Flow Hedging and Revealed Beliefs

Emanuele Luzzi*

Preliminary draft: July 2026

Abstract

This paper measures how the flow risk created by delegated management distorts what portfolios reveal about beliefs. Active funds earn fees on assets under management, and the flows that move those assets share a common component. Managers hedge this flow risk through their holdings, and the hedge is state-dependent: it strengthens with the conditional variance of the common flow factor. In US fund holdings, the rotation away from high-flow-beta stocks appears within fund and quarter as lagged flow variance rises, responds to flow risk rather than to aggregate volatility, and is an order of magnitude weaker in index funds facing the same flows. Because the hedge responds to flow risk rather than to expected returns, reading holdings as beliefs misattributes it as pessimism, and the misattribution scales with flow variance: naive holdings-implied expectations look most bearish when flow risk peaks. A structural framework recovers beliefs from holdings net of the estimated hedge; a first pass puts the misattribution near 0.7 percent per year. Part of what holdings reveal as pessimism is hedging.

Keywords: fund flows; flow hedging; mutual fund holdings; subjective beliefs; belief recovery.
JEL codes: D84, G11, G12, G23.

1 Introduction

In standard portfolio theory, an investor who expects higher returns holds more of the risky asset, with a sensitivity pinned down by risk aversion (Merton, 1969). The data fall short of this benchmark. Households allocate to equities with little regard to their stated return expectations (Giglio et al., 2021), and professional asset managers, who publish their expectations, move their allocations by a fraction of what the same theory implies (Dahlquist and Ibert, 2024). Either subjective beliefs matter less for portfolios, and hence for prices, than the survey literature assumes, or something stands between a belief and a trade. What that something is remains unsettled, and for a structural reason: a portfolio is one observation generated by three unknowns, the manager's beliefs, her risk aversion, and the constraints she trades under, and holdings alone cannot separate them. Candidate frictions are not scarce, and some are measured: for households, part of the gap traces to inattention, adjustment costs, and taxes (Giglio et al., 2021). For a delegated manager the natural candidates are institutional, mandates, benchmarking, the liability side of the fund, and they are mostly unobserved: an explanation that invokes such a constraint after seeing the

*Università della Svizzera italiana and Swiss Finance Institute. Email: emanuele.luzzi@usi.ch. Comments welcome.

data can fit any discrepancy between beliefs and holdings, and an explanation that fits everything rules out nothing. Progress requires a friction specified before the fact, with a signature in holdings that neither beliefs nor slow-moving constraints share.

This paper proposes such a friction and measures it. Active equity funds charge fees on assets under management, so a manager's income moves with her fund's flows, and the flows of active funds share a common component (Dou et al., 2024). Building on their framework, I model a manager who hedges this common flow risk through her holdings: she tilts away from stocks whose returns co-move with the common flow factor, and, because the motive is precautionary, she does so when flow risk is elevated rather than permanently. The testable implication is an interaction rather than a level: the tilt on flow betas should scale with the conditional variance of the common flow factor. Unlike frictions that dampen how portfolios respond to beliefs, the hedge adds to holdings a component unrelated to the manager's views; it does not weaken the belief signal but biases it, by an amount the model ties to observables.

On 16 million fund–stock–quarter positions of US active equity funds over 1992–2024, the within-fund tilt on flow betas is indistinguishable from zero on average and rotates negative as lagged flow variance rises, by about 0.011 standard deviations of active weight per standard deviation of flow beta and flow risk. The rotation is specific to flow risk: matching interactions with market variance and the VIX do not carry it, even though those states move the size and liquidity margins of the same portfolios. It is stable across the measures of conditional flow risk I construct, and it is an order of magnitude weaker in index funds, which face the identical flow environment under a replication mandate. At two standard deviations of flow risk the tilt reaches -0.023, comparable in size to the aggregate active-sector tilt of Dou et al. (2024), which I replicate on my sample, though the two come from different designs.

Calibrating the hedging intensity to the reduced-form rotation by minimum distance and feeding it into a belief recovery in the style of Ghosh et al. (2024), run with and without the hedge, yields the paper's second, preliminary result: holdings read without the hedge understate the manager's expected market return by about 0.7 percent per year on average, from about 0.6 in calm flow-risk states to 0.9 in stressed ones, reaching one and a half to two percent per year in 2001–03, 2008–09, and 2020. A manager who underweights flow-exposed stocks to hedge outflow risk looks pessimistic to any method that reads holdings as beliefs alone, and most pessimistic when flow risk peaks. The estimates run at the aggregate level and I flag them as preliminary, but they hint that part of the apparent conviction gap motivating this literature is hedging.

The paper makes two contributions. First, starting from Dou et al. (2024), it estimates the conditional margin their manager's problem implies: the per-fund rotation on flow betas scaled by the conditional variance of the common flow factor. Their evidence establishes the unconditional tilt, its strengthening with fund activeness, and, through quasi-experiments around natural disasters and trade-war announcements, the response of tilts to shifts in a fund's own outflow risk; the margin here is the continuous, within-fund response of every fund's tilt to the conditional variance of the common factor, the complement their theory points to. The state-dependence also answers directly

the identification concern that dogs holdings-based inference: the slow-moving style, size, and constraint confounds that can mimic a cross-sectional flow-beta tilt do not rotate with lagged flow-factor variance, so conditioning on the flow-risk state differences them out in a way an unconditional sort cannot. Second, the paper connects the hedge to belief measurement in a model that nests the belief recovery of Ghosh et al. (2024), with observed holdings replacing their assumed portfolios and the flow factor entering the manager's kernel as a second pricing argument. Because the hedging term is an additive, quantifiable component of holdings, a belief read off holdings net of the hedge is the manager's view purged of the hedging tilt, and the bias sits on the revealed side of the stated-versus-revealed comparisons this literature runs, so correcting it is a precondition for attributing the pass-through gap. I run that recovery at the aggregate level and find the correction largest when flow risk peaks. The fund-level version, and the model extension in which a scarce risk budget generates the attenuation the frictionless model cannot, disciplined by the stated capital-market assumptions I observe for a subset of managers, are not yet implemented; Section 6 proposes them as the immediate next steps.

Related literature. The paper sits at the intersection of three strands. A growing literature measures investor beliefs directly and asks how they map into portfolios: survey expectations and their properties (Greenwood and Shleifer, 2014), survey evidence for households (Giglio et al., 2021), the stated capital-market assumptions of asset managers (Dahlquist and Ibert, 2024), the return expectations of pension funds and their consultants (Andonov and Rauh, 2022; Andonov et al., 2026), and structural approaches that estimate beliefs and preferences jointly (Egan et al., 2026); when markets are inelastic, how much of a belief reaches the portfolio matters directly for prices (Gabaix and Koijen, 2021). A second strand studies fund flows and their consequences: aggregate flows co-move with market returns (Warther, 1995), flows respond to performance (Berk and Green, 2004) and force liquidity-motivated trading (Edelen, 1999) that is costly (Coval and Stafford, 2007; Lou, 2012), they interact with the scale of active management (Pástor et al., 2015), and they share a common component (Goetzmann et al., 2000; Ferson and Kim, 2012); the asset-pricing implications of flow risk are developed by Vayanos and Woolley (2013) and Savov (2014), and Dou et al. (2024) show that managers hedge the common component and markets price it. The idea that exposure to client flows constrains delegated capital goes back to Shleifer and Vishny (1997), whose arbitrageurs pull back from positions when flows turn against them; here the hedge is measured in holdings and its effect on inferred beliefs quantified. A third reads beliefs off portfolios and prices: Shumway et al. (2011) invert holdings, Ghosh et al. (2024) recover beliefs and preferences from asset prices, and Koijen and Yogo (2019) estimate demand systems in which unobserved demand remains a residual. My contribution is to use the second strand's friction to discipline the third strand's inversion, with the first strand's stated beliefs as the eventual validation.

Outline. The rest of the paper is organized as follows. Section 2 develops the model and its testable content. Section 3 describes the data and measurement. Section 4 presents the evidence on the conditional flow hedge. Section 5 calibrates the hedging intensity and reports the first-pass

belief recovery. Section 6 concludes.

2 Model

2.1 Environment

The economy runs one period, from t to $t+1$, with states drawn under the objective measure P_0. There are N stocks with excess returns $r^e_{t+1} \in \mathbb{R}^N$, conditional mean $\mu_t = \mathbb{E}^{P_0}_t[r^e_{t+1}]$ and covariance $\Sigma_t = \mathrm{Var}^{P_0}_t[r^e_{t+1}]$; write $v_t = \mathrm{diag}(\Sigma_t)$. A common fund-flow factor f_{t+1}, the first principal component of fund-flow shocks, drives the systematic component of flows; β^f_j denotes stock j's loading of returns on f, and $\sigma^2_{f,t} = \mathrm{Var}_t[f_{t+1}]$ its conditional variance.

Funds are indexed by l and each is run by a provider $p(l)$. A fund is observed through its holdings: the portfolio-weight vector $\varphi^{(l)}_t \in \mathbb{R}^N$ and the portfolio return $r^{(l)}_{W,t+1}$ are data, not parameters. This is the central difference from Ghosh et al. (2024), who lack holdings and, in their baseline, assume stylized corner portfolios for the investors whose beliefs they recover. The manager holds a subjective belief $P^{(l)}_\Pi$ about returns; its first two moments are observed for funds run by providers that publish capital-market assumptions (CMAs), which report an equity-premium mean μ^{CMA}_p and volatility σ^{CMA}_p. Stated beliefs never enter the recovery of the belief from holdings; their role is to test the model's predictions about the wedge between stated views and holdings, and they acquire an identifying role in the extension discussed in Section 6.

2.2 The manager's problem

Fund l manages assets $A^{(l)}_t$, which grow with the portfolio's gross return and change with net flows,

$$A^{(l)}_{t+1} = A^{(l)}_t\big(R^{(l)}_{W,t+1} + \Phi^{(l)}_{t+1}\big), \qquad \Phi^{(l)}_{t+1} = \mathbb{E}_t\big[\Phi^{(l)}_{t+1}\big] + b^{(l)} f_{t+1} + \varepsilon^{(l)}_{t+1}, \tag{1}$$

where $\mathbb{E}_t[\Phi^{(l)}_{t+1}]$ is the predictable component of the fund's flow, $b^{(l)}$ is the loading of its flow innovation on the common factor, and $\varepsilon^{(l)}_{t+1}$ is idiosyncratic. The manager is paid a fee proportional to end-of-period assets and has constant absolute risk aversion $\gamma^{(l)}$ over fee income. Because fee income inherits the flow loading $b^{(l)}$, the factor f enters her payoff as background risk that her stock positions do not span, and this is what gives her a motive to hedge.

With jointly normal returns and flows the manager's problem is the mean-variance program, and its first-order condition is

$$\varphi^{(l)}_t = \frac{1}{\gamma^{(l)}}\,\Sigma^{-1}_t\Big(\mu^{\Pi,(l)}_t - \tfrac{1}{2}v_t\Big) - \theta^{(l)}\,\Sigma^{-1}_t\,\mathrm{Cov}_t\big(r^e_{t+1}, f_{t+1}\big), \qquad \mu^{\Pi,(l)}_t \equiv \mathbb{E}^{P^{(l)}_\Pi}_t[r^e_{t+1}]. \tag{2}$$

The first term is the speculative demand of an investor with risk aversion $\gamma^{(l)}$ and subjective expected returns $\mu^{\Pi,(l)}_t$; the belief enters here and nowhere else. The second term tilts the portfolio away from stocks whose returns covary with flows, with intensity $\theta^{(l)}$. Equation (2) is the mean-variance analogue of the delegated-manager problem of Dou et al. (2024) (their equations 3.12–

3.13), stated here under the manager's subjective belief, which their model does not carry; their Theorem 1 gives the cross-sectional implication that active funds' tilts covary negatively with flow betas. In their baseline model with homogeneous active funds the hedging intensity is common across managers by symmetry; I allow it to vary, and I read the belief off the portfolio rather than assuming it.

The hedge is conditional, and this determines what the model predicts in holdings. With stable loadings, $\mathrm{Cov}_t(r^e_{j,t+1}, f_{t+1}) = \beta^f_j \, \sigma^2_{f,t}$: the hedge is not a permanent tilt on flow betas but a rotation that scales with the conditional variance of the flow factor, and it responds to flow risk specifically, since aggregate volatility enters the hedging term nowhere. The estimating implication of (2) is therefore the interaction of flow betas with conditional flow-factor variance, not an unconditional sort. An unconditional tilt need not appear if flow risk is low in the average state, and imposing constant $\sigma^2_{f,t}$ is a restriction the model does not make.

Solving (2) for the subjective mean gives the inversion

$$\mu_t^{\Pi,(l)} = \gamma^{(l)} \Sigma_t \varphi_t^{(l)} + \tfrac{1}{2} v_t + \gamma^{(l)} \theta^{(l)} \, \mathrm{Cov}_t\big(r^e_{t+1}, f_{t+1}\big). \tag{3}$$

A belief read off holdings while ignoring flow hedging omits the last term and is biased by exactly that amount. Because the covariance is conditional, so is the bias, largest in high flow-risk states.

2.3 The stochastic discount factor and recovery

The portfolio rule (2) is the demand of an agent whose stochastic discount factor is exponentially affine in her wealth return and the flow factor,

$$M^{(l)}_{t+1} \propto \big(R^{(l)}_{W,t+1}\big)^{-\gamma^{(l)}} e^{-\gamma^{(l)}\theta^{(l)} f_{t+1}}, \tag{4}$$

normalized to price the riskless asset. The first term is the power kernel of Ghosh et al. (2024), exponentially affine in the log wealth return, and prices exposure to the fund's own return; the second is the flow-hedging tilt, which raises marginal value in states of adverse flows; $\theta^{(l)} = 0$ recovers their kernel exactly. Note that the flow loading in the kernel is $\gamma^{(l)}\theta^{(l)}$, not $\theta^{(l)}$: curvature scales every exposure in marginal utility, while it cancels out of the hedge demand in (2). Holdings satisfy the conditional Euler equation $\mathbb{E}^{P^{(l)}_\Pi}[M^{(l)}_{t+1} \, r^e_{i,t+1} \mid \mathcal{I}_t] = 0$ for each traded test asset i, and this restriction is what the recovery solves: under conditional normality it solves in closed form and returns the inversion (3); without normality it is solved by the smoothed empirical likelihood of Ghosh et al. (2024), of which the inversion is the Gaussian special case. Relative to Ghosh et al. (2024) the recovery is unchanged; what changes is its input. Observed holdings replace their corner-portfolio assumption, and the flow factor enters the kernel as a second pricing argument.

Identification assigns each parameter its moment. Risk aversion $\gamma^{(l)}$ is set by profile likelihood, as in Ghosh et al. (2024): it is the curvature at which the beliefs that satisfy the Euler equations deviate least from the objective distribution. The hedging intensity $\theta^{(l)}$ is identified from the conditional quantity moment implied by (2): the loading of holdings tilts on flow betas scaled

by conditional flow-factor variance. The threat to this moment is not style or size, which are slow-moving and do not rotate with conditional flow risk, but any margin that does; liquidity management in particular responds to the same state variable, so liquidity exposures and their interactions with the flow-risk state enter the specification alongside the flow-beta margin. The belief is the residual in (3), so any tilt misattributed to $\theta^{(l)}$ passes one for one into the recovered belief. Section 4 estimates the reduced form of the θ moment; Section 5 calibrates θ to it and runs the recovery.

2.4 Testable content

The model's testable content at this stage rests on holdings alone: one prediction, and one quantity the recovery produces. The prediction is the conditional hedge: funds rotate away from high-flow-beta stocks when the conditional variance of the common flow factor is elevated. Its refutable features are that the rotation operates through the interaction of flow betas with lagged flow-factor variance rather than unconditionally, and that it responds to flow risk specifically rather than to aggregate volatility, which separates it from generic risk management: matching interactions with market volatility or the VIX should be absent even where those states move other portfolio margins.

The quantity is the decomposition of holdings-implied beliefs. Setting $\theta = 0$ in (3) reads a fund's hedging tilt as if it were a view: a manager who underweights high-flow-beta stocks to hedge looks bearish on them, and the term $\gamma^{(l)}\theta^{(l)}\,\mathrm{Cov}_t(r^e_{\mathrm{mkt}}, f)$ is exactly the size of the misreading. Because the hedge is conditional, the misreading is a time series, largest in high flow-risk states; Section 5 measures it.

3 Data and Measurement

The sample combines quarterly mutual-fund equity holdings from a three-source splice, Thomson-Reuters S12 through 2008Q2, CRSP Mutual Fund holdings to 2019Q2, and SEC N-PORT filings from 2019Q3 (Schwarz and Potter, 2016), with fund characteristics and net flows from the CRSP mutual-fund database and stock-level data from CRSP. The universe contains active, open-end, US domestic-equity funds: index funds and ETFs are identified by CRSP flags and name screens and retained as a separate passive arm, and funds must hold at least 70 percent of assets in equities, disclose a holdings book, exceed $15 million in lagged assets, and survive at least eight quarters, with incubation observations removed following Evans (2010). Fund identifiers, share-class aggregation, and the full filter battery are given in the Appendix. The active estimation sample spans 1992Q1–2024Q4 and contains 16,013,812 fund–stock–quarter positions across 4,495 funds; the passive arm contributes 15,835,889 positions across 1,460 index funds facing the same flow environment under a replication mandate (Table 1).

The common flow factor f follows Dou et al. (2024): fund-level flow shocks, purged of performance and lagged flows, are value-weighted within five size groups, the group series are linearly detrended, and f is the standardized first principal component of the result, which accounts for

Table 1: Sample overview

Active-equity fund panel (fund–quarters)	215,352
Active-equity funds	4,636
Sample window	1992Q1–2024Q4
Holdings source (fund–quarters): S12 / CRSP / N-PORT	88,214 / 97,614 / 29,524
Median holdings per fund–quarter	56
Estimation sample: positions (fund–stock–quarters)	16,013,812
Estimation sample: funds	4,495
Passive arm: positions / funds	15,835,889 / 1,460
Common flow factor: PC1 variance share (months)	80.5% (419)
Position-level active weight: cross-sectional sd	1.65 pp
Position-level active weight: median $\lvert aw \rvert$	0.45 pp
Stock flow beta β^f: sd (p10, p90)	0.034 (-0.003, $+0.072$)

Sample construction is described in Section 3, with the full filter battery in the Appendix. The estimation sample is the subset of the active-equity panel surviving the merge of holdings with lagged stock flow betas, the market benchmark, and stock controls. The bottom panel reports raw moments of the holdings panel merged with flow betas: active weights in percentage points of the fund's equity book; flow betas in monthly return per unit of the (unit-variance) flow factor.

80.5 percent of their variance over 419 months of flow data spanning 1990–2024; passive funds are excluded from the factor, which identifies the common client shock to active funds specifically. Stock j's flow beta β_j^f is the loading of its monthly excess returns on f, estimated on 36-month rolling windows and lagged one quarter. The conditional variance of the flow factor, $\sigma_{f,t}^2$, is the trailing 12-month variance of the flow shock; robustness uses the trailing volatility, a GARCH$(1,1)$ conditional variance, and 3- and 6-month windows, and the VIX placebo uses the monthly mean of the daily CBOE VIX close. All state variables are dated to the prior quarter-end. The dependent variable is the position-level active weight: the stock's share of the fund's equity book minus its weight in the market benchmark, the value-weighted CRSP common-stock universe excluding financials and utilities, as in Dou et al. (2024). Active weights, betas, and state variables are standardized in the regressions, so coefficients are in standard-deviation units; Table 1 reports the raw moments that convert them back.

4 The Conditional Flow Hedge

4.1 Specification

The hedging term of (2), with stable loadings, makes the estimating implication an interaction:

$$aw_{j,t}^{(l)} = \alpha_t^{(l)} + \beta_1 \, \beta_{j,t-1}^f + \beta_2 \left(\beta_{j,t-1}^f \times \sigma_{f,t-1}^2\right) + \Gamma' X_{j,t-1} + \varepsilon_{j,t}^{(l)}, \tag{5}$$

where $aw_{j,t}^{(l)} = \varphi_{j,t}^{(l)} - w_{j,t}^{\text{mkt}}$ is fund l's active weight in stock j, $\alpha_t^{(l)}$ is a fund×quarter fixed effect, and $X_{j,t-1}$ contains the stock's market beta, log size, and Amihud illiquidity, together with their interactions with $\sigma_{f,t-1}^2$ in the interacted specifications. Subtracting the market weight, a

stock characteristic common across funds, leaves the hedging prediction of (2) intact and folds the benchmark term into the stock controls. The fund×quarter fixed effect absorbs everything that varies at the fund-time level, including the level of $\sigma^2_{f,t}$ itself and any portfolio-wide response to it; identification comes from within-portfolio variation in how a given fund, in a given quarter, allocates across stocks with different flow betas. The model's prediction is $\beta_2 < 0$; β_1 is the tilt at the mean flow-risk state, left unrestricted by the model, and empirically it is indistinguishable from zero, so the hedge operates in the upper tail of the flow-risk distribution rather than on average. Standard errors are two-way clustered by stock and quarter; because $\sigma^2_{f,t}$ is a single time series, the interaction is identified off time variation across quarters, and the by-quarter clustering dimension is the binding one.

4.2 The state-dependent tilt

Table 2 reports the main result. Column (1) estimates the unconditional specification at the fund level: within fund and quarter, the active-weight tilt on flow betas is +0.0019 and indistinguishable from zero. The permanent tilt that a static reading of the hedging demand would predict is absent from individual portfolios. Column (2) adds the interaction with lagged flow-factor variance: the coefficient is −0.0105 ($t = -3.80$), so funds rotate away from high-flow-beta stocks when common flow risk is elevated, while the level remains indistinguishable from zero.

Column (3) runs the discrimination test within a single regression, adding the interaction of flow betas with lagged market variance. The flow-risk rotation strengthens to −0.0130 ($t = -4.20$); the market-variance interaction is positive (+0.0037, $t = 2.8$). Generic risk management would rotate portfolios off high-flow-beta stocks in every high-volatility state, a negative market-variance interaction; the estimate carries the opposite sign, so the flow-beta rotation is not generic volatility management. The model's prediction for this coefficient is zero and the small positive estimate is off-model; the discrimination is read off its sign, not its magnitude. Column (4) shows the result is not a financial-crisis artifact, and column (5) reports the mandated-non-response arm: estimated on index funds, which face the same common flow shocks but must replicate their benchmarks, the interaction is −0.0007 ($t = -2.1$), one-fifteenth of the active-fund magnitude and not exactly zero; the flag-and-name index set includes enhanced-index products with limited discretion. The rotation is an order of magnitude weaker where a replication mandate constrains the response. Figure 1 shows the estimated relation by flow-risk state: the conditional mean of the active weight against the flow beta is close to flat in the calm quintiles and slopes downward in the stressed ones, with per-quintile slopes from +0.006 to −0.034.

4.3 The aggregate tilt

Table 3 places the within-fund result against the aggregate fact of Dou et al. (2024), replicated on my sample. Conditional on stock size and illiquidity, the aggregate active sector underweights high-flow-beta stocks, with a tilt of −0.019 standard deviations per standard deviation of flow beta (column 2), stable under flat-TNA weighting (column 3). Column 1 shows the tilt is not visible

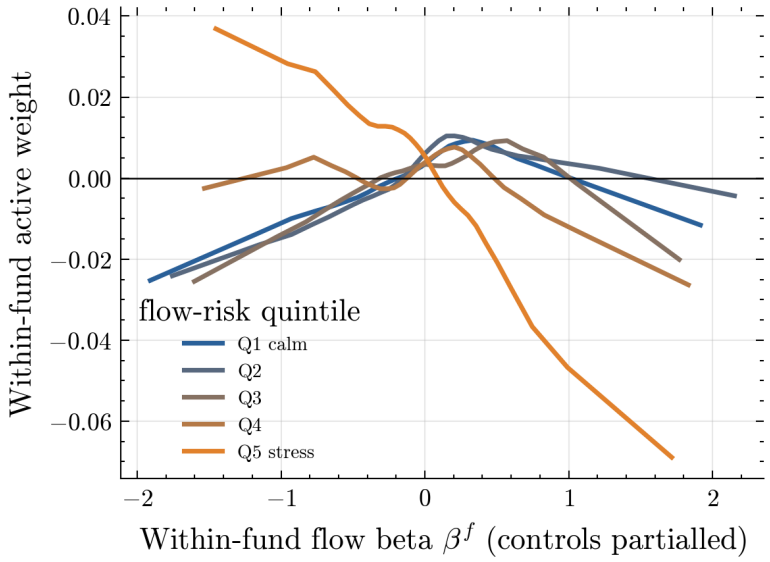


Figure 1: Active weight against flow beta, by flow-risk state. The conditional mean of the position-level active weight against the stock's lagged flow beta, estimated separately within five quintiles of the trailing twelve-month variance of the common flow factor (Q1 calm to Q5 stress), 1992–2024. Both variables are standardized, demeaned within fund×quarter, and residualized on market beta, log size, Amihud illiquidity, and the size and illiquidity interactions with the flow-risk state, matching the specification of Table 2, column (2). Each line is a LOWESS fit through quantile-bin means of the residualized data. The per-quintile least-squares slopes range from +0.006 in Q1 to −0.034 in Q5.

without the size control: the raw partial correlation is positive, because flow betas covary with the size and liquidity characteristics that active funds tilt on for unrelated reasons, variation that the within-fund design of Table 2 conditions out. The aggregate tilt and the within-fund level of Table 2, column (1), are different objects: the aggregate design pools across funds, so composition across funds can produce a sector tilt while the average within-fund level is zero, and the two estimates are not in tension.

4.4 Robustness and placebos

Table 4 varies the flow-risk state variable and adds the VIX placebo. The rotation is stable across the trailing 12-month volatility (−0.0093), the trailing variance (−0.0105), and a GARCH(1, 1) conditional variance (−0.0117), so the result does not depend on how conditional flow risk is measured. Column (4) interacts flow betas with the VIX alongside the flow-risk interaction: the flow-risk rotation is unchanged (−0.0095, $t = -3.69$) while the VIX interaction is statistically indistinguishable from zero (+0.0006, $t = 0.38$). The null is not for lack of power: in the same regression the VIX moves the size and illiquidity margins of the same portfolios significantly (Appendix Table A.2). Aggregate risk states rotate other dimensions of fund portfolios; only flow risk rotates the flow-beta dimension. Shortening the flow-variance window to six and three months leaves the interaction stable (−0.0093 and −0.0086, both significant at one percent; Appendix Table A.3).

Table 2: The conditional flow hedge

	(1) Baseline	(2) × Flow var	(3) + Market var	(4) Ex 2008–09	(5) Passive
Flow beta β^f	+0.0019	−0.0021	−0.0041*	−0.0031	+0.0020***
	(0.0021)	(0.0022)	(0.0024)	(0.0023)	(0.0006)
$\beta^f \times$ flow variance		−0.0105***	−0.0130***	−0.0110***	−0.0007**
		(0.0028)	(0.0031)	(0.0028)	(0.0004)
$\beta^f \times$ market variance			+0.0037***		
			(0.0013)		
Stock controls	Yes	Yes	Yes	Yes	Yes
Controls × flow variance	No	Yes	Yes	Yes	Yes
Controls × market variance	No	No	Yes	No	No
Fund × quarter FE	Yes	Yes	Yes	Yes	Yes
Observations	16,013,812	16,013,812	16,013,812	14,986,895	15,835,889
Funds	4,495	4,495	4,495	4,493	1,460
Within R^2	0.0061	0.0064	0.0065	0.0064	0.0164

Position-level active weight on lagged stock flow beta and its interaction with the lagged trailing 12-month variance of the common flow factor. The unit of observation is a fund–stock–quarter; columns (1)–(4) use active funds, column (5) the passive arm of index funds identified by CRSP index/ETF flags and name screens. Stock controls are market beta, log size, and Amihud illiquidity; interacted specifications include their interactions with the flow-risk state. Dependent variable, betas, and state variables standardized; state variables dated to the prior quarter-end. Two-way clustered (stock, quarter) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Full estimates in Appendix Table A.1.

4.5 Magnitudes

The estimates are in standard-deviation units, best read as a gradient in the flow-risk state. At the mean state the flow-beta gradient of active weights is the level coefficient, −0.0021 and indistinguishable from zero. One standard deviation above mean flow risk the gradient is $-0.0021 - 0.0105 = -0.0126$; two standard deviations above, −0.0231, roughly three times the size of the corresponding market-beta tilt (−0.0077) and, on my sample, comparable to the −0.019 aggregate tilt I replicate from Dou et al. (2024), though the two are estimated on different designs. In raw units, using the moments in Table 1: at two standard deviations of flow risk, moving from a p10 to a p90 flow-beta stock (a 2.2 standard-deviation move in β^f) lowers a stock's active weight by about 0.08 percentage points, roughly one fifth of the median position's absolute active weight of 0.45 percentage points. Per position the effect is modest; the economically relevant object is the coordinated rotation across the high-flow-beta names in a book whose median length is 56 stocks, and the tracking error that rotation consumes, which is the input to the risk-budget extension discussed in Section 6. The simultaneous rotation of the illiquidity margin (Appendix Tables A.1 and A.2) indicates that funds also shift toward liquid names in the same states, consistent with flow risk being managed jointly through the flow-beta and liquidity margins, a joint estimation the structural recovery requires.

Table 3: The aggregate active-weight tilt on flow betas

	(1) β^{mkt} only	(2) + size + Amihud	(3) Flat-TNA weight
Flow beta β^f	$+0.0177^{***}$	-0.0189^{***}	-0.0186^{***}
	(0.0058)	(0.0052)	(0.0053)
Market beta, size, Amihud	β^{mkt} only	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	347,052	347,052	347,052
Within R^2	0.0023	0.0549	0.0545

Aggregate active-fund active weight on stock flow beta, following Dou et al. (2024), eq. 2.6. The unit of observation is a stock–quarter on their Table-3 union sample: a stock enters in quarter t if the aggregate active portfolio held it in any of the trailing nine quarters, with unheld names carrying an active weight of minus the market weight; betas are lagged to the prior quarter-end; column (3) weights funds equally rather than by TNA. Flow betas and controls are standardized. Two-way clustered (stock, quarter) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Full estimates in Appendix Table A.4.

Table 4: Robustness to the flow-risk measure and the VIX placebo

	(1) Vol (12m)	(2) Variance	(3) GARCH	(4) + VIX
$\beta^f \times$ flow risk	-0.0093^{***}	-0.0105^{***}	-0.0117^{**}	-0.0095^{***}
	(0.0025)	(0.0028)	(0.0045)	(0.0026)
$\beta^f \times$ VIX				$+0.0006$
				(0.0015)
Flow beta level, stock controls	Yes	Yes	Yes	Yes
Controls $\times$ flow risk	Yes	Yes	Yes	Yes
Controls $\times$ VIX	No	No	No	Yes
Fund $\times$ quarter FE	Yes	Yes	Yes	Yes
Observations	16,013,812	16,013,812	16,013,812	16,013,812
Within R^2	0.0064	0.0064	0.0064	0.0065

Position-level active weight on lagged stock flow beta interacted with common-flow risk, where the flow-risk state is the trailing 12-month realized volatility (1), the trailing 12-month variance (2), or a GARCH$(1,1)$ conditional variance (3) of the common flow factor; column (4) uses the trailing volatility and additionally interacts flow beta, log size, and Amihud with the monthly VIX. Active funds only. Conventions as in Table 2. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Full estimates in Appendix Table A.2.

5 From the Hedge to Beliefs: A First Pass

This section calibrates the hedging intensity θ to the reduced-form rotation, feeds it into the kernel (4), and recovers the belief about the market return that the aggregate active portfolio implies, with and without the hedge. The exercise is minimal: one representative type, the aggregate active sector, and the market direction of beliefs only. Its purpose is to measure the size and state-dependence of the belief misreading that Section 2 predicts, not to deliver the full cross-section of fund-level parameters.

5.1 Calibrating the hedging intensity

The regression coefficient in Table 2 is not θ itself: the model's hedge demand is $-\theta\,\Sigma_t^{-1}(\beta^f\sigma^2_{f,t})$, so the coefficient mixes θ with the projection of the Σ_t^{-1}-transformed tilt onto the flow-beta regressor. I therefore calibrate θ by exactly identified minimum distance, passing the model through the same regression as the data. I construct the model tilt at $\theta = 1$, $\tilde{w}_t = -\Sigma_t^{-1}\big(\beta^f\sigma^2_{f,t}\big)$, using a two-factor covariance model $\Sigma_t = B_tF_tB_t' + D_t$ with market and flow-factor loadings B_t, trailing factor covariance F_t, and diagonal idiosyncratic variance D_t, all in raw units; I then run regression (5) with $\tilde{w}_t$ as the dependent variable on the same sample. The data regression, in raw units, gives $\hat{\pi} = -0.0086$ (clustered $t = -3.8$); the model tilt gives $c = -0.810$ per unit of θ; linearity of the demand in θ then identifies

$$\hat{\theta} = \hat{\pi}/c = 0.0106,$$

with a standard error of 0.0028 conditional on the covariance model. In portfolio terms, the estimate says funds carry about one percent of the unit-intensity hedge tilt $\tilde{w}_t$; the belief correction in the next subsection converts it into return units. Two caveats attach: the magnitude of $\hat{\theta}$ inherits the factor structure of Σ_t, since the hedge target $\beta^f\sigma^2_{f,t}$ is itself a factor exposure, and the standard error treats c as known. The sign and significance of $\hat{\theta}$, which is what the decomposition below chiefly uses, come directly from the reduced form.

5.2 The recovery

The type is the aggregate active portfolio: the asset-weighted monthly net return of the active funds in the sample, a median of 1,879 funds per month over 1992–2024. Its test assets are those of the model: the market excess return and the portfolio's own excess return. States are described by the trailing twelve-month market return, the current flow factor, and its conditional variance; the objective conditional distribution is a kernel estimate over these states, and the recovered belief is the minimal tilt of that distribution under which the kernel (4) prices both test assets, computed with the smoothed empirical likelihood estimator of Kitamura et al. (2004) as adapted to belief recovery by Ghosh et al. (2024). Curvature is set by profile likelihood at $\theta = 0$, giving $\hat{\gamma} = 4.0$, a conventional magnitude; re-profiling with the hedge in the kernel moves the estimate one grid step, to 3.5, on a nearly flat profile likelihood. Curvature is held fixed across runs so that the only

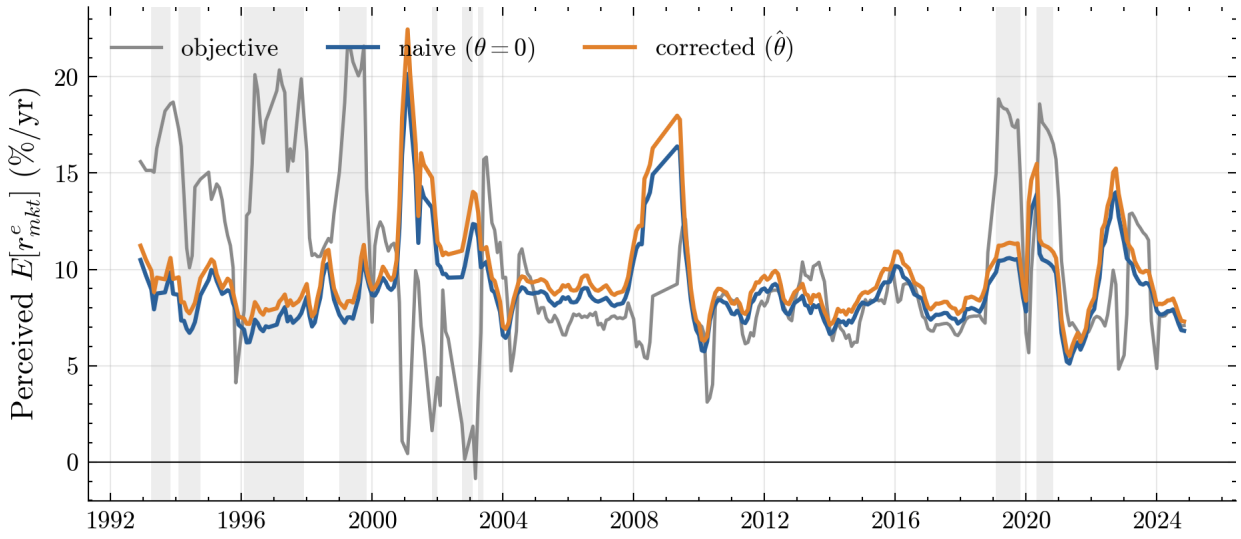


Figure 2: Recovered expected market return, with and without the flow hedge. The perceived expected excess market return $\mathbb{E}[r^e_{\text{mkt}}]$ of the aggregate active portfolio (the asset-weighted monthly net return of the sample's active funds, a median of 1,879 funds per month), monthly 1992–2024, annualized in percent per year, smoothed with a three-month centered moving average. The gray line is the objective conditional expectation, the kernel estimate over the conditioning states (trailing twelve-month market return, flow factor, and flow-factor variance). The blue line is the naive recovery, which sets $\theta = 0$ in the kernel (4); the orange line is the corrected recovery, which uses the calibrated $\hat{\theta} = 0.0106$ with kernel loading $\gamma\theta$; both use the same $\hat{\gamma} = 4.0$. Shaded bands mark months in the top quintile of the trailing twelve-month variance of the common flow factor.

difference between the two recoveries is the hedge: the naive belief sets $\theta = 0$, reading the entire portfolio as speculative demand, and the corrected belief uses $\hat{\theta}$ in the kernel loading $\gamma\theta$. Pricing errors at the solution are numerically zero in both runs.

Figures 2 and 3 report the result. The correction is positive throughout, as the sign of the hedge dictates: a manager who underweights flow-exposed stocks to hedge flow risk looks more pessimistic about the market than she is, so pricing the hedge out raises the recovered belief. On average the naive reading understates the expected market return by about 0.7 percent per year. The correction is state-dependent in the way (3) predicts: it averages 0.6 percent per year in the calmest quintile of flow-risk states and 0.9 percent in the most stressed quintile, reaches one and a half to two percent per year in 2001–03, 2008–09, and 2020, and its time series correlates at 0.38 with the model's closed-form expression $\gamma\theta\ \text{Cov}_t(r^e_{\text{mkt}}, f)$ for the misreading. The recovery also shows recovered expectations sitting above objective ones in crashes: a near-fully-invested equity book can only be rationalized by optimism, because with the cash margin unobserved the inversion cannot register a market-downturn view, a bias that Shumway et al. (2011) flag and set aside. The hedge correction does not resolve that level effect; it measures the part of apparent conviction that is flow hedging, which is largest when flow risk peaks.

A 0.6 to 0.9 percent per year wedge is comparable to the average time-series standard deviation of managers' stated US equity-premium expectations, 1.01 percent per year, and one-third to one-

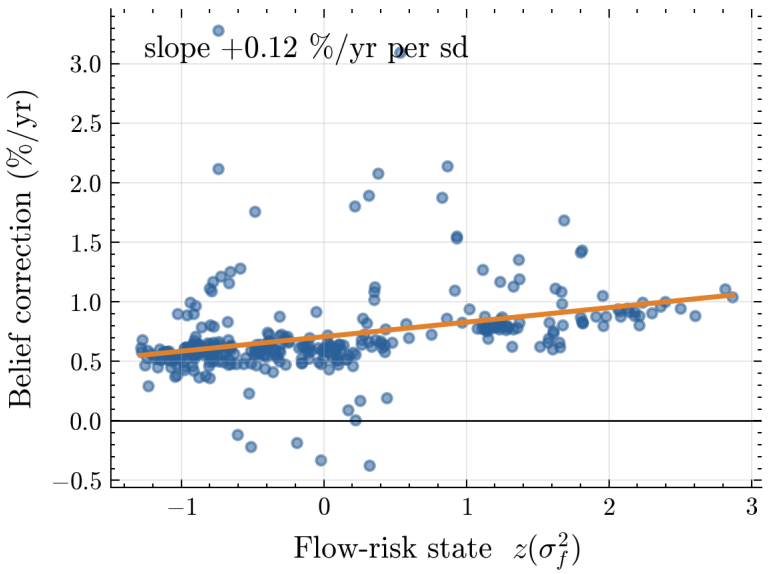


Figure 3: The belief correction against the flow-risk state. Each point is a month, 1992–2024. The vertical axis is the corrected minus the naive recovered expected market return of Figure 2, annualized, in percent per year. The horizontal axis is the trailing twelve-month variance of the common flow factor, standardized over the sample. The line is the least-squares fit; its slope is +0.12 percent per year per standard deviation of the flow-risk state.

half of the cross-sectional standard deviation, 1.75 percent, that Dahlquist and Ibert (2024) report. It is not the whole conviction gap, and the remainder is the object of the extension discussed in Section 6. Because this first pass runs at the aggregate, it averages over the fund-level heterogeneity in $\theta^{(l)}$ that the model allows; the cross-sectional version, which the fund-level rotation in Table 2 licenses, is where the correction should concentrate among high-flow-risk funds.

6 Conclusion

This paper set out to explain why portfolios respond to beliefs less than theory implies, and it delivers one measured friction and a first consequence. Active mutual funds hedge systematic flow risk state by state: within fund and quarter, active weights rotate away from high-flow-beta stocks as the conditional variance of the common flow factor rises; the rotation is specific to flow risk, stable across risk measures, and an order of magnitude weaker in index funds facing the same flows under a replication mandate. The conditional design is what the model prescribes, and it is what identifies the hedge against the slow-moving style and size confounds an unconditional sort cannot escape. Feeding the calibrated hedge into a belief recovery run with and without it yields the paper's second, preliminary result: holdings-implied expectations of the market return appear to understate the manager's view by about 0.7 percent per year on average, and by the most in high flow-risk states. A method that reads holdings as beliefs alone misreads hedging as pessimism, and the misreading peaks with flow risk.

The recovery pools the active sector into a single representative portfolio, so it averages over the fund-level heterogeneity the reduced form licenses; the near-term plan is fund-level hedging intensities and heterogeneous corrected beliefs. The calibration's standard errors are conditional on the covariance model behind the model tilt, and the magnitude of the hedging intensity inherits that model's factor structure; a bootstrap of the full pipeline and a richer factor model are the immediate estimation work. And the recovered beliefs are so far checked only against the model's own closed form; the external test comes next: for managers whose providers publish capital-market assumptions, the corrected belief can be compared with the stated one, a discipline holdings-based recoveries have generally lacked. Stated beliefs then take on a second role in the model extension the reduced form motivates. The frictionless model cannot generate attenuation of the belief pass-through, because the hedge and the speculative demand are additive; attenuation requires the two to compete for a scarce risk budget, as under a penalty on benchmark-relative risk (Roll, 1992), and such a penalty is observationally equivalent in holdings to higher risk aversion, so separating preference from penalty requires the external belief measure the stated assumptions provide. The extension's prediction is that pass-through declines, within provider, with the budget the hedge consumes. The fact this paper contributes, a per-fund flow hedge that turns on when flow risk is high, is the input each of these steps requires.

References

Aleksandar Andonov and Joshua D. Rauh. The return expectations of public pension funds. *Review of Financial Studies*, 35(8):3777–3822, 2022.

Aleksandar Andonov, Spencer J. Couts, Andrei S. Gonçalves, Johnathan Loudis, and Andrea Rossi. Subjective beliefs and the portfolio allocations of institutional investors. Working paper, 2026.

Jonathan B. Berk and Richard C. Green. Mutual fund flows and performance in rational markets. *Journal of Political Economy*, 112(6):1269–1295, 2004.

Joshua Coval and Erik Stafford. Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 86(2):479–512, 2007.

Magnus Dahlquist and Markus Ibert. Equity return expectations and portfolios: Evidence from large asset managers. *Review of Financial Studies*, 37(6):1887–1928, 2024.

Winston Wei Dou, Leonid Kogan, and Wei Wu. Common fund flows: Flow hedging and factor pricing. *Journal of Finance*, 2024. Forthcoming.

Roger M. Edelen. Investor flows and the assessed performance of open-end mutual funds. *Journal of Financial Economics*, 53(3):439–466, 1999.

Mark Egan, Alexander MacKay, and Hanbin Yang. What drives variation in investor portfolios?

Estimating the roles of beliefs and risk preferences. *Review of Financial Studies*, 2026. Advance Access, published online April 21, 2026.

Richard B. Evans. Mutual fund incubation. *Journal of Finance*, 65(4):1581–1611, 2010.

Wayne E. Ferson and Min S. Kim. The factor structure of mutual fund flows. *International Journal of Portfolio Analysis and Management*, 1(2):112–143, 2012.

Xavier Gabaix and Ralph S. J. Koijen. In search of the origins of financial fluctuations: The inelastic markets hypothesis. Working Paper 28967, National Bureau of Economic Research, 2021.

Anisha Ghosh, Arthur G. Korteweg, and Qing Xu. Recovering heterogeneous beliefs and preferences from asset prices. *Journal of Finance*, 2024. Forthcoming.

Stefano Giglio, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. Five facts about beliefs and portfolios. *American Economic Review*, 111(5):1481–1522, 2021.

William N. Goetzmann, Massimo Massa, and K. Geert Rouwenhorst. Behavioral factors in mutual fund flows. Working Paper 00-14, Yale International Center for Finance, 2000.

Robin Greenwood and Andrei Shleifer. Expectations of returns and expected returns. *Review of Financial Studies*, 27(3):714–746, 2014.

Yuichi Kitamura, Gautam Tripathi, and Hyungtaik Ahn. Empirical likelihood-based inference in conditional moment restriction models. *Econometrica*, 72(6):1667–1714, 2004.

Ralph S. J. Koijen and Motohiro Yogo. A demand system approach to asset pricing. *Journal of Political Economy*, 127(4):1475–1515, 2019.

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

Robert C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics*, 51(3):247–257, 1969.

Ľuboš Pástor, Robert F. Stambaugh, and Lucian A. Taylor. Scale and skill in active management. *Journal of Financial Economics*, 116(1):23–45, 2015.

Richard Roll. A mean/variance analysis of tracking error. *Journal of Portfolio Management*, 18 (4):13–22, 1992.

Alexi Savov. The price of skill: Performance evaluation by households. *Journal of Financial Economics*, 112(2):213–231, 2014.

Christopher G. Schwarz and Mark E. Potter. Revisiting mutual fund portfolio disclosure. *Review of Financial Studies*, 29(12):3519–3544, 2016.

Andrei Shleifer and Robert W. Vishny. The limits of arbitrage. *Journal of Finance*, 52(1):35–55, 1997.

Tyler Shumway, Maciej B. Szefler, and Kathy Yuan. The information content of revealed beliefs in portfolio holdings. Working paper, University of Michigan and London School of Economics, 2011.

Dimitri Vayanos and Paul Woolley. An institutional theory of momentum and reversal. *Review of Financial Studies*, 26(5):1087–1145, 2013.

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

Appendix

A.1 Data construction

Holdings come from a three-source splice: Thomson-Reuters S12 filings through 2008Q2, the CRSP Mutual Fund holdings database from 2008Q3 to 2019Q2, and SEC N-PORT filings from 2019Q3. On funds covered by more than one source, the implied portfolio weights agree closely, which is what licenses the splice. Funds are identified by MFLINKS `wficn`, and share classes are aggregated to the fund, with returns TNA-weighted across classes and assets summed. The active universe keeps open-end US domestic-equity funds that hold at least 70 percent of assets in equities (or carry an equity objective code in the S12 era), disclose a holdings book, exceed $15 million in lagged assets, and appear for at least eight quarters; incubation observations are removed following Evans (2010), and net flows are winsorized at the 1st and 99th percentiles. Index funds and ETFs, identified by CRSP flags and name screens, are excluded from the active universe and form the passive arm. The stock universe is CRSP common stocks (share codes 10 and 11), and the market benchmark is the value-weighted CRSP common-stock universe excluding financials and utilities, as in Dou et al. (2024).

The common flow factor follows Dou et al. (2024). Monthly fund-level flows are computed from assets and returns, and flow shocks are the residuals from a regression of flows on lagged fund performance and lagged flows. The shocks are value-weighted within five fund-size groups, each group series is linearly detrended, and the factor is the standardized first principal component of the five series, which accounts for 80.5 percent of their variance over 419 months spanning 1990–2024; passive funds are excluded throughout. Stock flow betas are slopes of monthly excess returns on the factor over rolling 36-month windows; market betas use rolling 60-month windows; Amihud illiquidity is computed quarterly from daily data. Betas are lagged one quarter, and every state variable, the trailing 3-, 6-, and 12-month volatility and variance of the flow factor, the GARCH$(1,1)$ conditional variance, the trailing market variance, and the monthly mean of the daily VIX close, is dated to the prior quarter-end.

A.2 Tables

The tables below report the full coefficient vectors behind the body tables.

Table A.1: The conditional flow hedge: full estimates

	(1) Baseline	(2) × Flow var	(3) + Market var	(4) Ex 2008–09	(5) Passive
Flow beta β^f	$+0.0019$	-0.0021	-0.0041^{*}	-0.0031	$+0.0020^{***}$
	(0.0021)	(0.0022)	(0.0024)	(0.0023)	(0.0006)
$\beta^f \times$ flow variance		-0.0105^{***}	-0.0130^{***}	-0.0110^{***}	-0.0007^{**}
		(0.0028)	(0.0031)	(0.0028)	(0.0004)
$\beta^f \times$ market variance			$+0.0037^{***}$		
			(0.0013)		
Market beta	-0.0085^{**}	-0.0077^{**}	-0.0075^{**}	-0.0073^{**}	-0.0024^{**}
	(0.0033)	(0.0033)	(0.0033)	(0.0033)	(0.0011)
Log size	$+0.0646^{***}$	$+0.0614^{***}$	$+0.0623^{***}$	$+0.0611^{***}$	$+0.0804^{***}$
	(0.0155)	(0.0159)	(0.0161)	(0.0159)	(0.0050)
Amihud illiq.	-0.0045^{**}	-0.0056^{***}	-0.0082^{***}	-0.0116^{***}	$+0.0030^{***}$
	(0.0022)	(0.0015)	(0.0014)	(0.0020)	(0.0006)
Log size × flow variance		-0.0125^{***}	-0.0101	-0.0122^{**}	-0.0191^{***}
		(0.0045)	(0.0061)	(0.0047)	(0.0037)
Amihud × flow variance		-0.0074^{***}	-0.0067^{***}	-0.0046^{***}	-0.0009^{**}
		(0.0013)	(0.0013)	(0.0013)	(0.0004)
Log size × market variance			-0.0048		
			(0.0037)		
Amihud × market variance			$+0.0027^{***}$		
			(0.0008)		
Fund × quarter FE	Yes	Yes	Yes	Yes	Yes
Observations	16,013,812	16,013,812	16,013,812	14,986,895	15,835,889
Within R^2	0.0061	0.0064	0.0065	0.0064	0.0164

Full coefficient vectors for Table 2. The unit of observation is a fund–stock–quarter; the dependent variable is the position-level active weight, the stock's share of the fund's equity book minus its weight in the market benchmark. Regressors are the stock's flow beta, market beta, log size, and Amihud illiquidity, lagged one quarter, together with their interactions with the trailing 12-month variance of the common flow factor (columns 2–5) and with the trailing market variance (column 3); state variables are dated to the prior quarter-end. All specifications include fund×quarter fixed effects; the dependent variable, betas, and state variables are standardized. Columns (1)–(3) use the active-fund sample, column (4) excludes 2008–09, and column (5) uses the passive arm of index funds identified by CRSP flags and name screens. Two-way clustered (stock, quarter) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.2: Robustness to the flow-risk measure and the VIX placebo: full estimates

	(1) Vol (12m)	(2) Variance	(3) GARCH	(4) + VIX
Flow beta β^f	-0.0016	-0.0021	-0.0039	-0.0020
	(0.0022)	(0.0022)	(0.0033)	(0.0022)
$\beta^f \times$ flow risk	-0.0093^{***}	-0.0105^{***}	-0.0117^{**}	-0.0095^{***}
	(0.0025)	(0.0028)	(0.0045)	(0.0026)
$\beta^f \times$ VIX				$+0.0006$
				(0.0015)
Market beta	-0.0076^{**}	-0.0077^{**}	-0.0077^{**}	-0.0077^{**}
	(0.0033)	(0.0033)	(0.0033)	(0.0033)
Log size	$+0.0615^{***}$	$+0.0614^{***}$	$+0.0593^{***}$	$+0.0618^{***}$
	(0.0158)	(0.0159)	(0.0171)	(0.0159)
Amihud illiq.	-0.0048^{***}	-0.0056^{***}	-0.0062^{***}	-0.0096^{***}
	(0.0017)	(0.0015)	(0.0016)	(0.0017)
Log size $\times$ flow risk	-0.0125^{***}	-0.0125^{***}	-0.0135	-0.0112^{***}
	(0.0040)	(0.0045)	(0.0137)	(0.0040)
Amihud $\times$ flow risk	-0.0076^{***}	-0.0074^{***}	-0.0085^{***}	-0.0060^{***}
	(0.0013)	(0.0013)	(0.0024)	(0.0011)
Log size $\times$ VIX				-0.0054^{***}
				(0.0019)
Amihud $\times$ VIX				$+0.0019^{***}$
				(0.0004)
Fund $\times$ quarter FE	Yes	Yes	Yes	Yes
Observations	16,013,812	16,013,812	16,013,812	16,013,812
Within R^2	0.0064	0.0064	0.0064	0.0065

Full coefficient vectors for Table 4. The unit of observation is a fund–stock–quarter on the active-fund sample; the dependent variable is the standardized position-level active weight. The flow-risk state is the trailing 12-month realized volatility of the common flow factor in column (1), its trailing 12-month variance in column (2), and a GARCH$(1,1)$ conditional variance in column (3); column (4) uses the trailing volatility and adds interactions of flow beta, log size, and Amihud illiquidity with the monthly mean of the daily VIX close. All specifications include the stock controls, their interactions with the flow-risk state, and fund$\times$quarter fixed effects; variables are standardized and state variables dated to the prior quarter-end. Two-way clustered (stock, quarter) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.3: Robustness to the flow-variance window

	(1) 3m	(2) 6m	(3) 12m
Flow beta β^f	$+0.0003$	-0.0006	-0.0021
	(0.0022)	(0.0022)	(0.0022)
$\beta^f \times$ flow variance	-0.0086^{***}	-0.0093^{***}	-0.0105^{***}
	(0.0027)	(0.0028)	(0.0028)
Market beta, size, Amihud	Yes	Yes	Yes
Controls $\times$ flow variance	Yes	Yes	Yes
Fund $\times$ quarter FE	Yes	Yes	Yes
Observations	16,013,812	16,013,812	16,013,812
Within R^2	0.0062	0.0063	0.0064

Position-level active weight on lagged flow beta and its interaction with the trailing variance of the common flow factor computed over 3-, 6-, and 12-month windows, on the active-fund sample. All specifications include the stock controls (market beta, log size, Amihud illiquidity), their interactions with the flow-variance state, and fund×quarter fixed effects; variables are standardized and state variables dated to the prior quarter-end. Two-way clustered (stock, quarter) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.4: The aggregate active-weight tilt on flow betas: full estimates

	(1) β^{mkt} only	(2) + size + Amihud	(3) Flat-TNA weight
Flow beta β^f	$+0.0177^{***}$	-0.0189^{***}	-0.0186^{***}
	(0.0058)	(0.0052)	(0.0053)
Market beta	$+0.0388^{***}$	$+0.0477^{***}$	$+0.0458^{***}$
	(0.0098)	(0.0107)	(0.0106)
Log size		-0.2456^{***}	-0.2449^{***}
		(0.0475)	(0.0474)
Amihud illiq.		-0.0326^{**}	-0.0326^{**}
		(0.0140)	(0.0140)
Quarter FE	Yes	Yes	Yes
Observations	347,052	347,052	347,052
Within R^2	0.0023	0.0549	0.0545

Full coefficient vectors for Table 3: the aggregate active-sector tilt following Dou et al. (2024), eq. 2.6. The unit of observation is a stock–quarter on their Table-3 union sample: a stock enters in quarter t if the aggregate active portfolio held it in any of the trailing nine quarters, unheld names carry an active weight of minus the market weight, and betas are lagged to the prior quarter-end; column (3) weights funds equally rather than by TNA. Flow betas and controls are standardized, and all specifications include quarter fixed effects. Two-way clustered (stock, quarter) standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.