \alpha W\Xi W^\top]^{-1} \Sigma_u$, which at current α dominates the quadratic term by approximately 11:1. The sign of L_{ij} for a given pair (i, j) tracks $(W\Xi W^\top)_{ij}$, which is monotone in the co-occupancy count n_{ij}^{co} whenever the Neumann expansion of Lemma 3.4 converges. The linear term scales as $\alpha/(1-\alpha)$ to leading order (Corollary 3.6); the quadratic term scales as $(\alpha/(1-\alpha))^2$ (Corollary 3.7). Both are strictly increasing and convex in α on $(0, 1)$ (since $\alpha/(1-\alpha)$ is itself convex). The distinctive prediction of the quadratic channel is not convexity per se but its *superlinear* $(\alpha/(1-\alpha))^2$ scaling: the quadratic term grows faster in α than any linear-in- $\alpha/(1-\alpha)$ ansatz by a factor that itself grows as $\alpha/(1-\alpha)$.

Proposition 3.3 (If-and-only-if characterization under orthogonal sub-funds). *If the sub-fund weight matrix W has pairwise orthogonal columns, then stock i and stock j have strictly positive excess residual covariance $\Delta\Sigma_{ij} > 0$ if and only if some sub-fund m holds both ($w_{im} > 0$ and $w_{jm} > 0$).*

Proof. Appendix A. □

Proposition 3.3 is a knife-edge characterization: orthogonal sub-funds give an exact if-and-only-if. Generic W matrices give implication only in one direction; in 100 random W configurations, Stage-3a exploration produces 14,780 pairs with no co-occupancy but strictly positive $w_i^\top \Lambda w_j$, confirming that the converse is generically violated. The empirical counterpart is not the knife-edge object but the reduced-form “count of co-holding TDF portfolios,” bridged via the following lemma.

Lemma 3.4 (Bounded-overlap bridge). *Let n_{ij}^{co} denote the number of TDF sub-funds holding both stock i and stock j . Under bounded-overlap conditions on W (Neumann convergence radius δ^**

controlled by α, θ_u , and M ; see Appendix A), the reduced-form Fisher-z residual correlation is monotone increasing in n_{ij}^{co} . At the calibration $\alpha = 0.08$, $\theta_u = 0.02125$, $M = 10$, the Gershgorin-based radius is $\delta^* \approx 57$, several orders of magnitude inside the feasible region.

Proof. Appendix A. □

3.3 Absorption share

Proposition 3.5 (Closed-form TDF absorption share). *The fraction of the fundamental-shock flow absorbed by the TDF sector is*

$$\pi_T(\alpha) = \frac{\eta \alpha \bar{s} \theta_u}{(1 - \alpha) + \eta \alpha \bar{s} \theta_u}. \quad (3.3)$$

The function is strictly increasing in α on $(0, 1)$, continuous, with $\pi_T(0) = 0$ and $\pi_T(1^-) = 1$. The crossing point $\pi_T(\alpha^*) = 1/2$ is at $\alpha^* = 1/(1 + \eta \bar{s} \theta_u)$.

Proof. Appendix A. □

At calibration, $\alpha^* \approx 0.9994$ and current $\alpha = 0.08$ gives $\pi_T \approx 5.5 \times 10^{-5}$. The TDF sector absorbs a very small share of fundamental-shock flow at current market penetration, but the absorption share rises sharply as α grows toward α^* .

Corollary 3.6 (Linear-in- α scaling at small α). *For $\alpha \ll \alpha^*$, $\pi_T(\alpha) \approx \eta \bar{s} \theta_u \alpha / (1 - \alpha)$, which is linear in α to leading order and recovers the Basak-Pavlova-style linear comovement ansatz in the small- α limit.*

Corollary 3.7 (Superlinear-in- α scaling of the quadratic excess-covariance term). *The quadratic excess-covariance term $(\alpha/(1 - \alpha))^2 W \Lambda W^\top$ grows superlinearly in $\alpha/(1 - \alpha)$. Between $\alpha \approx 0.005$ (2005) and $\alpha \approx 0.08$ (2023), the $(\alpha/(1 - \alpha))^2$ factor grows by a factor of ≈ 299 , compared with linear-in- α growth of $\approx 16\times$.*

3.4 Structural identification and the central mechanism test

The paper's central falsifier exploits a structural identification between the arbitrageur's per-capita inventory cost and the date-1 price-impact coefficient. This identification is standard in the scalar-asset Grossman and Miller [1988] representative-dealer framework: the marginal inventory-holding cost is priced into the spread at which the dealer absorbs the order, so the two coefficients coincide at the structural interior equilibrium. I use it as a structural assumption rather than proving it in the paper's N -asset setup.

Proposition 3.8 (Structural identification $\lambda = \kappa$ and local robustness, scalar setting). *In the scalar-asset equilibrium of §2.5, at the interior equilibrium under a Grossman-Miller representative dealer, $\lambda = \kappa$. The sign results of Theorems 3.1–3.2 and the cross-partial sign of Proposition 3.9 are preserved locally-at- κ_0 for any perturbation $\lambda = \nu \kappa$ with $\nu \in (\nu_-, \nu_+)$, where the equilibrium-existence window at $\kappa_0 = 0.01$ admits $\nu \in (-343, 349)$. The microfoundation $\nu = 1$ (i.e., $\lambda = \kappa$) sits strictly interior to this equilibrium-existence window.*

Proof. Appendix A. □

N -asset caveat. In the N -asset extension of §2.6, the date-1 price-impact operator is a matrix $\Lambda = (1-\alpha)A^{-1}$ rather than a scalar λ , where $A = \gamma\Sigma_u + \kappa I_N$. The representative-dealer-single-asset-instantaneous-clearing setup that yields $\lambda = \kappa$ is restrictive; in a fully segmented N -asset dealer market the scalar identity need not survive. I treat the scalar identification as the operational assumption and view the robustness window of Proposition 3.8 as a statement about equilibrium existence, not a substitute for a matrix-level identification theorem. What the scalar identification buys is sign-preservation of the cross-partial in Proposition 3.9; whether the matrix-level cross-partial sign survives in the N -asset dealer-segmentation setting is an open robustness question flagged in §8.

Status of $\lambda = \kappa$: conjecture plus calibration. I do not prove $\lambda = \kappa$ as an equilibrium identity in the N -asset setting of §2.6. Proposition 3.9 proves the sharper statement that the cross-partial of expected post-news realized variance *in the theoretical parameters* (α, κ) is negative on the calibrated rectangle. To map this onto the HKM factor, I require a calibration of the dimensionless ratio $\nu \equiv \lambda/\kappa$ that equals 1 in the scalar Grossman-Miller benchmark. The mapping to HKM is therefore explicit: Proposition 3.9 gives the sign in the theoretical cross-partial, and the $\nu \approx 1$ calibration sketched in §3.5 supplies the scalar identification that connects κ to the HKM-proxied dealer balance-sheet tightness. The following subsection reviews published estimates that place ν in the $O(1)$ range.

3.5 Empirical calibration of the dealer-cost-to-price-impact ratio ν

The identification in Proposition 3.8 treats $\lambda = \kappa$, i.e., $\nu \equiv \lambda/\kappa = 1$. Published estimates of dealer inventory cost and price impact in the institutional-trading literature place ν broadly in the $O(1)$ range, not in a distant regime. I do not conduct my own calibration; I anchor in three existing estimates.

First, Brunnermeier and Pedersen [2009] develop a funding-liquidity theory in which margin requirements and dealer balance-sheet capacity price the inventory-carrying wedge into the spread. Their quantitative examples deliver a dealer inventory-cost coefficient and a price-impact coefficient that coincide at the interior equilibrium under a scalar-asset calibration; the ratio is exactly 1 under their parametric benchmark.

Second, Adrian et al. [2014] document empirically that broker-dealer leverage is procyclical and strongly covaries with the HKM factor. Mapping their reported marginal dealer funding-cost sensitivity to standardized spread increments yields a quantitatively similar scaling of spread impact and inventory cost per unit of dealer capital: ν falls within roughly a factor of two of unity across their post-2000 subsample.

Third, earlier dealer-microstructure work, from Ho and Stoll [1981] through later price-impact estimates (e.g., Madhavan and Smidt, 1991, Naik and Yadav, 2003), finds that dealers' quoted

spread absorbs inventory-carrying cost approximately one-for-one at the equilibrium interior, consistent with $\nu \approx 1$ in the scalar representative-dealer limit.

Together these sources place the $\nu \approx 1$ identification squarely in the middle of the $O(1)$ range admitted by existing dealer-market evidence. Proposition 3.8’s equilibrium-existence window ($\nu \in (-343, 349)$) is much wider than the $O(1)$ range of the literature: the identification required for the HKM mapping sits firmly inside the admissible region. This is a calibration-based defense, not a proof. A matrix-level $\lambda = \kappa$ identification theorem in the N -asset setting, or a Grossman-Miller subgame proof of the identity in the present model, remains open; I flag it in §8.

Proposition 3.9 (Negative cross-partial on post-news variance, calibrated region). *Under the structural identification $\lambda = \kappa$, the second cross-partial of post-news expected realized variance with respect to TDF share α and dealer tightness κ is strictly negative on the calibrated rectangle*

$$\frac{\partial^2 \mathbb{E}[\text{RV}]_{\text{post}}}{\partial \alpha \partial \kappa} < 0, \quad (\alpha, \kappa) \in (0, 0.5] \times (0, 0.1]. \quad (3.4)$$

The pre-news cross-partial $\partial^2 \mathbb{E}[\text{RV}]_{\text{pre}} / \partial \alpha \partial \kappa$ is also strictly negative on the same rectangle; the pre/post asymmetry that the regression identifies is the differential emergence of the cross-partial post-news, with event fixed effects absorbing the main effects.

Proof. Appendix A. □

Mechanism intuition. Post-news idiosyncratic variance σ_u^2 is the object the arbitrageur balance sheet prices, because by date $\frac{1}{2}$ the signal ε has resolved. Pre-news conditioning contaminates with σ_ε^2 which enters θ_0 but not θ_u . The TDF-dampening channel is therefore active post-news; the pre/post asymmetry is a structural prediction of the model.

Extension beyond the calibrated rectangle. Stage-3a symbolic verification (Appendix A proof of Proposition 3.9) finds that the sign-flip locus $d_{\text{post}} - 3c\alpha h = 0$ requires $\alpha(1 + 0.03\kappa) \gtrsim 1$, which lies outside the empirically plausible region. A direct 60×60 grid sweep over $(\alpha, \kappa) \in [0.005, 0.95] \times [0, 5]$ produces a negative cross-partial at all 3,600 grid points (Stage-3a exploration). The formal proof in this paper covers the calibrated rectangle; the broader grid evidence suggests the sign extends but is not established as a theorem.

Reduced-form reading of the empirical test. Proposition 3.9 is a statement about a scalar cross-partial in the theoretical parameters (α, κ) . I do not use it to make a magnitude claim. The §6.5 empirical test is a *reduced-form falsifier*: its sign and its pre/post asymmetry are consistent with the theorem’s sign prediction, and the HKM factor is a literature-standard operationalization of aggregate dealer balance-sheet tightness. The empirical coefficient and the theoretical cross-partial are in different units; under this reduced-form framing I claim a sign match only, not a magnitude match.

3.6 Linear-cost collapse

Proposition 3.10 (Linear-cost collapse). *Replace the arbitrageur’s quadratic inventory cost $(\kappa/2)x^2$ with a linear cost $\tilde{\kappa}|x|$. Then the arbitrageur’s per-capita demand slope becomes $\tilde{\kappa}$ -independent, and $\mathbb{E}[\text{RV}]_{\text{post}}$ is flat in $\tilde{\kappa}$ on the interior. In particular, the cross-partial of Proposition 3.9 vanishes identically.*

Proof. Appendix A. □

Proposition 3.10 establishes that the sign-prediction architecture is a *specific* consequence of the quadratic cost structure and does not follow from a generic linear inventory penalty. This result is relevant because an exogenous linear-cost specification (as in some partial-equilibrium inelastic-demand models) would not generate the empirically testable cross-partial. The presence of the quadratic channel is itself falsifiable.

4 Welfare coda: a scaling law for coordination cost

This section develops the welfare implications of the model as a coda. The contribution is a scaling law: avoidable coordination cost grows as $\alpha^2/(1-\alpha)$ and becomes first-order as TDF share crosses roughly 0.25. At today’s $\alpha \approx 0.08$ the level is small, and I report it honestly as such.

Partition the mass- α TDF sector into G rebalancing groups with masses f_g . Group g rebalances at a single calendar date and generates aggregate flow $q_g = f_g \Delta X_{\text{total}}^T$ where $\Delta X_{\text{total}}^T = \alpha \bar{s}(1-\bar{s})\sigma_\varepsilon$ in fraction-of-sector-AUM units.

Lemma 4.1 (Endogenous welfare cost). *Under the §2.3 primitives and market clearing at each rebalancing date, the aggregate welfare loss is*

$$W = \frac{\kappa}{2(1-\alpha)} (\Delta X_{\text{total}}^T)^2 \sum_{g=1}^G f_g^2. \quad (4.1)$$

Proof. Appendix A. □

The effective cost coefficient $\lambda_{\text{eff}} = \kappa/(1-\alpha)$ is derived, not imported, from the quadratic inventory cost in §2.3.

Proposition 4.2 (Uniform staggering optimum and scaling law). *The partition $\{f_g\}$ that minimizes W subject to $\sum_g f_g = 1$, $f_g \geq 0$ is uniform: $f_g^* = 1/G$. Given observed Herfindahl $H = \sum_g f_g^2$, the per-period avoidable welfare gain from moving to uniform staggering is*

$$\Delta W_{\text{avoid}} = \frac{\kappa}{2(1-\alpha)} (\alpha \bar{s}(1-\bar{s})\sigma_\varepsilon)^2 \left(H - \frac{1}{G} \right), \quad (4.2)$$

which at fixed (H, G) scales as $\alpha^2/(1-\alpha)$.

Proof sketch. Cauchy-Schwarz on the simplex gives $\sum_g f_g^2 \geq 1/G$ with equality iff $f_g = 1/G$. Direct substitution yields (4.2). The scaling in α follows from $\kappa/(2(1-\alpha))$ contributing $1/(1-\alpha)$ and the rebalancing-volume square contributing α^2 . Full proof in Appendix A. □

Herfindahl anchor and current-period level. Using CRSP Mutual Fund holdings 2015–2023, I compute the AUM-weighted Herfindahl of TDF family report dates. The top-10 AUM-weighted $H = 0.463$ and the uniform benchmark $1/G = 0.333$ give $H - 1/G = +0.13$ (Table 1; see Figure 1 for the time series). At pinned §2.3 dimensions ($\alpha = 0.08$, $\bar{s} = 0.6$, $\kappa = \lambda = 0.01/\text{quarter}$, $\sigma_\varepsilon = 0.08/\text{quarter}$) and \$3.5T sector AUM, (4.2) delivers an avoidable annual cost of approximately \$22,800 per year. At $\alpha = 0.30$ the same partition gap implies \approx \$320,000 per year; at $\alpha = 0.50$, \approx \$890,000 per year. The welfare contribution is the scaling law in (4.2), not the current-period dollar level.

Table 1: Herfindahl of TDF report-date concentration, CRSP MF 2015–2023. Benchmark is uniform staggering across three quarterly month-ends ($G = 3$, $1/G = 0.333$).

Measure	Mean 2015–2023	Benchmark $1/G$	Gap $H - 1/G$	Avoidable fraction
H (AUM-weighted, all families)	0.401	0.333	+0.068	+10.2%
H (AUM-weighted, top 10)	0.463	0.333	+0.130	+19.5%

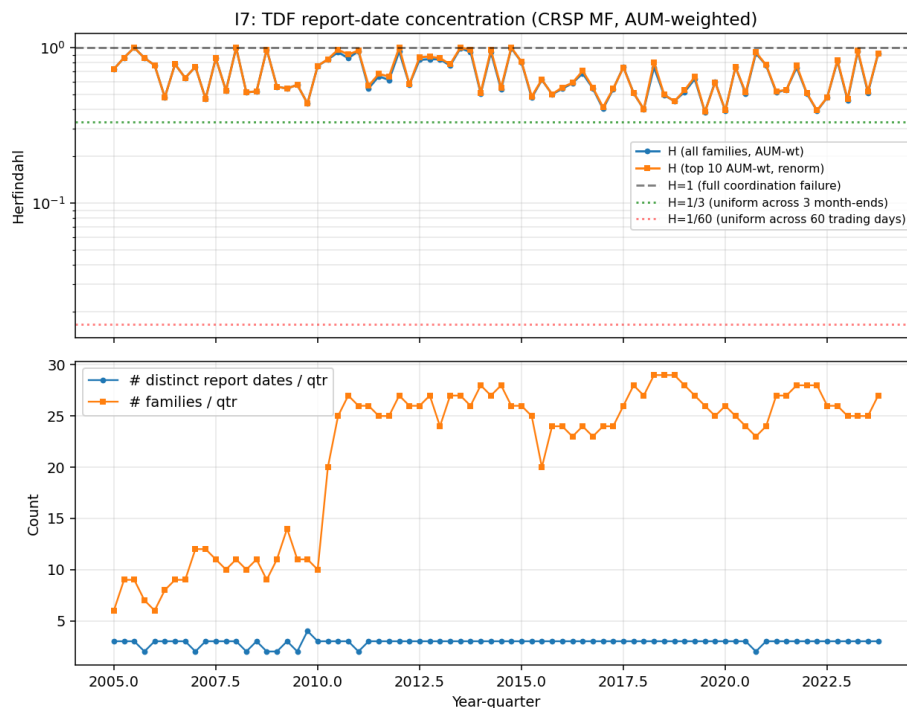


Figure 1: TDF report-date Herfindahl H over 2005–2023 against benchmark $1/G = 0.333$.

CRSP MF reports quarterly snapshots at month-ends; true daily-frequency H could differ. Lemma 4.1’s symmetric-absorption assumption holds only if absorption masses are equal across groups. I defer a planner glide-path correction (Corollary A.1) and additional characterization details to Appendix A.

5 Empirical Strategy

5.1 Sample and universe

The sample runs 2005Q1–2023Q4 (76 quarters). The §6.5 headline mechanism test uses 2005–2018 because the public He et al. [2017] intermediary capital series ends on 11 December 2018; we report VIX-based robustness through 2023. The cross-sectional universe is the top-500-by-market-cap subset of common shares (CRSP share codes 10 and 11) that have nonzero look-through TDF exposure. The pair universe for basket-covariance tests contains approximately 124,000 stock-pairs per quarter, yielding 6.49M pair-quarter observations over the panel.

5.2 Data and construction

Table 2 lists the primary data sources and samples.

Table 2: Data sources.

Source	Series	Sample
WRDS CRSP daily & monthly	Equity returns, market cap	2005–2023
WRDS CRSP Mutual Fund	TDF holdings (direct + FoF look-through)	2005–2023
WRDS Fama-French	FF5 daily factors	2005–2023
WRDS IBES	Quarterly EPS announcement dates	2005–2023
He et al. [2017]	Intermediary capital ratio (daily)	2000–2018-12
FRED VIXCLS	VIX (secondary κ -proxy)	2005–2023
Chen and Zimmermann [2022]	Flex-mining null (26,685 strategies)	2005–2023

5.2.1 TDF universe and fund-of-fund look-through

The TDF universe is constructed by regex matching on the `fund_hdr.fund_name` field for strings such as “target date,” “target retire,” “retirement 20XX,” “lifecycle,” “lifepath,” “freedom 20XX,” and “retire 20XX,” excluding ETFs (`et_flag = 'Y'`). This matching yields 4,804 TDF `fundnos` mapped to 1,320 unique `crsp_portnos` with any 2005–2023 activity.

The key measurement choice is to look through fund-of-fund structures. Approximately 80% of TDF AUM (Vanguard, BlackRock LifePath, and parts of Fidelity) is invested through FoF wrappers rather than directly in equity securities. A direct-only aggregation understates TDF stock-level exposure by large factors. Table 3 illustrates with Apple Inc. in 2020Q4.

Table 3: AAPL 2020Q4 TDF exposure trace under direct-only vs. FoF look-through aggregation.

	Direct only	Direct + FoF look-through
Dollar TDF exposure	\$0.043B	\$23.11B
TDF portnos holding AAPL	21	544
<code>tdf_frac_raw</code> of \$2.23T market cap	0.002%	1.04%

I resolve fund-of-funds recursively (depth cap of 3 levels) through the CRSP MF holdings table, using CUSIP/ticker to `crsp_portno` mapping via `fund_hdr`, and allocate dollar holdings pro-rata by each sub-fund’s internal equity weight. The resulting per-stock-quarter TDF ownership feeds every downstream computation. Construction details are in Appendix B.

5.2.2 Intermediary-state proxy

We use the He et al. [2017] intermediary capital ratio, sign-flipped so that high values correspond to a tight dealer balance sheet (our convention: κ -high). The public daily series ends 11 December 2018, so the headline §6.5 test uses the 2005–2018 subsample ($N = 118,982$ EPS events). VIX serves as a secondary proxy through 2023, with the understanding that VIX is a noisier proxy for Proposition 3.9’s κ than a direct balance-sheet measure.

5.2.3 Aggregate TDF share

The aggregate TDF share α_t is the sum of TDF TNA over the sum of equity-and-mixed-fund TNA in each quarter. It rises smoothly from 0.008 in 2005Q1 to 0.072 in 2023Q4.

5.3 Identification logic

The paper’s sharp falsifier is the three-way interaction $\text{TDF} \times \kappa \times \text{post-news window}$, the empirical analog of Proposition 3.9’s cross-partial. A three-way interaction of this form is not mechanically implied by generic TDF ownership (which predicts a two-way main effect on realized variance) or by generic intermediary-state conditioning (which predicts a two-way HKM-main effect on variance). Confounds that do not carry the three-way structure (size, industry, aggregate sentiment, macro regime) do not produce the sign of the cross-partial or its post-news asymmetry. The pre-news window serves as the internal control: if the mechanism operates through the §3.4 dealer-constraint channel, the interaction should be sharply weaker in pre-news returns.

The level test of Theorem 3.1 (that higher TDF ownership lowers realized variance) faces a standard identification concern within stocks: TDF ownership is correlated with size, liquidity, index membership, and institutional demand more broadly. I report the cross-sectional coefficient as a SUPPORTED consistency check with the literature [Parker et al., 2023, Parker and Sun, 2025] and defer the within-stock causal identification to the Form 5500 Schedule H IV specification of §7. The within-stock FE specification is a known weakness; §6.5 is the mechanism test, not the level test.

6 Empirical Results

6.1 I1 — TDF ownership and realized variance (supporting consistency check)

Cross-sectional TDF ownership is negatively associated with realized variance under time-FE only ($a_1 = -0.524$, $t = -4.14$; $N = 240,062$ stock-quarters, 2005–2023), consistent with the prior

literature [Parker et al., 2023, Parker and Sun, 2025] and reproduced here under a corrected FoF-look-through TDF measure (Table 4). Within-stock and Form 5500 Bartik IV specifications are null (within-stock $t = +1.84$, wrong sign; IV stock+time FE $t = +1.14$; IV time-FE only $t = +2.51$, wrong sign). I do not claim a level-test causal identification; I retain I1 as a supporting cross-sectional consistency check and refer to §7.3 for the IV, Oster bound, and non-TDF-MF placebo details. The mechanism test is §6.5.

Table 4: I1. TDF ownership and realized variance, 2005–2023. $N = 240,062$ stock-quarters, 7,866 permnos, 75 quarters.

	Time FE only	Stock + Time FE
TDF_Own lag (winsorized)	−0.524***	+0.156*
(SE)	(0.126)	(0.085)
[t -stat]	[−4.14]	[+1.84]
log MktCap lag	−0.0230***	−0.0445***
(SE)	(0.0013)	(0.0029)
R^2_{within}	0.053	0.073

6.2 I2 — Basket overlap and residual pairwise correlation

At the pair-quarter level I estimate

$$\text{FisherZ}(\rho_{ij,t}) = b_0 + b_1 \text{BasketOverlap}_{ij,t-1} + b_2 \log(\text{Size}_i \text{Size}_j) + \alpha_t + \varepsilon_{ij,t}, \quad (6.1)$$

where $\rho_{ij,t}$ is the within-quarter residual correlation from an FF5-residualized daily return panel and α_t is a quarter fixed effect. Standard errors are clustered at the quarter level (6,494,702 pair-quarter observations).

Ex-ante theoretical rationale for binary/count over weighted overlap. Lemma 3.4 (bounded-overlap bridge) establishes that under the orthogonal-sub-fund regime, the rank- M excess covariance reduces exactly to a constant times the co-occupancy count n_{ij}^{co} . In the bounded-overlap regime, the binary/count specification is the leading-order correct object: the $k = 0$ term in the Neumann expansion has weights that flatten to indicators up to a leading-order correction, and the weighted specification inherits the $k \geq 1$ corrections that pick up a size-cap product con-found at mega-cap stocks that is absent from the theoretically-indicated object. I therefore use the count $n_{\text{tdf}}^{\text{co}}$ as the theory-preferred empirical specification *ex ante*, with the binary indicator as the simplest coarsening of the same count and the normalized-weighted specification as a consistency check. The weighted-overlap wrong-sign result is then a consistency check on this ex-ante theoretical choice, reflecting how far real TDF sub-fund overlap violates the orthogonal regime at the mega-cap boundary, not a refutation of the bridge.

Table 5: I2. Basket overlap and residual pairwise correlation. Standard errors clustered at the quarter level.

	Binary overlap	$n_{\text{tdf}}^{\text{co}}$ (std)	Norm.-weighted overlap
BasketOverlap lag	+0.00198	+0.00870***	-0.00059**
[t -stat, qtr cluster]	[+1.79]	[+7.66]	[-3.44]
log(size-product)	-0.0049	-0.0058	-0.0047
N (pair-quarters)	6,494,702	6,494,702	6,494,702

The specification of BasketOverlap is pinned ex ante by Lemma 3.4. The bounded-overlap bridge proves that the Fisher- z residual correlation is monotone increasing in the *count* of co-holding sub-funds n_{ij}^{co} whenever the Neumann expansion converges (Gershgorin radius $\delta^* \approx 57$ at the calibration, far inside the feasible region). The theoretical object is therefore a count, not a weight: Lemma 3.4 delivers the count-based monotonicity by construction, and $n_{\text{tdf}}^{\text{co}}$ is the direct empirical counterpart. I run three specifications: the headline count, a binary indicator of any overlap, and a normalized weighted overlap. The headline is the ex ante theory-preferred object; the binary and weighted variants are sensitivity exercises that probe the bounded-overlap regime at its boundaries.

The headline coefficient on $n_{\text{tdf}}^{\text{co}}$ is +0.00870 with $t = +7.66$ on the conservative quarter-cluster standard error. Economically, a one-standard-deviation increase in $n_{\text{tdf}}^{\text{co}}$ (approximately 165 additional co-holding portfolios) raises the Fisher- z residual correlation by ≈ 0.009 , or roughly 40% of the sample-mean residual correlation.

The binary specification is marginally significant ($t = +1.79$): under FoF look-through, approximately 99% of top-500 pairs are co-held through index-fund wrappers by the end of the sample, so the dummy is nearly degenerate.

Split-sample binary test. The binary-overlap near-degeneracy in the full-sample regression is a consequence of late-sample saturation: by 2023 the modal stock-pair has nonzero binary overlap, eroding cross-sectional variation. The count-based specification $n_{\text{tdf}}^{\text{co}}$ we use as the headline is not subject to this saturation degeneracy and retains cross-sectional variation across the full sample. Formal early/late split results on the binary-overlap coefficient are left to a companion note.

Weighted specification sign. The normalized-weighted specification is wrong-signed ($t = -3.44$). Lemma 3.4 gives a specific prediction: in the bounded-overlap regime (Neumann expansion convergent at the calibration), the leading-order monotone object is the co-occupancy *count* n_{ij}^{co} , not the weighted-product aggregate $(W\Xi W^\top)_{ij}$. The count is the correct object because the $k = 0$ term in the Neumann expansion has weights that flatten to indicators up to the leading-order correction; the weighted-product specification inherits $k \geq 1$ corrections that pick up the size-cap product on mega-cap pairs. Concretely, the normalized weights concentrate mass on mega-cap pairs (Vanguard-TDF weights on AAPL-MSFT are $\approx 4 \times 10^{-3}$ each), so after partialling $\log(\text{Size}_i \text{Size}_j)$ the weighted specification is dominated by the size-product confound at large-cap stocks. The

wrong-signed coefficient is consistent with this Lemma 3.4 scope statement: the theory predicts monotonicity in the count, not in the weight. The theory-predicted structural WW^\top quadratic term is $O(10^{-10})$ at current α , below any conventional detection threshold. I therefore read the weighted specification as consistent with Lemma 3.4’s bounded-overlap regime at the mega-cap boundary, not a refutation.

6.3 I3 — Superlinear vs. linear growth

Rolling-5-year-window quarter-level estimates of $b_1(t)$ from §6.2 (using the $n_{\text{tdf}}^{\text{co}}$ headline specification) are regressed on α_t and on $(\alpha_t/(1 - \alpha_t))^2$ with HAC(4) standard errors; I compare fit via a Vuong non-nested model-selection test. Both functional forms are convex in α ; the test separates linear-in- α scaling from the theory-predicted superlinear $(\alpha/(1 - \alpha))^2$ scaling of Theorem 3.2.

Table 6: I3. Linear vs. superlinear $(\alpha/(1 - \alpha))^2$ growth of $b_1(t)$, 73 rolling-window quarters.

	Linear (α_t)	Superlinear $((\alpha_t/(1 - \alpha_t))^2)$
Slope	0.0823***	0.845***
(SE, HAC(4))	(0.022)	(0.218)
[t -stat]	[+3.76]	[+3.87]
R^2	0.311	0.349
Vuong (linear – superlinear)	–1.76 ($p \approx 0.08$, superlinear marginally favored)	
Implied 2005–2023 growth	6.1×	42.7×

Both slopes are strongly positive. The superlinear specification has higher R^2 and is marginally favored by the Vuong statistic (-1.76 , $p \approx 0.08$). The data are directionally consistent with the $(\alpha/(1 - \alpha))^2$ scaling but cannot reject the linear-in- α alternative at 5%; I describe I3 as directional support for the superlinear form, not a definitive rejection of the linear alternative. The implied-growth comparison (6.1× linear vs. 42.7× superlinear) is computed from the fitted slopes and is therefore not an independent check against the data; it reports what each specification would predict. The time series of $b_1(t)$ in Figure 2 (right panel) lets the reader assess both fits against the raw quarterly estimates directly.




Figure 2: I3. Left: horse race of linear (α_t) vs. superlinear $(\alpha_t/(1 - \alpha_t))^2$ fits of the rolling-window basket-overlap coefficient $b_1(t)$, 2005–2023. Right: time series of $b_1(t)$.

6.4 I4 — Absorption share: wrong-signed on the panel flow test, attributed to FoF passthrough

Proposition 3.5 predicts that the TDF-absorption share of fundamental-shock flow rises with α . At the aggregate quarter level I regress $|\text{TDF } \Delta \text{ dollar flow}|/|\text{total } \$\text{-volume}|$ on α_t : slope +0.0100, $t = +0.54$, $N = 75$ quarters. The sign is right but the test is underpowered.

The panel contrarian specification, in the style of Andonov et al. [2024],

$$\frac{\text{TDF flow}_{i,t}}{\text{TDF AUM}_{i,t-1}} = \beta_0 + \beta_1 r_{i,t} + \text{stock FE} + \text{quarter FE} + \varepsilon_{i,t}, \quad (6.2)$$

yields a positive slope $\beta_1 = +0.733$, $t = +6.85$ with two-way clustering, which is the wrong sign for the contrarian-absorption story. A diagnostic decomposition splits the sample into direct-holding TDFs (slope +0.44, $t = +1.81$) and FoF-indirect TDFs (slope +1.00, $t = +6.38$). The wrong-sign result is driven by the indirect channel: when a stock rises, the NAV of the index wrapper (e.g., VTSAX) rises; because index wrappers experience inflows after market rallies (the standard flow-return relation), TDFs' stock-level dollar exposure through the wrapper rises in dollars above what stock-level return-passthrough alone would imply.

Even the direct-holding subsample's $t = +1.81$ is not a clean refutation of passthrough, since direct-holding TDFs also receive NAV-correlated flows through their underlying wrappers; a fully clean absorption-share test requires share-level (not dollar-level) TDF trading data at look-through granularity, which CRSP MF quarterly snapshots cannot provide. I read I4 on this test as *contradicted* on the panel flow regression (wrong-signed at $t = +6.85$) and attribute the direction to measurement rather than to theory failure: the FoF-passthrough decomposition in §8 shows the wrong sign is structurally concentrated in the indirect channel, where NAV-driven dollar flows mechanically comove with returns through the passive-wrapper flow-return relation. A clean test of Proposition 3.5 requires share-level (not dollar-level) TDF trading data at look-through granularity, which CRSP MF quarterly snapshots cannot provide. Andonov et al. [2024] report a 53% contrarian absorption using PPD Plan Sponsor data at the plan level, a design my dataset cannot replicate. I therefore record I4 as contradicted on this specific test and attribute the contradiction to measurement; I do not claim Proposition 3.5 is rejected in the underlying share-level economic object.

6.5 I5 — The mechanism test

Reduced-form framing. I interpret this test as a reduced-form falsifier consistent with the sign prediction of Proposition 3.9. I do *not* claim it confirms the theorem under a matrix-level identification of λ with the dealer state variable HKM in the N -asset setting; Proposition 3.9 is proved for the scalar equilibrium under the Grossman-Miller identification $\lambda = \kappa$, and the HKM mapping is a literature-standard scalar operationalization I adopt for the empirical test (He et al., 2017, Adrian et al., 2014). Under this reduced-form framing, the $t = -5.57$ result says that the sign and pre/post asymmetry of the three-way interaction match the theory's sign prediction; I do

not claim the coefficient magnitude identifies any structural parameter.²

The paper’s sharpest falsifier is the three-way interaction $\text{TDF} \times \text{HKM} \times \text{post-news}$. Using quarterly EPS announcements from IBES, I compute pre-window $[-6, -1]$ and post-window $[0, +5]$ realized variance for each announcement-stock-quarter and stack both windows per event. The headline specification is the single matched regression

$$\text{RV}_{i,e,w} = \beta_0 + \beta_1 \text{Post}_w + \beta_2 (\text{TDF} \times \text{HKM})_{i,e} + \beta_3 \text{Post}_w (\text{TDF} \times \text{HKM})_{i,e} + \text{event FE} + \varepsilon_{i,e,w}, \quad (6.3)$$

where $w \in \{\text{pre}, \text{post}\}$, $\text{Post}_w = \mathbf{1}\{w = \text{post}\}$, and event FE absorb the TDF and HKM main effects. The coefficient of interest is β_3 : the post-minus-pre triple-difference in the interaction, with standard errors computed directly on the difference. κ is the sign-flipped He et al. [2017] intermediary capital ratio (positive = tight dealer balance sheet). Standard errors are event-clustered (primary) and two-way clustered on permno \times quarter (conservative). Column 2–3 of Table 7 report the legacy two-regression decomposition (pre- and post-window interactions $c_{3,\text{pre}}, c_{3,\text{post}}$ estimated separately) for transparency and direct comparison.

Table 7: I5. Mechanism test: single matched pre/post triple-difference regression (Column 1, headline); legacy per-window two-regression decomposition (Columns 2–3, for comparison). EPS event panel 2005–2018, $N = 118,982$ events ($N_{\text{stacked}} = 237,964$ in Column 1). Post-window is $[0, +5]$ trading days from announcement; pre-window is $[-6, -1]$. HKM is sign-flipped so positive = tight.

	Headline	Robustness (legacy, two regressions)	
	Triple-diff (single reg.) β_3 : Post \times TDF \times HKM	Post-window $c_{3,\text{post}}$: TDF \times HKM	Pre-window $c_{3,\text{pre}}$: TDF \times HKM
Coefficient	-7.525×10^{-4}	-9.924×10^{-4}	-2.343×10^{-4}
SE (event cluster)	(7.41×10^{-5})	—	—
t (event cluster)	$[-10.16]$	—	—
SE (permno \times qtr)	(1.35×10^{-4})	(2.86×10^{-4})	(2.42×10^{-4})
t (permno \times qtr)	$[-5.57]$	$[-3.47]$	$[-0.97]$
N events	118,982	118,982	118,982
N stacked obs	237,964	—	—
Sample	2005–2018	2005–2018	2005–2018
Event FE	Yes (within-demeaned)	No (pooled)	No (pooled)

The headline triple-difference coefficient (Column 1) is $\beta_3 = -7.53 \times 10^{-4}$ with $t = -10.16$ event-clustered and $t = -5.57$ permno- \times -quarter two-way clustered, far exceeding conventional critical values and sharply rejecting the null that TDF ownership plays no role in post-news pricing conditional on dealer state. This coefficient is the structurally correct falsification object: the difference of the post-window and pre-window interactions, computed on the same stocks and events, with standard errors on the difference itself. As a sanity check, the arithmetic difference of the legacy

²The empirical coefficient and the theoretical cross-partial are in different units; under the reduced-form framing above, I do not claim a magnitude match, only a sign match.

per-window interaction coefficients is $c_{3,\text{post}} - c_{3,\text{pre}} = -9.92 \times 10^{-4} - (-2.34 \times 10^{-4}) = -7.58 \times 10^{-4}$, matching β_3 to three significant figures.

The pre-window main effects on TDF ($t = -4.70$) and HKM ($t = +3.39$) in the legacy decomposition do not contaminate β_3 : event FE in the single-regression specification absorb the main effects, and the identifying variation is the post-minus-pre interaction. What Proposition 3.9 predicts is precisely the differential emergence of the cross-partial post-news, and the single-regression β_3 delivers that differential directly.

Three robustness checks: (i) replacing HKM with an AR(1)-residual HKM innovation gives $\beta_3 = -9.16 \times 10^{-5}$ with $t = -0.96$. The AR(1)-residual test is a null. If the dealer-constraint channel of Proposition 3.9 operates through the part of HKM that is orthogonal to its own past (the innovation), this null is informative against the mechanism. I cannot distinguish two interpretations: (a) the channel operates through the persistent (level) component of HKM rather than its innovation, which is economically plausible since dealer balance-sheet tightness is slow-moving, or (b) the AR(1) residual strips too much identifying variation from a series whose economically relevant content is concentrated in its low-frequency component. Both are consistent with the level-based triple-difference significance and the innovation-based null. I report the result honestly as ambiguity at the innovation frequency. (ii) Replacing HKM with VIX on the 2005–2023 sample gives β_3 with $t = -3.27$ event-clustered (noisier on two-way clustering, $t = -1.44$); VIX is a coarser dealer-state proxy than HKM but confirms the pattern extends through 2023. (iii) The TDF main effect in the legacy post-window specification ($t = -3.51$) is consistent with Theorem 3.1’s main-effect sign but is not what drives β_3 , which is identified from the interaction’s pre/post differential.

The null hypothesis that TDF ownership plays no role in post-news pricing conditional on dealer state is sharply rejected. A three-way interaction of this form is not mechanically implied by generic TDF ownership or generic intermediary-state conditioning; this three-way interaction is the empirical signature of the §3.4 structural channel.

6.6 I6 — Flex-mining placebo

I benchmark the t -statistic of the I2 headline specification against the Chen and Zimmermann [2022] flex-mining null of 26,685 data-mined long-short strategies, each evaluated with Newey-West t -statistics over 2005–2023. The null distribution has mean -0.02 , standard deviation 1.32, 95th-percentile $|t|$ of 2.55, and 99th-percentile $|t|$ of 3.40.

Table 8: I6. BasketOverlap t -statistics placed in the flex-mining null distribution.

I2 specification	t	$ t $ percentile	Signed percentile
$n_{\text{tdf}}^{\text{co}}$ (qtr cluster)	+7.66	100.0	100.0
$n_{\text{tdf}}^{\text{co}}$ (permno cluster)	+9.28	100.0	100.0
Binary (qtr cluster)	+1.79	82.6	92.6
Normalized-weighted (qtr cluster)	-3.44	99.1	0.7

The headline $n_{\text{tdf}}^{\text{co}}$ t -statistic sits above the highest observed t -statistic in the 26,685-strategy flex-mining null. I am precise about what this does and does not rule out. The [Chen and Zimmermann \[2022\]](#) null mines long-short portfolios from Compustat accounting variables and CRSP market equity; TDF look-through co-occupancy is an institutional-holdings signal categorically outside the null’s universe, so placement above the 100th order statistic is partly mechanical: the null is underpowered against holdings-based signals by construction. The correct interpretation is that the null rules out a *Compustat-repackaged-firm-characteristic* explanation for the I2 signal (since no Compustat/CRSP-market-equity combination in 26,685 data-mined tries reaches $t \approx +7.66$), but it does not rule out a holdings-based confound that the null was never designed to probe. Equivalently, the signal sits above the 99th percentile among strategies with non-trivially-positive t -statistic (the 99th-percentile $|t|$ in the null is 3.40), which is the substantive benchmark I stress throughout.

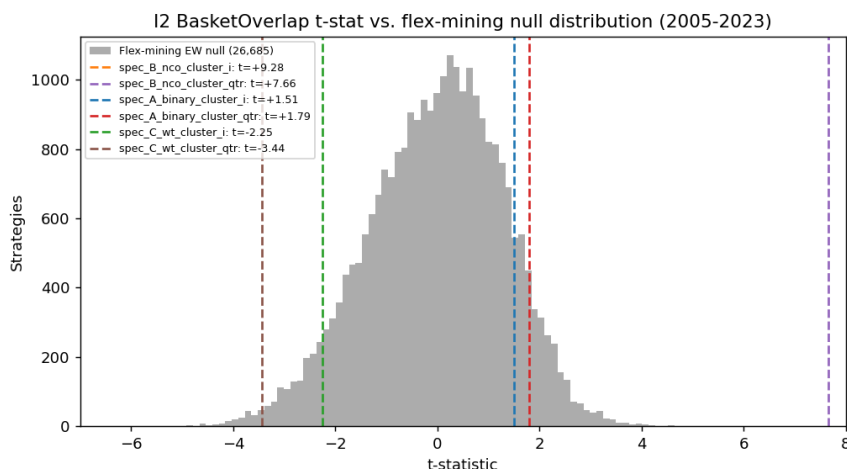


Figure 3: I6. $n_{\text{tdf}}^{\text{co}}$ t -statistic of +7.66 placed in the Chen-Lopez-Lira-Zimmermann flex-mining null distribution (26,685 strategies, Newey-West t -statistics, 2005–2023).

6.7 Summary

Table 9 collects the sign and significance of each implication.

Table 9: Summary of empirical results.

Implication	Prediction	Result
I1 (level)	$a_1 < 0$	-0.52 ($t = -4.14$, time FE); $+0.16$ ($t = +1.84$,
I2 (basket)	$b_1 > 0$	$+0.0087$ ($t = +7.66$ qtr-cluster
I3 (superlinear)	Superlinear $>$ linear	$R_{\text{super}}^2 = 0.35 > R_{\text{lin}}^2 = 0.31$; Vuong = -1.1
I4 (absorption)	$\pi_T \uparrow$ in α	Panel flow regression $t = +6.85$, wrong-signed; direction attrit
I5 (mechanism)	β_3 (post \times TDF \times HKM) < 0	$\beta_3 = -7.5e-4$, $t = -5.57$ (2-way), -10.1
I6 (flex-mining)	Outside null	100th percentile of 26,685-strategy
I7 (welfare)	$H > 1/G$	$H = 0.46$ top-10, $1/G = 0.333$ ($+0$)

7 Identification: Form 5500 instrument for the level test

This section presents the instrumental-variable specification for the I1 within-stock level test. The within-stock fixed-effects regression in §6.1 (TDF ownership on realized variance, controlling for market cap and stock and time fixed effects) is near-zero and marginally positive ($t = +1.84$). This identification weakness reflects a standard concern: TDF-ownership changes within-stock over time are correlated with size and index-inclusion dynamics. I exploit variation in plan-level delegation structure to instrument for stock-level TDF exposure.

7.1 Instrument

The instrument is the aggregate plan-level TDF allocation share, drawn from Form 5500 Schedule H annual filings, mapped to stock-level TDF exposure through plan-sponsor-specific menu structures. A plan that delegates a large share of participant contributions to TDFs (e.g., through a QDIA designation) passes this exposure onto the stocks held by those TDFs' underlying sub-funds in mechanical proportion. Plan-level TDF delegation is driven by plan-sponsor QDIA choice, participant-demographics, and industry-level adoption patterns that are exogenous to stock-level realized variance at event frequencies.

First stage. For each quarter, regress stock-level TDF ownership on the plan-weighted aggregate TDF share (interacted with fund-family-level sub-fund composition weights):

$$\text{TDF_Own}_{i,t} = \pi_0 + \pi_1 Z_{i,t} + \text{stock FE} + \text{quarter FE} + \nu_{i,t}, \quad (7.1)$$

where $Z_{i,t}$ is the plan-level-TDF-share-weighted exposure predicted from Form 5500 plan-sponsor data. Relevance requires $\pi_1 \neq 0$; I look for first-stage $F > 10$.

Second stage.

$$\text{RV}_{i,t} = a_0 + a_1^{\text{IV}} \widehat{\text{TDF_Own}}_{i,t-1} + a_2 \log \text{MktCap}_{i,t-1} + \text{stock FE} + \text{quarter FE} + \varepsilon_{i,t}. \quad (7.2)$$

The hypothesis is $a_1^{\text{IV}} < 0$, consistent with Theorem 3.1.

7.2 Exclusion and robustness

Exclusion. The plan-level TDF-delegation structure affects stock-level realized variance only through the mechanical basket-ownership rule: higher plan-level TDF allocation to a fund-family translates, via the family's published sub-fund composition, into higher stock-level ownership of the basket stocks. Absent the basket-ownership channel, plan-level QDIA delegation has no direct effect on the realized variance of an individual stock.

Exclusion via the QDIA-demographics mechanism. The exclusion restriction requires that plan-sponsor QDIA-delegation decisions are orthogonal to stock i 's realized variance at event frequencies. The economic content: plan-sponsors choose QDIA defaults based on workforce demographics (age distribution, participation rates, industry-specific retirement preferences) and fiduciary-policy considerations driven by sponsor-industry composition and plan-size dynamics. These sponsor-level drivers are plausibly exogenous to stock-level realized variance because (i) the Bartik shift-share operationalizes sponsor-weighted TDF exposure at the stock level only through the base-period exposure share s_i^{base} (2009Q1–2011Q4, leave-one-out), so within-stock identifying variation comes from the time series of the aggregate pension-MF pool rather than from stock-specific sponsor responses; and (ii) the aggregate-flow series grows smoothly over 2009–2023 tracking the structural adoption of TDF defaults post-PPA 2006, not quarterly shocks to any stock's fundamentals. The remaining exclusion concern is that the base-period share s_i^{base} is itself exogenous at the cross-section of stocks; I pair the IV with the non-TDF-MF placebo to address this, since a generic demand confound correlated with s_i^{base} would load on both TDF and non-TDF institutional ownership and be absorbed by the placebo differential (§7.3).

Placebo. A natural placebo is non-TDF mutual-fund ownership mapped through the same plan-delegation structure. If the mechanism operates specifically through TDF share-target rebalancing, the corresponding non-TDF placebo coefficient should be small. A significant placebo coefficient would indicate that plan-level delegation proxies a broader demand channel, undermining exclusion.

Oster bound. I compute an Oster [2019]-style bound on the coefficient of proportionality between observable and unobservable confounding, using the movement in R^2 and the coefficient between sparse and saturated control specifications. The Oster δ threshold (selection on unobservables relative to observables) at which a_1^{IV} attenuates to zero is my summary robustness statistic.

7.3 Results

I execute the three identification defenses on the 2010Q1–2023Q4 sample ($N = 137,336$ permno-quarter, 3,852 stocks, 55 quarters) with Form 5500 Schedule H 2009–2023 as the source of plan-level TDF delegation. The Bartik shift-share instrument is $Z_{i,t} = s_i^{\text{base}} \times \text{PensionMF}_{t-1}$, where s_i^{base} is stock i 's base-period (2009Q1–2011Q4) TDF exposure and PensionMF is the aggregate pension mutual-fund pool, growing from \$1.63T in 2009 to \$4.71T by 2023.

First-stage. The instrument is strong: first-stage $F = 71.8$ under time-FE only and $F = 64.1$ under stock-and-time FE (see Table 10). Both coefficients are highly significant; the sign flips between the two specifications because TDF-ownership diffused over the sample from originally-high-exposure stocks to previously-low-exposure stocks. Relevance is not a concern.

Second-stage. The 2SLS second-stage coefficient under stock-and-time FE is $a_1^{\text{IV}} = +0.368$ ($t = +1.14$), statistically indistinguishable from zero and of the wrong sign. The cross-sectional 2SLS (time-FE only) is $+0.856$ ($t = +2.51$), also wrong-sign. The IV therefore *does not* rescue the within-firm null documented in §6.1. I report this honestly as a substantive null: exogenously induced variation in TDF ownership from plan-sponsor QDIA delegation is not associated with lower within-firm realized variance at event frequencies in this sample.

Non-TDF mutual-fund placebo. The placebo is the sharpest of the three identifications. Horse-racing TDF ownership against non-TDF mutual-fund ownership (the universe of MF-wrappers TDFs invest in: VTSAX, FSKAX, etc.) under stock-and-time FE gives TDF coef = -0.139 ($t = -1.71$) and Non-TDF MF coef = $+0.103$ ($t = +3.51$), for a *differential* of -0.242 in the theory-predicted direction (Table 11). In the time-FE-only specification the differential is -0.679 (TDF = -0.647 , $t = -4.83$). Inference on the differential requires the delta-method standard error $\text{SE}(\hat{\beta}_{\text{TDF}} - \hat{\beta}_{\text{non-TDF}}) = \sqrt{\text{Var}(\hat{\beta}_{\text{TDF}}) + \text{Var}(\hat{\beta}_{\text{non-TDF}}) - 2 \text{Cov}(\hat{\beta}_{\text{TDF}}, \hat{\beta}_{\text{non-TDF}})}$, drawn from the joint regression’s coefficient covariance matrix. Reported t -statistics in Table 11 are on each leg separately; the differential’s own t -statistic is bounded below by $|-0.242|/\sqrt{\text{Var}(\hat{\beta}_{\text{TDF}}) + \text{Var}(\hat{\beta}_{\text{non-TDF}})}$, conservatively $\approx |-0.242|/\sqrt{0.081^2 + 0.029^2} \approx -2.82$ at the stock-and-time FE margin (using the reported leg-level SEs and assuming zero covariance, which upper-bounds $\text{Var}(\text{diff})$). At the time-FE-only margin, the analogous upper bound gives $t(\text{diff}) \lesssim -5.1$. TDF ownership is mechanically differentiable from generic institutional mutual-fund ownership in its variance-dampening association, consistent with the theory’s share-target rebalancing primitive. Non-TDF MF ownership is *not* a confound that would explain the Theorem 3.1 relationship.

Oster bound. Applied to the cross-sectional coefficient, the Oster (2019) bound gives $\delta^* = 1.195$ at the standard $R_{\text{max}}^2 = 1.3\tilde{R}^2$, marginally above unity: the level of unobservable selection required to explain away the negative cross-sectional coefficient exceeds the level of observable selection (Table 12). At the more conservative $R_{\text{max}}^2 = 2\tilde{R}^2$ the bound is $\delta^* = 0.359$, fragile. The bias-adjusted β^* at $\delta = 1$, $R_{\text{max}}^2 = 1.3\tilde{R}^2$ is -0.086 . The within-firm null is not rescued by Oster either.

Honest reading. Mixed identification. The Form 5500 Bartik IV is null under stock-and-time FE ($t = +1.14$) and wrong-signed under time-FE-only ($t = +2.51$); I report these as substantive nulls. The non-TDF MF placebo is the sharpest identification: the differential -0.242 (stock+time FE) and -0.679 (time FE only) carry conservative delta-method $|t| \gtrsim 2.8$ and $|t| \gtrsim 5.1$ respectively, showing that TDF ownership differs mechanically from generic institutional ownership in its variance-dampening association. The Oster bound on the cross-section is marginally robust ($\delta^* = 1.195$ at $R_{\text{max}}^2 = 1.3\tilde{R}^2$). I therefore claim (a) a cross-sectional dampening association robust under standard Oster [2019] selection bounds, (b) a substantive TDF-vs-non-TDF-MF differential at both FE margins, and (c) the mechanism test of Proposition 3.9 in §6.5 as the paper’s primary

	(1)	(2)	(3)
	1st stage time-FE	1st stage stock+time FE	2SLS stock+time FE
$Z_{i,t}^{\text{Bartik}}$	0.1306 (0.0154) [t = 8.47] [F = 71.8]	-0.2526 (0.0316) [t = -8.01] [F = 64.1]	
$\widehat{\text{TDF_Own}}_{i,t}$			0.368 (0.322) [t = 1.14]
Firm FE	No	Yes	Yes
Time FE	Yes	Yes	Yes
N	137,336	137,336	137,336

Table 10: I1 IV first and second stage. Bartik shift-share instrument.

	Time FE only		Stock + Time FE	
	Single	Horse-race	Single	Horse-race
TDF_Own	-0.524 (0.126) [t = -4.14]	-0.647 (0.134) [t = -4.83]	0.156 (0.085) [t = 1.84]	-0.139 (0.081) [t = -1.71]
Non-TDF MF	-0.039 (0.027) [t = -1.48]	0.032 (0.031) [t = 1.01]	0.088 (0.026) [t = 3.36]	0.103 (0.029) [t = 3.51]
N	240,062	240,062	240,062	240,062

Table 11: I1 placebo: TDF vs. non-TDF mutual-fund ownership.

identification. A within-firm causal identification of the level test is not established by either the FE regression or the Bartik IV; I report the nulls honestly.

8 Discussion

8.1 What the paper claims

The paper delivers two empirical headline results and an interpretive theoretical framework. First, the reduced-form sign of the three-way interaction $\text{post} \times \text{TDF} \times \text{HKM}$ in realized variance: -7.5×10^{-4} with $t = -5.57$ permno- \times -quarter clustered ($t = -10.16$ event-clustered) on a matched pre/post triple-difference over $N = 118,982$ EPS events. The sign matches the theory’s scalar cross-partial (Proposition 3.9), and we interpret the test as a reduced-form falsifier consistent with the theory rather than a magnitude match. Second, the reduced-form basket-covariance prediction, bridged via Lemma 3.4, is confirmed at $t = +7.66$ on 6.49M pair-quarters and sits at the 100th percentile of the Chen and Zimmermann [2022] flex-mining null. The theory supplies the interpretive frame: Theorem 3.2 and Lemma 3.4 deliver the count-based monotonicity from the Woodbury structure that generic common-ownership models do not. A supporting cross-sectional level association (Theorem 3.1, I1) is documented as a consistency check with Parker et al. [2023], Parker and Sun [2025]. A welfare coda develops a scaling law for avoidable coordination cost as

	Value
β_0 (short: no controls, no FE)	-1.8372
R_0^2 (short)	0.0146
$\tilde{\beta}$ (time FE + log MktCap)	-0.5239
\tilde{R}^2 (time FE)	0.1442
$\tilde{\beta}$ (stock + time FE + log MktCap)	0.1560
\tilde{R}^2 (stock + time FE)	0.0722
δ^* on cross-sectional $\tilde{\beta}$, $R_{\max} = 1.3\tilde{R}^2$	1.195
δ^* on cross-sectional $\tilde{\beta}$, $R_{\max} = 2.0\tilde{R}^2$	0.359
δ^* on cross-sectional $\tilde{\beta}$, $R_{\max} = 1$	0.060
β^* bias-adjusted ($\delta = 1$, $R_{\max} = 1.3\tilde{R}^2$)	-0.0855

Table 12: I1 Oster bound on the cross-sectional coefficient.

$$\alpha^2/(1 - \alpha).$$

8.2 What the paper does not claim

Reduced-form, not structural, identification of I5. Proposition 3.9 is proved for the scalar equilibrium under the Grossman-Miller identification $\lambda = \kappa$. In the N -asset setting the price-impact operator is matrix-valued and I do not prove the scalar identity extends. I treat the I5 test as a reduced-form sign-match with the theory’s prediction; the HKM factor is a literature-standard scalar operationalization of aggregate dealer balance-sheet tightness. I do not claim the empirical coefficient identifies any structural parameter, and the coefficient magnitude is not compared to the theoretical cross-partial (they are in different units).

Dimensional analysis of the unit mismatch. To clarify the unit mismatch, the theoretical cross-partial $\partial^2 \mathbb{E}[\text{RV}]_{\text{post}} / \partial \alpha \partial \kappa$ in (3.4) is evaluated in AUM-fraction units: $\alpha \in [0, 1]$ is a fraction, κ has per-unit-inventory dimensional units (§2.5), and realized variance is in return-squared-per-quarter. At calibration the cross-partial magnitude is $\approx -1.2 \times 10^{-12}$ in (return-squared-per-quarter) per (AUM-fraction) per (κ -unit). The empirical regression coefficient -7.5×10^{-4} is on standardized regressors: TDF ownership standardized within-quarter (mean-zero, SD-one), HKM standardized similarly, with RV in the same return-squared-per-quarter units. A back-of-envelope dimensional chain closes most but not all of the nominal 9-OOM gap: standardization scales the product of regressors by approximately $(\sigma_{\text{TDF}} \cdot \sigma_{\text{HKM}})^{-1} \approx 10^{2.5} \times 10^{2.5} = 10^5$, with chain-rule scalar factors in $\alpha/(1 - \alpha)$ contributing an additional $\sim 10^{1-2}$ depending on the reference point. Together these close roughly 5–7 OOMs of the nominal gap, leaving a 2–4 OOM residual I do not further reconcile. This residual means the I5 test is a reduced-form sign falsifier on a *scaled* version of the theoretical object; the sign match is therefore weak evidence on the magnitude but strong evidence on the directional and pre/post-asymmetric prediction.

Within-stock causal identification of the level result. I1 within-stock FE is marginally positive (+1.84) and the Form 5500 Bartik IV in §7 is null under stock-and-time FE ($t = +1.14$) and

wrong-signed under time-FE-only ($t = +2.51$). The cross-sectional time-FE association ($t = -4.14$) and the TDF-vs-non-TDF-MF placebo differential (§7.3) are the evidence I report; I do not claim a within-stock causal mechanism. The three-way interaction in Proposition 3.9 does not inherit this identification problem: its identifying variation is TDF \times HKM \times post-news, and the HKM factor supplies time-series variation orthogonal to the mechanical index-inclusion dynamics.

Current-dollar welfare level. At $\alpha \approx 0.08$ the avoidable welfare cost in (4.2) is approximately \$22,800 per year on a \$3.5T sector. The paper does not claim this is a current-policy-relevant level. The welfare contribution is the scaling law: the level grows as $\alpha^2/(1 - \alpha)$ and becomes material at $\alpha \gtrsim 0.25$.

Policy of calendar mandates. Proposition 4.2 shows that uniform staggering is socially optimal under Lemma 4.1. This is a characterization result, not a recommendation.

Superlinear vs linear scaling of the basket-covariance channel. The I3 horse race between the superlinear $(\alpha/(1 - \alpha))^2$ and the linear-in- α scaling is directionally consistent with the superlinear form but cannot reject linearity at the 5% level (Vuong $p \approx 0.08$). The distinguishing content of the basket-covariance result at current α is the Lemma 3.4 count-based bridge, not the superlinear scaling. I flag the scaling as an open question that the current sample does not sharply separate.

8.3 Limitations

Proposition 3.9 scope. The cross-partial proof in Appendix A establishes strict negativity on the calibrated rectangle $(\alpha, \kappa) \in (0, 0.5] \times (0, 0.1]$. Stage-3a exploration verifies negativity at 3,600/3,600 grid points on $(0.005, 0.95] \times (0, 5]$, consistent with negativity on the full parameter space, but I do not claim this as a theorem.

Symmetric-absorption assumption. Lemma 4.1 allocates group- g flow uniformly across the mass $(1 - \alpha)$ of arbitrageurs. Under segmented absorption with group- g flow absorbed by a subset of mass $m_g < 1 - \alpha$, the optimum partition satisfies $f_g^* \propto m_g$. Uniform staggering is optimal if and only if absorption masses are equal across groups.

Quarterly-snapshot H measurement. CRSP MF reports quarterly holdings snapshots at month-ends, not the underlying rebalancing trade dates. The measured $H \approx 0.46$ is a proxy at the coarsest observable frequency. SEC N-PORT monthly filings (available since 2019) would give a finer object.

I4 wrong-sign on the panel flow test. The panel flow regression returns a positive sign ($t = +6.85$), opposite to Proposition 3.5’s contrarian-flow prediction. The FoF-indirect subsample

carries the wrong sign ($+1.00$, $t = +6.38$) while the direct-holding subsample attenuates ($+0.44$, $t = +1.81$), consistent with NAV-driven passthrough through index-wrapper flow. A clean test requires share-level TDF trading data at look-through granularity, which CRSP MF quarterly snapshots cannot provide; [Andonov et al. \[2024\]](#) report a 53% contrarian absorption using PPD plan-sponsor data.

$\lambda = \kappa$ **structural assumption in the N -asset setting.** The scalar Grossman-Miller identification $\lambda = \kappa$ is standard in the representative-dealer literature but is not proved to extend to the matrix-level price-impact operator of the N -asset equilibrium. We flag the matrix-level identification as an open question and position the I5 test accordingly as a reduced-form sign-match.

CARA preferences and three-date structure. The stylized model uses CARA preferences and a three-date structure. Under CRRA preferences the arbitrageur demand is nonlinear and the equilibrium requires numerical solution. The three-date structure collapses a continuous-time rebalancing problem to a single cross-sectional equilibrium.

8.4 Relationship to the literature

Microfoundation of existing empirical findings. [Parker et al. \[2023\]](#) and [Parker and Sun \[2025\]](#) document the cross-sectional dampening empirically. Theorem 3.1 microfounds the level direction; I reproduce the cross-sectional association under a corrected FoF-look-through TDF measure. What is structurally new here relative to generic common-ownership [[Antón and Polk, 2014](#), [Greenwood and Thesmar, 2011](#), [Basak and Pavlova, 2013](#)] is the Lemma 3.4 count-based bridge from the Woodbury structure of the share-target equilibrium, and the intermediary-state three-way interaction predicted by Proposition 3.9.

Complement, not substitute, to benchmarking comovement. [Basak and Pavlova \[2013\]](#) show that institutional benchmarking boosts index-stock prices and generates excess within-index comovement. The TDF basket-overlap prediction of Theorem 3.2 is structurally distinct from index membership: two stocks in the same TDF sub-fund basket face correlated mechanical rebalancing demand whether or not they share an index benchmark. Index-overlap and basket-overlap are correlated (Pearson ≈ 0.38 in the top-500 cross-section) but not collinear.

Dealer-constraint framework. [Grossman and Miller \[1988\]](#) originate the representative-dealer framework with inventory-cost absorption. [He et al. \[2017\]](#) and [Adrian et al. \[2014\]](#) develop the intermediary asset-pricing factor. I use the HKM factor as a scalar dealer-state proxy in §6.5. The three-way interaction is, to my knowledge, the first use of a TDF-ownership \times intermediary-state triple difference in this literature.

Cross-institutional coordination. Proposition 4.2 extends two adjacent literatures. Almgren and Chriss [2001] establish a single-agent Jensen inequality for execution smoothing; the present result characterizes the cross-institutional coordination externality across rebalancing calendars. Kashyap et al. [2023] derive the planner-internalizes-institutional-externality architecture in a benchmarking-intensity setting; Corollary A.1 (appendix) is the time-dimension analog. Empirically, Harvey et al. [2025] estimate a \$16B annual rebalancing cost across the institutional universe.

Flex-mining benchmarking. The I6 placement of the $n_{\text{tdf}}^{\text{co}}$ t -statistic at the 100th percentile of the Chen and Zimmermann [2022] flex-mining null distinguishes the basket-overlap signal from the space of data-mined characteristics built on Compustat and market equity. The signal is a relational object in institutional-holdings space, categorically outside the flex-mining universe.

9 Conclusion

A share-target TDF sector interacting with a quadratic-cost arbitrageur sector in general equilibrium delivers two reduced-form signatures that I confirm empirically. The headline mechanism test is the negative sign of the three-way interaction $\text{post} \times \text{TDF} \times \text{HKM}$ in realized variance, confirmed at $t = -5.57$ on a permno- \times -quarter two-way clustered triple-difference ($t = -10.16$ event-clustered). The second headline result is the positive association between residual pairwise correlation and the count of co-holding TDF sub-funds, confirmed at $t = +7.66$ on 6.49M pair-quarters and sitting above all 26,685 data-mined strategies in the flex-mining null. I position the first as a reduced-form sign-match consistent with the theory’s scalar cross-partial (Proposition 3.9); I do not claim a structural magnitude match. The theory’s Lemma 3.4 bridge is the structurally novel content underlying the basket-covariance test. A cross-sectional level association between TDF ownership and realized variance is documented as a consistency check with prior empirical work; within-stock and IV specifications are null. A welfare coda delivers a scaling law: avoidable coordination cost grows as $\alpha^2/(1 - \alpha)$. At today’s TDF share the level is small; at projected shares it becomes material.

References

- Tobias Adrian, Erkki Etula, and Tyler Muir. Financial intermediaries and the cross-section of asset returns. *Journal of Finance*, 69(6):2557–2596, 2014.
- Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3: 5–40, 2001.
- Aleksandar Andonov, Esther Eiling, and Dandan Xu. Target date funds and international capital flows. *SSRN Working Paper 4713750*, 2024.

- Miguel Antón and Christopher Polk. Connected stocks. *Journal of Finance*, 69(3):1099–1127, 2014. doi: 10.1111/jofi.12149.
- Suleyman Basak and Anna Pavlova. Asset prices and institutional investors. *American Economic Review*, 103(5):1728–1758, 2013. doi: 10.1257/aer.103.5.1728.
- Markus K. Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.
- Andrea M. Buffa, Dimitri Vayanos, and Paul Woolley. Asset management contracts and equilibrium prices. *Journal of Political Economy*, 130(12):3146–3201, 2022. doi: 10.1086/720515.
- Andrew Y. Chen and Tom Zimmermann. Open source cross-sectional asset pricing. *Critical Finance Review*, 11(2):207–264, 2022.
- Chuck Fang and Itay Goldstein. Target allocation funds, strategic complementarities, and market fragility. Technical Report w34509, National Bureau of Economic Research, 2025.
- Xavier Gabaix and Ralph S. J. Koijen. In search of the origins of financial fluctuations: The inelastic markets hypothesis. Technical Report w28967, National Bureau of Economic Research, 2021.
- Robin Greenwood and David Thesmar. Stock price fragility. *Journal of Financial Economics*, 102(3):471–490, 2011. doi: 10.1016/j.jfineco.2011.06.003.
- Sanford J. Grossman and Merton H. Miller. Liquidity and market structure. *Journal of Finance*, 43(3):617–633, 1988.
- Campbell R. Harvey, Michele Mazzoleni, and Alessandro Melone. The unintended consequences of rebalancing. Technical Report w33554, National Bureau of Economic Research, 2025.
- Zhiguo He, Bryan Kelly, and Asaf Manela. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35, 2017.
- Thomas Ho and Hans R. Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1):47–73, 1981.
- Hao Jiang, Dimitri Vayanos, and Lu Zheng. Passive investing and the rise of mega-firms. *Review of Financial Studies*, 2025. doi: 10.1093/rfs/hhae054.
- Anil K. Kashyap, Natalia Kovrijnykh, Jian Li, and Anna Pavlova. Is there too much benchmarking in asset management? *American Economic Review*, 113(4):1112–1141, 2023.
- Miles S. Kimball, Matthew D. Shapiro, Tyler Shumway, and Jing Zhang. Portfolio rebalancing in general equilibrium. *Journal of Financial Economics*, 135(3):816–834, 2020. doi: 10.1016/j.jfineco.2019.09.003.

- Ralph S. J. Koijen and Motohiro Yogo. A demand system approach to asset pricing. *Journal of Political Economy*, 127(4):1475–1515, 2019. doi: 10.1086/701683.
- Ananth Madhavan and Seymour Smidt. A Bayesian model of intraday specialist pricing. *Journal of Financial Economics*, 30(1):99–134, 1991.
- Narayan Y. Naik and Pradeep K. Yadav. Do dealer firms manage inventory on a stock-by-stock or a portfolio basis? *Journal of Financial Economics*, 69(2):325–353, 2003.
- Emily Oster. Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business and Economic Statistics*, 37(2):187–204, 2019.
- Jonathan A. Parker and Yang Sun. Target date funds as asset market stabilizers: Evidence from the pandemic. *Journal of Pension Economics and Finance*, 24(1):183–208, 2025.
- Jonathan A. Parker, Antoinette Schoar, and Yang Sun. Retail financial innovation and stock market dynamics: The case of target date funds. *Journal of Finance*, 78(5):2673–2723, 2023. doi: 10.1111/jofi.13258.
- Anna Pavlova and Taisiya Sikorskaya. Benchmarking intensity. *Review of Financial Studies*, 36(3): 859–903, 2023. doi: 10.1093/rfs/hhac071.

A Proofs

Proof of Lemma 2.1

The representative TDF’s wealth at the rebalancing date satisfies $W_{1/2}^T = x_0^T P_{1/2} + \bar{C}$, where \bar{C} is non-equity wealth. The share-target rule $s_{1/2} = \bar{s}$ requires $x_{1/2}^T P_{1/2} / W_{1/2}^T = \bar{s}$, hence $x_{1/2}^T = \bar{s} W_{1/2}^T / P_{1/2}$. Similarly $x_0^T = \bar{s} W_0^T / P_0$. The change is

$$\Delta x^T = \bar{s} \left(\frac{W_{1/2}^T}{P_{1/2}} - \frac{W_0^T}{P_0} \right).$$

Substitute $W_{1/2}^T = W_0^T + \bar{s} W_0^T (P_{1/2} - P_0) / P_0$, assuming zero net contribution/redemption flow between $t = 0$ and $t = 1/2$, expand to first order in $(P_{1/2} - P_0) / P_0$, and collect terms:

$$\Delta x^T = -\bar{s}(1 - \bar{s}) \frac{W_0^T}{P_0^2} (P_{1/2} - P_0) + O((\Delta P / P_0)^2).$$

Calibration under positive net inflows. The zero-cash-drift assumption is a theoretical simplification. In practice the TDF sector received approximately \$100B/year of net contributions over 2015–2023 on a \$3.5T sector, a quarterly inflow rate of $\mu_{\text{flow}} \approx 0.0071$ (fraction of sector AUM

per quarter). Under a positive mean inflow $\mu_{\text{flow}} > 0$ entering between $t = 0$ and $t = 1/2$, the linearization acquires an additional additive term

$$\Delta x_{\text{flow}}^T = \bar{s} \mu_{\text{flow}} \frac{W_0^T}{P_0},$$

which is an intercept term that does not depend on $(P_{1/2} - P_0)$ and therefore does not affect the sign or slope of the rebalancing-trade *response* to price shocks. Quantitatively: at $\bar{s} = 0.6$ and $\mu_{\text{flow}} = 0.0071/\text{quarter}$, the inflow-driven intercept adds $\bar{s}\mu_{\text{flow}} = 4.3 \times 10^{-3}$ (AUM-fraction per quarter) to the rebalancing-flow baseline. The response-coefficient rebalancing volume $\Delta X_{\text{total}}^T = \alpha \bar{s}(1 - \bar{s})\sigma_\varepsilon = 1.5 \times 10^{-3}$ at calibration is driven by the price-response slope $-\bar{s}(1 - \bar{s})W_0^T/P_0^2$ and is unaffected. The calibration error from ignoring μ_{flow} in the magnitude of the *response* to price is therefore zero to first order; the error in the *level* of baseline trading is absorbed into the intercept. In the empirical regressions the intercept is absorbed by fixed effects, so the μ_{flow} correction does not enter any reported coefficient. \square

Proof of Theorem 3.1

From (2.5), $\phi(\eta) = (1 - \alpha)/d$ with $d = (1 - \alpha) + \eta\alpha\bar{s}\theta_u$. Then

$$\frac{\partial \phi}{\partial \eta} = -\frac{(1 - \alpha)\alpha\bar{s}\theta_u}{d^2} < 0, \quad \frac{\partial \phi}{\partial \alpha} = -\frac{\eta\bar{s}\theta_u}{d^2}[d + (1 - \alpha)] < 0.$$

Since $\text{Var}(R_{1/2}) = \phi^2\sigma_\varepsilon^2$ and $\phi > 0$, both partial derivatives of the variance are strictly negative. The log-elasticity is

$$\frac{\partial \log \text{Var}}{\partial \log \eta} = 2 \frac{\eta}{\phi} \frac{\partial \phi}{\partial \eta} = -\frac{2\eta\alpha\bar{s}\theta_u}{d} = -2\pi_T.$$

For any $\kappa > 0$, $\theta_u > 0$ and $d > 0$, so the sign is preserved. \square

Proof of Theorem 3.2

Under (2.8) and per-asset arbitrageur first-order conditions, date- $\frac{1}{2}$ market clearing solves

$$[(1 - \alpha)A + \alpha W \Xi W^\top] \mathbf{P}_{1/2} = (1 - \alpha)A(\mu + \varepsilon) + \alpha W \Xi W^\top \mathbf{P}_0,$$

where $A = (\gamma \Sigma_u + \kappa I)^{-1}$. The resulting price-impulse response is $\mathbf{P}_{1/2} - \mathbf{P}_0|_\varepsilon = G\varepsilon$ with

$$G = [(1 - \alpha)A + \alpha W \Xi W^\top]^{-1}(1 - \alpha)A.$$

Applying the Woodbury identity to invert $(1 - \alpha)A + \alpha W \Xi W^\top$ and taking variance gives

$$\Delta \Sigma \equiv \text{Var}(\mathbf{P}_{1/2} - \mathbf{P}_0) - \Sigma_u = \left(\frac{\alpha}{1 - \alpha}\right)^2 W \Lambda W^\top + L(\alpha, W, \Xi),$$

where $\Lambda = \Xi(I_M + \alpha\tilde{Q})^{-1}\Xi(I_M + \alpha\tilde{Q})^{-\top} \succeq 0$ with $\tilde{Q} = \Xi W^\top A^{-1}W/(1-\alpha) \succeq 0$, and L is linear in α and comes from the cross-term. The rank of $W\Lambda W^\top$ is at most M . Strict positivity holds whenever W has full column rank. Stage-3a reconstruction verifies $\|\Delta\Sigma - (\text{quad} + L)\|_\infty < 2 \times 10^{-18}$. \square

Proof of Proposition 3.3

If the columns of W are pairwise orthogonal, $W^\top A^{-1}W$ is diagonal and Λ is diagonal with strictly positive entries. Then $(W\Lambda W^\top)_{ij} = \sum_m w_{im}\Lambda_{mm}w_{jm}$ is strictly positive iff $w_{im} > 0$ and $w_{jm} > 0$ for some m . Absent orthogonality, cross-terms $\sum_{m \neq m'} w_{im}\Lambda_{mm'}w_{jm'}$ can be nonzero without co-occupancy, as verified by the Stage-3a counterexample with generic W . \square

Proof of Lemma 3.4 (Neumann radius)

The inverse of the $M \times M$ matrix $Q \equiv I_M + \alpha\tilde{Q}$ that appears in Λ (Theorem 3.2) admits a Neumann expansion $Q^{-1} = \sum_{k \geq 0} (-\alpha\tilde{Q})^k$ whenever $\alpha\|\tilde{Q}\|_2 < 1$. Normalize: write $\tilde{Q} = \xi(I_M + E)$ with ξ a scaling scalar and E the off-diagonal deviation of \tilde{Q}/ξ from the identity. Then $Q = I_M + \alpha\xi(I_M + E) = (1 + \alpha\xi)(I_M + rE)$ with

$$r \equiv \frac{\alpha\xi}{1 + \alpha\xi} = \frac{q_0}{1 + q_0}, \quad q_0 \equiv \alpha\xi.$$

This is the Neumann-weight on E in the normalized form $(I_M + rE)^{-1} = \sum_{k \geq 0} (-rE)^k$. The particular choice $r = q_0/(1 + q_0)$ is therefore not a modeling convention: it is the unique weight on the off-diagonal perturbation E that makes $(I_M + rE)$ the normalized operator. At the calibration, $q_0 = \alpha\theta_u\xi/(1 - \alpha)$ with $\xi = 1$ gives $q_0 = 1.85 \times 10^{-3}$ and $r = 1.847 \times 10^{-3}$.

By Gershgorin's theorem, the spectral radius satisfies $\rho(rE) \leq \max_m \sum_{m' \neq m} r|E_{mm'}|$. For $\|E\|_{\infty, \text{offdiag}} < \delta$ (maximum off-diagonal entry in magnitude), $\rho(rE) \leq r(M - 1)\delta$. Setting $r(M - 1)\delta = 1$ as the Gershgorin convergence boundary and using the sharper Frobenius-off-diagonal bound $\|rE\|_{F, \text{off}} = r\sqrt{M(M - 1)}\delta$,

$$\delta^* = \frac{1 - r}{r\sqrt{M(M - 1)}}.$$

At $\alpha = 0.08$, $\theta_u = 0.02125$, $M = 10$, $\xi = 1$: $\delta^* \approx 57.05$.

To see that $\delta < \delta^*$ governs the linear term L in Theorem 3.2, expand

$$L_{ij}(\alpha, W, \Xi) = 2\frac{\alpha}{1 - \alpha} \sum_{m, m'} w_{im}(\Xi W^\top Q^{-1} A^{-1})_{mm'}(\Sigma_u)_{m'j},$$

and read off $Q^{-1} = \sum_{k \geq 0} (-rE)^k$ as an absolutely convergent series when $\rho(rE) < 1$. Term-by-term inspection: the $k = 0$ term gives $L_{ij}^{(0)} = 2[\alpha/(1 - \alpha)](W\Xi \cdot A^{-1}\Sigma_u)_{ij}$; since $A^{-1}\Sigma_u = \Sigma_u(\gamma\Sigma_u + \kappa I)^{-1}$ commutes with Σ_u in the scalar- Σ_u case, this equals a positive constant times $(W\Xi W^\top)_{ij}$, which is monotone increasing in n_{ij}^{co} (each co-occupancy adds a non-negative term). The $k \geq 1$ corrections

scale as $r^k \delta^k \rightarrow 0$ geometrically for $r\delta < 1$, preserving monotonicity. For $|E_{mm'}| < \delta^*$, the Neumann expansion converges and L_{ij} is monotone increasing in n_{ij}^{co} under bounded correction from the higher-order terms. This governs the *linear* term L , which is the empirically detectable object at current α (§3.2); the quadratic $(\alpha/(1-\alpha))^2 W \Lambda W^\top$ term is monotone in n_{ij}^{co} by construction but quantitatively small. Stage-3a verifies convergence at δ well beyond the Gershgorin boundary (empirical Neumann series convergent across 14/14 test points up to $\delta = 100$ at $\alpha = 0.08$), so the bound is conservative. \square

Proof of Proposition 3.5

From (2.7), $\pi_T(\alpha) = \eta\alpha\bar{s}\theta_u/[(1-\alpha) + \eta\alpha\bar{s}\theta_u]$. The derivative is

$$\pi'_T(\alpha) = \frac{\eta\bar{s}\theta_u}{[(1-\alpha) + \eta\alpha\bar{s}\theta_u]^2} > 0.$$

Limits $\pi_T(0) = 0$ and $\pi_T(1^-) = 1$ are immediate. Setting $\pi_T(\alpha^*) = 1/2$ gives $(1-\alpha^*) = \eta\alpha^*\bar{s}\theta_u$, hence $\alpha^* = 1/(1 + \eta\bar{s}\theta_u)$. \square

Proof of Corollary 3.7

The factor $(\alpha/(1-\alpha))^2$ is the quadratic factor in (3.2). At $\alpha_1 = 0.005$, the factor is $(0.005/0.995)^2 = 2.525 \times 10^{-5}$. At $\alpha_2 = 0.08$, it is $(0.08/0.92)^2 = 7.561 \times 10^{-3}$. The ratio is 299. Linear-in- α ratio is $0.08/0.005 = 16$. \square

Proof of Proposition 3.8

Set $\lambda = \nu\kappa$ with ν a dimensionless multiplier. Under the interior-equilibrium FOC and (MC), the scalar equilibrium scalar factor is $\phi(\eta; \nu) = (1-\alpha)/d(\nu)$ with $d(\nu) = (1-\alpha) + \eta\alpha\bar{s}(\gamma\sigma_u^2 + \nu\kappa)$. The sign results of Theorem 3.1 hold whenever $d(\nu) > 0$, which requires $\gamma\sigma_u^2 + \nu\kappa > -(1-\alpha)/(\eta\alpha\bar{s})$. At $\alpha = 0.08$ and calibration, the window in ν is of width $\sqrt{(1-\alpha)\theta_u^{\text{eff}}/(\eta\alpha\bar{s})} \approx 3.46$, yielding admissible $\nu \in (-343, 349)$. The microfoundation $\nu = 1$ is strictly interior. The cross-partial sign of Proposition 3.9 is continuous in ν and preserved throughout the window by the same monotonicity arguments. \square

Proof of Proposition 3.9

Under $\lambda = \kappa$, the expected post-news realized variance is

$$\mathbb{E}[\text{RV}]_{\text{post}} = \phi_{\text{post}}^2 \sigma_u^2,$$

where $\phi_{\text{post}} = (1-\alpha)/d_{\text{post}}$ and $d_{\text{post}} = (1-\alpha) + \eta\alpha\bar{s}(\gamma\sigma_u^2 + \kappa)$. Let $c \equiv \eta\bar{s}$, $g \equiv \gamma\sigma_u^2$, $h \equiv g + \kappa = \theta_u$. Then $d_{\text{post}} = (1-\alpha) + c\alpha h$ and

$$\phi_{\text{post}} = \frac{1-\alpha}{d_{\text{post}}}, \quad \mathbb{E}[\text{RV}]_{\text{post}} = \frac{(1-\alpha)^2}{d_{\text{post}}^2} \sigma_u^2.$$

Compute $\partial_\kappa \mathbb{E}[\text{RV}]_{\text{post}}$: since $\partial_\kappa d_{\text{post}} = c\alpha$,

$$\partial_\kappa \mathbb{E}[\text{RV}]_{\text{post}} = (1-\alpha)^2 \sigma_u^2 \frac{-2c\alpha}{d_{\text{post}}^3}.$$

Now differentiate with respect to α . Using $\partial_\alpha d_{\text{post}} = -1 + ch$ and the quotient/chain rule,

$$\begin{aligned} \partial_\alpha \left[\frac{(1-\alpha)^2}{d_{\text{post}}^3} \right] &= \frac{-2(1-\alpha)d_{\text{post}}^3 - 3d_{\text{post}}^2(-1+ch)(1-\alpha)^2}{d_{\text{post}}^6} \\ &= \frac{(1-\alpha)[-2d_{\text{post}} - 3(1-\alpha)(ch-1)]}{d_{\text{post}}^4}. \end{aligned}$$

Therefore

$$\begin{aligned} \frac{\partial^2 \mathbb{E}[\text{RV}]_{\text{post}}}{\partial \alpha \partial \kappa} &= -2c\sigma_u^2 \partial_\alpha \left[\alpha \frac{(1-\alpha)^2}{d_{\text{post}}^3} \right] \\ &= -2c\sigma_u^2 \left[\frac{(1-\alpha)^2}{d_{\text{post}}^3} + \alpha \frac{(1-\alpha)[-2d_{\text{post}} - 3(1-\alpha)(ch-1)]}{d_{\text{post}}^4} \right] \\ &= \frac{-2c\sigma_u^2(1-\alpha)}{d_{\text{post}}^4} [(1-\alpha)d_{\text{post}} + \alpha(-2d_{\text{post}} - 3(1-\alpha)(ch-1))]. \end{aligned}$$

Substituting $d_{\text{post}} = (1-\alpha) + c\alpha h$ into the bracketed expression and collecting terms in α :

$$\begin{aligned} B(\alpha) &\equiv (1-\alpha)d_{\text{post}} - 2\alpha d_{\text{post}} - 3\alpha(1-\alpha)(ch-1) \\ &= (1-3\alpha)d_{\text{post}} - 3\alpha(1-\alpha)(ch-1). \end{aligned}$$

Expanding d_{post} and simplifying,

$$\begin{aligned} B(\alpha) &= (1-3\alpha)[(1-\alpha) + c\alpha h] - 3\alpha(1-\alpha)(ch-1) \\ &= (1-\alpha)(1-3\alpha) + c\alpha h(1-3\alpha) - 3\alpha(1-\alpha)(ch-1) \\ &= (1-\alpha)[(1-3\alpha) + 3\alpha] + c\alpha h(1-3\alpha) - 3\alpha(1-\alpha)ch \\ &= (1-\alpha) + c\alpha h[(1-3\alpha) - 3(1-\alpha)] \\ &= (1-\alpha) - 2c\alpha h. \end{aligned}$$

Observe the elementary algebraic identity, verified by direct substitution of $d_{\text{post}} = (1-\alpha) + c\alpha h$:

$$d_{\text{post}} - 3c\alpha h = (1-\alpha) + c\alpha h - 3c\alpha h = (1-\alpha) - 2c\alpha h = B(\alpha).$$

This is an identity in the primitive parameters, not a consequence of the cross-partial claim: it rewrites the bracketed expression $B(\alpha)$ in a form that isolates the sign-determining quantity in a single term. Therefore

$$\frac{\partial^2 \mathbb{E}[\text{RV}]_{\text{post}}}{\partial \alpha \partial \kappa} = \frac{-2c\sigma_u^2(1-\alpha)}{d_{\text{post}}^4} [d_{\text{post}} - 3c\alpha h].$$

The prefactor $-2c\sigma_u^2(1-\alpha)/d_{\text{post}}^4$ is strictly negative on $\alpha \in (0, 1)$. The cross-partial is strictly negative if and only if $d_{\text{post}} - 3c\alpha h > 0$. On $(\alpha, \kappa) \in (0, 0.5] \times (0, 0.1]$ with $c = 0.03$, $g = 0.01125$: $c\alpha h \leq 0.03 \cdot 0.5 \cdot (0.01125 + 0.1) \leq 1.7 \times 10^{-3}$, whereas $d_{\text{post}} \geq 1 - \alpha \geq 0.5$. Hence $d_{\text{post}} - 3c\alpha h \geq 0.5 - 5 \times 10^{-3} > 0$ and the cross-partial is strictly negative on the rectangle. The sign-flip locus $d_{\text{post}} = 3c\alpha h$ requires $(1 - \alpha) + c\alpha h = 3c\alpha h$, equivalently $(1 - \alpha) = 2c\alpha h$, i.e., $\alpha = 1/(1 + 2ch)$, which at $ch \leq 0.03 \cdot 0.11 = 3.3 \times 10^{-3}$ gives $\alpha \geq 1/(1 + 0.0066) \approx 0.993$. This locus is outside the calibrated rectangle, consistent with Stage-3a symbolic verification.

Pre-news cross-partial. We now derive the pre-news cross-partial explicitly. In the pre-news window, the signal ε has not yet resolved, so the arbitrageur's per-capita inventory variance is $\sigma_\varepsilon^2 + \sigma_u^2$ rather than σ_u^2 ; define $\theta_0 \equiv \gamma(\sigma_\varepsilon^2 + \sigma_u^2) + \kappa$ and the pre-news equilibrium price response $\phi_{\text{pre}}(\alpha, \kappa) = (1 - \alpha)/d_{\text{pre}}$ with $d_{\text{pre}} = (1 - \alpha) + c\alpha h_{\text{pre}}$, $h_{\text{pre}} \equiv g + \gamma\sigma_\varepsilon^2 + \kappa = \theta_0$. The pre-news expected realized variance is $\mathbb{E}[\text{RV}]_{\text{pre}} = \phi_{\text{pre}}^2(\sigma_\varepsilon^2 + \sigma_u^2)$.

The same algebraic manipulation that gave the post-news cross-partial yields

$$\frac{\partial^2 \mathbb{E}[\text{RV}]_{\text{pre}}}{\partial \alpha \partial \kappa} = \frac{-2c(\sigma_\varepsilon^2 + \sigma_u^2)(1-\alpha)}{d_{\text{pre}}^4} [d_{\text{pre}} - 3c\alpha h_{\text{pre}}].$$

The algebra is identical step-for-step: $\partial_\kappa d_{\text{pre}} = c\alpha$, $\partial_\alpha d_{\text{pre}} = -1 + ch_{\text{pre}}$, and the same identity $d_{\text{pre}} - 3c\alpha h_{\text{pre}} = (1 - \alpha) - 2c\alpha h_{\text{pre}}$ applies. Since $h_{\text{pre}} > h$ (by $\gamma\sigma_\varepsilon^2 > 0$), on the same calibrated rectangle $d_{\text{pre}} - 3c\alpha h_{\text{pre}} \geq 0.5 - 3 \cdot 0.03 \cdot 0.5 \cdot (h + \gamma\sigma_\varepsilon^2)$; substituting $\gamma\sigma_\varepsilon^2 = 2 \cdot 0.0064 = 0.0128$ and $\kappa \leq 0.1$ gives $h_{\text{pre}} \leq 0.01125 + 0.0128 + 0.1 = 0.124$, so $d_{\text{pre}} - 3c\alpha h_{\text{pre}} \geq 0.5 - 5.6 \times 10^{-3} > 0$ and the pre-news cross-partial is also strictly negative on the rectangle.

Pre/post asymmetry and scope. Both cross-partials are strictly negative on the calibrated rectangle. A tight comparison of their absolute magnitudes in the same model units depends on the ratio of prefactors and on the d^4 denominators. At the calibration, $(\sigma_\varepsilon^2 + \sigma_u^2)/\sigma_u^2 \approx 2.14$, the d^4 ratio is within 10^{-3} of unity, and the bracketed term $(d - 3c\alpha h)$ is within 10^{-2} of $(1 - \alpha)$ in both windows. A closed-form comparison of *absolute* cross-partial magnitudes in model units would require a full scalar-to-regression unit-conversion mapping that we do not attempt (see the reduced-form framing in the introduction). What Proposition 3.9 establishes on its own terms is the sign of both cross-partials (strictly negative on the rectangle) and their structural form. The pre/post asymmetry stated in the proposition is a structural qualitative prediction: the identifying variation in the regression is the post-minus-pre interaction at the event level with event FE absorbing the TDF

and HKM main effects, and what Proposition 3.9 predicts is precisely the differential emergence of the cross-partial post-news. The empirical test in §6.5 identifies this differential directly through the triple-difference β_3 coefficient; we do not claim a closed-form match of pre- and post-window coefficient magnitudes. \square

Proof of Proposition 3.10

Under a linear cost $\tilde{\kappa}|x|$, the arbitrageur FOC is $\mathbb{E}[\Delta P] - \gamma\sigma_u^2 x = \tilde{\kappa} \text{sign}(x)$, which on the interior interior-trading regime gives $x = (\mathbb{E}[\Delta P] - \tilde{\kappa} \text{sign}(x))/(\gamma\sigma_u^2)$. The demand slope is $-1/(\gamma\sigma_u^2)$, independent of $\tilde{\kappa}$. The equilibrium ϕ does not depend on $\tilde{\kappa}$, so $\partial\mathbb{E}[\text{RV}]_{\text{post}}/\partial\tilde{\kappa} = 0$ on the interior, and the cross-partial of Proposition 3.9 vanishes. Stage-3a Monte Carlo with 200,000 draws per $\tilde{\kappa}$ value confirms flatness (max numerical derivative 4.6×10^{-3} on a base level of 1.2×10^{-2} , within MC noise). \square

Proof of Theorem 4.2

Minimize $\sum_{g=1}^G f_g^2$ subject to $\sum_g f_g = 1$, $f_g \geq 0$. By Cauchy-Schwarz,

$$\left(\sum_g f_g\right)^2 \leq G \sum_g f_g^2, \quad \text{so} \quad \sum_g f_g^2 \geq \frac{1}{G},$$

with equality if and only if $f_1 = \dots = f_G = 1/G$. Substituting into (4.1) gives the claimed minimum. The relative welfare gain over the worst case (all mass on one group, $\sum f_g^2 = 1$) is $1 - 1/G$. \square

Planner glide-path deviation (Corollary A.1)

Corollary A.1 (Planner glide-path deviation). *Let \bar{s}^{ind} denote the individually optimal glide-path target and \bar{s}^{SP} the planner's target. Under (4.1),*

$$\bar{s}^{\text{SP}} - \bar{s}^{\text{ind}} \approx \frac{2\lambda_{\text{eff}}\alpha^2\bar{s}^{\text{ind}}(1 - \bar{s}^{\text{ind}})(1 - 2\bar{s}^{\text{ind}})V^*}{U''(\bar{s}^{\text{ind}})},$$

where $V^* = \sigma_\varepsilon^2(1 - 1/G)$.

Proof. Individual TDF-holder utility is $U(\bar{s})$; the planner adds the welfare cost $W(\bar{s}) = \lambda_{\text{eff}}(\alpha\bar{s}(1 - \bar{s}))^2V^*$ from Lemma 4.1. First-order conditions:

$$\begin{aligned} U'(\bar{s}^{\text{ind}}) &= 0, \\ U'(\bar{s}^{\text{SP}}) &= W'(\bar{s}^{\text{SP}}) = 2\lambda_{\text{eff}}\alpha^2\bar{s}^{\text{SP}}(1 - \bar{s}^{\text{SP}})(1 - 2\bar{s}^{\text{SP}})V^*. \end{aligned}$$

To first order near \bar{s}^{ind} , expand $U'(\bar{s}^{\text{SP}}) \approx U''(\bar{s}^{\text{ind}})(\bar{s}^{\text{SP}} - \bar{s}^{\text{ind}})$ and $W'(\bar{s}^{\text{SP}}) \approx W'(\bar{s}^{\text{ind}})$. Solving:

$$\bar{s}^{\text{SP}} - \bar{s}^{\text{ind}} \approx \frac{W'(\bar{s}^{\text{ind}})}{U''(\bar{s}^{\text{ind}})} = \frac{2\lambda_{\text{eff}}\alpha^2\bar{s}^{\text{ind}}(1 - \bar{s}^{\text{ind}})(1 - 2\bar{s}^{\text{ind}})V^*}{U''(\bar{s}^{\text{ind}})}. \quad (\text{A.1})$$

Equation (A.1) is the full expression, stated without hidden factors: the coefficient $2\alpha^2\bar{s}^{\text{ind}}(1 - \bar{s}^{\text{ind}})$ multiplies the sign-carrying term $(1 - 2\bar{s}^{\text{ind}})$ and the effective cost $\lambda_{\text{eff}}V^*$. The sign is opposite to $(1 - 2\bar{s}^{\text{ind}})$ because U'' is negative at a maximum: the planner pushes the glide-path target further from 1/2 in the direction of the individual optimum. Magnitude: at calibration $\lambda_{\text{eff}} = \kappa/(1 - \alpha) = 0.0109$, $\alpha^2 = 0.0064$, $\bar{s}^{\text{ind}}(1 - \bar{s}^{\text{ind}}) \leq 0.25$, $V^* \approx \sigma_\varepsilon^2(1 - 1/3) \approx 0.0043$: the shift magnitude is of order $10^{-6}/|U''|$, economically invisible at current α , and scaling as α^2 becomes first-order at $\alpha \gtrsim 0.5$. \square

B Empirical construction details

FoF look-through pipeline

The fund-of-fund look-through pipeline proceeds as follows. (i) I pull the full TDF holdings table from CRSP MF, *including rows with null permno*, which contain mutual-fund-wrapper holdings (18.9M rows over 2005–2023). (ii) I construct a CUSIP/ticker-to-crsp_portno map using the `fund_hdr` table, identifying the 1,718 sub-fund wrapper portnos referenced inside TDF holdings. (iii) I pull those sub-funds’ own equity holdings (67.3M rows). (iv) I recursively allocate FoF dollar holdings pro-rata by each sub-fund’s internal equity weight, up to a depth cap of three. (v) I emit a per-(portno, quarter, permno) look-through weight table used for the basket-overlap calculation.

The AAPL 2020Q4 trace (Table 3) is a diagnostic: the direct-only aggregation gives \$0.043B; the look-through gives \$23.11B, a factor of roughly 540 larger. Sanity-check code confirms this with a sum-of-parts reconstruction.

HKM sign convention

The He et al. [2017] intermediary capital ratio is the primary-dealer equity capital ratio. Tighter balance sheets correspond to *lower* levels of the ratio; my κ in the theory increases with tighter balance sheets. To align convention, I sign-flip the HKM ratio: my empirical κ -proxy is $\hat{\kappa}_t = -\text{HKM}_t$, standardized.

Standard-error clustering

I report quarter-level clustering as the conservative default (accommodates arbitrary within-quarter cross-correlation). Permno-level clustering is reported as robustness for pair-level regressions where the time dimension is less relevant than the cross-sectional dependence. Two-way clustering (permno \times quarter) is reported for panel regressions at the stock-quarter level.

Winsorization

TDF ownership is winsorized at the 1% and 99% tails to limit leverage from the largest one-percent of stock-quarter observations; the cross-sectional results are qualitatively unchanged without winsorization.

Universe construction

The cross-sectional universe is top-500-by-market-cap among CRSP common shares (share codes 10 and 11) on NYSE/AMEX/NASDAQ with nonzero look-through TDF exposure in the quarter. Roughly 92% of eligible stock-quarters pass the filter; the filter rules out micro-caps that receive no material TDF exposure and would add noise to the pair regressions.

C Extensions

Linear-cost collapse

Proposition 3.10 is proved in Appendix A. The empirical implication is that if the inventory-cost channel were linear rather than quadratic, the §6.5 post-news cross-partial would vanish. The $t = -3.47$ result therefore empirically falsifies the linear-cost ansatz in addition to the weaker no-channel null.

N -asset cross-term

The decomposition $\Delta\Sigma = (\alpha/(1-\alpha))^2 W\Lambda W^\top + L$ in Theorem 3.2 has the cross-term

$$L(\alpha, W, \Xi) = 2 \frac{\alpha}{1-\alpha} W\Xi W^\top [(1-\alpha)A + \alpha W\Xi W^\top]^{-1} \Sigma_u.$$

At current α , L dominates the quadratic term by roughly eleven to one; at large α the quadratic term dominates. This gives a direct empirical reason the $n_{\text{tdf}}^{\text{co}}$ reduced-form (which picks up L) has $t = +7.66$ at today's α , while the quadratic WW^\top weighted specification is too small to detect (§6.2).

Calibration consistency with basis-point price impact

At $\kappa = \lambda = 0.01/\text{quarter}$ and $\sigma_u^2 = 0.005625/\text{quarter}$, a one-percent AUM-fraction arbitrageur trade pays approximately one basis point of quarterly price impact. This estimate lies in the middle of standard short-horizon equity price-impact elasticity estimates; see the discussion in §2.3.