Reaching for Habitats: A Stabilizing Channel of Monetary Policy Transmission

Jialiang Lou*

London School of Economics

August 4, 2026

Abstract

This paper shows a broad class of institutional investors hold persistent duration habitats and actively rebalance to maintain them. Monetary shocks reprice long bonds more than short ones, pushing portfolio durations off target and triggering trades against the repricing. This forms a stabilizing channel of monetary policy transmission, Reaching for Habitats, opposite to the amplification channels emphasized in the literature. Using a comprehensive panel of institutional bond holdings, I show that investor types from mutual funds and insurers to pensions and variable annuities actively offset most of their price-induced duration drift within a few quarters, both in Treasuries and corporate bonds. I construct a bond-level measure of exposure to duration-rebalancing demand and find that around FOMC announcements, bonds with higher exposure to strong rebalancers exhibit dampened price reactions to monetary surprises. The gap grows with maturity, builds for two weeks, and closes by the next meeting. A long-short portfolio on this exposure earns an annualized alpha above 3% when held for two weeks after each FOMC meeting. A preferred-habitat model with a soft penalty on duration deviations explains the channel and decomposes the rebalancing into a habitat-anchoring and a mean-variance margin. Estimated fund by fund, habitat anchoring accounts for nearly 90% of the rebalancing. The calibrated model matches the empirical regression on the scale of the long-end dampening: per 100 basis points of monetary shock, the channel dampens the ten-year yield response by about 6(5) basis points in Treasuries (corporates).

*PhD Candidate in Finance, London School of Economics. Email: j.lou6@lse.ac.uk. I am deeply grateful to Dimitri Vayanos, Christian Julliard, Walker Ray, Manav Chaudhary, Christopher Polk, Ian Martin, Kathy Yuan, Cameron Peng, Peter Kondor and Martin Oehmke for their valuable comments and suggestions. All remaining errors are my own.

1 Introduction

Monetary policy is transmitted to the real economy largely through interest rates, and a long-standing puzzle is why short rate policy shocks move long-term rates as much as they do. Hanson and Stein (2015) show that a 100 basis point increase in the two-year nominal yield on an FOMC announcement day moves the ten-year nominal forward rate by 45 basis points, with almost all of it, 42 basis points, falling on the real rate rather than on expected inflation. Such a strong response at the long end is hard to reconcile with standard New Keynesian models, in which nominal rigidities have unwound a decade out, leaving monetary policy close to neutral for real rates at that horizon. The pattern is robust across samples and shock measures (Bauer and Swanson, 2023; Hillenbrand, 2025).

Hanson and Stein interpret this excess sensitivity through the term premium, and their account is a demand story called "Reaching for Yield": when short rates fall, income-oriented investors such as insurers and pension funds buy long-term bonds to keep portfolio yield up, compressing term premia and pushing long rates further in the direction of the shock. Because the break-even-inflation component does not significantly respond to the shock at this horizon, the nominal and real forwards nearly move one-for-one this far out. However, not every account of the puzzle is demand-based; Nakamura and Steinsson (2018), for one, trace the long-rate response to a Fed information effect that revises expectations about fundamentals rather than term premia. Within the demand view, a growing literature identifies forces that move long rates by reshaping who bears duration risk: wealth effects on risk-averse arbitrageurs (Kekre, Lenel, and Mainardi, 2024), granular reallocations in Treasury demand (Jansen, Li, and Schmid, 2025), and flow-induced scaling by bond mutual funds (Fang, 2025). What these channels share is a direction: the trading they identify pushes long-term prices the same way as the shock.

This excess sensitivity, however, does not last. The initial excess response of long-term forward rates holds for the first couple of weeks, then unwinds over the next two to three months (Figure 1, left panel). Hanson and Stein (2015) note this mean reversion and attribute it to term-premium dynamics, but leave the demand-side mechanism implicit. The reversal could have more than one source: part is the natural dissipation of announcement-window pressure as prices revert toward fundamentals, and part could be an active stabilizing demand that this paper identifies, present from the announcement itself and building over the following days as

investors trade against the price-induced drift.

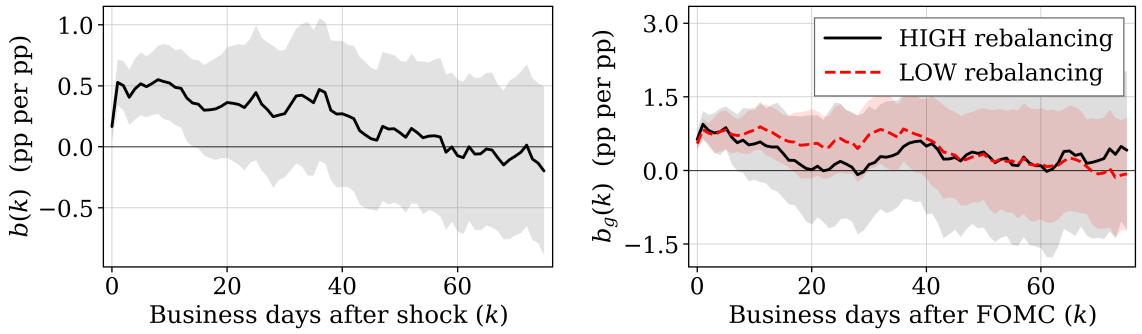
This paper identifies one such stabilizing demand force, which I call "Reaching for Habitats". Institutions hold a preferred portfolio duration, and a monetary shock disturbs it mechanically. Because bonds of different maturity reprice by different scales, any held-fixed portfolio's duration drifts away from its pre-shock level: a rate cut normally drifts it up, a rate hike drifts it down. An investor anchored to a target duration zone trades back toward it, selling[1] long-term bonds just as the shock has bid their prices up more than short-term, and buying them when it has knocked them down more than short-term. This duration-restoring demand leans against the initial price move and stabilizes the transmission of monetary policy to long rates.

Its fingerprint is already visible in the raw cross-section. Dividing a nine-to-eleven-year Treasury cohort by the average rebalancing strength of the institutions holding each bond, the bonds held by strong rebalancers show visibly smaller yield responses over the two months following each FOMC than those held by weak rebalancers, the same window over which the aggregate reversal accumulates (Figure 1, right panel).[2] Because the stabilizing demand and the amplifying forces are set off by the same shock and act over overlapping windows, the two are hard to disentangle directly in the aggregate. The analysis therefore works from the cross-section, the differential response of bonds whose holders rebalance more versus less strongly, and recovers the channel's magnitude by combining the regression spread with the model calibration, rather than reading it off the aggregate response directly. "Rebalancing strength" is itself a fund-level measure that I construct in the sections below; here it serves only to motivate.

The empirical analysis begins with showing that a broad class of institutional investors holds persistent, sticky portfolio duration. Combining the Morningstar mutual-fund and eMAXX institutional holdings data, I assemble a comprehensive panel of institutional bond holdings covering the major types of U.S. institutions over more than two decades. I classify investors by account type: mutual funds, insurers, pensions, variable annuities, hedge funds, and closed-end funds, among others. Each account is treated as a separate fund, so that the analysis isolates ac-

[1] Here "selling" need not mean submitting a sale or a short order. It denotes a *relative* duration demand that resets the portfolio back toward its habitat, and can be implemented without selling any long-term bond, for instance by directing net new purchases toward shorter maturities. I refer to this relative tilt as net selling of duration: for prices, what matters is the relative demand across maturities, which exerts the same directional pressure along the curve and in the cross-section as an outright sale would.

[2] The aggregate panel estimates $\Delta f^{10y}_{t-1,t+k} = a(k) + b(k)\,\Delta y^{2y}_{t-1,t+1} + \varepsilon_{t,k}$ separately at each horizon k on FOMC dates, with Newey–West HAC lag $\lceil 1.5k/25 \rceil$ following the spacing of FOMC meetings. The cross-sectional panel sorts the nine-to-eleven-year-maturity Treasury cohort into within-duration-tertile HIGH and LOW rebalancing-strength bins and runs the same regression on each group's cohort-mean yield change (post-2010, $N = 114$ events).



(a) Aggregate response of the 10y forward

(b) Cross-section by rebalancing strength

Figure 1: **The Hanson–Stein Reversal and Its Cross-Sectional Fingerprint.** Both panels regress a long-rate change on the change in the 2-year nominal yield over the FOMC announcement window, $\Delta y^{2y}_{t-1,t+1}$, and trace the resulting coefficient as the horizon k lengthens. Panel (a) replicates Hanson and Stein (2015): the dependent variable is the cumulative change in the 10-year-ahead instantaneous forward rate from the day before the announcement to k trading days after, $\Delta f^{10y}_{t-1,t+k}$, so the plotted coefficient $b(k)$ measures how much of a 2-year shock still passes through to the 10-year forward k days later (1995–2024 FOMC sample, Federal Reserve GSW Svensson curve). The pass-through peaks within the first weeks and then reverses over the following months. Panel (b) runs the same regression within a 9–11-year Treasury cohort, with the dependent variable the cohort-mean yield change of bonds sorted, within duration tertiles, into HIGH and LOW rebalancing-strength bins (post-2010); the high-rebalancer cohort's response is smaller and unwinds sooner. Shaded bands are 95% confidence intervals.

counts with economically distinct duration mandates that share a common manager. Two descriptive facts characterise this universe. First, portfolio duration is highly persistent: in both Treasury and corporate-bond markets, nearly 90% of funds remain in the same duration bucket from one quarter to the next, with AR(1)-style persistence around 0.96–0.99 at the one-year horizon and above 0.80 even at five years, essentially universal across the institution types in the panel. Second, this stickiness is not passive: decomposing quarterly fund-level duration changes into a passive component (the mechanical duration drift that price moves alone induce on a held-fixed portfolio) and an active component (the trading-induced residual), the active component systematically and significantly offsets the passive drift, with cumulative three-quarter offsets of a one-year passive drift around 81% in Treasury and 54% in Corporate, observed across the institution types in the panel.

The tightness of these duration anchors varies across institution groups: life insurers and bond mutual funds sit at the strongly-anchored end, hedge funds at the discretionary end, and the rest in between, with the dominant driver for each group detailed in Table 1 of Section 2. The paper does not investigate which driver is operative for any given fund, since each is studied in its own institutional literature. It

takes stable duration targets as given and asks how the resulting rebalancing trades transmit monetary policy to long-term interest rates.

To identify who supplies this rebalancing demand, I estimate for every fund a rebalancing sensitivity $\widehat{\beta}^{\text{reb}}_{f,t}$ for each quarter: the slope of the fund's active duration changes on its passive price-induced deviations, estimated on the fund's own history and shrunk toward an institution-type prior. A value of zero means the fund lets duration drift with prices; minus one means every unit of passive drift is fully offset within the quarter. The resulting panel is centered near full rebalancing and preserves wide heterogeneity across funds.

If duration rebalancing transmits monetary policy along the yield curve, bonds disproportionately held by strong rebalancers should react less to FOMC shocks than bonds held by weak rebalancers. Aggregating the fund-level rebalancing sensitivities into a bond-level signal makes this testable: $\text{rebal_str}_{i,t}$ is the ownership-weighted rebalancing intensity of a bond's holders, a shift-share construction with predetermined par-value ownership shares as exposures and the real-time fund sensitivities as shifts (Borusyak, Hull, and Jaravel, 2022). A bond inherits a high rebal_str because its holders happen to rebalance strongly on average historically. Section 5.1 develops the identification argument.

The FOMC announcement window is the natural laboratory for this test: the policy-news shock has relatively clean high-frequency identification (Nakamura and Steinsson, 2018; Bauer and Swanson, 2023), and it mechanically pushes every multi-bond portfolio off its target duration, opening a synchronised rebalancing trigger across the institutional universe. Regressing cumulative bond returns over post-FOMC windows of length k trading days on the shock interacted with $\text{rebal_str}_{i,t-1}$, the interaction is significantly positive: bonds with higher rebalancing-strength exposure react less to FOMC surprises. The dampening grows along the duration spectrum, switches on at the announcement, builds over roughly two weeks, and closes over the remainder of the inter-meeting window. A differential that opens with the shock and closes before the next meeting is the signature of transitory demand pressure rather than information or a mechanical price artifact, and it is tradeable: a long-short portfolio sorted on rebal_str around FOMC events earns Newey–West annualized alphas of 2.5–4.1% over the first two post-FOMC weeks and nothing when held across the full inter-meeting window.

The channel can be modelled by micro-founding a continuous-time preferred-habitat model (Vayanos and Vila, 2021). Habitat investors are mean-variance op-

timizers who, beyond trading off risk against return, pay a soft quadratic penalty δ whenever their portfolio duration deviates from an investor-specific target; they trade against risk-averse arbitrageurs, and the bond market clears. Three results organize the rest of the paper. First, the stabilizing rebalancing arises even under pure mean-variance optimization. A mean-variance investor wants to hold a certain amount of interest-rate risk; when a rate cut lifts long-bond prices, its existing long positions come to carry more of that risk than it wants, so it sells some of them, leaning against the price rise, even with no duration target at all ($\delta = 0$). The cut also raises long bonds' expected return, which pulls the other way, but in the empirically relevant region the risk-rebalancing motive dominates and the investor sells on net. The soft duration anchor ($\delta > 0$) strengthens the pull back toward target. Stabilization is thus a property of ordinary portfolio choice, not of hard-mandate compliance. Second, each fund's rebalancing splits cleanly into a mean-variance margin (rational re-optimization of duration risk) and a habitat margin (the pull back toward the target), with the relative weight set by how tightly the fund is anchored, a split I estimate fund by fund in the data. Third, all of this reaches prices through a single channel. By selling long bonds into the rally, habitat investors hand duration risk to the risk-averse arbitrageurs just as the rate cut has swollen the dollar amount of that risk in the market; bearing more of it, arbitrageurs demand a higher term premium, which pushes long yields back up and blunts their initial fall.

How much of the rebalancing is habitat? Active rebalancing could reflect habitat anchoring, the pull back toward a fund-specific duration target, or rational mean-variance re-optimisation as expected returns move with the shock. The model nests both and identifies the mix fund by fund through the habitat-pull strength ω_f; estimated from each fund's own trading, habitat anchoring carries roughly ninety percent of the rebalancing in both markets. The split matters beyond decomposition, because it pins the model's calibration. With the mean reversion taken from the yield curve and the rebalancing intensity from the holdings panel, neither touching the pricing regression, the model predicts the dampening the FOMC cross-section should display, and at the ten-year maturity the model puts the dampening at roughly six basis points per 100 basis point shock in Treasuries and five in corporates, inside the regression's confidence bands. Two identification chains that share no data land on the same magnitude.

The stabilization is partial, not complete. The rebalancing channel coexists with the amplification channels documented in recent work; both are set off by the same

shock, and what reaches prices is their net. The data pin down the timing of the stabilizing side: the cross-sectional dampening switches on at the announcement and accumulates over the following two weeks as portfolios trade back toward their targets. When the amplifying flows act is not something this paper measures, though the event studies in that literature place much of their action inside the announcement window itself. Identification here does not depend on settling that question: the dampening is identified from cross-sectional variation in exposure to rebalancing demand, and requires no stand on the two forces' relative timing or size.

The paper's contribution is to document, identify, and quantify the stabilizing side of this ledger. A broad cross-section of institutions holds persistent duration habitats and actively restores them; the behavior is strong enough to show up in the cross-section of bond prices around FOMC announcements; and a canonical preferred-habitat model, calibrated without touching the pricing regression, reproduces the magnitude of the cross-sectional dampening. In doing so the paper also supplies an empirical micro-foundation for the downward-sloping habitat demand that the preferred-habitat literature has assumed since Vayanos and Vila (2021): the anchoring strength that generates the slope is measured fund by fund, and it is the dominant margin.

Related Literature. This paper speaks to four strands of literature on how monetary policy is transmitted through intermediaries and how demand shapes the term structure.

Amplification channels in monetary policy transmission. A large literature identifies forces that amplify the response of long-term yields to monetary surprises. Hanson and Stein (2015) attribute the excess sensitivity of long-term real forward rates to the reaching-for-yield behavior by income-oriented investors. Kekre, Lenel, and Mainardi (2024) formalize a financial-accelerator amplification channel via arbitrageur wealth effects. Jansen, Li, and Schmid (2025) embed estimated investor demand functions into a Vayanos–Vila (2021) equilibrium and find that monetary tightening raises term premia because granular Treasury investors with high cross-maturity elasticities rebalance toward shorter Treasuries, forcing risk-averse arbitrageurs to absorb additional duration risk. Fang (2025) and Kuong, O'Donovan, and Zhang (2024) emphasize a flow-induced amplification channel: bond funds experience large outflows after tightening shocks and scale down holdings, contracting credit supply more than traditional investors would. Lu and Wu (2025) study cross-asset-class rebalancing, in which mandate-bound institutions sell equities to rebuild target allocations after

a shock and amplify the equity reaction. See also Harvey, Mazzoleni, and Melone (2025) on stock–bond reallocation, and Selgrad (2023) on QE-induced cross-market reallocation within fixed income. The channel in this paper runs on a different margin and in the opposite direction: within-bond duration rebalancing that leans against the initial price move and partially undoes it over the days following the announcement, complementing rather than competing with these cross-asset amplification channels and making the net long-rate response reflect the relative magnitude of the two forces. Identification of monetary surprises themselves has been refined by Gürkaynak, Sack, and Swanson (2005), Nakamura and Steinsson (2018), Bauer and Swanson (2023), and Cieslak and Schrimpf (2019); Hillenbrand (2025) documents a secular Fed influence on long rates; Hanson, Lucca, and Wright (2021), and Boehm and Kroner (2025) study related transmission margins.

Preferred-habitat theory and demand-driven term structure. The preferred-habitat view dates back to Culbertson (1957) and Modigliani and Sutch (1966) and has been formalized in modern equilibrium models in which maturity-specific investors with downward-sloping demand interact with risk-averse arbitrageurs (Vayanos and Vila, 2021). The framework has been applied to government debt policy (Guibaud, Nosbusch, and Vayanos, 2013), exchange rates (Gourinchas, Ray, and Vayanos, 2025; Greenwood, Hanson, Stein, and Sunderam, 2023), swap markets (Hanson, Malkhozov, and Venter, 2024), and quantitative easing (Ray, Droste, and Gorodnichenko, 2024); empirical work has linked supply and demand shocks to term-structure movements within this framework (Greenwood and Vayanos, 2014; D'Amico and King, 2013; Greenwood, Hanson, and Vayanos, 2024). Closest to this paper, Jansen, Li, and Schmid (2025) and Chaudhary, Fu, and Zhou (2026) estimate granular demand functions across institutional sectors of the Treasury market in the Koijen–Yogo (2019) demand-system tradition. Relative to that work, I do not estimate a full demand system; instead, I extract a single, theoretically motivated dimension of investor behavior, the active offsetting of passive duration drift, and show that variation across bonds in the exposure to this behavior delivers cross-sectional predictions about FOMC-window returns that the demand-system approach is not designed to address. What's more, while the downward-sloping demand is a foundational primitive of preferred-habitat theory, this paper provides direct micro-level evidence on its empirical counterpart and on the price impact it generates around monetary events. Concurrent work by Ding and Czech (2026) examines the arbitrageur side of the same framework using UK regulatory dealer data, complementing this paper's focus

on the habitat-investor side.

Institutional duration management and ALM-driven trading. A long-standing literature documents that life insurers and pension funds practice asset–liability management and adjust the duration of their bond portfolios in response to interest-rate movements (Domanski, Shin, and Sushko, 2017; Ozdagli and Wang, 2019; Wenning and Li, 2025; Li, 2025). Closest to this paper, Ozdagli and Wang (2019) introduce an "active duration adjustment" decomposition for U.S. life insurers. My setting extends this work along three margins: I cover a broad cross-section of U.S. institutional bond investors rather than life insurers alone (and find qualitatively similar offsetting across nearly all sectors); I take the active–passive decomposition from a fund-level holdings statistic to a bond-level cross-sectional asset-pricing test via the shift-share signal $\text{rebal_str}_{i,t}$; and I provide a model-based decomposition into mean-variance and habitat components, showing that mean-variance optimization with even a soft duration anchor is sufficient to generate stabilizing rebalancing, so that the empirical finding does not hinge on identifying the precise institutional friction (mandate, benchmark, regulatory capital, ALM) that produces the duration target. The negative-convexity "hunt-for-duration" direction documented in Domanski, Shin, and Sushko (2017) for German insurers in 2014 is shown in Appendix Table 14 to be a single-episode and single-jurisdiction phenomenon.

Demand-based asset pricing and market inelasticity. The inelastic-markets hypothesis of Gabaix and Koijen (2021) implies that systematic rebalancing flows can have large price effects. The demand-system approach of Koijen and Yogo (2019) has been extended to Treasuries (Chaudhary, Fu, and Zhou, 2026; Jansen, Li, and Schmid, 2025), corporate bonds (Chaudhary, Fu, and Li, 2022; Coppola, 2025; Siani, 2025), money markets (Doerr, Eren, and Malamud, 2023), and international markets (Jiang, Krishnamurthy, and Lustig, 2022). Bogousslavsky (2016) and Peng and Wang (2025) show that mechanical, infrequent, or factor-based rebalancing can generate predictable return autocorrelation and demand pressure. Methodologically, my rebal_str measure has the structure of a shift-share (Bartik) instrument, in the spirit of Borusyak, Hull, and Jaravel (2022) and Goldsmith-Pinkham, Sorkin, and Swift (2020), with predetermined ownership shares as exposures and fund-level rebalancing sensitivities as shifts. This paper identifies a specific source of that inelasticity, duration rebalancing, and links it to monetary policy transmission through the cross-section of bond returns around FOMC events.

The rest of the paper is organized as follows. Section 2 describes the data. Sec-

tion 3 develops the equilibrium model. Section 4 documents duration habitats and active rebalancing across institutional investors. Section 5 tests for stabilization in the cross-section of bond returns around FOMC shocks. Section 6 decomposes the rebalancing into its mean-variance and habitat margins and confronts the calibrated model with the cross-sectional estimates. Section 7 concludes.

2 Data

I combine portfolio-level holdings of broad coverage of U.S. institutional bond investors with bond-level price and characteristic data and high-frequency FOMC monetary policy surprises. The panel runs from early 2000 through end-2024.

2.1 Bond returns and characteristics

For Treasury securities, I obtain prices, returns, and characteristics (maturity, duration, convexity, coupon, on-the-run status, amount outstanding, bid–ask spread) from the CRSP Treasury databases. I exclude Treasury bills, callable Treasuries, and TIPS from the main analysis, whose pricing is confounded by convenience yields, callability, and inflation risk.

For corporate bonds, I obtain transaction-level prices from TRACE Enhanced and issue-level characteristics (coupon, maturity, ratings, issue size, callability) from Mergent FISD. Issuer-level financial information is taken from Compustat. I follow Dickerson, Mueller, and Robotti (2023) and the Open Source Bond Asset Pricing project for the cleaning of TRACE records and the construction of daily and monthly bond returns, bond analytics (duration, convexity, yield-to-maturity, credit spreads), and bond-month liquidity measures. The corporate-bond sample is filtered at two levels. For the holdings-based analysis (Section 4), which measures fund-level duration habitats and rebalancing, I retain the broad universe of fixed-cash-flow, non-distressed corporate bonds in fund holdings, to maximize coverage across institution types. For the FOMC-window evidence and asset-pricing tests (Section 5), I additionally restrict to liquid investment-grade-eligible bonds: face value outstanding of at least $250 million, monthly turnover of at least 5%, monthly trading volume of at least $50 million, and a zero-trade-day ratio of at most 0.5.[3] The rebalancing measures are thus estimated on the broad institutional universe, while the return tests run on bonds whose price reactions to FOMC shocks are reliably measured.

[3]The $50 million volume threshold matches the inclusion criterion of the Bloomberg US Aggregate Bond Index.

Daily and monthly broad-market indices for Treasuries and corporates come from the WRDS Bond Database and Bloomberg. Daily Gürkaynak–Sack–Wright (GSW) yield-curve factors used as Treasury risk factors are from the Federal Reserve.

2.2 Institutional holdings

I assemble a comprehensive panel of institutional bond holdings by combining two complementary data sources, each with its own coverage advantage.

Mutual funds: Morningstar. I obtain quarterly holdings of U.S. open-end mutual funds from the Morningstar Global Databases for 1990–2024Q1.[4] I organise the Morningstar mutual fund universe into three analysis categories: "Taxable Bond MF" and "Allocation MF" are kept as their own groups, and "Other MF" aggregates all remaining open-end U.S. mutual fund types (equity, international, municipal-bond, commodity, and insurer- or pension-affiliated funds), retained whenever they hold material Treasury or corporate-bond positions. I exclude money-market funds, whose holdings are largely restricted to T-bills with convenience-yield dynamics outside this paper's scope.

Other institutional investors: Refinitiv eMAXX. I obtain quarterly portfolio-level holdings of non-mutual-fund institutional bond investors from Refinitiv eMAXX, which compiles fixed-income holdings across the U.S. institutional fixed-income universe over 1998Q2–2024Q4. Each holding in eMAXX carries a FUND_CLASS code that identifies the mandate-bearing sub-account: life-insurance general account, P&C general account, separate account funding variable annuities, defined-benefit pension trust, hedge-fund vehicle, closed-end fund, sovereign or central-bank account, and so on. I keep the FUND_CLASS taxonomy throughout rather than aggregating to the parent-firm level, following Bretscher, Hsu, Schmid, and Tamoni (2026); sub-accounts within the same firm can have economically distinct duration mandates (a life insurer's general account vs. its separate accounts), and the FUND_CLASS cut isolates these.[5]

[4] Morningstar's coverage of U.S. mutual fund holdings is comprehensive from the early 1990s onward. The 2024Q1 cutoff reflects the latest vintage available at the time of analysis. For funds reporting at monthly frequency, I keep the quarter-end observation.

[5] Mapping the 35 raw FUND_CLASS codes into research groups gives nine eMAXX-side categories used in the analysis (Life, P&C, and Other Insurers; Variable Annuity; Pension; Hedge Fund; Closed-End MF; Sovereign; Other Institutional); combined with the three Morningstar mutual-fund groups (Taxable Bond MF, Allocation MF, Other MF) the panel has twelve groups in total. Classifying at the parent-firm level would instead pool economically distinct sub-accounts (for instance, 22% of a firm-level investment-manager bucket is externally managed life-insurance separate-account money), which the FUND_CLASS cut avoids. The combined sample restricts to U.S.-domiciled funds whose bond portfolios contain at least 10% Treasuries (or 10% corporate bonds) and at least five Treasury

The groups differ in why they hold a duration preference. I take these targets as given rather than modelling their origin; Table 1 summarises the arrangement behind each, from liability matching at life insurers to benchmark tracking at bond mutual funds.

Table 1: **Sources of Duration Preferences by Institution Group.** For each group, the institutional arrangement that pins down its target duration. Groups follow the eMAXX FUND_CLASS classification, which assigns each holding to the mandate-bearing sub-account rather than to the parent firm.

Institution Group	Where the duration target comes from
Life Insurers	Life policies and annuities create long, predictable liabilities; asset duration is managed to match them, and capital regulation reinforces the match.
P&C Insurers	Claims arrive sooner and less predictably, so asset duration is kept short to intermediate.
Variable Annuities	Separate accounts track policyholder-chosen sub-funds; the insurer sets no duration target of its own.
Pensions	Asset duration is matched to benefit obligations, with trustees setting funding glide paths.
Taxable Bond MFs	The prospectus states a duration range and the portfolio tracks a bond index such as the Bloomberg US Aggregate.
Allocation MFs	A stated asset mix keeps the fixed-income sleeve close to its benchmark's duration.
Hedge Funds	Risk limits but no duration target; positioning is opportunistic.
Other groups	Other insurers, other and closed-end mutual funds, sovereign accounts, and residual institutional mandates; targets are mandate-specific and each group is individually small in the panel.

Coverage and external validity. The combined panel covers the actively managed institutional segment of the U.S. bond market. The panel captures about 17% of the combined Treasury and investment-grade corporate market's dollar duration (market value times duration) over the post-2010 window: 38% of investment-grade corporates, 7% of Treasuries, and roughly 28% of the market beyond twelve years (Appendix G; Section 5.1). The holders the panel misses are large, chiefly foreign investors and the Federal Reserve in Treasuries. Three considerations contain the concern. First, the tests are cross-sectional. What moves the price spread is the gap in rebalancing demand between bonds held by strong rebalancers and bonds held

holdings or ten corporate-bond holdings in a given quarter, excluding the rare short positions reported by some hedge funds. Throughout the paper, f indexes a separately managed portfolio (the eMAXX FUND_CLASS-level identifier or the Morningstar fund identifier) rather than its parent firm. Because each holding is classified by mandate, bond portfolios run by hedge-fund, bank, or asset-management firms enter the panel under the mandate they actually serve (most often a mutual-fund or insurance sub-account); the resulting institutional panel is dominated by insurance accounts (Life and P&C together about 85% of institutional par value), and the hedge-fund, sovereign, and other-institutional groups are populated but small, retained for completeness rather than as representative samples of those universes.

by weak ones, so the panel need not be a representative draw of the market; an unobserved holder distorts the comparison only if its own FOMC-window trading lines up, bond by bond, with the rebalancing signal. Second, the missing holders have essentially no rebalancing demand to contribute, so they are unlikely to line up with it. Foreign official reserves hold Treasuries for safety and liquidity, revise their allocations at the horizon of years rather than days, and face no pressure to trade when a rate shock moves their portfolio's duration. The Federal Reserve is not a portfolio investor at all: it buys and sells on pre-announced programs whose amounts and maturities are set by policy, not chosen to manage duration. The construction of rebal_str handles such holders correctly, because ownership shares are measured against total amount outstanding: any holder the panel does not see enters the signal exactly as if its rebalancing sensitivity were zero, which for these holders is close to the truth. Consistent with this, Chaudhary, Fu, and Zhou (2026) find that since 2008 Treasury yield movements in shock episodes have been driven by domestic investors, the segment the panel does observe. Third, coverage is deepest exactly where the channel operates: beyond twelve years the panel holds close to three of every ten dollars of the market's duration risk. The aggregate calibration in Section 6 should nonetheless be read as a scaling of the cross-sectional estimates rather than a market-wide accounting identity.

2.3 Monetary policy surprises

FOMC announcement dates are obtained from the Federal Reserve Board. This paper uses the high-frequency monetary policy surprises proposed by Nakamura–Steinsson (2018) and refined by Bauer and Swanson (2023). The NS–BS FOMC shock data is publicly available and updated through 2024 at the Federal Reserve Bank of San Francisco.

Table 2 summarizes the data sources.

3 Model

I develop a continuous-time preferred-habitat model with soft duration anchoring that microfounds the downward-sloping habitat demand of Vayanos and Vila (2021) from observable portfolio behavior, generates closed-form predictions about long-end yield-response dampening, and provides direct cross-sectional empirical hypotheses. Full derivations and proofs are in Appendix J.

Data	Sources
Bond Returns and Characteristics	
Treasury prices, returns, characteristics	CRSP Treasury
Corporate-bond transactions and characteristics	TRACE Enhanced; Mergent FISD; Open Source Bond Asset Pricing
Issuer financials	Compustat
Treasury and credit market indices	WRDS Bond Database; Bloomberg
GSW yield-curve factors	Federal Reserve
Institutional Holdings (1990–2024)	
Open-end mutual funds	Morningstar Global Databases
Insurance (Life / P&C / Other), variable-annuity separate accounts, pensions, hedge funds, closed-end funds, sovereign accounts, other institutional	Refinitiv eMAXX (FUND_CLASS sub-account level)
Monetary Policy and Macro Expectations	
FOMC meeting dates	Federal Reserve Board
High-frequency monetary policy surprises	FRBSF; Nakamura and Steinsson (2018); Bauer and Swanson (2023)

Table 2: **Data Sources.** Bond prices, returns, and characteristics combine the CRSP Treasury database with TRACE Enhanced and Mergent FISD for corporates. Open Source Bond Asset Pricing supplies the standard cleaning and additional features. Institutional holdings combine Morningstar (U.S. open-end mutual funds, 1990–2024) with Refinitiv eMAXX (life and P&C insurers, variable-annuity separate accounts, pensions, hedge funds, closed-end funds, sovereign accounts, and other institutional mandates, 1998–2024), classified at the FUND_CLASS sub-account level (Bretscher, Hsu, Schmid, and Tamoni, 2026). FOMC dates are obtained from the Federal Reserve Board; high-frequency monetary surprises follow Nakamura and Steinsson (2018) and Bauer and Swanson (2023), available through the Federal Reserve Bank of San Francisco.

3.1 Setup

Time is continuous. The short rate r_t follows an OU process under the physical measure:

$$dr_t = \kappa_r(\bar{r} - r_t)\,dt + \sigma_r\,dB_t, \tag{1}$$

with mean-reversion speed $\kappa_r > 0$, long-run mean $\bar{r}$, and volatility $\sigma_r > 0$. There is a continuum of zero-coupon bonds indexed by remaining maturity $\tau \in (0, T]$. Bond prices are exponential-affine in the short-rate deviation,

$$P(t,\tau) = \exp\big[-A(\tau)(r_t - \bar{r}) - C(\tau)\big]\,M_t(\tau), \tag{2}$$

where $A(\tau)$ is the *effective duration loading*, $C(\tau)$ is a deterministic level term, and $M_t(\tau)$ is a maturity-specific idiosyncratic martingale capturing bond-level non-systematic risk.

Supply. Bond supply is fixed in face value: each maturity has a constant outstanding face value $\bar{N}(\tau)$. This is to mimick the fact that short-term interest rate policy is not

linked with central-banks long-term bonds operation, and isolating the short-term policy effect from other non-conventional monetary policies like QE, which could confound the clear illustration of the channel. Thus, the market value of supply, $S_t(\tau) = \bar{N}(\tau)P(t,\tau)$, therefore contracts mechanically when rates rise (long-bond prices fall).

Steady-state snapshot. The rebalancing analysis takes a snapshot at the steady state $r_t = \bar{r}$ with $M_t(\tau) = 1$ and applies a rate shock Δr. The level of $M_t(\tau)$ rescales prices uniformly and independent of every rebalancing object, so the normalisation is without loss of generality.

Agents. There are two types of agents: Arbitrageurs and Preferred-habitat (PH) investors, both of whom are CARA mean–variance optimizers over instantaneous dollar returns.

Arbitrageurs hold dollar positions $x_t(\tau)$ in τ year bonds and solve

$$\max_{\{x_t(\tau)\}} \int_0^T x_t(\tau)\big(\mu(t,\tau) - r_t\big)\, d\tau \;-\; \frac{\gamma_A}{2}\,\Sigma_{A,t}, \tag{3}$$

where $\gamma_A > 0$ is absolute risk aversion and $\Sigma_{A,t} \equiv \sigma_r^2 V_{A,t}^2$ is the factor variance of dollar returns, with $V_{A,t} \equiv \int_0^T x_t(u)A(u)du$ the dollar-duration exposure. Arbitrageurs run diversified books, so their idiosyncratic exposure is not penalised (Assumption 1 below).

Preferred-habitat (PH) investors form a continuum indexed by $i \in I = [0,1]$, normalised to total mass $|I| = 1$. Each investor has a personal target duration D_i^* and a personal duration-penalty intensity $\delta_i > 0$, while they are common in $(\gamma_H, \bar{W}, \sigma^{\mathrm{id}}(\tau))$. Investor i holds dollar positions $y_t^i(\tau)$ for maturity i, bears idiosyncratic risk ($\Sigma_{H,t}^i = \sigma_r^2(V_{H,t}^i)^2 + \int_0^T y_t^i(\tau)^2\sigma^{\mathrm{id}}(\tau)^2 d\tau$), and solves the mean–variance problem augmented by a quadratic penalty for deviations of portfolio Macaulay duration $D_{\mathrm{port},t}^i$ from D_i^*:

$$\max_{\{y_t^i(\tau)\}} \int_0^T y_t^i(\tau)\big(\mu(t,\tau) - r_t\big)d\tau \;-\; \frac{\gamma_H}{2}\Sigma_{H,t}^i \;-\; \frac{\delta_i}{2}\big(D_{\mathrm{port},t}^i - D_i^*\big)^2\bar{W}, \tag{4}$$

with $D_{\mathrm{port},t}^i = (1/\bar{W})\int_0^T y_t^i(\tau)\,\tau\,d\tau$ and $\bar{W}$ a constant per-investor wealth scale (Assumption 2). The mass-1 normalisation makes integrals over i both averages and aggregates ($|I| = N$ would rescale parameters in the standard CARA sense without changing any pricing object). The penalty stands in for whatever anchors a fund's

duration—mandates, benchmark tracking, regulatory limits, or preference; the model is agnostic about the source.

The model treats bonds as differentiated by duration alone; bond-level non-systematic variation, including credit risk for corporates, is absorbed into the idiosyncratic martingale $M_t(\tau)$. This is the standard simplification in the preferred-habitat tradition—Vayanos and Vila (2021) note that habitats "could alternatively concern bonds that differ in liquidity or in the type of issuer," and recent empirical work treats maturity as the habitat dimension while handling credit through controls (Selgrad, 2023; Greenwood, Hanson, and Vayanos, 2024; Coppola, 2025; Cavaleri, 2023, extends the framework to joint duration–credit pricing). Empirically, non-duration variation is handled by sample selection and controls: the corporate sample is restricted to highly liquid investment-grade bonds (Section 2); credit, liquidity, and issuer controls are interacted with the FOMC surprise in every cross-sectional regression (Section 5.1); the long-short factor model risk-adjusts with corporate-market and credit-spread factors (Section 5.2); and σ_0^{id} is calibrated higher in Corporate (0.040) than in Treasury (0.010). Therefore, the "mean-variance" channel in Section 6.1 is the duration-driven MV channel, not a multi-factor loading on credit or liquidity premia.

3.2 Key assumptions

There are three assumptions structuring the model.

Assumption 1 (Asymmetric idiosyncratic risk bearing)**.** *PH investors hold concentrated portfolios and bear idiosyncratic risk: $\sigma^{id}(\tau) > 0$ enters $\Sigma^i_{H,t}$. Arbitrageurs run diversified books: their portfolio variance is the factor variance $\sigma_r^2 V_{A,t}^2$, and their first-order conditions carry no idiosyncratic term.*

Idiosyncratic risk is carried only by preferred-habitat investors, so that their portfolio choice does not degenerate through a decoupling of Macaulay duration from factor duration. The arbitrageur is the Vayanos–Vila (2021) arbitrageur, compensated for factor risk only, which makes idiosyncratic risk unpriced in equilibrium exactly; the model inherits this setting unchanged. [6]

Assumption 2 (Constant sector wealth)**.** *The PH sector's reference scale $\bar{W}$ is held constant in time: fund inflows and outflows roughly balance, keeping each fund's*

[6]The assumption is not only a normalisation: if neither sector priced idiosyncratic risk, Macaulay duration and factor duration would decouple—the mandate could be met costlessly through the shape of holdings—and rebalancing would leave prices unchanged (Appendix J, Lemma 2). The channel therefore operates through, and scales with, the idiosyncratic risk borne by the habitat sector.

wealth at a constant level.

This standard simplification isolates the model from the confounding channel of fund flows; moreover, under CARA preferences optimal dollar positions in risky assets are independent of wealth, so the FOCs and equilibrium pricing implications are unchanged from a fully wealth-dynamic specification.

Assumption 3 (Local linearisation around the steady state)**.** *For small deviations of r_t from $\bar{r}$, the aggregate equilibrium objects Θ_t (effective-duration supply), $\mathcal{D}_{dev,t}$ (PH portfolio duration deviation), and $S_t(\tau)$ (bond market value) are linear in $(r_t - \bar{r})$.*

This is the standard affine-pricing approximation in the term-structure literature; the precise definitions of these objects appear in Appendix J.

3.3 Equilibrium drift and the pricing ODE

Three conditions determine the equilibrium drift. The arbitrageur's first-order condition prices systematic risk only (Assumption 1),

$$\mu(t,\tau) - r_t \;=\; \gamma_A \sigma_r^2 \, A(\tau) \, V_{A,t}, \tag{5}$$

with $V_{A,t} \equiv \int_0^T x_t(u) A(u)\, du$ the arbitrageur's factor exposure.

Preferred Habitat investor i's first-order condition adds compensation for idiosyncratic risk and the habitat wedge,

$$\mu(t,\tau) - r_t \;=\; \gamma_H \sigma_r^2 \, A(\tau) \, V_{H,t}^i \;+\; \gamma_H \, \sigma^{\mathrm{id}}(\tau)^2 \, y_t^i(\tau) \;+\; \delta_i \, \tau \left(D^i_{\mathrm{port},t} - D_i^* \right), \tag{6}$$

where $V_{H,t}^i \equiv \int_0^T y_t^i(u) A(u)\, du$ is investor i's factor exposure.

The market clears maturity by maturity,

$$x_t(\tau) + y_t(\tau) \;=\; S_t(\tau), \qquad \tau \in (0, T]. \tag{7}$$

Because every condition is linear, integrating (6) over i leaves only sector total values, while the same heterogeneity generates the cross-sectional dispersion in demands and rebalancing intensities exploited empirically (Appendix J.4.1). Weighting (7) by $A(\tau)$ and integrating over $\tau \in (0, T]$ gives $V_{A,t} + V_{H,t} = \Theta_t$, where $\Theta_t \equiv \int_0^T S_t(u) A(u)\, du$ is the aggregate effective-duration supply and $V_{H,t} \equiv \int_0^T y_t(u) A(u)\, du$ is the habitat sector's *factor duration*. Substituting into (5),

$$\mu(t,\tau) - r_t = \lambda_t \, A(\tau), \qquad \lambda_t = \gamma_A \, \sigma_r^2 \big(\Theta_t - V_{H,t} \big), \tag{8}$$

so arbitrageurs, the marginal pricer, link the factor risk with the market price of risk, which also carries γ_A. Substituting the Itô drift of the affine bond price into (8) and matching coefficients of $(r_t - \bar{r})$ yields the ODE for $A(\tau)$.

Proposition 1 (Fundamental pricing ODE)**.** *Under Assumptions 1–3, the effective duration loading satisfies*

$$\boxed{A'(\tau) + \kappa_\eta\, A(\tau) = 1, \qquad A(0) = 0,} \tag{9}$$

with the unique closed-form solution

$$A_\eta(\tau) = \frac{1 - e^{-\kappa_\eta \tau}}{\kappa_\eta}, \qquad \kappa_\eta \equiv \kappa^* + \eta_r. \tag{10}$$

The effective mean reversion κ_η can be decomposed into a no-rebalancing benchmark κ^* and a rebalancing increment η_r,

$$\kappa^* \equiv \kappa_r + \gamma_A \sigma_r^2\, \Omega_0 = \kappa_r + \frac{\sigma_r^2\, \Xi}{\tau_{agg}}, \qquad \eta_r \equiv \gamma_A \sigma_r^2 \big(\Omega - \Omega_0\big) \;>\; 0, \tag{11}$$

where $\Omega \equiv -\partial(\Theta_t - V_{H,t})/\partial r_t|_{r_t=\bar r}$ measures how strongly the duration that arbitrageurs must absorb varies with the short rate, Ω_0 is its value when habitat investors behave as pure mean–variance investors ($\delta_i \equiv 0$), $\Xi \equiv \int_0^T \bar S(s) A(s)^2\, ds$ is the corresponding swing in total duration supply, and $\tau_{agg} \equiv 1/\gamma_A + 1/\gamma_H$. The sign $\eta_r > 0$ is unconditional; η_r rises with each δ_i and vanishes as $\sigma^{id} \to 0$ (Appendix J, Proposition 7).

Equation (9) is the canonical Vayanos–Vila pricing equation. Rebalancing raises the effective mean reversion from κ^* to $\kappa_\eta = \kappa^* + \eta_r$, so the entire dampening profile below is pinned by (κ^*, η_r).

The objects in (11) have a simple reading.

Step 1: a rate cut creates extra duration to hold. Supply is fixed in face value. When rates fall, every bond's price rises, so the market value of duration outstanding grows: $\Theta_t \approx \bar\Theta - \Xi\,(r_t - \bar r)$. Per unit of rate cut, an extra Ξ of duration must end up on someone's balance sheet.

Step 2: without duration habitats, the two sectors split it. If habitat investors cared only about risk and return, each sector would take a share of the extra duration in proportion to its risk tolerance. The arbitrageurs' share is

$$\Omega_0 = \Xi \cdot \frac{1/\gamma_A}{\tau_{\text{agg}}},$$

which is the standard supply term in κ^*.

Step 3: with duration habitats, arbitrageurs take more than their share, so the term premium must move more. The rate cut has already pushed PH investors' portfolio duration above target, so habitat investors would be selling duration, not

buying it. Arbitrageurs, consequently, end up holding more than the proportional share: $\Omega > \Omega_0$. To make them willing, the term premium has to rise more after a cut and fall more after a hike, which is translated into a mean-reversion increment:

$$\eta_r = \gamma_A \sigma_r^2 (\Omega - \Omega_0) > 0.$$

The sign needs no side conditions. Appendix J (Proposition 7) gives Ω in closed form and shows η_r grows with the penalties δ_i and vanishes as $\sigma^{\text{id}} \to 0$: when reshaping the portfolio is free, the mandate no longer constrains how much interest-rate risk the sector holds (Lemma 2).

3.4 Stabilization of the long-end yield response

The closed-form solution (10) implies that rebalancing dampens long-rate sensitivities relative to the Vasicek benchmark $A_{BM}(\tau) = (1 - e^{-\kappa^*\tau})/\kappa^*$, the loading that obtains at $\eta_r = 0$.[7]

Proposition 2 (Dampened yield response: a hump-shaped profile)**.** *For any $\eta_r > 0$ and $\kappa^* > 0$, the effective duration is strictly below benchmark at every maturity, $A_\eta(\tau) < A_{BM}(\tau)$, and the implied dampening of the yield response per unit short-rate shock,*

$$D(\tau) \;=\; \frac{A_{BM}(\tau) - A_\eta(\tau)}{\tau}, \tag{12}$$

is hump-shaped*. It vanishes as $\tau \to 0$ and as $\tau \to \infty$, and attains a single interior peak.*

The proof is in Appendix J. The hump has economic meaning. At the short end, prices are pinned by the policy rate, and arbitrageurs' capacity is ample relative to any rebalancing flow. At the long end there is little left to stabilize: with mean reversion, a transitory short-rate move barely reaches thirty-year yields in the first place. Although the price impact could be larger (yields times duration), the yield response itself is small. Section 6.2 documents this hump in the cross-section, with regression-implied magnitudes in line with the calibrated model.

[7] Vayanos and Vila (2021) describe the same phenomenon as a "flatter yield-response function" (the function $\tau \mapsto \partial y(t,\tau)/\partial r_t$ flatter against τ). We use "dampening" or "stabilization of the long-end yield response" to avoid potential confusion with level yield-curve flattening: in our model, an unexpected rate cut still produces a steepening of the level yield curve (long yields move less than the short rate), and the channel makes long yields move less still than the Vasicek benchmark would predict. The two terminologies refer to the same mathematical object, $\partial y(t,\tau)/\partial r_t = A(\tau)/\tau$, being pushed lower at intermediate τ.

3.5 Heterogeneity and the rebalancing mechanism

We now move from the aggregate equilibrium to the individual PH investor, where the cross-sectional predictions live.

Proposition 3 (Steady-state portfolio duration)**.** *For an investor with type (D_i^*, δ_i), the steady-state (pre-shock) optimal portfolio duration is the convex combination*

$$D_{port}^{i,old} \;=\; (1-\omega(\delta_i))\, D_{port}^{MV} + \omega(\delta_i)\, D_i^*, \qquad \omega(\delta_i) \in [0,1], \tag{13}$$

where D_{port}^{MV} is the common duration all PH investors would hold if every $\delta_i = 0$, and $\omega(\delta_i)$ is the habitat pull strength, rising from $\omega(0) = 0$ to $\omega \to 1$ as $\delta_i \to \infty$.

3.5.1 Response to a short-rate shock

Now consider an unexpected change in the short rate, Δr. Three quantities characterise investor i's response:

1. $D_{\text{port}}^{i,\text{old}}$: pre-shock optimal duration, equation (3).
2. $D_{\text{port}}^{i,\text{passive}}$: the no-trade duration—positions held fixed, prices repriced by Δr. To first order (Appendix J, eq. (78)),

   $$D_{\text{port}}^{i,\text{passive}} - D_{\text{port}}^{i,\text{old}} \;\approx\; -\Delta r \cdot \text{Cov}^{i,\text{old}}\big(A(\tau),\, \tau\big),$$

   where $\text{Cov}^{i,\text{old}}(A, \tau)$ is the maturity covariance under investor i's pre-shock portfolio weights, positive for any portfolio since A is increasing in τ.
3. $D_{\text{port}}^{i,\text{new}}$: the post-shock optimum from re-solving the FOC at the new state.

Define the model-implied *passive duration deviation* and *active duration rebalancing*:

$$\mathcal{D}ev_t^i \equiv D_{\text{port}}^{i,\text{passive}} - D_{\text{port}}^{i,\text{old}}, \tag{14}$$

$$\mathcal{R}eb_t^i \equiv D_{\text{port}}^{i,\text{new}} - D_{\text{port}}^{i,\text{passive}}. \tag{15}$$

Proposition 4 (Sign-opposite rebalancing)**.** *In the empirically relevant parameter region where the price-revaluation effect on long-bond positions dominates the equilibrium rise in expected excess returns, $\mathcal{D}ev_t^i$ and $\mathcal{R}eb_t^i$ have opposite signs: when prices push the portfolio's duration up, the investor rebalances down, and vice versa. The investor-specific rebalancing slope*

$$\beta_i^{reb} \equiv \frac{\mathcal{R}eb_t^i}{\mathcal{D}ev_t^i} \;\in\; \big(-1,\, \beta^{MV}\big],$$

with $\beta_i^{reb} \to -1$ as $\delta_i \to \infty$ (perfect rebalancing) and $\beta_i^{reb} = \beta^{MV} < 0$ at $\delta_i = 0$ (mean-variance baseline floor). β_i^{reb} is monotonically decreasing in δ_i.

Even a pure mean-variance investor rebalances partway back ($\beta^{\text{MV}} < 0$); the duration penalty pushes the slope toward -1. The magnitude of β_i^{reb} therefore measures anchoring intensity without taking a stand on its source. Figure 2 visualises this structure on the calibrated grid; the empirical AUM-weighted aggregates sit near $(\omega, \beta^{\text{reb}}) = (0.90, -0.71)$ in Treasury and $(0.90, -0.79)$ in Corporate.

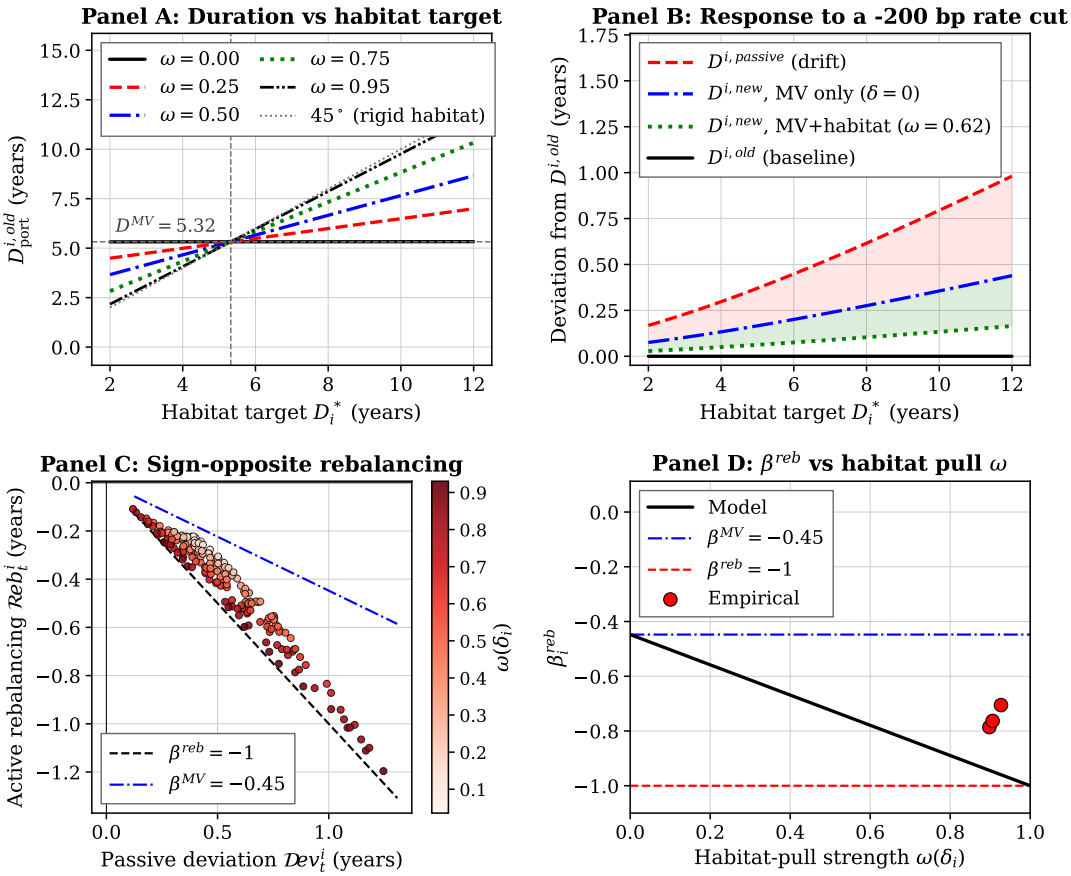


Figure 2: **Heterogeneous PH Duration Response to a Short-Rate Shock.** Numerical evaluation of equation (3) and Proposition 4 on the calibrated grid (Appendix I), varying the habitat-pull strength $\omega(\delta) \in [0, 1]$ ($\omega = 0$: MV-only; $\omega = 1$: rigid habitat); the empirical distribution of $\widehat{\omega}_f$ by institution group is in Figure 18. Panel (a): $D^{i,\text{old}}_{\text{port}}$ versus D_i^* across ω. Panel (b): rate-cut response ($\Delta r = -200$ bps) at $\omega = 0.625$, split into MV and habitat channels (using the approximation $\text{Cov}^{i,\text{old}}(A, \tau) \approx D^{i,\text{old}}_{\text{port}} A(D^{i,\text{old}}_{\text{port}})$). Panel (c): $\mathcal{R}eb_t^i$ against $\mathcal{D}ev_t^i$ for $n = 200$ heterogeneous funds, with reference lines at $\beta^{\text{reb}} = -1$ and an illustrative $\beta^{\text{MV}} = -0.45$. Panel (d): β^{reb} versus ω, with red markers at the empirical AUM-weighted aggregates $\approx (0.93, -0.71)$ Treasury and $(0.90, -0.79)$ Corporate.

3.6 Empirical hypotheses

The model yields four testable predictions.

- **H1 (persistent habitats).** Heterogeneous targets generate persistent dispersion in portfolio durations.
- **H2 (sign-opposite rebalancing).** Funds would partially rebalance back their price-induced passive duration drifts, with the slope of active rebalancing on passive deviation between -1 and 0.
- **H3 (cross-sectional pricing).** Bonds held by stronger rebalancers react less to monetary surprises.
- **H4 (maturity and transitory pressure).** The dampening concentrates at mid-to-long maturities and would show return reversal or spread closure, reflecting transitory demand rather than information.

4 Duration Habitats and Rebalancing

This section establishes the paper's two behavioral hypotheses, H1 and H2. Institutional bond investors hold persistent duration habitats, and they actively trade against the passive duration drift that price moves induce in their portfolios.

4.1 Persistent duration habitats

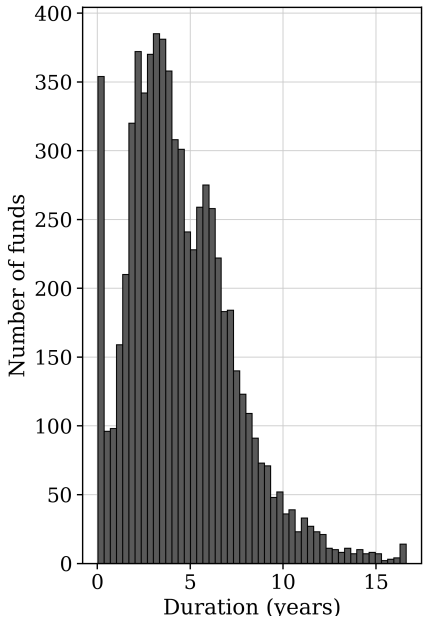
For Treasury and corporate bonds, I obtain Macaulay duration for each asset from CRSP and Mergent FISD. For each investor fund f in quarter t, portfolio duration is the value-weighted sum of asset-level durations:

$$Duration_{f,t} = \sum_{i \in f} w_{i,f,t}\, Duration_{i,t}, \tag{16}$$

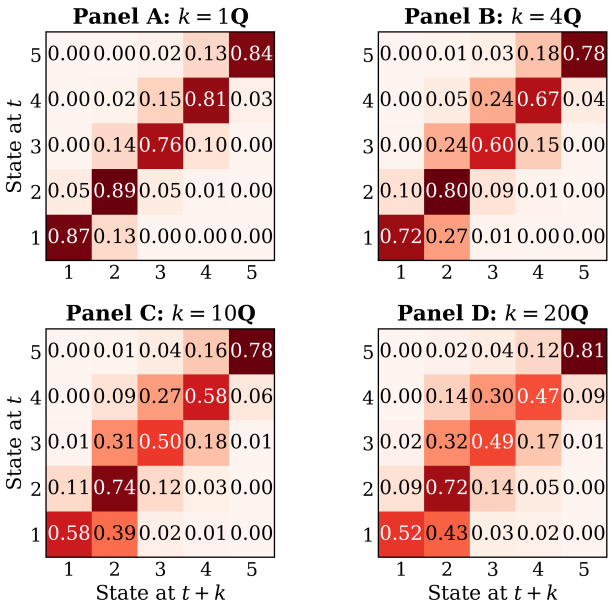
where i indexes the assets held in portfolio f and $w_{i,f,t}$ is the market-value weight.

I group portfolios into five duration buckets, $[0, 2]$, $(2, 5]$, $(5, 7]$, $(7, 12]$, $(12, \infty)$ years, and for each horizon $k \in \{1, 4, 10, 20\}$ quarters estimate the 5×5 transition matrix of bucket membership from t to $t + k$ by empirical frequencies.

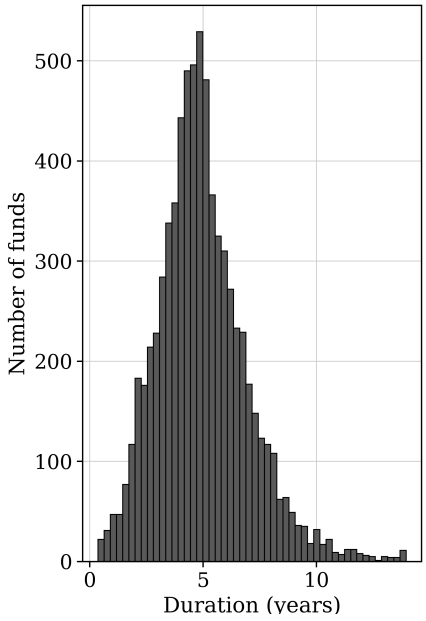
Figure 3 reports the pooled duration distribution and the transition matrices. The Treasury distribution is multi-modal—a spike inside $[0, 2]$ years, broad mass in $(2, 5]$ and $(5, 7]$, a thin right tail—while the Corporate distribution concentrates near 4–5 years. The transition matrices are strongly diagonal: 80–95% of institutions stay in their bucket after one quarter, and even at twenty quarters most mass remains at, or drifts back toward, each group's modal bucket rather than spreading to the extremes.



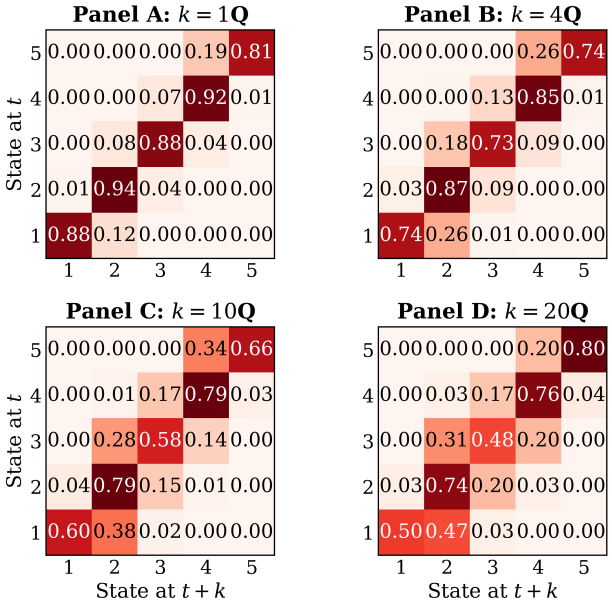
(a) Treasury: Fund Duration Distribution



(b) Treasury: Average Transition Matrices



(c) Corporate Bond: Fund Duration Distribution



(d) Corporate Bond: Average Transition Matrices

Figure 3: **Duration Distribution and State Transition Matrix: All Institutions Pooled.** The figure pools all institution types covered in the data. The top row shows results for the Treasury market and the bottom row for the corporate bond market. Panels (a) and (c) plot the distribution of average institution-level duration (Treasury and Corporate Bond, respectively); the x-axis is portfolio duration in years and the y-axis is the number of institution–quarter observations. Panels (b) and (d) plot the average state-probability transition matrix at horizons $k = 1, 4, 10, 20$ quarters; rows index the duration state at time t and columns index the state at $t+k$, with cell color encoding transition probability. Across both markets, the distributions display visible peaks corresponding to short-, mid-, and long-duration habitats, and the transition matrices are strongly diagonal: roughly 90% of institutions remain in the same duration bucket after one quarter, and even after five years ($k = 20$) most institutions still occupy their original habitat.

Stickiness holds for every institution group separately (Appendix A), though its form varies: mutual funds look like a cross-section of many habitats—multi-modal with persistent cross-fund dispersion—while insurers and pensions look like a single type-wide habitat concentrated at mid-durations. Both patterns are consistent with anchoring to a stable target D_i^*.

I also regress fund duration on its own k-quarter lag:

$$Dur_{i,t} = \alpha + \beta_k \, Dur_{i,t-k} + \varepsilon_{i,t}. \tag{17}$$

with standard errors clustered at the fund level. Table 3 reports β_k close to one at every horizon: one-year persistence is 0.96 in Treasury and 0.99 in Corporate, and five-year persistence remains 0.82 and 0.89. Duration habitats are stable investor attributes, as H1 predicts.

Table 3: **Persistence of Fund Duration.** This table reports the persistence of fund-level portfolio duration from pooled OLS regressions of $Dur_{i,t}$ on $Dur_{i,t-k}$ at horizons $k = 1, 4, 8, 16, 20$ quarters, as in equation (17), estimated separately for the Treasury and corporate-bond markets. Standard errors clustered at the fund level are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

	Dependent variable: $Duration_{i,t}$				
	$k = 1$	$k = 4$	$k = 8$	$k = 16$	$k = 20$
Panel A: Treasury market					
$Duration_{i,t-k}$	0.962^{***}	0.898^{***}	0.850^{***}	0.814^{***}	0.816^{***}
	(0.00228)	(0.00641)	(0.01080)	(0.01905)	(0.02416)
Observations	67,925	57,897	47,329	32,196	26,428
R^2	0.917	0.779	0.677	0.581	0.549
Panel B: Corporate bond market					
$Duration_{i,t-k}$	0.987^{***}	0.954^{***}	0.913^{***}	0.888^{***}	0.886^{***}
	(0.000689)	(0.00227)	(0.00446)	(0.00900)	(0.01076)
Observations	162,090	147,839	130,620	100,992	88,284
R^2	0.963	0.870	0.765	0.639	0.598

The persistence coefficients translate into an economically fast pull back to habitat. Estimating the same regression with time fixed effect, capturing deviations from each fund's time-average duration, we can converting the horizon-k coefficient into an implied AR(1) half-life, a duration deviation halves in roughly four quarters in Treasury and seven in Corporate, stable across estimation horizons and across institution groups (Appendix B).

4.2 Active Duration rebalancing

Persistence alone does not distinguish an actively maintained habitat from inertia: maturity shortening, differential price moves, and issuer turnover all drift portfolio duration around mechanically at quarterly frequency. The model predicts that habitat investors would partially offset such drifts. This section tests this prediction.

For each institutional account f and quarter t, I construct a counterfactual passive portfolio that holds the $t-1$ position fixed and lets the price and duration update freely. Define the passive weight of asset i in fund f as:

$$w_{i,f,t}^{pass} = \frac{w_{i,f,t-1}(1 + ret_{i,t})}{1 + R_{f,t}^{pass}}, \tag{18}$$

where $ret_{i,t}$ is the asset's return and $R_{f,t}^{pass}$ is the hypothetical passive portfolio return assuming no trading:

$$R_{f,t}^{pass} \equiv \sum_{i \in f} w_{i,f,t-1}\, ret_{i,t}. \tag{19}$$

The passive duration of account f is then

$$Dur_{f,t}^{pass} = \sum_{i \in f} w_{i,f,t}^{pass} \cdot Dur_{i,t}, \tag{20}$$

while the realised duration $Dur_{f,t} = \sum_{i \in f} w_{i,f,t}\, Dur_{i,t}$ is observed directly from the holdings panel. From these I form two signals, the passive duration deviation accumulated since $t-1$ and the active rebalancing within t:

$$Dur_{f,t}^{\text{dev}} = Dur_{f,t}^{\text{pass}} - Dur_{f,t-1} \tag{21}$$

$$Dur_{f,t}^{\text{reb}} = Dur_{f,t} - Dur_{f,t}^{pass} \tag{22}$$

If institutions actively offset the mechanical drift toward their habitat target, $Dur_{f,t-l}^{\text{dev}}$ and $Dur_{f,t}^{\text{reb}}$ should take opposite signs.

I regress active rebalancing on the current and two lagged passive deviations, so that the summed coefficients measure the fraction of passive drift offset within three quarters:

$$Dur_{f,t}^{\text{reb}} = \alpha + \sum_{l=0}^{2} \beta_l\, Dur_{f,t-l}^{\text{dev}} + \sum_i \gamma_i Control_{i,f,t} + \lambda_t + \varepsilon_{f,t} \tag{23}$$

Controls comprise current and lagged log assets under management ($\log TNA_{f,t}$, $\log TNA_{f,t-1}$). The main specification absorbs quarter fixed effects λ_t, with standard

errors clustered at the institution level. As a robustness check I report a two-way fixed-effects specification that adds account fixed effects μ_f.[8] Per-sector regressions are run separately within each institution type under the same specification.

Table 4 reports the estimates. In both markets, current and lagged passive deviations significantly negatively predict active rebalancing. Summing the three lags, Treasury institutions offset -0.81 of the passive drift within three quarters and corporate institutions -0.54; the two-way FE column leaves the totals essentially unchanged (-0.81 and -0.50). The significance holds across institution types: $\hat{\beta}_0$ is negative and significant for every sizeable group in both markets, and the cumulative offsets are economically substantial, ranging from about -0.4 to -1.4 across groups and markets.

4.3 Fund-Level Rebalancing Propensity

To bring the channel to bond pricing I construct a fund-level rebalancing propensity. I construct it in two stages: an expanding-window OLS on each fund's history of (passive deviation, active rebalancing) pairs, and a Bayesian shrinkage of the fund-quarter OLS estimate toward a within-institution-type prior.

Stage 1: Expanding-window OLS. For each fund f and quarter t, I estimate the slope of active rebalancing on passive deviation by OLS with no intercept:

$$Dur^{\text{reb}}_{f,\tau} = \beta^{\text{reb}}_{f,t}\, Dur^{\text{dev}}_{f,\tau} + \varepsilon_{f,\tau}, \qquad \tau \in \mathcal{T}_{f,t}. \tag{24}$$

where $\mathcal{T}_{f,t} = \{\tau \leq t-1\}$ is the history of fund-quarters. Intercepts are set to zero, because the level-offset relationship itself is the habitat-anchoring signal (Section 4.2), and an intercept would absorb it. I require at least 3 historical observations, and accounts with fewer receive the type prior directly in Stage 2.

Stage 2: Within-type prior and Bayesian posterior. Within each estimation quarter t, I form the prior empirically from the cross-section of fund-level OLS estimates. I compute the within-institution-type mean and variance from estimates of funds with

[8] Fund fixed effects would demean each fund's deviation and rebalancing by its own averages, so the slope would be identified only from quarters in which a fund's drift is unusually large relative to its own norm. But the fund-level average pairing is itself evidence of habitat defence: a fund with a long-maturity target, for example, sees its portfolio duration shorten mechanically every quarter as its bonds age, and must buy duration every quarter to stay put—a persistently negative deviation matched with persistently positive rebalancing. The main specification keeps this component; the two-way FE column shows the estimates barely move when it is removed.

Table 4: **Sector-Level Duration Rebalancing.** The table reports estimates of equation (23), regressing active duration rebalancing on the current and two lagged passive duration deviations. Each row is a separate regression. The two pooled rows estimate over all institution types within the market; the main spec absorbs quarter fixed effects λ_t and the two-way FE row adds account fixed effects μ_f. Each by-type row estimates (23) within a single institution group under the main spec. Controls comprise $\log TNA_{f,t}$ and $\log TNA_{f,t-1}$. The $\sum_l \widehat{\beta}_l$ column reports the sum of the three coefficients, the cumulative fraction of passive drift offset within three quarters (— where the lag coefficients are not estimated). Standard errors are clustered at the institution level (omitted for compactness). ***, **, * denote significance at the 1%, 5%, and 10% levels.

	$\widehat{\beta}_0$	$\widehat{\beta}_{l1}$	$\widehat{\beta}_{l2}$	$\sum_l \widehat{\beta}_l$	Obs.	Controls
Panel A: Treasury market						
Pooled (Time FE)	-0.233^{***}	-0.371^{***}	-0.211^{***}	-0.815	73,558	Yes
Pooled (Two-way FE)	-0.244^{***}	-0.371^{***}	-0.195^{***}	-0.810	73,558	Yes
By institution type (Time FE):						
Taxable Bond MFs	-0.350^{***}	-0.394^{***}	-0.193^{***}	-0.937	18,300	Yes
Allocation MFs	-0.580^{***}	-0.465^{***}	-0.313^{*}	-1.358	2,582	Yes
Other MFs	-0.372^{***}	-0.363^{***}	-0.216^{*}	-0.951	6,385	Yes
Insurance: Life	-0.360^{***}	-0.393^{***}	-0.330^{***}	-1.083	9,370	Yes
Insurance: P&C	-0.277^{***}	-0.436^{***}	-0.436^{***}	-1.149	30,016	Yes
Insurance: Other	-0.258^{**}	-0.534^{***}	-0.125	-0.917	3,169	Yes
Variable Annuity	-0.286	—	—	—	3,548	Yes
Closed-End MF	-0.804	—	—	—	168	Yes
Panel B: Corporate bond market						
Pooled (Time FE)	-0.430^{***}	-0.039^{***}	-0.076^{***}	-0.545	169,703	Yes
Pooled (Two-way FE)	-0.418^{***}	-0.026^{***}	-0.060^{***}	-0.504	169,703	Yes
By institution type (Time FE):						
Taxable Bond MFs	-0.504^{***}	-0.042	-0.073^{***}	-0.619	42,011	Yes
Allocation MFs	-0.647^{***}	-0.062	-0.067	-0.776	3,717	Yes
Other MFs	-0.374^{***}	0.022	-0.055^{**}	-0.407	9,954	Yes
Insurance: Life	-0.342^{***}	-0.096^{***}	-0.084^{***}	-0.522	42,729	Yes
Insurance: P&C	-0.577^{***}	-0.064^{***}	-0.077^{***}	-0.718	58,052	Yes
Insurance: Other	-0.485^{***}	-0.023	-0.084^{*}	-0.592	6,089	Yes
Variable Annuity	-0.388^{***}	—	—	—	5,689	Yes
Closed-End MF	-0.750^{***}	—	—	—	1,441	Yes

at least 10 historical observations:

$$\bar{\beta}_{g,t} = \text{mean}\big(\hat{\beta}^{\text{ols}}_{f,t} \ : \ f \in g,\ n_{f,t} \geq 10 \big), \tag{25}$$

$$\hat{\sigma}^2_{g,t} = \text{var}\big(\hat{\beta}^{\text{ols}}_{f,t} \ : \ f \in g,\ n_{f,t} \geq 10 \big), \tag{26}$$

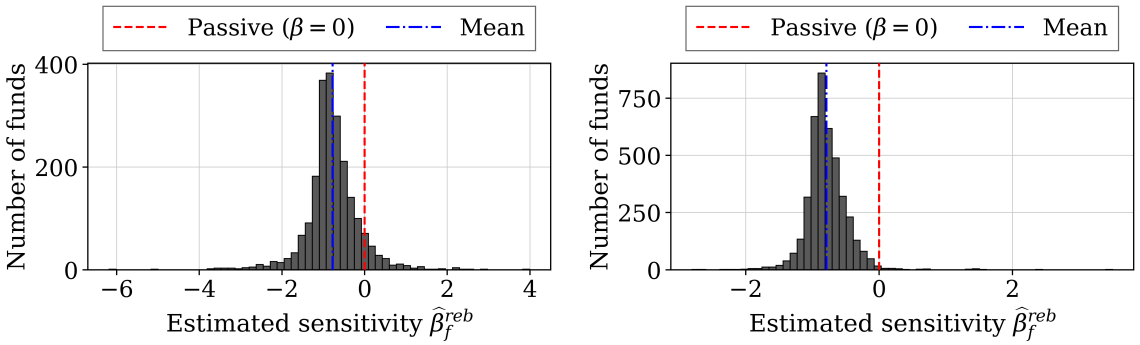
where g indexes the institution types (twelve groups). The prior is type-specific: it pulls each account toward funds with the same account type, preserving the cross-type heterogeneity. With $\hat{\beta}^{\text{ols}}_{f,t} \sim \mathcal{N}(\beta^{\text{reb}}_{f,t}, \hat{se}^2_{f,t})$ and $\beta^{\text{reb}}_{f,t} \sim \mathcal{N}(\bar{\beta}_{g(f),t}, \hat{\sigma}^2_{g(f),t})$, the

posterior mean is:

$$\widehat{\beta}^{\text{reb}}_{f,t} = \frac{d_{f,t}\,\hat{\beta}^{\text{ols}}_{f,t} + p_{g,t}\,\bar{\beta}_{g(f),t}}{d_{f,t} + p_{g,t}}, \qquad d_{f,t} = \frac{1}{\hat{se}^2_{f,t}}, \quad p_{g,t} = \frac{1}{\hat{\sigma}^2_{g(f),t}}. \tag{27}$$

The shrinkage weight is endogenous: long-history accounts keep their estimates, while sparse-history accounts are pulled to their type mean prior. The prior uses only pre-t data, so the panel of $\widehat{\beta}^{\text{reb}}_{f,t}$ is free of look-ahead bias.

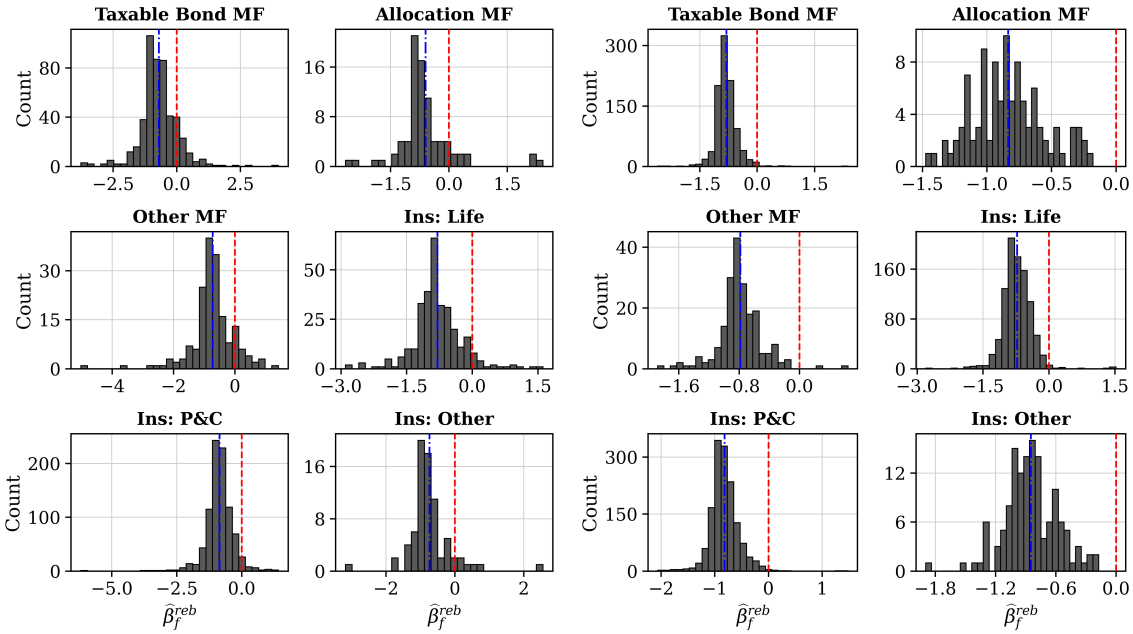
Figure 4 reports the cross-sectional distribution of $\widehat{\beta}^{\text{reb}}_{f,t}$ over all fund-quarters. In both markets the distribution is centered well below zero, with a clear mode near -1, the full-rebalancing point at which active rebalancing exactly offsets the passive duration drift on average—the bulk of the mass sits in the $[-1, 0)$ range that H2 predicts. A small minority of accounts have positive sensitivities; these are accounts that, on average, entend rather than offset the passive drift. Figure 5 further shows the distribution by institution group. The negative-mass concentration is broad-based: every group's modal mass sits at or below zero, with within-group medians clustering roughly between -0.5 and -1 in both markets. Some groups produce visibly bell-shaped distributions while others do not, but the location of the mass below zero is consistent across types.



(a) Treasury (b) Corporate Bond

Figure 4: **Rebalancing Sensitivity Distribution.** Distribution of posterior fund-level rebalancing sensitivities $\widehat{\beta}^{\text{reb}}_{f,t}$ from (27), pooled across all fund-quarters. Left: Treasury market. Right: corporate bond market. The vertical dashed red line marks $\beta = 0$ (passive); the blue dash-dotted line marks the cross-sectional mean. Most accounts have negative sensitivities, with a mode near -1 corresponding to full rebalancing.

I additionally rank accounts within each quarter by $\widehat{\beta}^{\text{reb}}_{f,t}$ into ascending quintiles, from strong rebalancers (G1) to weak (G5). Table 5 shows the quintile means ascend monotonically in both markets and every major institution type populates all five quintiles. This indicates the heterogeneity in rebalancing preference is prevailing



(a) Treasury market

(b) Corporate bond market

Figure 5: **Posterior Rebalancing Sensitivity Distribution by Institution Group.** Histograms of $\widehat{\beta}_{f,t}^{\text{reb}}$ pooled across fund-quarters, separately for each institution group in each market. Vertical dashed red line marks $\beta = 0$; blue dash-dotted line marks the within-group mean. Group β distributions are concentrated at or below zero across all groups, with cross-type heterogeneity in both location and dispersion.

across investors. Comparatively, the weakest Treasury quintile averages only -0.12—a genuine non-rebalancer tail—while the weakest Corporate quintile still averages -0.53, so the Corporate market contains only rebalancers of different strengths.

5 Monetary Shock Stabilization around FOMC Shocks

Section 4 establishes that bond institutions hold duration habitats and actively defend them. This section asks whether the rebalancing leaves a footprint in prices, using FOMC announcements as the laboratory. FOMC events suit the test for two reasons: the high-frequency surprise is a cleanly identified shock (Nakamura and Steinsson, 2018; Bauer and Swanson, 2023), and it mechanically forces every multi-bond portfolio off its target duration at the same moment, opening a fresh deviation $Dur_{f,t}^{\text{dev}}$ at every fund.

Funds with stronger rebalancing tendencies close that gap more aggressively in the days that follow, so bonds owned by such funds receive more of the offsetting demand pressure. Holding bond characteristics fixed, the cross-sectional differential between

Table 5: **Fund-Quarter Distribution by Institution Type and Sensitivity Quintile.** Cells in the body report counts of fund-quarter observations within each institution type × sensitivity-quintile cell. Rows are institution types; columns are ascending sensitivity quintiles, from G1 (most negative posterior $\widehat{\beta}^{\text{reb}}_{f,t}$, strong rebalancers) to G5 (least negative, weak rebalancers). The rightmost column reports each institution type's fund-quarter total and pooled mean $\widehat{\beta}^{\text{reb}}$. The bottom row reports the pooled mean $\widehat{\beta}^{\text{reb}}$ within each quintile (across all institution types). Pensions have the smallest fund-quarter counts and noisier within-cell statistics.

	G1 (strong)	G2	G3	G4	G5 (weak)	Total	Mean $\widehat{\beta}^{\text{reb}}$
Panel A: Treasury market							
Taxable Bond MF	8,206	6,190	6,031	7,771	8,400	36,598	−0.71
Allocation MF	1,463	844	1,182	1,544	1,641	6,674	−0.70
Other MF	1,581	1,807	1,811	2,220	2,383	9,802	−0.66
Insurance: Life	4,715	6,088	7,215	6,735	5,158	29,911	−0.70
Insurance: P&C	18,236	19,522	17,539	16,296	16,024	87,617	−0.72
Insurance: Other	922	1,410	1,967	1,002	966	6,267	−0.74
Pension	0	78	0	0	0	78	−0.91
Mean $\widehat{\beta}^{\text{reb}}$	−1.21	−0.92	−0.76	−0.55	−0.12	—	—
Panel B: Corporate bond market							
Taxable Bond MF	11,342	14,925	16,115	14,867	8,941	66,190	−0.83
Allocation MF	2,178	1,410	1,522	1,383	1,122	7,615	−0.87
Other MF	1,977	3,000	3,520	2,939	1,202	12,638	−0.85
Insurance: Life	12,897	10,767	10,996	11,996	18,847	65,503	−0.79
Insurance: P&C	19,038	19,604	17,884	19,664	19,335	95,525	−0.82
Insurance: Other	3,086	2,432	1,894	1,207	1,336	9,955	−0.88
Pension	0	0	72	0	4	76	−0.82
Mean $\widehat{\beta}^{\text{reb}}$	−1.06	−0.92	−0.84	−0.75	−0.53	—	—

bonds held by strong versus weak rebalancers is the signature of the stabilising channel; the model predicts the differential is positive (H3), concentrates at longer maturities, and should show return reversals (H4).

5.1 Rebalancing strength and bond-return reactions

I aggregate fund-level posterior sensitivities $\widehat{\beta}^{\text{reb}}_{f,t}$ from Section 4.3 into a continuous, ownership-weighted bond-level measure, *rebalancing strength*. For bond i on quarter t,

$$\text{rebal_str}_{i,t} \;=\; -\sum_{f} \widehat{\beta}^{\text{reb}}_{f,t}\, \frac{H_{i,f,t}}{S_{i,t}}, \tag{28}$$

where the sum is over every fund f holding bond i at the most recent quarter-end before t, $H_{i,f,t}$ is fund f's face-value holding on bond i, and $S_{i,t}$ is the bond i's amount outstanding, so $H_{i,f,t}/S_{i,t}$ is the ownership share. The minus sign maps stronger rebalancers (more negative $\widehat{\beta}^{\text{reb}}$) to larger $\text{rebal_str}_{i,t}$. The measure has

a Bartik shift-share structure (Borusyak, Hull, and Jaravel, 2022): predetermined ownership shares are the exposures, fund-level sensitivities the shifts.

A natural worry is selection: funds with strong rebalancing habits might simply hold bonds that would react less to FOMC shocks anyway. Three properties of the construction address it. First, the shifts are predetermined: $\widehat{\beta}^{\text{reb}}_{f,t}$ is estimated from each fund's pre-t history alone, so it cannot reflect the post-announcement returns it is later asked to explain. Second, the shifts are fund-level traits, not institution-type labels: every major type spans all five sensitivity quintiles (Table 5), so $\text{rebal_str}_{i,t}$ does not proxy for type-level bond preferences such as life insurers tilting long or P&C insurers tilting short. Third, the measure captures who holds a bond, not how much of it is institutionally held; the ownership level is controlled directly in the pricing tests. The maintained assumption is then that, conditional on observable bond and issuer characteristics, rebalancing strength is unrelated to other determinants of FOMC-window returns.

Implied from the theory, bonds who are more exposed to rebalancing force should show dampened price reaction to monetary surprises. Therefore, I design the following regression specification, with bond and event-date fixed effects:

$$\begin{aligned} Ret_{i,[t,t+k]} = \alpha + \beta_0\, FOMC_t^{\text{Shock}} + \beta_1\, FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t} + \beta_3\, \text{rebal_str}_{i,t} \\ + \mathbf{X}'_{i,t}\delta + \gamma_i + \theta_t + \varepsilon_{i,t}, \end{aligned} \tag{29}$$

where γ_i absorbs time-invariant bond characteristics and θ_t absorbs the aggregate FOMC effect, so the level $\beta_0\, FOMC_t^{\text{Shock}}$ would be absorbed by θ_t and not appear separately. The control vector $\mathbf{X}_{i,t}$ contains bond and issuer-level characteristics together with their interactions with $FOMC_t^{\text{Shock}}$, listed in Table 6. For Treasuries, these are the standard bond characteristics; for corporates, I add variables about credit risks and issuers' financial signals (Bao, Pan, and Wang, 2011; Dickerson, Mueller, and Robotti, 2023; Fang, 2025). Total institutional ownership $\text{own}_{i,t}$ and its shock interaction enter the control set as well. Continuous variables are winsorised at the 1% tails.

Without time fixed effect, β_0 is supposed to be negative, so the dampened price reaction would predict a positive β_1. β_0 would be absorbed when including time fixed effects. However, β_1 would still be hypothesized to be positive.

The regression includes FOMC events from January 2010 onward, the point at which coverage of the rebalancing-strength signal becomes stable in both markets (Appendix H). For Treasuries I exclude T-bills, TIPS, and callable issues, whose

Table 6: **Bond and Issuer Controls in the FOMC-Window Regression.** The table lists the controls that enter the regression in equation (29) for each market. All variables enter both as a level and as their interaction with $FOMC_t^{\text{Shock}}$. Continuous variables are winsorised at the 1% tails. The level $FOMC_t^{\text{Shock}}$ itself is absorbed by event-date fixed effects.

Category	Treasury	Corporate
Bond rate sensitivity	duration; convexity; yield; coupon	duration; log convexity; yield spread; coupon
Bond liquidity / structure	on-the-run; bid–ask; log issue size	log Amihud illiquidity; log issue size
Issuer financial state	—	log assets; leverage (debt / assets); cash / assets
Credit risk	—	rating dummies AA, A, BBB (AAA = base)
Total institutional ownership	$\text{own}_{i,t}$	$\text{own}_{i,t}$
Fixed effects	bond, event-date	bond, event-date
Standard errors	bond-clustered	bond-clustered

prices also load on safety demand, expected inflation, and negative convexity—separate stories from the channel studied here. For corporates, because TRACE prices of thinly traded bonds are dominated by quote noise (Bao, Pan, and Wang, 2011; Dickerson, Mueller, and Robotti, 2023), I keep investment-grade bonds with price $\geq$ \$5, yield-to-maturity < 1, monthly volume $\geq$ \$50M, face value $\geq$ \$250M, and a zero-trading-days ratio ≤ 0.5: large, liquid IG corporates whose post-FOMC reactions are reliably observable. The main specification uses the no-crisis sample, dropping the COVID liquidity event (February–May 2020) and the 2022 rate-hike cycle—14 months out of 15 years, leaving about 92% of post-2010 FOMC observations. Both windows featured fire sales and liquidity disruptions that overwhelm arbitrageurs' risk-bearing capacity and would confound the demand channel of interest; restricting to normal times is standard in this literature (e.g., Li, 2025, removes the Global Financial Crisis and March 2020 on the same grounds).[9]

Table 7 reports main results across different specifications. Spec 1 is the pooled-OLS version with only the standard duration and ownership controls and without

[9]The excluded windows are not simply dropped: Appendix D reports estimates on the crisis-inclusive samples (Appendix Table 15). The corporate estimate stays positive and significant even within the eight stress-window FOMC dates (0.68 to 1.22 across horizons), while the Treasury interaction flips sign there (-1.6 to -2.4): Treasuries are the market's safe and collateral asset, and in episodes of dealer balance-sheet strain the fire-sale and dash-for-cash forces documented by Kekre, Lenel, and Mainardi (2024), Fang (2025), and Siriwardane, Sunderam, and Venter (2025) dominate rebalancing demand. A continuous version makes the same point without a sample split: interacting the shock and $\text{rebal_str}_{i,t}$ with standardised pre-announcement ΔVIX volatility on the full panel gives -0.320 ($t = -2.48$) in Treasury and -0.064 ($t = -2.36$) in IG corporate: the dampening is near its no-crisis value in calm conditions and declines as stress rises. The channel is a normal-times stabilizer; in crises, Treasuries' safe-asset role takes over.

fixed effects; Spec 2 adds the remaining bond and issuer controls; Spec 3 is the headline equation (29) with bond and event-date fixed effects.

Table 7: **Rebalancing-Strength Effect on Bond-Return Reactions to FOMC Shocks: Specification Progression.** Estimates of $\widehat{\beta}_1$ on $FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t}$ in equation (29) and its pooled-OLS variants, on the no-crisis sample (post-2010 FOMC events, excluding the COVID liquidity event of February–May 2020 and the 2022 rate-hike cycle; Appendix D analyses those windows). *Spec 1* is the pooled OLS baseline (no fixed effects, level $FOMC_t^{\text{Shock}}$ included) with only the standard duration and ownership controls; *Spec 2* adds the full bond and (for corporates) issuer controls of Table 6, still pooled OLS; *Spec 3* adds bond and event-date fixed effects (equation (29), the headline). Bond-clustered standard errors in parentheses. *, **, *** denote significance at the 10%, 5%, 1% levels. R^2 is the within-R^2; the fixed effects in Spec 3 can drive it negative, a standard PanelOLS quirk that does not affect cluster-robust inference.

	Spec	$\widehat{\beta}_1$ on $FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t}$	Bond controls	Fixed Effects	R^2	N
Panel A: Treasury market						
$Ret_{t,t+2}$	1	2.021*** (0.549)	—	—	0.123	29,581
	2	2.018*** (0.496)	Yes	—	0.181	29,581
	3	1.076*** (0.334)	Yes	Bond & event-date	−0.571	29,581
$Ret_{t,t+4}$	1	2.029*** (0.675)	—	—	0.069	29,581
	2	2.091*** (0.647)	Yes	—	0.102	29,581
	3	1.603*** (0.507)	Yes	Bond & event-date	−0.066	29,581
$Ret_{t,t+6}$	1	1.976*** (0.695)	—	—	0.070	29,581
	2	2.077*** (0.672)	Yes	—	0.104	29,581
	3	1.825*** (0.529)	Yes	Bond & event-date	0.005	29,581
Panel B: Corporate bond market						
$Ret_{t,t+2}$	1	1.598*** (0.125)	—	—	0.098	104,583
	2	1.913*** (0.127)	Yes	—	0.117	103,983
	3	0.732*** (0.145)	Yes	Bond & event-date	−0.150	103,983
$Ret_{t,t+4}$	1	1.461*** (0.142)	—	—	0.083	104,583
	2	1.717*** (0.144)	Yes	—	0.098	103,983
	3	0.527*** (0.160)	Yes	Bond & event-date	−0.050	103,983
$Ret_{t,t+6}$	1	2.233*** (0.165)	—	—	0.076	104,583
	2	2.551*** (0.169)	Yes	—	0.090	103,983
	3	0.567*** (0.175)	Yes	Bond & event-date	0.054	103,983

Stricter pooled-OLS specifications without the duration-channel control are reported in Appendix Table 16 as a transparency check. Crisis-inclusive samples are analysed in Appendix D.

As is shown in table 7, $\widehat{\beta}_1$ is significantly positive at the 1% level across all three specifications and both markets.[10] The regression conclusion is essentially unchanged when the full shock-interacted controls are added (Spec 2) and survives bond and event-date fixed effects (Spec 3), with the effect building increasingly over the days after the announcement for both Treasuries and corporate bonds.

To further illustrate the timing of this dampening effect, I conduct an event-study analysis: I estimate equation (29) with cumulative-return windows $Ret_{i,[t-k,t-1]}$ and $Ret_{i,[t,t+k]}$ for each $k \in \{1, 2, \ldots, 10\}$, and Figure 6 plots the resulting $\widehat{\beta}_{1,k}$ profile relative to event day 0 for both markets. There shows no pre-trend: Pre-announcement

[10]Two further checks: selection on unobservables runs the wrong way (adding the shock-interacted controls *raises* $\widehat{\beta}_1$, and the Oster (2019) identified set excludes zero in both markets), and Adão–Kolesár–Morales standard errors at the fund level (Borusyak, Hull, and Jaravel, 2022) leave the headline t-statistics at 2.5 in Treasury and 3.0 in IG corporate.

coefficients are flat and indistinguishable from zero in both markets (all $|t| < 1.7$), and the dampening effect starts to be significant on only from the shock. From day 0 onward, $\widehat{\beta}_{1,k}$ rises and accummulates, stabilzing at around 1.8 in Treasuries by $k = 5$–6 and builds to 1.1–1.4 in corporates by the second week. This documents that rebalancing flows accumulating gradually rather than executing instantaneously. Preferred Habitat investors are facing many portfolio adjustment costs and it could take several days to update investment decisions. Therefore, it should be expected that the rebalancing flows are gradually executed over several days, which is consistent with the intertia flow of institutions documented in the literature (e.g., Vayanos and Woolley, 2013; Dickerson, Mueller, and Robotti, 2023; Fang, 2025; Chaudhary, Fu and Li, 2026).

Heterogeneity by maturity. The rebalancing channel should concentrate where arbitrage capacity is most strained, say, the mid to long end of each market. To test this, I augment Spec 3 of equation (29) with a triple interaction between the FOMC shock, rebalancing strength, and bond duration measured as a deviation from its sample mean:

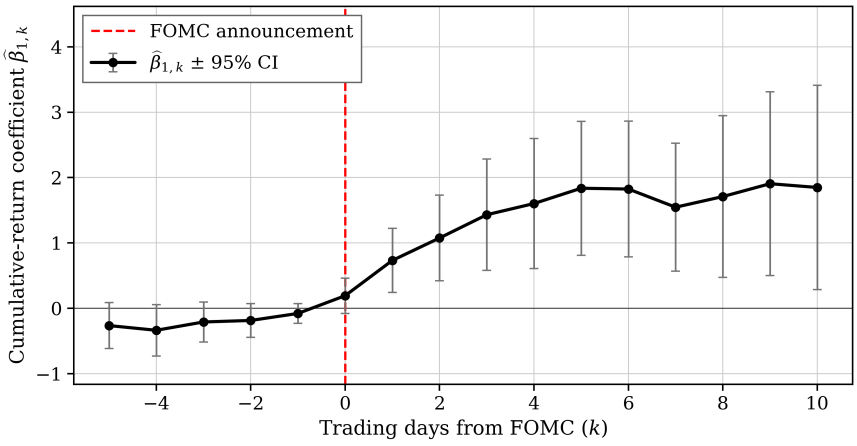
$$\begin{aligned} Ret_{i,[t,t+k]} = \alpha_i + \delta_t + \beta_1 \, FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t} + \beta_2 \, FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t} \times (D_{i,t} - \bar{D}) \\ + \beta_3 \, \text{rebal_str}_{i,t} + \beta_4 \, \text{rebal_str}_{i,t} \times (D_{i,t} - \bar{D}) + \mathbf{X}'_{i,t}\delta + \varepsilon_{i,t}, \end{aligned} \tag{30}$$

where $D_{i,t}$ is bond i's modified duration at $t-1$, $\bar{D}$ is the sample-mean duration on the no-crisis panel. $\widehat{\beta}_1$ measures the dampening at the average duration; $\widehat{\beta}_2$ measures how the dampening grows per additional year of bond duration. A positive $\widehat{\beta}_2$ delivers the long-end concentration prediction without any binning choice.

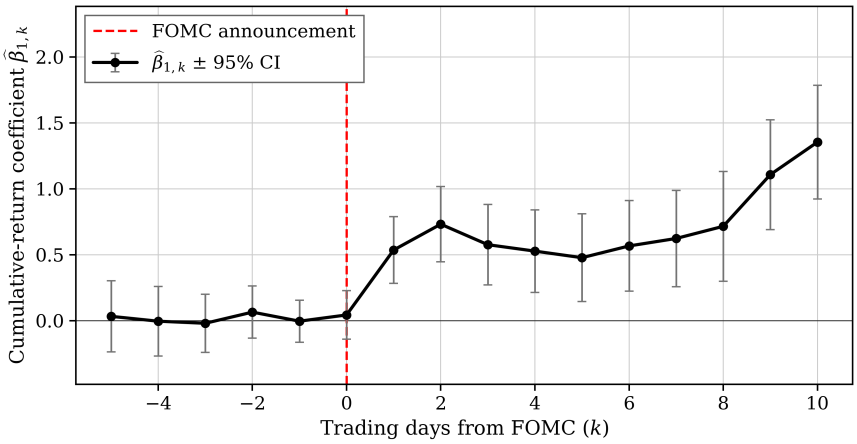
$\widehat{\beta}_2$ is positive and significant at every horizon in both markets. In the meantime, the dampening increases with bond duration. To read the magnitude at a given maturity, the table also reports the implied dampening $\widehat{\beta}_1 + \widehat{\beta}_2 \, (D - \bar{D})$ at $D \in \{3, 7, 10, 15\}$ years, with standard errors from the delta method. Every implied coefficient is positive and significant at the 1% level, and the effect at $D = 15$ years is roughly twice that at $D = 3$ years. The maturity profile would also be the empirical counterpart that Section 6.2 confronts with the calibrated model's dampening, where it shows the two agree with each other.

5.2 Long-short portfolio and the round trip

If the dampening reflects transient rebalancing pressure rather than information, the return spread between high- and low-rebal_str bonds should reverse or even close



(a) Treasury



(b) Corporate Bond

Figure 6: **Event Study of $FOMC^{\text{Shock}}$ × rebal_str, No-Crisis Sample.** Each point plots $\widehat{\beta}_{1,k}$ from equation (29) on cumulative returns $Ret_{i,[t,t+k]}$ for $k \geq 0$ and $Ret_{i,[t-k,t-1]}$ for $k < 0$. Vertical bars are 95% confidence intervals from bond-clustered standard errors. The dashed vertical line marks the FOMC announcement date.

before the next meeting as the pressure fades and arbitrage capital moves in. I trade on this spread directly. On each FOMC date I sort bonds into quintiles by $\text{rebal_str}_{i,t}$, measured from holdings published before the announcement,[11] and form a value-weighted long-short conditional on the shock: on a cut I am long the lowest-rebal_str quintile and short the highest, whose rise is dampened; on a hike the legs flip. Either way the position collects the dampening differential as it opens.

[11]The signal at each formation date uses only holdings whose report date precedes it by at least 90 calendar days, matching the publication lags of mutual-fund N-PORT filings (60 days) and statutory insurance filings compiled by eMAXX (45–90 days). Every FOMC observation in the sample therefore trades on publicly available holdings. Results under contemporaneous quarter-end signals are essentially identical: alphas move by less than half a percentage point, with no sign changes in any headline cell.

Table 8: **Heterogeneity of the Rebalancing-Strength Effect by Duration.** Estimates of $\widehat{\beta}_1$ on $FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t}$ and $\widehat{\beta}_2$ on $FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t} \times (D_{i,t} - \bar{D})$ from the augmented Spec 3 of equation (30). Bond and event-date fixed effects, full bond and issuer controls of Table 6 with their FOMC-shock interactions, bond-clustered standard errors. No-crisis sample. $\widehat{\beta}_1$ is the dampening at the average duration $\bar{D}$; $\widehat{\beta}_2$ is the gradient per additional year. Implied effects $\widehat{\beta}_1 + \widehat{\beta}_2(D - \bar{D})$ at $D \in \{3, 7, 10, 15\}$ reported below, with delta-method significance. ***, **, * denote significance at the 1%, 5%, 10% levels.

	Treasury market			Corporate bond market		
	$k = 2$	$k = 4$	$k = 6$	$k = 2$	$k = 4$	$k = 6$
Headline coefficients						
$\widehat{\beta}_1$ (at $D = \bar{D}$)	0.877^{***}	1.312^{***}	1.552^{***}	1.023^{***}	0.738^{***}	0.788^{***}
	(0.298)	(0.452)	(0.475)	(0.148)	(0.160)	(0.178)
$\widehat{\beta}_2$ (per +1y)	0.052^{**}	0.075^{**}	0.070^{**}	0.042^{***}	0.030^{***}	0.032^{***}
	(0.021)	(0.029)	(0.030)	(0.006)	(0.006)	(0.008)
Implied dampening $\widehat{\beta}_1 + \widehat{\beta}_2(D - \bar{D})$ *at:*						
$D = 3$ y	0.764^{***}	1.149^{***}	1.399^{***}	0.843^{***}	0.608^{***}	0.652^{***}
$D = 7$ y	0.970^{***}	1.447^{***}	1.680^{***}	1.011^{***}	0.729^{***}	0.779^{***}
$D = 10$ y	1.125^{***}	1.671^{***}	1.890^{***}	1.138^{***}	0.820^{***}	0.875^{***}
$D = 15$ y	1.383^{***}	2.044^{***}	2.241^{***}	1.348^{***}	0.971^{***}	1.034^{***}
$\bar{D}$ (sample mean)	5.18	5.18	5.18	7.28	7.28	7.28
Observations	29,581	29,581	29,581	103,983	103,983	103,983
Bond & event-date FE	Yes	Yes	Yes	Yes	Yes	Yes
Full controls of Table 6	Yes	Yes	Yes	Yes	Yes	Yes

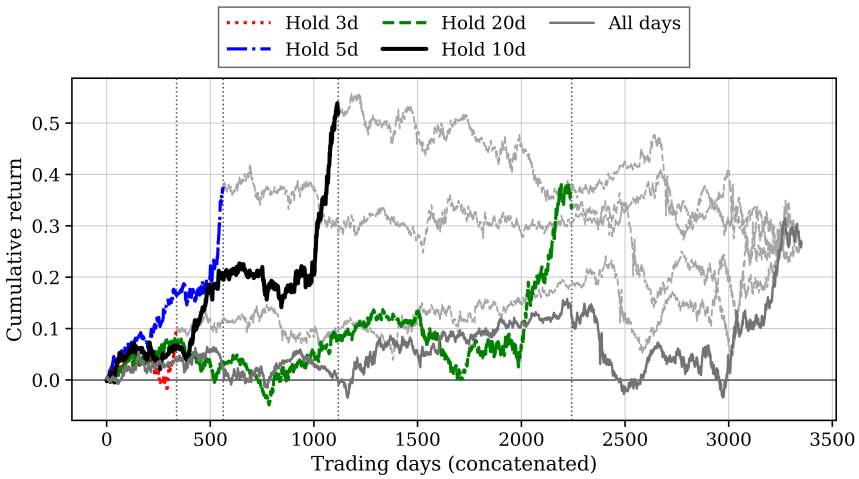
The position starts at $t + 1$, once the shock is public, and is deployed two ways: hold-and-close holds for $k \in \{3, 5, 10, 20\}$ trading days and harvests only the opening of the spread; hold-and-reverse flips the position at day k and holds to the next FOMC, additionally harvesting the close.[12] For Treasuries, I work with raw returns; for corporate bonds, I work with bond-level duration-adjusted returns, the standard treatment in the corporate-bond literature (Nozawa, 2017; Dickerson, Mueller, and Robotti, 2023; Kelly, Palhares, and Pruitt, 2023; Dickerson, Julliard, and Mueller, 2026).[13]

Figure 7 shows the round trip directly: returns accrue through the first one-to-two weeks after each FOMC (coloured segments) and then revert toward zero over the rest of the inter-meeting period (light grey). Hold-and-reverse, which rides the second leg, earns most at $k = 5$ and $k = 10$, the horizons that flip closest to the spread's peak. Cumulative returns by holding horizon for both variants are in Appendix F

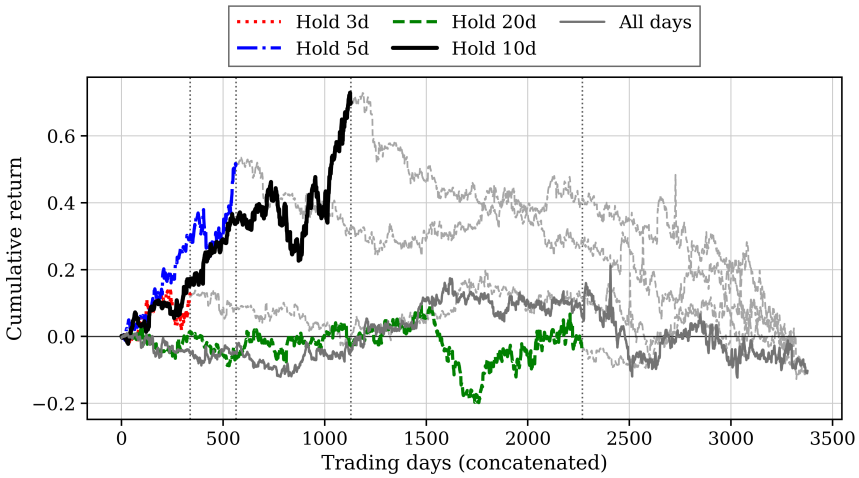
[12] The annualised alpha in Table 9 ($\bar{r} \times 252$ on calendar-time) is the realised return of a full-year deployment for hold-and-reverse and of a $\approx$ 50-active-day deployment for hold-and-close. The Sharpe ratio is computed on the active-day subseries and annualised by $\sqrt{252}$, giving a deployment-invariant measure comparable across variants.

[13] The adjustment is $r_{i,t}^{\text{adj}} = r_{i,t}^{\text{corp}} - \left(-D_{i,t}^{\text{mod}}\, \Delta y_{i,t}^{\text{matched}}\right)$, where $\Delta y_{i,t}^{\text{matched}}$ is the change in the Gürkaynak–Sack–Wright (GSW) Treasury yield interpolated at the bond's modified duration, following van Binsbergen, Nozawa, and Schwert (2025).

(Figures 19 and 20).



(a) Treasury



(b) Corporate (duration-adjusted)

Figure 7: **Active vs. residual inter-meeting window, hold-and-close rebal_str long-short.** Coloured segments accumulate the return through the k-day active window, with line styles denoting the holding horizon ($k = 10$ in solid black); light grey segments continue through the rest of the inter-meeting period. The reversion in the grey segment motivates the hold-and-reverse variant.

Table 9 reports the alphas. Following the standard bond-pricing literature (Nozawa, 2017; Dickerson, Mueller, and Robotti, 2023), Treasury tests use the yield-curve shocks, the duration-carry factor, and changes in MOVE as control factors; corporate tests add the corporate market return, credit-spread changes, and changes in VIX. I report the static factor alpha and a conditional alpha in which each factor is interacted with the direction of the FOMC shock.[14]

[14]Conditional regression: $r_t^{LS} = \alpha^{\text{cond}} + \mathbf{b}'(s_t \cdot \mathbf{f}_t) + \varepsilon_t$, with $s_t = \text{sign}(FOMC^{\text{Shock}}_{T(t)})$ on active days for hold-and-

Alpha results also show the peak-and-reverse pattern. The naive alpha is significantly positive at the one-to-two-week horizons, fades by $k = 20$, and dies in the "All days" column, in both markets and both strategies. The conditional alpha is much smaller but significant at $k = 10$ for Treasury under both strategies and at $k = 5$–10 for corporate hold-and-reverse strategy. What's more, the active-day Sharpe ratios at $k = 5$ are 2.31 and 1.98 for hold-and-close deployment, quite high relative to the roughly 0.5 of the U.S. equity market; hold-and-reverse earns larger alphas at lower Sharpe ratios. The pervasive pattern presents in every test: the spread opens after the announcement and fully closes by the next meeting, the signature of transitory demand pressure rather than information.

Note that everything in this section, like the panel regressions, identifies relative responses of high- versus low-rebal_str bonds within an inter-meeting window. The dampening is a transitory difference in the path, not a permanent wedge, as the reversal shows: bonds more exposed to rebalancing demand travel less of the aggregate round trip, while bonds less exposed overshoot and come back. Stabilization means a smaller round trip in prices, and the model says the same: the κ-shift dampens the impact response while long-run values remain pinned by fundamentals.

5.3 Which institution moves the yield?

The rebalancing strength construction (28) aggregates rebalancing exposures across institution groups. I follow GolWdsmith-Pinkham, Sorkin, and Swift (2020) and decompose the pooled $\widehat{\beta}_1$ into a weighted sum of group-specific estimates,

$$\widehat{\beta}_1 \;\approx\; \sum_g \widehat{\alpha}_g \, \widehat{\beta}^{\mathrm{JI}}_{1,g}, \qquad \widehat{\alpha}_g = \frac{\widehat{\mathrm{Cov}}(Z_g, Z_{\mathrm{agg}})}{\widehat{\mathrm{Var}}(Z_{\mathrm{agg}})},$$

where $Z_g \equiv FOMC_t^{\mathrm{Shock}} \times \mathrm{rebal_str}_{g,i,t}$ is the group-specific shifter (with $\mathrm{rebal_str}_{g,i,t}$ built using only group-g ownership shares), Z_{agg} is its sum across groups, and $\widehat{\beta}^{\mathrm{JI}}_{1,g}$ is the coefficient from estimating (29) with Z_g replacing Z_{agg}. The weight $\widehat{\alpha}_g$ sums to one, measuring how much group g's shifter co-moves with the pooled shifter.

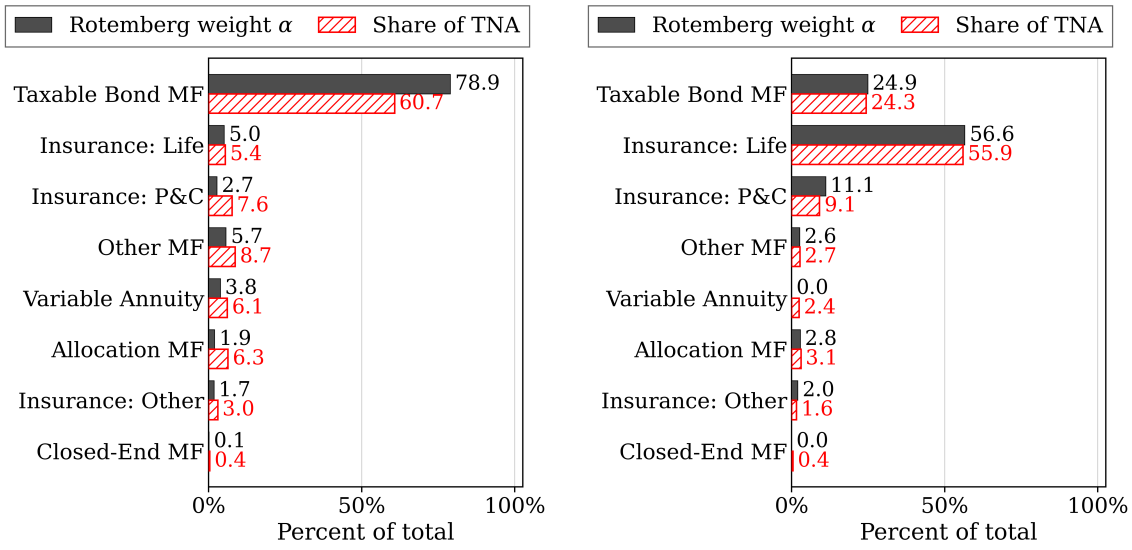
In Treasury, taxable bond mutual funds carry 79% of the $\widehat{\beta}_1$; in Corporate, the weight shifts to life insurers (57%) and bond mutual funds (25%) (Figure 8). Some might concern that whether the channel only works in one large group. Two checks confirm that the result does not stand or fall with any one holder class. First, rebuilding the Treasury signal after removing mutual funds entirely delivers a coefficient of

close, additionally flipping sign on reverse-phase days for hold-and-reverse, where $T(t)$ is the FOMC event whose position is held at t.

Table 9: **rebal_str Long-Short Portfolio Alphas: Naive and Conditional.** Annualised realised PnL (%/yr, $\bar{r} \times 252$ on calendar-time daily series), Newey–West t-statistics in parentheses. *Naive*: unconditional factor regression. *Conditional*: Ferson–Schadt (1996) variant with each factor interacted with the position direction s_t. Factor specs: GSW + DUR + ΔMOVE for Treasury raw returns; MKTB + CRF + GSW + ΔMOVE + ΔVIX for Corporate duration-adjusted returns (DUR dropped because the bond return is duration-stripped). "All days" is $t+1$ through two days before the next FOMC (not applicable for the reverse variant). Sharpe is computed on the active-day return subseries, annualised by $\sqrt{252}$. Sample: 2010-01 onward. */**/***: 10%/5%/1% significance.

	$k=3$	$k=5$	$k=10$	$k=20$	All days
Panel A: Treasury, raw returns					
Hold-and-close					
$\widehat{\alpha}^{\text{naive}}$ (t)	0.75 (1.38)	2.45*** (3.47)	3.28*** (3.52)	2.37* (1.85)	1.99 (1.35)
$\widehat{\alpha}^{\text{cond}}$ (t)	0.18 (1.11)	0.30 (1.63)	0.53** (2.36)	0.44 (1.22)	0.32 (0.72)
R^2 naive / cond	0.01 / 0.86	0.02 / 0.87	0.03 / 0.88	0.04 / 0.87	0.11 / 0.87
Sharpe	1.12	2.31	1.59	0.60	0.33
Hold-and-reverse					
$\widehat{\alpha}^{\text{naive}}$ (t)	−0.49 (−0.33)	2.91** (1.99)	4.56*** (3.03)	2.75* (1.77)	—
$\widehat{\alpha}^{\text{cond}}$ (t)	−0.06 (−0.13)	0.19 (0.43)	0.83** (2.04)	0.56 (1.39)	—
R^2 naive / cond	0.07 / 0.87	0.04 / 0.87	0.03 / 0.87	0.02 / 0.87	—
Sharpe	−0.08	0.49	0.75	0.45	—
Panel B: Corporate, duration-adjusted returns					
Hold-and-close					
$\widehat{\alpha}^{\text{naive}}$ (t)	1.02 (0.95)	3.36*** (2.64)	4.09** (2.46)	0.55 (0.28)	0.02 (0.01)
$\widehat{\alpha}^{\text{cond}}$ (t)	−0.29 (−0.31)	0.98 (0.94)	1.18 (0.86)	−2.11 (−1.36)	−2.97 (−1.64)
R^2 naive / cond	0.00 / 0.26	0.01 / 0.30	0.01 / 0.28	0.02 / 0.29	0.03 / 0.29
Sharpe	0.84	1.98	1.26	0.00	−0.04
Hold-and-reverse					
$\widehat{\alpha}^{\text{naive}}$ (t)	2.01 (0.90)	6.70*** (2.79)	8.16*** (3.14)	1.07 (0.42)	—
$\widehat{\alpha}^{\text{cond}}$ (t)	2.90 (1.55)	5.24*** (2.63)	5.47** (2.57)	−1.09 (−0.52)	—
R^2 naive / cond	0.02 / 0.29	0.01 / 0.29	0.01 / 0.29	0.01 / 0.29	—
Sharpe	0.23	0.70	0.90	0.04	—

the same size (Appendix Table 13), and also significant in 10% though a bit weaker, because what remains varies less across bonds. Second, the estimate does not lean on a few big funds. Counting each fund's contribution to the identifying variation, the effective number of independent funds behind the estimate is about 3,400 in Treasury and 840 in IG corporate, and the five largest funds together only account for 1.1% and 6.6% of the total. The coefficient is identified from the behavior of thousands of funds acting independently, which is also an indicator that shift-share estimates are well behaved (Borusyak, Hull, and Jaravel, 2022).



(a) Treasury (b) Corporate Bond

Figure 8: **Rotemberg Decomposition Weights for rebal_str.** Horizontal bars (bottom axis) report $\widehat{\alpha}_g$, the share of identifying variation in $FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t}$ attributable to institution group g on the no-crisis sample; weights sum to one. The dashed line (top axis) overlays the average quarterly holdings TNA of each group, which scales the cross-sectional ownership concentration in each market. The figure makes explicit that Rotemberg weights track where ownership of the rebalancing-strength signal is concentrated, rather than which groups participate in the channel.

6 Decomposing and Quantifying the Channel

The theory and empirics can interact in multiple ways. As indicated from the previous theory, the reduced-form $\widehat{\beta}_1$ in Section 5.1 could be risen from two potential forces, the mean-variance channel[15] and the habitat channel, both of which are implied by the model and can be calibrated to the data. What's more, the calibrated model can also quantify the dampening of the long-end, which, alternatively, could also be estimated from the regression.

6.1 Mean-variance and Habitat Channel of Rebalancing

By equation (3), each fund's steady-state optimal duration is a convex combination of the optimal mean-variance duration D_q^{MV} and the fund's own duration target D_f^* with weight $\omega_f \equiv \omega(\delta_f) \in [0, 1]$ standing for the relative pull strength of habitat. It

[15] "Mean-variance" here means the duration-driven MV channel of an investor optimising over interest-rate risk; non-duration variation in corporate-bond returns (credit, liquidity, issuer-specific premia) is absorbed by the liquid-IG sample restriction and the bond-level controls interacted with the FOMC surprise.

can be shown that the post-shock active rebalancing inherits the same structure,

$$\Delta D_q^f \;\equiv\; D_q^{f,\text{new}} - D_q^{f,\text{pass}} \;=\; \underbrace{(1-\omega_f)\big(D_q^{MV} - D_q^{f,\text{pass}}\big)}_{\text{MV channel}} \;+\; \underbrace{\omega_f\big(D_f^{*} - D_q^{f,\text{pass}}\big)}_{\text{habitat channel}}. \quad (31)$$

A higher ω_f means a higer share of the rebalancing can be accrued to the habitat anchoring incentive. At $\omega_f = 0$ the fund is a pure mean-variance optimizer; at $\omega_f = 1$ it is rigidly pinned to D_f^*. The mix ω_f determines the channel decomposition fund by fund.

Empirically, to estimate ω_f, I implement (31) as a per-fund regression of active rebalancing on the two gaps. Firstly, the MV optimum is estimated from the model's closed form expression:

$$D_q^{MV} \;=\; \frac{\lambda_q\, M_{\tau A}}{\gamma_H\, \bar{W}\,\big(\sigma_0^{\text{id}\,2} + \sigma_r^2\, M_{AA}\big)},$$

with $\kappa^* = 0.144$, $\sigma_r = 0.017$, $\sigma_0^{\text{id}} = 0.010$ (Treasury) and 0.040 (Corporate), $\gamma_H = 5$, and $\lambda_q = \text{ACMTP10}_q/A(10)$. Over the fund-panel window the optimum averages 2.6 years in Treasury and 1.6 in Corporate (5.4 and 4.6 at the long-history average term premium).[16] Quarter by quarter it swings with the term premium, from about +9 years to −4, while actual fund durations stay pinned between four and five and a half years; Appendix C plots the two series and explains why a negative optimum is exactly what myopic mean-variance logic prescribes once the term premium turns negative, as it does after 2015 (Sangvinatsos and Wachter, 2005). Funds visibly do not chase the mean-variance optimum, which is the fact the estimates below quantify. The habitat target D_f^* is proxied by the fund's expanding-window mean realised duration $\bar{D}_{f,q-1}$. Adding a partial-adjustment friction $\rho_f \in (0,1]$ to (31) and rearranging,

$$Dur_{f,q}^{\text{reb}} \;=\; \beta_f^{MV}\,\Delta_{f,q}^{MV} \;+\; \beta_f^{\text{hab}}\,\widehat{\Delta}_{f,q}^{\text{hab}} \;+\; \varepsilon_{f,q}, \quad (32)$$

with $\Delta_{f,q}^{MV} \equiv D_q^{MV} - Dur_{f,q}^{\text{pass}}$, $\widehat{\Delta}_{f,q}^{\text{hab}} \equiv \bar{D}_{f,q-1} - Dur_{f,q}^{\text{pass}}$, and the mapping $\beta_f^{MV} = \rho_f(1-\omega_f)$, $\beta_f^{\text{hab}} = \rho_f\,\omega_f$. The habitat-pull strength is therefore

$$\widehat{\omega}_f \;=\; \frac{\widehat{\beta}_f^{\text{hab}}}{\widehat{\beta}_f^{MV} + \widehat{\beta}_f^{\text{hab}}} \;\in\; [0,1], \quad (33)$$

[16] These averages use the supply-weighted moment convention ($\bar{S} = 1/T_{\max}$) of the estimation code. Under the unweighted convention of the appendix formulas and of Table 18, the Treasury average is unchanged and the Corporate average rises to about 6.2 years; the fund-level $\widehat{\beta}_f^{MV}$ absorbs the rescaling of the regressor $\Delta_{f,q}^{MV}$, so the downstream $\bar{\omega}$ and η_r are invariant to the choice (within 0.3%, verified on the raw-OLS aggregation).

clipped to the unit interval and invariant to ρ_f. Estimation requires a minimum of twelve quarterly observations per fund; results are aggregated AUM-weighted within institution group and pooled.[17]

Table 10 reports the AUM-weighted $\widehat{\omega}_g$ for each type of investors. The AUM-weighted $\bar{\omega}$ is 0.88 in Treasury and 0.89 in Corporate: habitat anchoring carries roughly 90% of the rebalancing strength in both markets. Both channels exist, but habitat dominates. This not only reveals the pervasive duration preference across different types of investors, but also implies that investors are more influenced by their own duration targets than by the market-wide mean-variance optimum, which has important implications for understanding the dynamics of bond markets and the impact of monetary policy on yields.

The group-level estimates also echo the institutional priors of Table 1: in Corporate the insurance groups anchor most strongly (0.90–0.91) while variable-annuity separate accounts sit lowest (0.74). Variable-annuity accounts follow the sub-fund portfolios their clients choose, so the insurer has no duration target of its own to defend. In Treasury, Life insurers sit lowest (0.80): a Life general account's long-duration anchor binds mainly in its much larger credit and structured-product book, not its Treasury sleeve.

Table 10: **Per-Group AUM-Weighted Habitat-Pull Strength $\widehat{\omega}_g$ and MV-Share.** For each institution group g, $\widehat{\omega}_g$ is the within-group AUM-weighted mean of the per-fund OLS slope ratio $\widehat{\omega}_f$ in equation (33); MV-share equals $1-\widehat{\omega}_g$. The POOLED row aggregates across all funds, AUM-weighted at the fund level. Sample: 2010-01-31 onward, ≥ 12 quarterly observations per fund.

	Treasury			Corporate		
Group	n	$\widehat{\omega}_g$	MV-share	n	$\widehat{\omega}_g$	MV-share
Insurance: Life	196	0.80	20%	781	0.91	9%
Insurance: P&C	593	0.90	10%	1,223	0.89	11%
Insurance: Other	60	0.91	9%	129	0.90	10%
Variable Annuity	114	0.92	8%	186	0.74	26%
Allocation MF	56	0.96	4%	80	0.89	11%
Other MF	158	0.87	13%	181	0.89	11%
Taxable Bond MF	390	0.86	14%	790	0.88	12%
Closed-End MF	5	0.98	2%	45	0.82	18%
POOLED (AUM-wt fund-level)	1,572	**0.88**	**12**%	3,415	**0.89**	**11**%

[17]The model defines ω_f through the deep parameter δ_f, which is interpreted as the strength of any force pulling fund f's portfolio toward a fund-specific target, whether an explicit mandate, a prospectus duration band, a slow-moving strategic view, or a universe restriction that makes the fund's effective MV optimum differ from the market-wide D_q^{MV}. The estimated $\widehat{\omega}_f$ measures the combined anchoring strength rather than distinguishing source.

6.2 Model and Regression-implied long-end dampening

Given the evidence of "Reaching for habitat", it's always important to ask about the quantity: How much of the long-end response does the channel dampen? There are two routes to address: model calibration, and regression revelation. The model route uses the estimated $\bar{\omega}$ and β^{MV} to calibrate the implied dampening; the regression route estimates the dampening maturity by maturity from the FOMC regression in Section 5. This section shows that the two routes show similar pattern and agree in the long-end.

The model in Section 3 showed that rebalancing dampens the yield response of maturity τ by

$$D(\tau) \;=\; \frac{A_{BM}(\tau) - A_{\eta}(\tau)}{\tau} \cdot 100\,\mathrm{bp}, \qquad A_{BM}(\tau) = \frac{1 - e^{-\kappa^* \tau}}{\kappa^*}, \quad A_{\eta}(\tau) = \frac{1 - e^{-\kappa_\eta \tau}}{\kappa_\eta}, \tag{34}$$

$D(\tau)$ can be explained as the reduction in the τ-maturity yield's response to a 100 bp short-rate shock. It is hump-shaped, vanishing at both ends and peaking at $\tau^\dagger \approx 1.79/\kappa^*$; the peak location depends on κ^* alone, and η_r sets only the height.[18] The model's implication therefore doesn't require asset pricing regressions.

The first, κ^*, comes from the short rate itself: an AR(1) gives 0.145 in Treasury and 0.177 in IG corporate.[19] The second is $\eta_r = \gamma_A \sigma_r^2(\Omega - \Omega_0)$ of equation (11), the extra state-sensitivity of the duration risk arbitrageurs absorb. Substituting the aggregated fund-level slopes $(\bar{\omega}, \beta^{\mathrm{MV}})$ into Ω gives the closed form (Appendix I)

$$\eta_r \;=\; \bar{\omega}\left(\beta^{\mathrm{MV}} + 1\right) f_H\, \kappa^*, \tag{35}$$

where f_H is the habitat sector's share of aggregate duration risk. Solving the mapping to its fixed point[20] gives $\eta_r^* = 0.030$ in Treasury and 0.031 in IG corporate, hence

[18] The 1.79 comes from the peak's first-order condition. Setting $D'(\tau) = 0$ and writing $x = \kappa^*\tau$, the condition reduces, to leading order in η_r, to $e^x = 1 + x + x^2$, whose unique positive root is $x^\dagger \approx 1.79$; hence $\tau^\dagger \approx 1.79/\kappa^*$ (Appendix J). The economics of the hump: at short maturities both loadings grow nearly one-for-one with maturity, so the gap between them, spread over τ years of yield, is negligible; at long maturities both loadings have flattened toward their ceilings, so the gap stops growing while it is averaged over ever more years. The trade-off is sharpest about 1.8 mean-reversion horizons out, and a larger η_r deepens the gap without moving where it peaks.

[19] An AR(1) on the one-year yield identifies the *physical* mean reversion; the loading A_{BM} uses the *risk-neutral* κ^*, which adds a term-premium contribution. A Vasicek fit to the term structure gives a higher risk-neutral mean reversion in Treasury ($\kappa^*_{\mathrm{YC}} = 0.227$, versus 0.145 physical, consistent with non-trivial Treasury term-premium dynamics) and an essentially identical one in IG corporate (0.172 versus 0.177). We report the physical estimate as a transparent single-series baseline; the risk-neutral κ^*_{YC} would lower and pull in the dampening hump (peak near 8y rather than 11y in Treasury), but the qualitative middle-maturity match is preserved under either calibration.

[20] The mapping is derived at the benchmark loading A_{BM}, but η_r itself changes the loading, and with it the supply moment $\Xi(\kappa) = \int_0^{T_{\max}} A(u;\kappa)^2\, \bar{S}(u)\, du$ that enters Ω. The self-consistent value therefore solves $\eta_r = \bar{\omega}(\beta^{\mathrm{MV}} + 1) f_H\, [\Xi(\kappa^* + \eta_r)/\Xi(\kappa^*)]\, \kappa^*$. I compute it by iterating from $A^{(0)} = A_{BM}$, updating η_r and the loading in turn; the iteration converges within fifteen iterations in both markets, with no damping needed (Appendix I, Figure 24).

$\kappa_\eta = 0.175$ and 0.208 [21]. This calibration shows the dampening peak to be near 11y in Treasury and 9y in IG corporate. Because neither input touches the cross-sectional regression, both the location and the height of the hump are predictions, not fits.

The regression side delivers the same object independently. The maturity-gradient specification of Section 5.1 estimates how much a bond's FOMC-window return is dampened per unit of rebalancing exposure at duration τ, namely $\widehat{\beta}_1 + \widehat{\beta}_2(\tau - \bar{D})$. Converting this per-bond estimate into the market-level yield dampening the model speaks about takes two steps:

$$D_{\text{reg}}(\tau) \;=\; \underbrace{\left[\widehat{\beta}_1 + \widehat{\beta}_2(\tau - \bar{D})\right]}_{\text{return dampening per unit exposure}} \;\times\; \underbrace{\text{Share}^H(\tau)}_{\text{habitat share of the market}} \;\times \frac{1}{\tau} \cdot 100\,\text{bp}. \quad (36)$$

The first term, $\widehat{\beta}_1 + \widehat{\beta}_2(\tau - \bar{D})$, answers a hypothetical question: how much would a bond's return be dampened if it were held entirely by unit-intensity rebalancers? The second term brings this back to the real market: at maturity τ, only a fraction $\text{Share}^H(\tau)$ of outstanding bonds sits with preferred-habitat investors, mutual funds and ETFs, insurers, pensions, and banks, so the market-wide dampening is the per-unit response times that fraction. The shares come from post-2010 quarterly means from Federal Reserve Z.1 sector holdings combined with MSPD and ICI maturity distributions, with quarterly variation propagated into the confidence bands; values and sources are in Appendix Table 11.[22] The division by τ converts a return dampening into a yield dampening.

Figure 9 overlays the model-implied dampening $D(\tau)$ on the regression-implied dampening $D_{\text{reg}}(\tau)$ with 95% confidence bands. For Treasury, the model and regression agree inside the CI band at all fourteen maturity points ($\tau \in [2, 15]$y). In IG corporate, it matches at eight of fourteen, every point from eight years out, spanning the middle-maturity peak where both objects identify the channel most cleanly.[23] Pattern and scale line up: the hump the model predicts is the hump the regression finds, at every maturity in Treasuries and from eight years out in corporates, the range that matters most for long-rate analysis. The match validates both sides at once. The dampening estimated by comparing bonds' price reaction round FOMC and scaled to the whole market, agree with the number the theoretical model implies

[21] TEvaluating the mapping just once at the benchmark loading, before any feedback, gives $\eta_r = 0.040$ (Treasury) and 0.037 (IG corporate) instead

[22] An alternative convention treats the coefficient as a response per unit of raw exposure and multiplies $\text{Share}^H(\tau)$ by the panel's average sensitivity; Appendix I reports the comparison under this more conservative scaling.

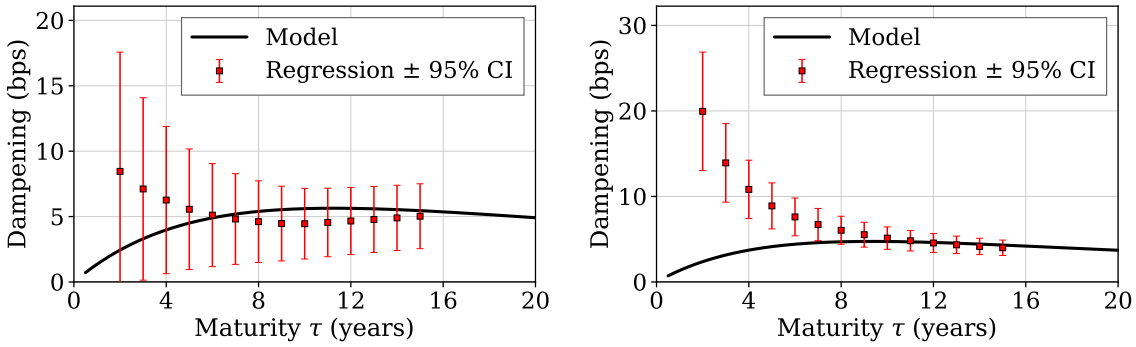
[23] The match survives the κ^* identification choice: a Vasicek yield-curve fit gives a higher risk-neutral κ^* in Treasury, which moves the predicted peak from roughly twelve to eight years without disturbing the middle-maturity agreement.

Table 11: **Habitat-Sector Ownership Shares by Maturity.** Share of each maturity segment's amount outstanding held by habitat-prone sectors, post-2010 averages. $\text{Share}^H(\tau)$, the sum of the four columns, is the market-coverage weight in equation (36); its quarterly variation over 2010–2024 enters the confidence bands of Figure 9 by the delta method.

	Treasury					IG Corporate				
Maturity	MF+ETF	Ins.	Pens.	Banks	Share^H	MF+ETF	Ins.	Pens.	Banks	Share^H
$<$ 2y	0.09	0.01	0.01	0.06	0.17	0.45	0.10	0.05	0.05	0.65
2–5y	0.14	0.02	0.02	0.07	0.25	0.36	0.20	0.08	0.04	0.68
5–7y	0.15	0.03	0.03	0.06	0.27	0.30	0.25	0.10	0.03	0.68
7–12y	0.16	0.05	0.04	0.05	0.30	0.27	0.30	0.10	0.02	0.69
$>$ 12y	0.13	0.16	0.12	0.05	0.46	0.19	0.42	0.12	0.01	0.74

Sector holdings from the Federal Reserve Financial Accounts (Z.1), quarterly 2010–2024: open-end mutual funds L.121, ETFs L.122, life and P&C insurers L.115–L.116, private and state and local pensions L.118–L.119, U.S.-chartered depository institutions L.111. Maturity allocation within sectors from ICI cross-sections (funds), NAIC Schedules D and DB (insurers), FFIEC Call Report Schedule RC-B (banks), and Bloomberg fixed-income ETF composition. Market totals by maturity from the Treasury Monthly Statement of the Public Debt (Treasuries) and Z.1 L.213 with ICI (IG corporates). The remainder of each segment is held by foreign investors, the Federal Reserve, money-market funds, dealers, and households.

only from holdings and yield curve.



(a) Treasury (b) IG Corporate

Figure 9: **Model-Implied versus Regression-Implied Dampening.** Each panel plots the model's hump-shaped dampening $D(\tau) = [A_{BM}(\tau) - A_\eta(\tau)]/\tau$ per 100 bp short-rate shock (solid), calibrated at the self-consistent η_r^* from the yield curve and the holdings-based habitat-pull panel, against the regression-implied $D_{\text{reg}}(\tau)$ of equation (36) with 95% confidence bands (points and bars); the bands include the quarterly variation of the Z.1 habitat shares. In Treasury the model lies inside the band at all fourteen maturity points; in IG corporate at eight of fourteen, every point from eight years out. The model peak, near $1.79/\kappa^*$, is pinned by κ^* alone, not fit to the dampening; the snapshot evaluation of the calibration mapping gives a visually identical model curve in the identified range. The corporate short-end divergence is the measurement issue discussed in the text.

The divergence in the corporate short end, below eight years, is from measurement. Short-maturity corporate moves load on forces the model deliberately leaves out, like credit risk dynamics and money-market conditions, and dividing a return effect by a small τ mechanically amplifies any such contamination.

How much of the Hanson–Stein puzzle does this buy back? Hanson and Stein (2015) put the excess sensitivity of the 10-year Treasury forward at 42 bp, so the comparable object is the forward-rate dampening $\Delta f(\tau) = e^{-\kappa^*\tau} - e^{-\kappa_\eta\tau}$. Table 12 reports the average forward dampening profile. It peaks at about six years and 6.9 bp per 100 bp shock in Treasury and five years and 5.8 bp in IG corporate. At the 10-year forward, the channel removes about 6 bp of the response in Treasury, undoing roughly 15% of the puzzle, and 4.5 bp in corporates. [24] Noted that, the 42 bp in Hanson and Stein (2015) should be a net number, amplification minus the dampening measured here. Decomposing and quantifying them is left to future works.

Table 12: **Forward-Rate Dampening Across Maturities.** Model-implied dampening of the τ-year-ahead instantaneous forward rate per 100 bp short-rate shock, $\Delta f(\tau) = e^{-\kappa^*\tau} - e^{-\kappa_\eta\tau}$, at the headline calibration. The forward profile peaks earlier than the yield profile: at 6.3 years (6.9 bp) in Treasury and 5.2 years (5.8 bp) in IG corporate.

τ (years)	2	5	7	10	12	15
Treasury (bp)	4.4	6.8	6.9	6.1	5.3	4.2
IG Corporate (bp)	4.2	5.8	5.6	4.5	3.7	2.6

Conversion between the two dampening objects: a τ-year yield is the average of the instantaneous forwards beneath it, so the yield dampening of Figure 9 is $D(\tau) = \frac{1}{\tau}\int_0^\tau \Delta f(u)\,du$. Equivalently, $\Delta f(\tau)$ is the maturity derivative of $\tau D(\tau)$. The forward numbers are the ones comparable to excess-sensitivity estimates stated for forward rates, such as Hanson and Stein (2015).

7 Conclusion

A central question in macro-finance is what determines the magnitude of monetary policy transmission to long-term interest rates. A large literature documents forces that amplify this transmission: reaching-for-yield (Hanson and Stein, 2015), arbitrageur balance-sheet effects (Kekre, Lenel, and Mainardi, 2024), granular Treasury demand reallocations (Jansen, Li, and Schmid, 2025), and flow-induced scaling by bond funds (Fang, 2025). This paper documents an opposing force, Reaching for Habitat: Institutional investors hold persistent duration habitats, actively offset the passive duration drift that rate moves induce in their portfolios, and thereby stabilize the transmission of monetary policy to long rates.

The evidence runs from behavior to pricing to magnitude. First, I document the portfolio duration is highly persistent, and funds offset most of a passive drift

[24] Hanson and Stein (2015) report this excess sensitivity as 45 bp in nominal and 42 bp in real terms at the ten-year-ahead forward, so it is almost entirely a real-rate, term-premium phenomenon with a negligible expected-inflation response at that horizon. The duration-rebalancing channel operates on the nominal term premium; because nominal and real forwards move nearly one-for-one this far out, the nominal dampening measured here maps onto the Hanson–Stein real-forward number without requiring TIPS-based estimation.

within three quarters, a behaviour shared across the institution types in the panel, mutual funds, insurers, pensions, and variable annuities alike. Around FOMC announcements, bonds held disproportionately more heavily by strong rebalancers show significantly dampened price reactions, stronger along maturities, building up for two weeks and reversing over the rest of the inter-meeting window, indicating transitory demand pressure rather than information. A preferred-habitat model in which preferred habitat investors possess mean-variance preference but pay a soft penalty for deviating from their duration targets micro-founds the channel, and the fund-level empirical decomposition attributes roughly ninety percent of the rebalancing to habitat anchoring, rest to mean-variance reoptimization. Calibrated from the yield curve and holdings, the model-implied dampening perfectly matches the regression-implied dampening at the maturities over mid-to-long ends. Quantitatively, the channel removes about 6 basis points of the ten-year forward's response per 100 basis point shock in Treasury and 4.5 in corporates; undoes roughly 15% of the 42 basis point excess sensitivity puzzle from Hanson and Stein (2015).

More broadly, the findings suggest that monetary transmission to long-term interest rates depends not only on expectations and risk premia but on the portfolio-management behavior of the investor base: the composition of holders and the distribution of mandate strengths among them shape how far and large monetary policy reaches along the curve.

8 References

Bauer, Michael D., and Eric T. Swanson. 2023. "A Reassessment of Monetary Policy Surprises and High-Frequency Identification." *NBER Macroeconomics Annual* 37: 87–155.

Bao, Jack, Jun Pan, and Jiang Wang. 2011. "The Illiquidity of Corporate Bonds." *The Journal of Finance* 66(3): 911–946.

Bretscher, Lorenzo, Po-Hsuan Hsu, Lukas Schmid, and Andrea Tamoni. 2026. "Heterogeneous Bond Investors and the Cross-Section of Corporate Bond Returns." Working paper / forthcoming *Review of Financial Studies*.

Cavaleri, Filippo. 2023. *A Preferred-Habitat Model with a Corporate Sector*. SSRN Scholarly Paper No. 4603103, Social Science Research Network. `https://doi.org/10.2139/ssrn.4603103`.

Chaudhary, Manav, Zhiyu Fu, and Jian Li. 2022. "Corporate Bond Multipliers: Substitutes Matter." *SSRN Scholarly Paper* No. 4292266, Social Science Research Network.

Chaudhary, Manav, Zhiyu Fu, and Haonan Zhou. 2026. "Anatomy of the Treasury Market: Who Moves Yields?" *Preprint*, SSRN.

Coppola, Antonio. 2025. "In Safe Hands: The Financial and Real Impact of Investor

Composition over the Credit Cycle." *The Review of Financial Studies* 38(8): 2275–2325.

Culbertson, J. M. 1957. "The Term Structure of Interest Rates." *The Quarterly Journal of Economics* 71(4): 485–517.

Doerr, Sebastian, Egemen Eren, and Semyon Malamud. 2023. "Money Market Funds and the Pricing of Near-Money Assets." *SSRN Electronic Journal* (ahead of print).

Wenning, Derek, and Ziang Li. "From Long to Short: How Interest Rates Shape Life Insurance Markets." Preprint, SSRN, 2025.

Fang, Chuck. n.d. *Monetary Policy Amplification through Bond Fund Flows*. Working paper.

Ferson, Wayne E., and Rudi W. Schadt. 1996. "Measuring Fund Strategy and Performance in Changing Economic Conditions." *The Journal of Finance* 51(2): 425–461.

Gabaix, Xavier, and Ralph S J Koijen. "Granular Instrumental Variables." Journal of Political Economy 132, no. 7 (2024): 2274–303.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu. 2020. "Empirical Asset Pricing via Machine Learning." *The Review of Financial Studies* 33(5): 2223–2273.

Gourinchas, Pierre-Olivier, and Walker Ray. n.d. *A Preferred-Habitat Model of Term Premia, Exchange Rates, and Monetary Policy Spillovers*. Working paper.

Greenwood, Robin, and Dimitri Vayanos. 2014. "Bond Supply and Excess Bond Returns." *The Review of Financial Studies* 27(3): 663–713.

Hanson, Samuel G., Aytek Malkhozov, and Gyuri Venter. 2024. "Demand-and-Supply Imbalance Risk and Long-Term Swap Spreads." *Journal of Financial Economics* 154: 103814.

Harvey, Campbell R., Michele G. Mazzoleni, and Alessandro Melone. n.d. *The Unintended Consequences of Rebalancing*. Empirical Asset Pricing 3, working paper.

Hillenbrand, Sebastian. "The Fed and the Secular Decline in Interest Rates." The Review of Financial Studies 38, no. 4 (2025): 981–1013. Macro Finance II. https://doi.org/10.1093/rfs/hhae089.

Jansen, Kristy A. E., Wenhao Li, and Lukas Schmid. n.d. *Granular Treasury Demand with Arbitrageurs*. Working paper.

Koijen, Ralph S. J., and Motohiro Yogo. 2019. "A Demand System Approach to Asset Pricing." *Journal of Political Economy* 127(4): 1475–1515.

Kelly, Bryan, Diogo Palhares, and Seth Pruitt. 2023. "Modeling Corporate Bond Returns." *The Journal of Finance* 78(4): 1967–2008.

Kuong, John Chi-Fong, James O'Donovan, and Jinyuan Zhang. 2024. "Monetary Policy and Fragility in Corporate Bond Mutual Funds." *Journal of Financial Economics* 161: 103931.

Li, Ziang. n.d. *Long Rates, Life Insurers, and Credit Spreads*. Working paper.

Lu, Xu, and Lingxuan Wu. "Monetary Transmission and Portfolio Rebalancing: A Cross-Sectional Approach." SSRN Electronic Journal, ahead of print, 2023.

Nozawa, Yoshio. 2017. "What Drives the Cross-Section of Credit Spreads?: A Variance Decomposition Approach." *The Journal of Finance* 72(5): 2045–2072.

Ozdagli, Ali, and Mihail Velikov. 2020. "Show Me the Money: The Monetary Policy Risk Premium." *Journal of Financial Economics* 135(2): 320–339.

Peng, Cameron, and Chen Wang. "Factor Rebalancing." SSRN Scholarly Paper No.

3327849. Social Science Research Network, March 7, 2023.

Ray, Walker, Michael Droste, and Yuriy Gorodnichenko. "Unbundling Quantitative Easing: Taking a Cue from Treasury Auctions." Journal of Political Economy 132, no. 9 (2024): 3115–72. https://doi.org/10.1086/729581.

Selgrad, Julia. n.d. *Testing the Portfolio Rebalancing Channel of Quantitative Easing*. Macro Finance II, working paper.

van Binsbergen, Jules H., Yoshio Nozawa, and Michael Schwert. 2025. "Duration-Based Valuation of Corporate Bonds." *The Review of Financial Studies* 38(1): 158–191.

Treynor, Jack L., and Kay K. Mazuy. 1966. "Can Mutual Funds Outguess the Market?" *Harvard Business Review* 44: 131–136.

Vayanos, Dimitri, and Jean-Luc Vila. 2021. "A Preferred-Habitat Model of the Term Structure of Interest Rates." *Econometrica* 89(1): 77–112.

Wilson, Robert. 1968. "The Theory of Syndicates." *Econometrica* 36(1): 119–132.

Hanson, Samuel G., and Jeremy C. Stein. "Monetary Policy and Long-Term Real Rates." *Journal of Financial Economics* 115, no. 3 (2015): 429–448.

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales. 2019. "Shift-Share Designs: Theory and Inference." *The Quarterly Journal of Economics* 134(4): 1949–2010.

Boehm, Christoph E., and Niklas Kroner. n.d. *The U.S., Economic News, and the Global Financial Cycle*. Working paper.

Bogousslavsky, Vincent. 2016. "Infrequent Rebalancing, Return Autocorrelation, and Seasonality." *The Journal of Finance* 71(6): 2967–3006.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-Experimental Shift-Share Research Designs." *The Review of Economic Studies* 89(1): 181–213.

Cieslak, Anna, and Andreas Schrimpf. 2019. "Non-Monetary News in Central Bank Communication." *Journal of International Economics* 118: 293–315.

D'Amico, Stefania, and Thomas B. King. 2013. "Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply." *Journal of Financial Economics* 108(2): 425–448.

Dickerson, Alexander, Christian Julliard, and Philippe Mueller. 2026. "The Co-Pricing Factor Zoo." Working paper.

Dickerson, Alexander, Philippe Mueller, and Cesare Robotti. 2023. "Priced Risk in Corporate Bonds." *Journal of Financial Economics* 150(2): 103707.

Ding, Sitong, and Robert Czech. 2026. "Who Moves the Long End? The Dynamics of Bond Market Segmentation." Working paper, London School of Economics and Bank of England.

Domanski, Dietrich, Hyun Song Shin, and Vladyslav Sushko. 2017. "The Hunt for Duration: Not Waving but Drowning?" *IMF Economic Review* 65(1): 113–153.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110(8): 2586–2624.

Hu, Grace Xing, Jun Pan, and Jiang Wang. 2013. "Noise as Information for Illiquidity." *The Journal of Finance* 68(6): 2341–2382.

Greenwood, Robin, Samuel G. Hanson, Jeremy C. Stein, and Adi Sunderam. 2023. "A Quantity-Driven Theory of Term Premia and Exchange Rates." *The Quarterly Journal of Economics* 138(4): 2327–2389.

Greenwood, Robin, Samuel G. Hanson, and Dimitri Vayanos. 2024. "Supply and Demand and the Term Structure of Interest Rates." Working paper.

Guibaud, Stéphane, Yves Nosbusch, and Dimitri Vayanos. 2013. "Bond Market Clienteles, the Yield Curve, and the Optimal Maturity Structure of Government Debt." *The Review of Financial Studies* 26(8): 1914–1961.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1(1): 55–93.

Hanson, Samuel G., David O. Lucca, and Jonathan H. Wright. 2021. "Rate-Amplifying Demand and the Excess Sensitivity of Long-Term Rates." *The Quarterly Journal of Economics* 136(3): 1719–1781.

Jiang, Zhengyang, Arvind Krishnamurthy, and Hanno Lustig. 2022. "Foreign Safe Asset Demand and the Dollar Exchange Rate." *The Journal of Finance* 76(3): 1049–1089.

Kekre, Rohan, Moritz Lenel, and Federico Mainardi. 2024. *Monetary Policy, Segmentation, and the Term Structure*. Working paper.

Modigliani, Franco, and Richard Sutch. 1966. "Innovations in Interest Rate Policy." *The American Economic Review* 56(1/2): 178–197.

Nakamura, Emi, and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *The Quarterly Journal of Economics* 133(3): 1283–1330.

Oster, Emily. 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business & Economic Statistics* 37(2): 187–204.

Ozdagli, Ali, and Zixuan (Kevin) Wang. 2019. *Interest Rates and Insurance Company Investment Behavior*. Working paper.

Sangvinatsos, Antonios A., and Jessica A. Wachter. 2005. "Does the Failure of the Expectations Hypothesis Matter for Long-Term Investors?" *The Journal of Finance* 60(1): 179–230.

Siani, Kerry. n.d. *Raising Bond Capital in Segmented Markets*. Working paper.

Siriwardane, Emil N., Adi Sunderam, and Gyuri Venter. 2025. "The Liquidity State of the Market and Monetary Transmission." Working paper.

A Appendix: Duration Habitats by Institution Group

This appendix replicates the duration distribution and state-probability transition matrix in Figure 3 separately for each institution group used in the analysis. As in the main text, every figure has the Treasury market in the top row and the corporate bond market in the bottom row; the left column reports the distribution of the institution-level portfolio duration (clipped at the 0.1% and 99.9% percentiles) and the right column reports the average state transition matrix at horizons $k = 1, 4, 10, 20$ quarters. The five duration buckets are $[0, 2]$, $(2, 5]$, $(5, 7]$, $(7, 12]$, $(12, \infty)$ years.

The institution groups follow the FUND_CLASS taxonomy described in Section 2: the Morningstar mutual-fund side contributes Taxable Bond MF, Allocation MF, and Other MF (equity, international, municipal bond, and commodity); the eMAXX side contributes Life Insurance, P&C Insurance, Other Insurance, Variable Annuity, Pension, Hedge Fund, Closed-End MF, Sovereign, and Other Institutional. Below I report six representative groups spanning the three main mandate families: open-end mutual funds, insurance and pension accounts, and closed-end vehicles. Distributions and transition matrices for the remaining groups (Allocation MF, Other MF, Other Insurance) are available upon request.

A common pattern emerges. At one quarter the transition matrix is strongly diagonal for every group, which already establishes the existence of a duration habitat at short horizons. At twenty quarters, several groups (most cleanly P&C and Other insurers in corporate bond) show that the long-horizon distribution collapses into one or two duration buckets regardless of where each institution started. This is not evidence against duration habitat; it is in fact the strongest evidence for it: the off-diagonal cells that look populated in any given quarter are largely snapshot outliers, positions an institution holds transiently away from its home bucket, which revert to the group's home habitat over multi-year horizons. The buckets into which probability mass converges are precisely the buckets where each group's unconditional distribution peaks, confirming that institutions in these groups have well-defined duration habitats and that the apparent dispersion at any cross-section overstates the long-run heterogeneity within each group.

Pensions (Figure 14). Pensions are the smallest group by institution count (46 Treasury and 31 corporate institutions), so distributional shapes are noisier. The corporate distribution is concentrated around 5 years with a long right tail; on corporate,

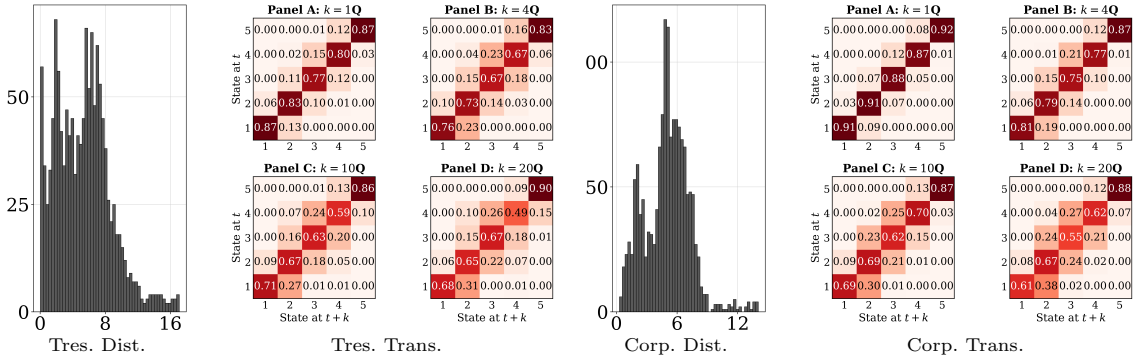


Figure 10: **Taxable Bond Mutual Funds.** Duration distribution and state transition matrix at $k = 1, 4, 10, 20$ quarters; Treasury (left two panels), Corporate (right two panels).

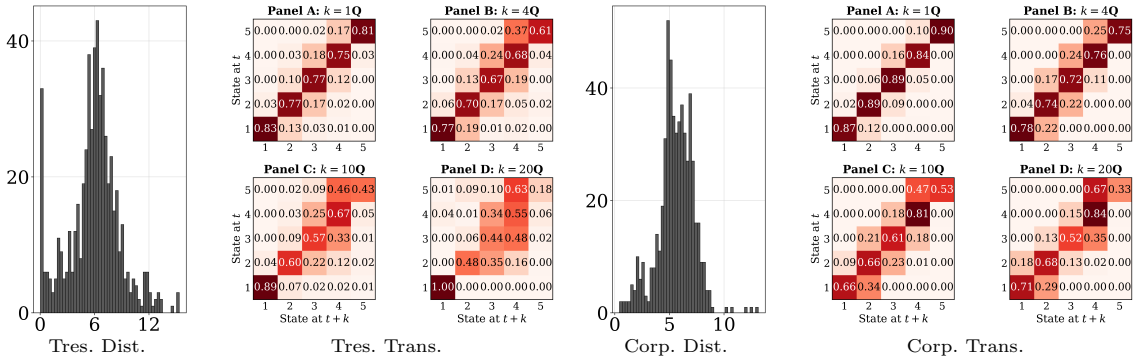


Figure 11: **Variable Annuity Separate Accounts.** Duration distribution and state transition matrix at $k = 1, 4, 10, 20$ quarters; Treasury (left two panels), Corporate (right two panels).

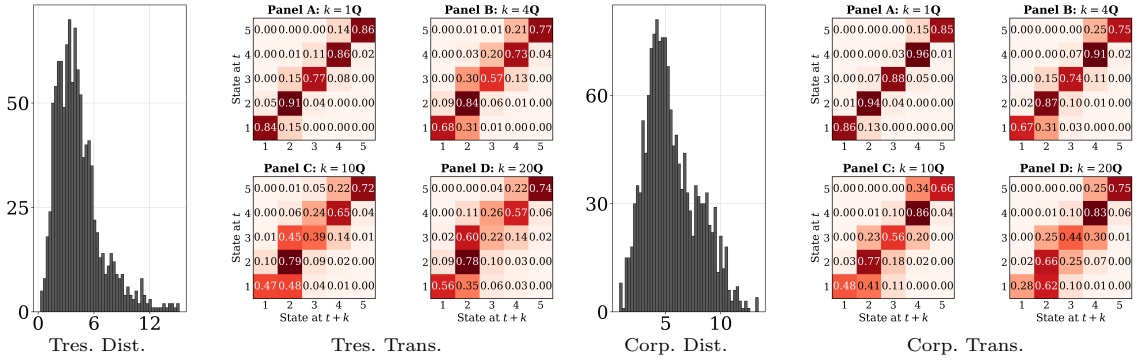


Figure 12: **Life Insurers.** Duration distribution and state transition matrix at $k = 1, 4, 10, 20$ quarters; Treasury (left two panels), Corporate (right two panels).

the 20-quarter transition matrix is sparse but still informative: state 2 $\rightarrow$ state 2 at 1.00, state 4 $\rightarrow$ state 4 at $\sim$ 65%, with the remainder drifting between adjacent buckets. On Treasury the 10- and 20-quarter matrices contain too few observations for a stable estimate (most cells are zero by construction) and should be read with

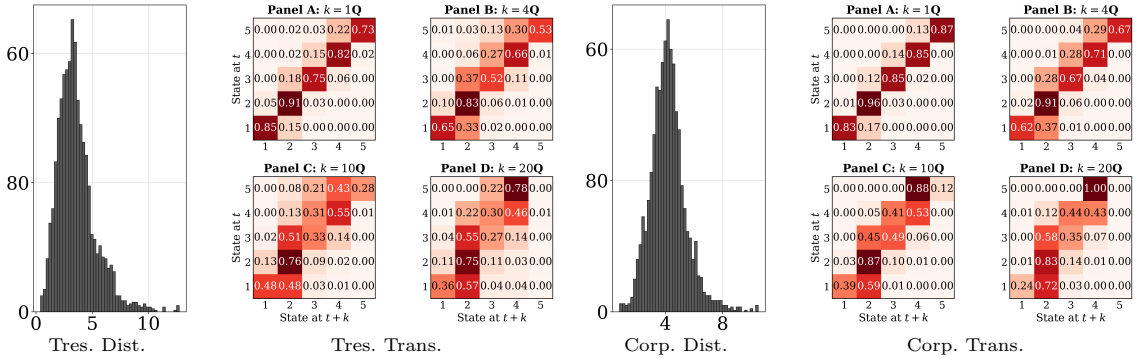


Figure 13: **P&C Insurers.** Duration distribution and state transition matrix at $k = 1, 4, 10, 20$ quarters; Treasury (left two panels), Corporate (right two panels).

caution.

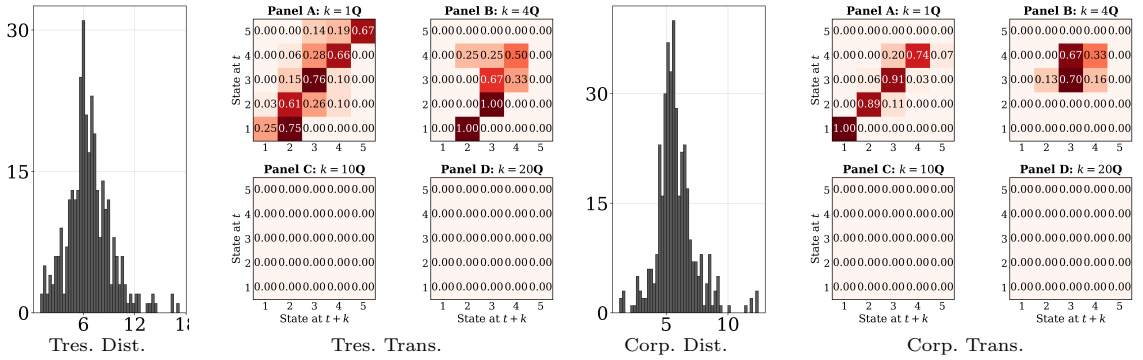


Figure 14: **Pensions.** Duration distribution and state transition matrix at $k = 1, 4, 10, 20$ quarters; Treasury (left two panels), Corporate (right two panels).

Other Institutional Mandates: Closed-End Funds and Residual Categories

Closed-End Mutual Funds (Figure 15). Closed-end bond funds carry permanent capital and typically run leveraged duration positions in narrowly defined mandates (e.g. municipal duration, high-yield duration). The corporate distribution peaks in the $(5, 10]$ bucket and the 20-quarter transition matrix is concentrated on the diagonal at the mandate duration, consistent with the closed-end-fund discipline of holding stated-mandate maturity without redemption pressure.

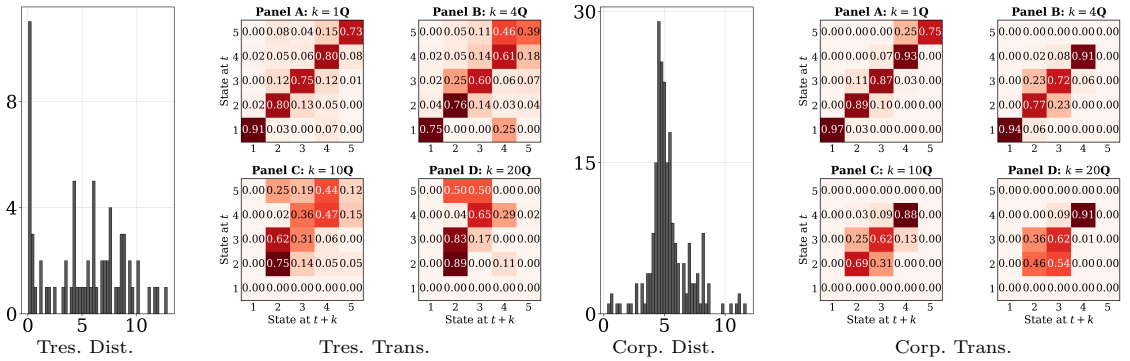


Figure 15: **Closed-End Mutual Funds.** Duration distribution and state transition matrix at $k = 1, 4, 10, 20$ quarters; Treasury (left two panels), Corporate (right two panels).

B Appendix: Implied Half-Lives of Duration Mean Reversion

This appendix details the half-life calculation summarised in Section 4.1. I estimate the AR(1)-style regression in deviations from each fund's time-average duration,

$$Dur_{i,t} - \bar{Dur}_i = \beta_k \left(Dur_{i,t-k} - \bar{Dur}_i \right) + \varepsilon_{i,t}, \tag{37}$$

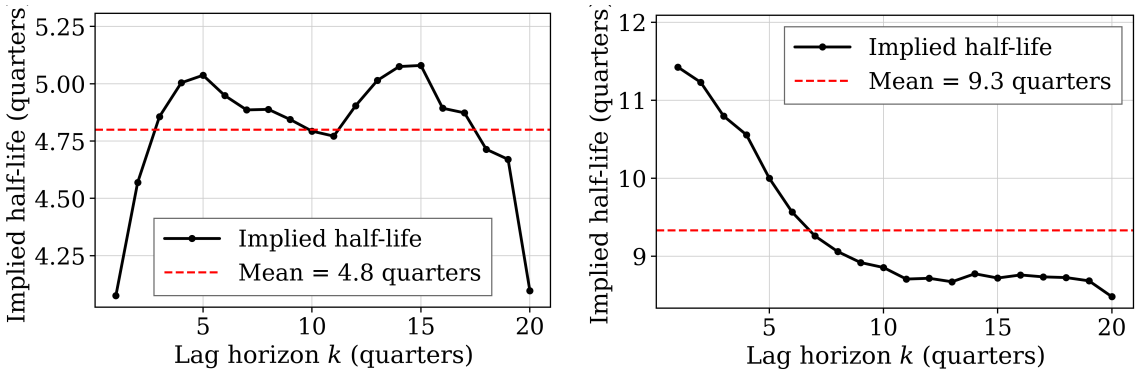
where $\bar{Dur}_i = \frac{1}{T_i} \sum_t Dur_{i,t}$. Interpreting the process as an underlying AR(1) $Dur_{i,t} = \rho\, Dur_{i,t-1} + u_{i,t}$ with $\beta_k \approx \rho^k$, the implied quarterly coefficient is $\hat{\rho}_k = \text{sign}(\hat{\beta}_k)\, |\hat{\beta}_k|^{1/k}$ and the implied half-life from the k-horizon regression is

$$h_{1/2,k} = -\frac{\ln 2}{\ln |\hat{\rho}_k|} = -\,k\, \frac{\ln 2}{\ln |\hat{\beta}_k|}. \tag{38}$$

Figure 16 reports $h_{1/2,k}$ for $k = 1, \ldots, 20$ quarters. The implied half-life fluctuates around four quarters in the Treasury market and seven quarters in the corporate bond market. Re-estimating within each of the twelve institution groups gives half-lives in a narrow range of roughly four to eight quarters throughout, so the fast pull back to habitat is a property of the institutional universe rather than of any single investor type.

C Appendix: Mechanism Decomposition Robustness

This appendix documents the cross-sectional distribution of $\widehat{\omega}_f$ by institution group, the calibration invariance of the decomposition to γ_H, and a Treasury-side robustness check that drops Taxable Bond Mutual Funds from the rebalancing-strength signal.



(a) Treasury (b) Corporate Bond

Figure 16: **Implied Half-Life of Duration Mean-Reversion.** Implied half-life $h_{1/2,k}$ from equation (38) for the Treasury (left) and corporate-bond (right) markets, plotted against the regression horizon k in quarters.

Calibration invariance to γ_H. The per-fund $\widehat{\omega}_f$ from the OLS slope ratio in equation (33) is exactly γ_H-invariant: under the no-intercept two-variable convention, both per-fund slopes scale as $1/\gamma_H$, so the ratio cancels. The per-group $\widehat{\beta}_g^{MV}$ is identified from $(\widehat{\omega}_f, \widehat{\beta}_f^{\text{reb}})$ pairs and inherits the rescaling, so the cohort MV-share is correspondingly invariant. I verify bit-for-bit invariance empirically across $\gamma_H \in \{2.5, 5, 10, 25\}$.

The mean-variance optimum over time. Figure 17 plots the quarterly D_q^{MV} against the cross-section of actual fund durations. Two features deserve comment. First, D_q^{MV} turns negative after 2015. This is not a pathology of the calibration: the optimum is proportional to the term premium, and when compensation for bearing duration risk is negative, a myopic mean-variance investor optimally holds negative duration exposure, shorting long bonds against short ones. This is the standard prescription of portfolio choice under time-varying bond premia (Sangvinatsos and Wachter, 2005), and negative-duration bond strategies exist as investable products. That essentially no fund follows the optimum into negative territory is precisely the anchoring the decomposition measures. Second, the common D_q^{MV} reflects the common-γ_H benchmark; heterogeneous risk aversion would rescale each fund's optimum by $1/\gamma_i$. A common rescaling leaves $\widehat{\omega}_f$ unchanged (the invariance above), and fund-specific level differences in the optimum are absorbed by the habitat-target proxy, which is itself a fund-specific level. What identifies the mean-variance margin is the common time variation in λ_q, not the level of the optimum.

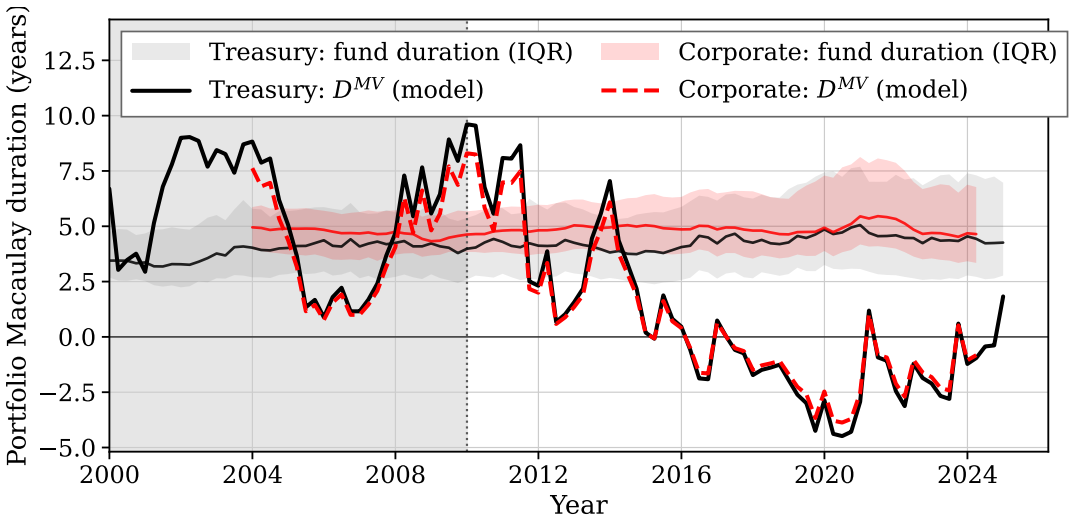


Figure 17: **The Mean-Variance Optimum versus Actual Fund Durations.** Quarterly model-implied mean-variance optimal duration D_q^{MV} (Treasury: solid black; IG corporate: dashed red) against the cross-section of actual fund durations (thin lines: median; shaded: interquartile range); the grey region marks the pre-2010 build-up of the estimation sample. D_q^{MV} moves with the term premium and turns negative after 2015, when compensation for duration risk does; actual fund durations stay between four and five and a half years throughout. The gap between the two lines is the visual counterpart of the habitat-share estimate $\bar{\omega} \approx 0.9$ below.

Distribution of $\widehat{\omega}_f$ by institution group. Figure 18 plots the cross-sectional distribution of $\widehat{\omega}_f$ within each institution group with measurable mass. Per-group AUM-weighted means range from 0.80 (Insurance: Life) to 0.98 (Closed-End MF) in Treasury, and from 0.74 (Variable Annuity) to 0.91 (Insurance: Life) in Corporate, with the insurance and closed-end groups at the top and the variable-annuity and taxable-bond-MF groups in the central band. The cross-group ranking is consistent with the institutional priors of Table 1.

Drop-Taxable-Bond-MF Treasury robustness (full grid). Table 13 reports the main mechanism regression coefficient $\widehat{\beta}_1$ on $FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t}$ for the pooled Treasury rebalancing-strength signal alongside the non-TBM aggregate $\text{rebal_str}_{i,t}^{\text{no-TBM}}$ defined in Section 5.3. The non-TBM signal is constructed by subtracting the Taxable-Bond-MF group's contribution from the pooled signal, retaining the other institution groups. The per-unit coefficient survives dropping TBMs at every horizon: the no-crisis $k = 4$ coefficient moves from 1.50 (pooled) to 1.44 (no-TBM), with similar magnitudes at $k = 2, 6$. Significance attenuates from the 1% level to the 10% level because the cross-sectional standard deviation of $\text{rebal_str}^{\text{no-TBM}}$ is only about 36% of the pooled signal: TBMs supply most of the cross-sectional ownership

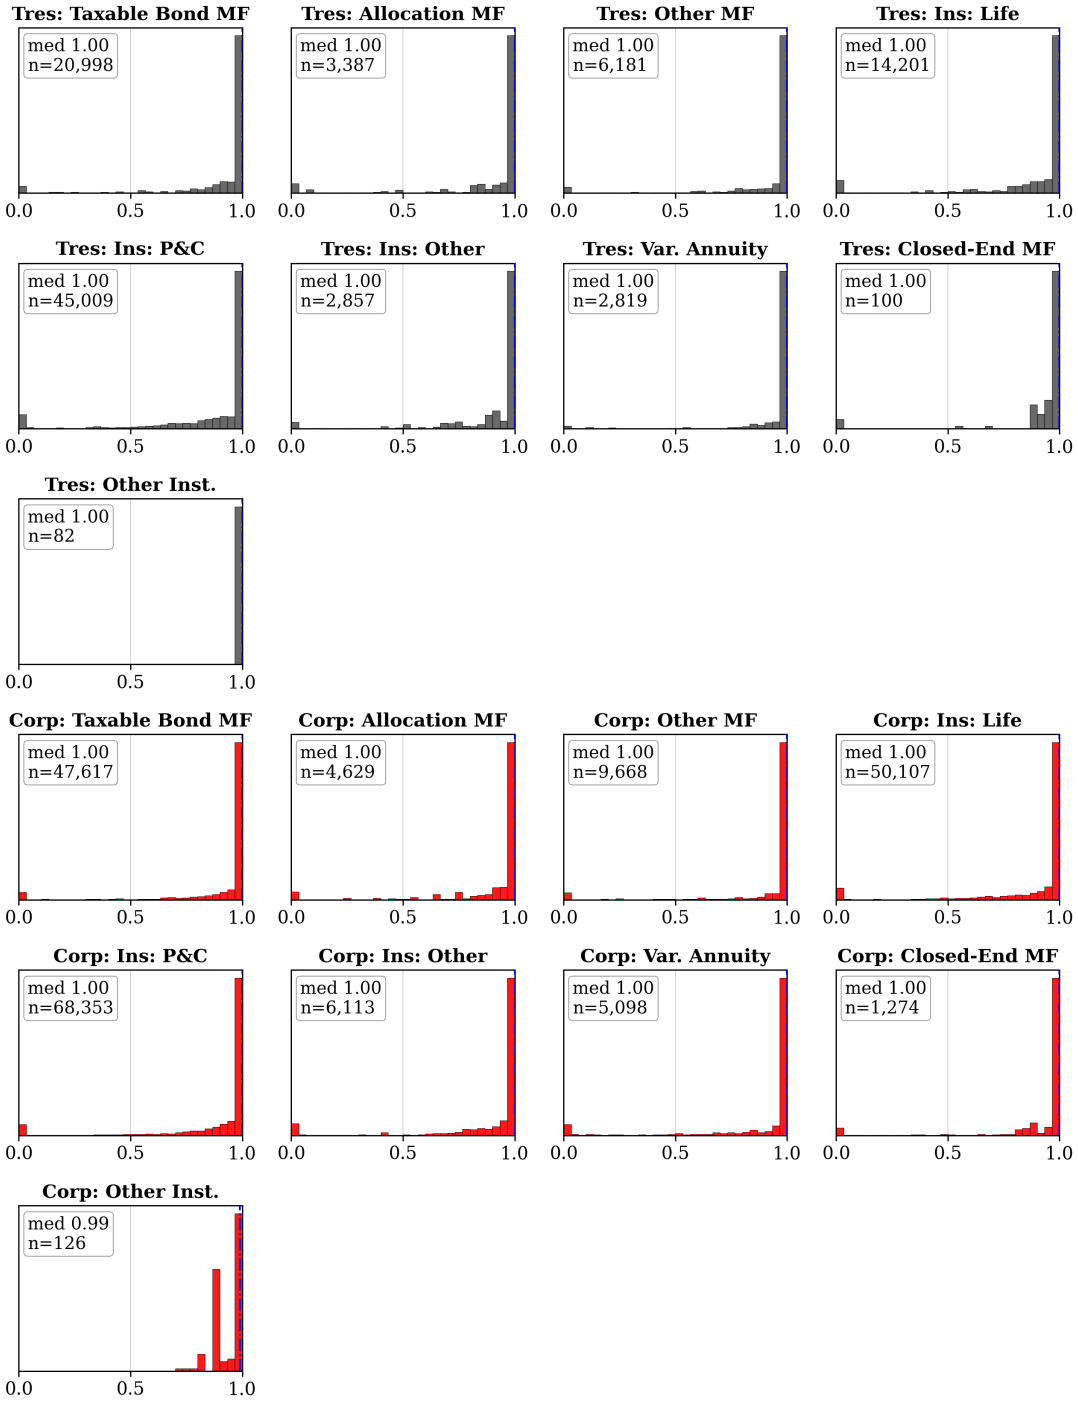


Figure 18: **Distribution of Habitat-Pull Strength $\widehat{\omega}_f$ by Institution Group.** Histograms within each of the twelve institution groups, pooling fund-quarters across the FOMC sub-sample. Vertical line marks the within-group median.

dispersion that powers identification, but the per-unit price-impact lives in the more habitat-anchored non-TBM segment.

Table 13: **Treasury Cross-Sectional Dampening: Pooled vs. Non-Taxable-Bond-MF Signal.** Estimates of $\widehat{\beta}_1$ on $FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t}$ from equation (29), comparing the pooled rebalancing-strength signal against the non-TBM aggregate. The non-TBM signal $\text{rebal_str}_{i,t}^{\text{no-TBM}} \equiv \text{rebal_str}_{i,t} - \text{rebal_str}_{i,t}^{\text{TBM}}$ excludes Taxable Bond Mutual Funds from the bond-level aggregation, retaining only Allocation MFs, Other MFs, insurers, variable-annuity accounts, pensions, hedge funds, closed-end funds, sovereign accounts, and other institutional mandates. Bond and event-date fixed effects, full controls, bond-clustered standard errors. Stars: * p¡0.10, ** p¡0.05, *** p¡0.01.

	Full sample			No-crisis sample		
Signal	$k=2$	$k=4$	$k=6$	$k=2$	$k=4$	$k=6$
$\text{rebal_str}_{i,t}$ (pooled)	−0.296	+0.491	+0.542	+1.000***	+1.504***	+1.711***
(t-stat)	(−1.02)	(+1.35)	(+1.46)	(+3.07)	(+3.02)	(+3.28)
$\text{rebal_str}_{i,t}^{\text{no-TBM}}$	+0.184	+1.080	+1.053	+1.203	+1.443*	+1.573*
(t-stat)	(+0.34)	(+1.55)	(+1.51)	(+1.63)	(+1.79)	(+1.93)

U.S. insurer rebalancing slopes across rate-cycle subsamples. Table 14 reports the cross-fund mean $\widehat{\beta}_f^{\text{reb}}$ within each insurance group, estimated separately on five subsamples: the full 2010–2024 sample; rate-cut quarters (period-level $\Delta\lambda_t$ in the bottom quartile); rate-hike quarters (top quartile); calendar year 2014 (the year of Domanski-Shin-Sushko's German hunt-for-duration episode); and the post-2020 sample (COVID + 2022 hike cycle). All thirty cells are negative; calendar year 2014 is the most negative subsample for five of six insurer-group × market combinations, consistent with U.S. insurers exhibiting near-perfect active rebalancing in the year that DSS document opposite-direction hunt-for-duration behavior in German insurers.

D Appendix: Crisis-Window Estimates

Why crisis windows interfere with identification. The full and stress-only samples include the COVID liquidity event of February–May 2020 and the 2022 rate-hike cycle, two episodes of severe dealer-intermediation breakdown and arbitrageur wealth shocks during which the cross-section of bond returns is dominated by amplification rather than rebalancing. Mechanically, in those windows:

- Levered intermediaries unwound Treasury inventory under funding stress (Kekre, Lenel, and Mainardi, 2024), producing large same-direction sales that move long-bond prices in the direction of the shock rather than against it.

Table 14: **U.S. Insurer Rebalancing Slopes Across Rate-Cycle Subsamples.** Cross-fund mean $\widehat{\beta}_f^{\text{reb}}$ for each insurance group, estimated separately on five subsamples. Within each subsample, fund-level $\widehat{\beta}^{\text{reb}}$ is the OLS slope of active rebalancing on passive deviation (equation (24), no intercept), aggregated across funds via cross-group mean. Year 2014 is the most negative subsample for five of six insurer-group × market combinations, consistent with U.S. insurers exhibiting near-perfect active rebalancing in the same year that Domanski, Shin, and Sushko (2017) document opposite-direction hunt-for-duration extension in German insurers.

	Full	Rate-cut	Rate-hike	Year 2014	Post-2020
Panel A. Treasury market					
Insurance: Life	−0.81	−0.69	−0.82	**−1.26**	−0.81
Insurance: Other	−0.75	−1.08	−0.69	−0.66	−0.76
Insurance: P&C	−0.72	−0.84	−0.78	**−0.95**	−0.76
Panel B. Corporate market					
Insurance: Life	−0.74	−0.72	−0.71	**−1.00**	−0.50
Insurance: Other	−0.78	−0.84	−0.75	**−1.07**	−0.59
Insurance: P&C	−0.80	−0.88	−0.76	**−1.01**	−0.73

Table 15: **State Dependence of the Rebalancing-Strength Effect.** Spec 3 estimates of $\widehat{\beta}_1$ on $FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t}$ (equation (29): bond and event-date fixed effects, full controls of Table 6) on the full post-2010 sample and on the stress-only sample, the eight FOMC dates inside February–May 2020 and March–December 2022. No-crisis estimates are in Table 7. Bond-clustered standard errors in parentheses. *, **, *** denote significance at the 10%, 5%, 1% levels.

	$k = 2$	$k = 4$	$k = 6$	N
Panel A: Treasury market				
Full sample	−0.300 (0.299)	0.521 (0.365)	0.576 (0.374)	32,067
Stress-only	-2.370^{***} (0.563)	-1.630^{***} (0.508)	-1.790^{***} (0.581)	2,486
Panel B: Corporate bond market				
Full sample	0.294^{**} (0.124)	0.322^{**} (0.135)	0.437^{***} (0.150)	113,655
Stress-only	0.680^{***} (0.249)	0.980^{***} (0.295)	1.220^{***} (0.390)	9,672

- Bond funds faced outflows that forced redemption-driven selling correlated with the shock sign (Fang, 2025), producing flow-induced amplification rather than habitat-driven dampening.
- Liquidity-screened corporate bonds were less affected than off-the-run Treasury notes, generating a cross-market asymmetry: the corporate-stress $\widehat{\beta}_1$ remains positive (the rebalancing channel survives in IG corporates) while the Treasury-stress $\widehat{\beta}_1$ flips sign in Spec 3.

E Appendix: Vanilla-Spec Robustness for the Cross-Sectional Regression

This appendix reports three pooled-OLS specifications that are stricter than Spec 1 of Table 7 and explains why they are not used as the headline baseline. The strict-vanilla specification keeps only the three core regressors ($FOMC_t^{\text{Shock}}$, $FOMC_t^{\text{Shock}} \times$ rebal_str$_{i,t}$, and rebal_str$_{i,t}$) with no other controls and no fixed effects. The two single-control variants additionally control for either ownership (own$_{i,t}$ and its FOMC interaction) or duration ($Duration_{i,t-1}$ and its FOMC interaction), but not both. The Spec 1 in the main Table 7 adds both ownership and duration controls and is the minimum economic baseline used in the paper.

The pattern in Table 16 explains why these stricter specifications cannot serve as the baseline. In the strict-vanilla column, the unconditional cross-sectional slope of FOMC-window returns on $FOMC_t^{\text{Shock}} \times$ rebal_str$_{i,t}$ is wrong-signed or insignificant in most Treasury cells: it is dominated by the standard duration channel (longer-duration bonds bear larger FOMC-window price reactions), and rebal_str$_{i,t}$ is mechanically correlated with duration through the holdings of habitat-anchored institutions. Adding only the ownership control (column +own) does not fix the duration confound and produces large, often opposite-signed estimates as $\widehat{\beta}_1$ instead absorbs the mechanical level co-movement with rebal_str$_{i,t}$. Adding only the duration control (column $+Duration$) over-corrects in the opposite direction: the duration interaction picks up most of the cross-sectional variance on its own, leaving an attenuated and often insignificant residual for $\widehat{\beta}_1$. Only when both ownership and duration are netted out together, the Spec 1 in Table 7, does $\widehat{\beta}_1$ stabilise at a positive, economically meaningful magnitude that survives the addition of the remaining controls (Spec 2) and fixed effects (Spec 3).

F Appendix: Hold-and-Reverse Cumulative Returns

Figure 20 plots the cumulative return of the hold-and-reverse deployment of the rebal_str long-short (Section 5.2), which flips the position at day k and holds the reversed position until the next FOMC, thereby collecting both the opening and the closing leg of the round trip. The $k = 5$ and $k = 10$ deployments—those that reverse closest to the spread's peak in Figure 19—accumulate the largest returns in both markets, while $k = 3$ (reversing before the spread has fully opened) and $k = 20$ (reversing after part of the retracement has already occurred) trail; the dashed

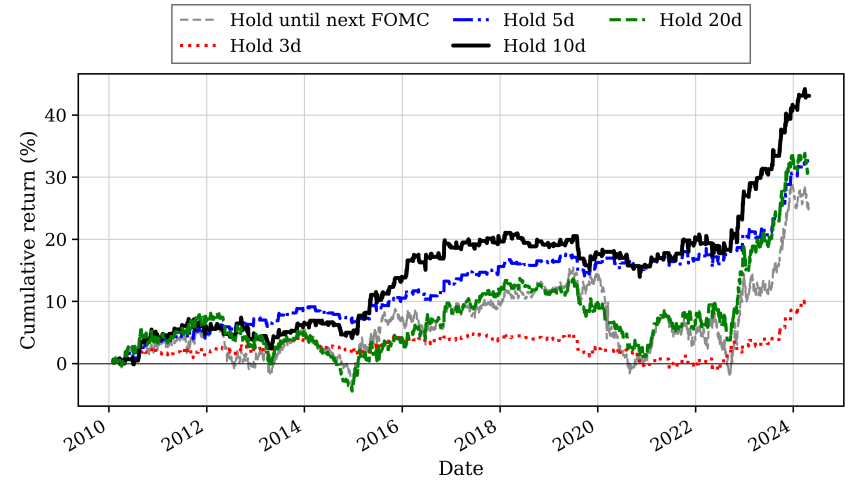
Table 16: **Vanilla-Spec Robustness for $\widehat{\beta}_1$ on $FOMC_t^{\text{Shock}} \times$ rebal_str$_{i,t}$.** Three pooled-OLS specifications stricter than Spec 1 of Table 7: (i) strict vanilla keeps only the three core regressors $FOMC_t^{\text{Shock}}$, $FOMC_t^{\text{Shock}} \times \text{rebal_str}_{i,t}$, rebal_str$_{i,t}$; (ii) +own adds own$_{i,t}$ and $FOMC_t^{\text{Shock}} \times \text{own}_{i,t}$; (iii) $+Duration$ instead adds $Duration_{i,t-1}$ and $FOMC_t^{\text{Shock}} \times Duration_{i,t-1}$. None of the three add fixed effects or any other controls. Standard errors clustered at the bond level in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Sample / Horizon	Strict vanilla	+own	$+Duration$
Panel A: Treasury market			
No-crisis sample			
$Ret_{t,t+2}$	-0.456^{***} (0.142)	3.123^{***} (0.599)	-0.106 (0.143)
$Ret_{t,t+4}$	-0.439^{***} (0.158)	4.431^{***} (0.786)	0.118 (0.164)
$Ret_{t,t+6}$	-0.619^{***} (0.168)	5.145^{***} (0.852)	0.045 (0.173)
Full sample			
$Ret_{t,t+2}$	-0.190^{*} (0.098)	-0.549 (0.359)	-0.092 (0.103)
$Ret_{t,t+4}$	-0.331^{***} (0.115)	3.247^{***} (0.465)	0.066 (0.114)
$Ret_{t,t+6}$	-0.503^{***} (0.129)	4.122^{***} (0.538)	-0.015 (0.121)
Stress-only sample			
$Ret_{t,t+2}$	0.225 (0.243)	-11.568^{***} (1.132)	-0.602^{***} (0.123)
$Ret_{t,t+4}$	0.696^{***} (0.164)	-4.310^{***} (0.720)	-0.730^{***} (0.166)
$Ret_{t,t+6}$	-0.876^{***} (0.218)	0.827 (0.593)	-0.945^{***} (0.199)
Panel B: Corporate bond market			
No-crisis sample			
$Ret_{t,t+2}$	0.108^{***} (0.010)	1.587^{***} (0.054)	0.105^{***} (0.010)
$Ret_{t,t+4}$	0.137^{***} (0.011)	2.234^{***} (0.061)	0.130^{***} (0.011)
$Ret_{t,t+6}$	0.179^{***} (0.013)	3.204^{***} (0.071)	0.171^{***} (0.012)
Full sample			
$Ret_{t,t+2}$	0.059^{***} (0.008)	1.233^{***} (0.045)	0.059^{***} (0.008)
$Ret_{t,t+4}$	0.096^{***} (0.010)	2.438^{***} (0.055)	0.085^{***} (0.009)
$Ret_{t,t+6}$	0.139^{***} (0.011)	3.508^{***} (0.065)	0.123^{***} (0.010)
Stress-only sample			
$Ret_{t,t+2}$	0.006 (0.036)	-1.808^{***} (0.210)	0.079^{**} (0.032)
$Ret_{t,t+4}$	0.348^{***} (0.044)	1.903^{***} (0.272)	0.068 (0.043)
$Ret_{t,t+6}$	0.151^{***} (0.051)	5.064^{***} (0.357)	-0.036 (0.052)

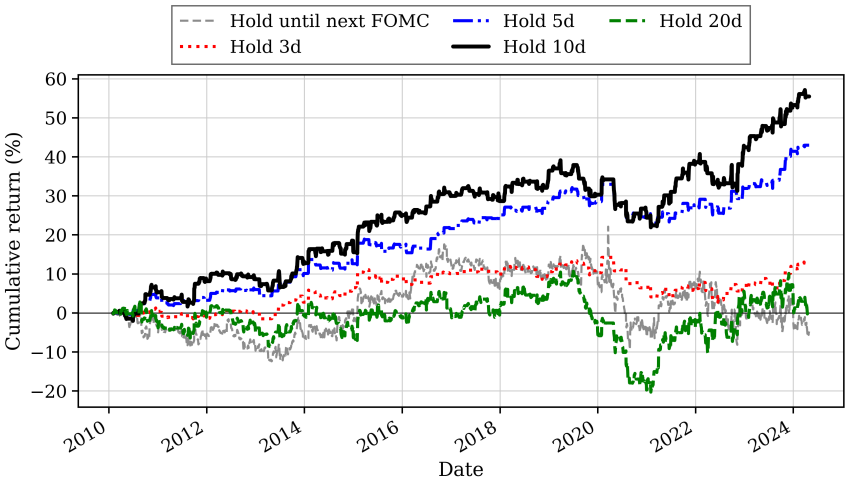
hold-until-next-FOMC benchmark stays near zero. The corresponding alphas are in Table 9.

G Appendix: Holdings Panel Coverage

This appendix reports the coverage of the combined eMAXX and Morningstar holdings panel relative to the U.S. bond market, as discussed in Section 2. Coverage is measured both by market value and by dollar-duration (market value times modified duration), the latter being the economically relevant metric for a duration-rebalancing channel because it weights each holding by its sensitivity to a rate shock.



(a) Treasury

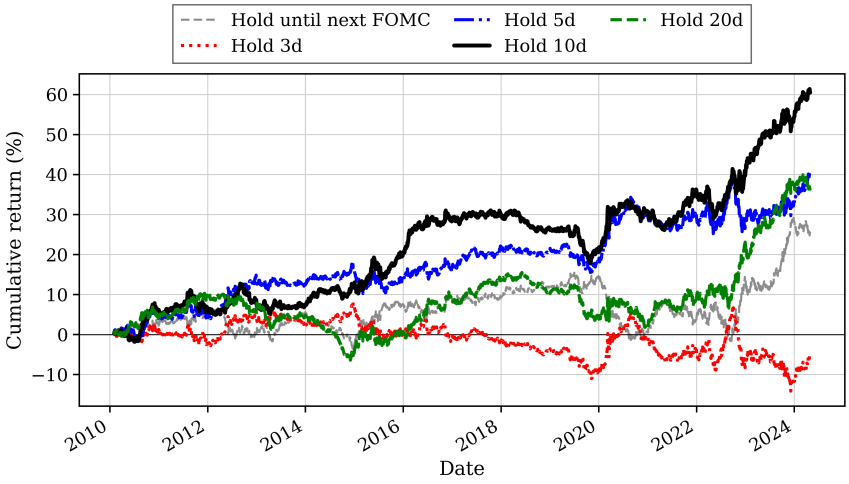


(b) Corporate (dur-adj.)

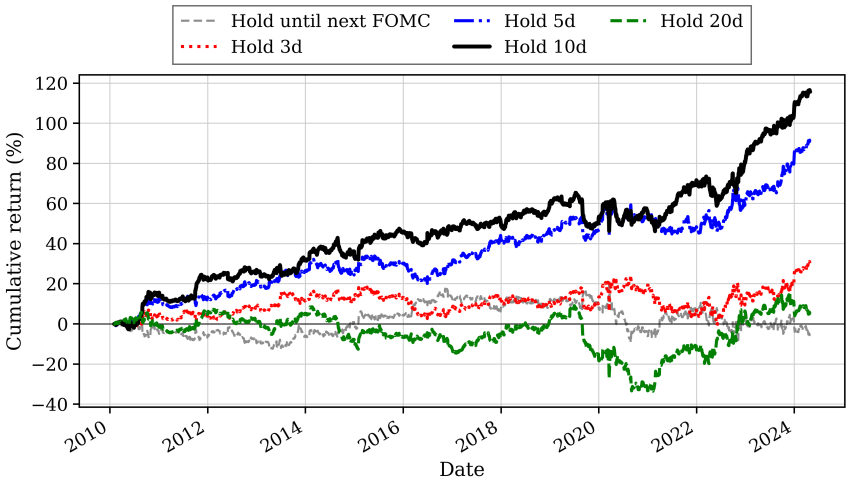
Figure 19: **Cumulative Return of the rebal_str Long-Short by Holding Horizon, Hold-and-Close Variant.** Position held from $t+1$ for $k \in \{3, 5, 10, 20\}$ trading days and then closed. Corporate returns are bond-level duration-adjusted (van Binsbergen, Nozawa, and Schwert, 2025); Treasury returns are raw. The dashed line is the "hold-until-next-FOMC" benchmark. The hold-and-reverse variant is reported in Table 9.

Averaged over the post-2010 window the panel holds about \$3.1T in market value and \$21T-years of dollar-duration. Panel A reports coverage by market segment; Panel B reports dollar-duration coverage by maturity, separately for Treasury and investment-grade corporate, which rises at the long end to about 28% for the two markets combined and 57% within investment-grade corporates, precisely where the channel concentrates (Section 5.1).

Figures 21 and 22 visualise the same coverage in two cuts: the quarterly time



(a) Treasury



(b) Corporate (dur-adj.)

Figure 20: **Cumulative Return of the rebal_str Long-Short, Hold-and-Reverse Variant.** Position entered at $t+1$, flipped after $k \in \{3, 5, 10, 20\}$ trading days, and held reversed until the next FOMC. Corporate returns are bond-level duration-adjusted (van Binsbergen, Nozawa, and Schwert, 2025); Treasury returns are raw. The dashed line is the "hold-until-next-FOMC" benchmark. The hold-and-close variant is in Figure 19; alphas for both variants are in Table 9.

series of market-value and dollar-duration coverage, which builds steadily through the 2000s and stabilises after 2010, and the dollar-duration coverage by maturity bin, which shows the long-end concentration reported in Panel B.

H Appendix: Posterior-Estimate Coverage Over Time

This appendix reports the time-series coverage of the fund-level posterior rebalancing sensitivities $\widehat{\beta}^{\text{reb}}_{f,t}$ defined in Section 4.3. Figure 23 plots the number of accounts with

Table 17: **Holdings Panel Coverage of the U.S. Bond Market (post-2010 average).** Coverage of the combined eMAXX and Morningstar panel by market value (MV) and by dollar-duration (DV = MV × modified duration), computed as the average of quarterly coverage ratios over the 2010Q1–2023Q4 window. Panel A is by market segment; Panel B is DV coverage by maturity bin, reported separately for Treasury and investment-grade corporate and for the two combined. Market references: U.S. marketable Treasury debt (SIFMA) and the investment-grade corporate market (Bloomberg US Aggregate Corporate Index), with maturity-bin reference weights from the Bloomberg index and Treasury MSPD composition. Reference magnitudes carry roughly ±10% relative uncertainty.

Panel A: Coverage by market segment (MV and DV)

Segment	MV coverage	DV coverage
Treasury	6.0%	6.9%
Investment-grade corporate	40.5%	37.9%
Combined Treasury + IG corporate	**14.9**%	**16.8**%

Panel B: Dollar-duration coverage by maturity

Maturity bin	Treasury	IG corporate	Combined
< 2 years	5.1%	38.5%	9.2%
2–5 years	7.5%	30.8%	14.4%
5–7 years	4.8%	32.3%	13.3%
7–12 years	4.7%	39.9%	15.4%
> 12 years	**12.2**%	**56.7**%	**28.0**%

Average of quarterly ratios over 2010Q1–2023Q4 (56 quarters); 2024Q2 onward excluded for partial holdings ingestion. A single end-of-sample snapshot understates the period-averaged coverage because late quarters have thinner ingestion.

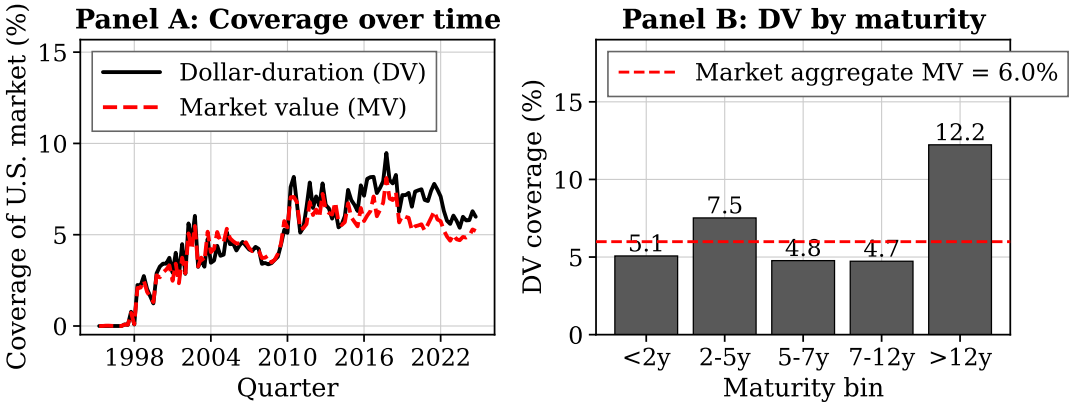


Figure 21: **Holdings Panel Coverage: Treasury.** Panel A: quarterly dollar-duration (solid black) and market-value (dashed red) coverage of the U.S. Treasury market, 1995–2024. Panel B: post-2010 average dollar-duration coverage by maturity bin, with the aggregate market-value coverage overlaid as a dashed reference line.

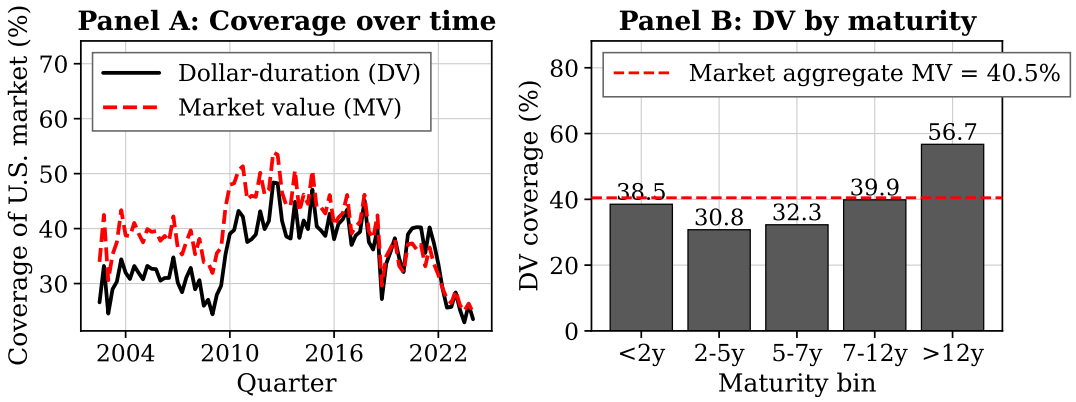
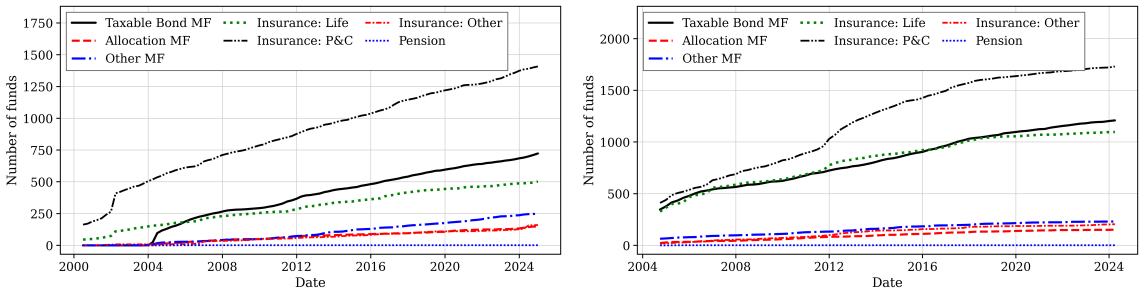


Figure 22: **Holdings Panel Coverage: Investment-Grade Corporate.** Panel A: quarterly dollar-duration (solid black) and market-value (dashed red) coverage of the investment-grade corporate market; the final two quarters, where holdings filings are still arriving, are omitted. Panel B: post-2010 average dollar-duration coverage by maturity bin, with the aggregate market-value coverage overlaid as a dashed reference line.

a non-missing $\widehat{\beta}^{\text{reb}}_{f,t}$ in a given quarter, by institution group. Coverage builds up in the early years as the within-type prior accumulates enough reliable observations to be informative, and stabilises by approximately 2010 in both markets. This is why the asset-pricing tests in Section 5 restrict to the post-2010 sample.



(a) Treasury market

(b) Corporate bond market

Figure 23: **Number of Accounts with Estimated Sensitivity Over Time, by Institution Group.** Each line tracks the number of fund-quarters with a non-missing posterior $\widehat{\beta}^{\text{reb}}_{f,t}$ in a given quarter. Coverage builds up over the early sample as the within-type prior accumulates reliable accounts and stabilises by around 2010.

I Appendix: Long-End Yield-Response Dampening Calibration

This appendix lists the structural constants entering the model-implied long-end yield-response dampening reported in Section 6.2 and derives the rebalancing incre-

ment η_r structurally from the habitat-pull panel.

I.1 Calibration of structural constants

Table 18 lists the structural constants entering the closed-form $A(\tau)$ in equation (10) and the supply-weighted moments $(M_{AA}, M_{\tau A}, M_{\tau\tau})$ used to compute the closed-form mean–variance optimum D_q^{MV} in Section 6.1.

Table 18: **Calibration Inputs for the Long-End Yield-Response Dampening.**

Symbol	Value	Source
κ^*	0.1445	AR(1) on quarterly ACMY01 ($\phi = 0.965$, $n = 260$
σ_r (annualized)	0.0172	AR(1) innovation $\times\sqrt{4}$
σ_0^{id}	0.010 (Treasury); 0.040 (Corporate)	idiosyncratic bond-return vol
γ_H	5.0	habitat-sector CARA risk aversion
$T_{\max}$	30 years	longest maturity in the bond panel
$M_{AA} = \int_0^{T_{\max}} A(u)^2\, du$	948.27 (Tres); 686.34 (Corp)	unweighted maturity moments, per-market κ^*
$M_{\tau A} = \int_0^{T_{\max}} u\, A(u)\, du$	2805.92 (Tres); 2363.05 (Corp)	same convention
$M_{\tau\tau} = \int_0^{T_{\max}} u^2\, du$	9000.00	κ-independent
$\Phi = M_{\tau\tau} - M_{\tau A}^2/(M_{AA} + {\sigma_0^{\mathrm{id}}}^2/\sigma_r^2)$	700.28 (Tres); 927.63 (Corp)	inverse of FOC curvature
$\bar{\omega}_{\mathrm{Tres}}$ (AUM-wt, 1,535 funds)	0.898	direct mean of per-fund OLS $\widehat{\omega}_f$
$\bar{\omega}_{\mathrm{Corp}}$ (AUM-wt, 3,332 funds)	0.887	direct mean of per-fund OLS $\widehat{\omega}_f$
η_r snapshot (Tres / Corp)	0.040 / 0.037	structural map $\bar{\omega}(\beta^{\mathrm{MV}}+1)f_H\kappa^*$ at $f_H = 0.5$, eq. (3
η_r^* fixed point (Tres / Corp)	0.030 / 0.031	self-consistent solution of the mapping; headline va
FOMC 1 SD shock	8 bps	Nakamura–Steinsson (2018); Bauer–Swanson (2023

The rebalancing increment η_r is mapped from the habitat-pull panel by a deliberately simple aggregation. The model's closed form (Appendix J, Proposition 7) expresses η_r through moment ratios and elasticities that are not separately observable at the fund level, so we use the heuristic map

$$\eta_r \;=\; \bar{\omega}\,(\beta^{\mathrm{MV}} + 1)\, f_H\, \kappa^*, \tag{39}$$

where $\bar{\omega}$ is the AUM-weighted habitat-pull strength, β^{MV} the mean-variance rebalancing baseline (Section 6.1), and $f_H \in (0, 1)$ the habitat sector's share of aggregate factor-duration risk. The map scales the benchmark absorption by the habitat share f_H and by the beyond-MV fraction of passive drift the sector reverses, $\bar{\omega}(\beta^{\mathrm{MV}} + 1) = |\beta^{\mathrm{reb}} - \beta^{\mathrm{MV}}|$. Two approximations are built in: it attributes the calibrated κ^* entirely to the supply-absorption term $\gamma_A \sigma_r^2 \Omega_0$ (i.e. treats κ_r as small relative to it, whereas the AR(1) estimate of κ^* blends the two), and it equates Macaulay-space rebalancing shares with factor-space elasticity reductions. The single free parameter f_H absorbs both. A benchmark of $f_H = 0.5$, an equal arbitrageur–habitat duration split as in standard Vayanos–Vila calibrations, gives $\eta_r = 0.040$ (Treasury) and 0.037 (IG corporate); the bracket $f_H \in [0.3, 0.7]$ gives $\eta_r \in [0.024, 0.057]$, over

which the mid-maturity match is preserved. None of these inputs uses the dampening regression, so the confrontation in Section 6.2 remains a genuine over-identification. The habitat-pull strength $\bar{\omega}$ is the AUM-weighted cross-fund mean of $\widehat{\omega}_f$ from the per-fund OLS two-variable regression in Section 6.1 (equation (33)).

Numerical fixed-point verification. The structural mapping in equation (39) evaluates the supply-elasticity moment $\Xi(\kappa) = \int_0^{T_{\max}} A(u;\kappa)^2 \, \bar{S}(u) \, du$ at the pre-shift benchmark κ^*, whereas in a fully self-consistent equilibrium Ξ is evaluated at the post-shift $\kappa_\eta = \kappa^* + \eta_r$, which itself depends on η_r. To verify that this feedback does not undo the calibration, I iterate the fixed point

$$\eta_r^{(k+1)} \;=\; \bar{\omega}\,(\beta^{\mathrm{MV}} + 1)\, f_H \, \frac{\Xi\big(\kappa_\eta^{(k)}\big)}{\Xi(\kappa^*)} \, \kappa^*, \qquad A^{(0)} = A_{BM},$$

so that the first iteration reproduces the structural snapshot and subsequent iterations impose self-consistency. Both markets converge in fewer than fifteen iterations, with no damping required, to $\eta_r^* = 0.030$ (Treasury) and 0.031 (IG corporate); Figure 24 traces the paths. The self-consistency correction shrinks η_r by about a quarter relative to the snapshot, in the direction the theory implies: a stronger effective mean reversion flattens the loading and shrinks the supply-elasticity moment. The converged values lie well inside the bracket $\eta_r \in [0.024, 0.057]$ spanned by $f_H \in [0.3, 0.7]$, over which the mid-maturity match of Section 6.2 is preserved. The paper's headline magnitudes use η_r^*; the snapshot evaluation brackets them from above, and in Treasury the self-consistent curve fits the regression slightly better, entering the confidence band at all fourteen maturity points against thirteen for the snapshot.

The intensity multiplier $R^B(\tau)$. The regression-implied dampening scales the per-signal interaction effect to market-wide intensity by $R^B(\tau) = \mathrm{rebal_str}(\tau)/\mathrm{coverage}_{DV}(\tau)$, the ratio of the holdings-weighted rebalancing-strength signal to the dollar-duration coverage in the maturity bin containing τ (bin midpoints $1, 3.5, 6, 9.5, 17$ years, linearly interpolated off-midpoint). The middle-maturity comparison is insensitive to this multiplier: perturbing R^B by $\pm 50\%$ leaves $D_{\mathrm{reg}}(10\mathrm{y})$ in $[5, 12]$ bps for Treasury and $[2, 6]$ bps for IG corporate, both bracketing the model values 7.3 and 5.6 bps. The short-end divergence, by contrast, is a pure multiplier artifact, as discussed in Section 6.2.

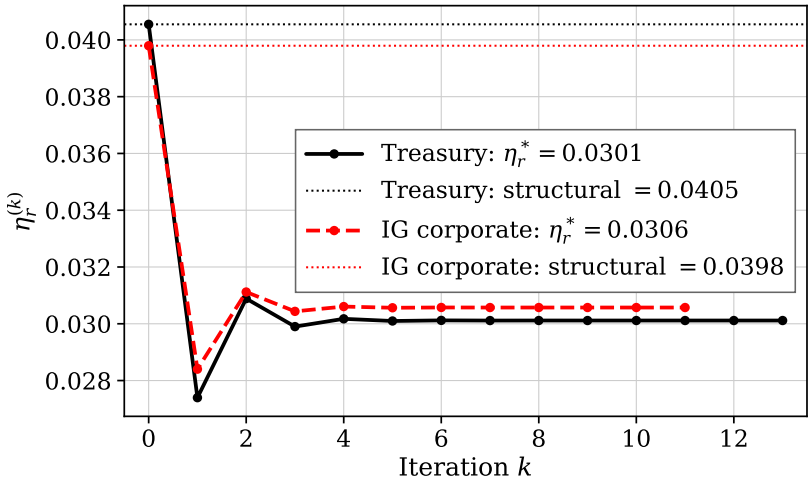


Figure 24: **Fixed-Point Convergence of the Rebalancing Increment η_r.** Iteration paths of $\eta_r^{(k)}$ for Treasury and IG corporate. Iteration 0 starts at the structural-mapping snapshot of equation (39) (dotted horizontal reference lines); the paths settle at the self-consistent values η_r^* reported in the legend within fifteen iterations.

I.2 Group-level empirical $\widehat{\beta}^{\text{reb}}$ distribution

Table 19 reports the group-level mean $\widehat{\beta}^{\text{reb}}$ within each institution group, computed from the fund-quarter posteriors of the Bayesian-shrinkage regression (24)–(27). Groups are sorted from least to most habitat-anchored within each market.

Table 19: **Group-Level Mean $\widehat{\beta}^{\text{reb}}$, by Institution Group.** Per-group rows report the cross-fund-quarter simple mean of the Bayesian-shrunk posterior $\widehat{\beta}^{\text{reb}}_{f,t}$, with each fund-quarter weighted equally. Groups are sorted from least to most habitat-anchored within each market. The summary block below additionally reports the AUM-weighted cross-fund average (the statistic that feeds the calibration in Section 6.2); the AUM-weighted average is computed over the subset of fund-quarters with valid TNA.

Treasury group	$\widehat{\beta}^{\text{reb}}$	Corporate group	$\widehat{\beta}^{\text{reb}}$
Other MF	−0.655	Insurance: Life	−0.787
Insurance: Life	−0.701	Pension	−0.816
Allocation MF	−0.701	Insurance: P&C	−0.816
Taxable Bond MF	−0.708	Taxable Bond MF	−0.830
Insurance: P&C	−0.722	Other MF	−0.846
Insurance: Other	−0.741	Allocation MF	−0.866
Pension	−0.912	Insurance: Other	−0.881
Cross-fund mean (full panel)	−0.708	*Cross-fund mean (full panel)*	−0.821
AUM-weighted mean	**−0.705**	*AUM-weighted mean*	**−0.786**
N fund-quarters	176,947		257,502

The cross-market asymmetry, Corporate $\widehat{\beta}^{\text{reb}}$ averaging more negative than Treasury (−0.79 AUM-weighted vs. −0.71), reflects institutional ownership composition.

Corporate ownership is more uniformly mandate-anchored across institution types: per-group means span a tight $[-0.88, -0.79]$ range in Corporate, while Treasury exhibits wider cross-group dispersion ranging from -0.66 (Other MF, the least mandate-anchored Treasury group) through -0.91 in Pension. Within Treasury, the most mandate-anchored groups are pension plans (whose Treasury holdings are predominantly liability-driven), the insurance groups, and discretionary mutual-fund vehicles whose duration positions vary more cross-sectionally.

I.3 Robustness of the dampening calibration

Alternative scaling of the regression-implied dampening. The baseline $D_{\text{reg}}(\tau)$ of equation (36) scales the regression coefficient by $\text{Share}^H(\tau)$, the habitat-sector share of the τ-maturity market, treating the average rebalancing intensity as embedded in the coefficient because the bond-level signal weights each holding by its fund's sensitivity. The alternative convention treats the coefficient as a response per unit of raw exposure and scales by $\text{Share}^H(\tau)$ times the panel's AUM-weighted average sensitivity (0.72 in Treasury, 0.79 in IG corporate; banks, outside the panel, proxied at the panel average). This more conservative scaling lowers D_{reg} by roughly a quarter and reduces the in-CI counts to 9/14 in Treasury and 6/14 in IG corporate. The two conventions answer different aggregation questions and bracket the market-wide magnitude; the middle-maturity agreement of the baseline, and the order-of-magnitude agreement of both, are the results the paper relies on. A second variant excludes banks, whose securities books are duration-managed against deposit franchises but sit outside the holdings panel. Dropping them lowers Share^H by five to six percentage points in Treasury and one to two in corporates, and lowers the regression-implied dampening by roughly a fifth in Treasury; the in-CI counts are unchanged at 14/14 and 8/14, so the comparison does not depend on where the habitat boundary is drawn.[25] Sector-by-maturity shares and their sources are in `habitat_share_sources`; the Z.1 quarterly variation of each share enters the confidence bands of Figure 9 by the delta method.

The dampening comparison of Section 6.2 rests on two market-specific inputs, κ^* and η_r, neither drawn from the dampening regression. To first order in η_r the hump's peak location is pinned by κ^* alone, at $\tau^\dagger \approx 1.79/\kappa^*$; η_r mainly scales the hump's height, moving the exact peak only modestly (from ≈ 12.4y at $\eta_r \to 0$ to ≈ 11.3y at

[25] Under the raw-exposure convention, measurement error in the fund-level sensitivities attenuates the observed signal and inflates the estimated coefficient by the inverse factor, so the coefficient must be paired with the *observed* average sensitivity rather than a full-rebalancing value of one; the two errors cancel in the product. Under the baseline convention the sensitivity never appears as a separate multiplier, so the issue does not arise.

$\eta_r = 0.03$ and ≈ 10.5y at the bracket top $\eta_r = 0.057$, all within the middle-maturity bins). Varying κ^* over the range delivered by the short-rate dynamics and η_r over the cross-fund $\widehat{\omega}_f$ range therefore leaves the middle-maturity agreement with the cross-sectional regression intact. A full (κ^*, η_r) sensitivity grid is reported in the replication materials.

J Appendix: Continuous-Time Model Derivations

This appendix presents the full continuous-time derivation, proofs, and approximations underlying the model summarized in Section 3.

J.1 Model Setup

We consider a frictional bond market in continuous time. Uncertainty is driven by a single systematic Brownian motion $(B_t)_{t\geq 0}$, which drives the short rate, and a family of idiosyncratic Brownian motions $\{Z_t(\tau)\}_{\tau\in(0,T]}$, mutually independent and independent of B_t, which generate bond-specific risk; all are defined on a filtered probability space satisfying the usual conditions.

J.1.1 Short rate and bond prices

The single aggregate state is the short rate r_t, an Ornstein–Uhlenbeck process under the physical measure $\mathbb{P}$:

$$dr_t = \kappa_r(\bar{r} - r_t)\,dt + \sigma_r\,dB_t, \tag{40}$$

with mean-reversion speed $\kappa_r > 0$, long-run mean $\bar{r}$, and volatility $\sigma_r > 0$. There is a continuum of zero-coupon bonds indexed by remaining maturity $\tau \in (0, T]$, with $P(t,\tau)$ the time-t price of one unit of the numéraire paid at $t+\tau$. The systematic part of the term structure is exponential-affine in $(r_t - \bar{r})$:

$$P^{\text{sys}}(t,\tau) \equiv \exp\big(-A(\tau)(r_t - \bar{r}) - C(\tau)\big), \tag{41}$$

where $A(\tau)$ is the loading of the log price on the short-rate deviation and $C(\tau)$ a deterministic level term, with $A(0) = C(0) = 0$ so that $P^{\text{sys}}(t,0) = 1$. $A(\tau)$ is the bond's *effective duration* with respect to the short-rate factor—a shock Δr_t changes the systematic log price by $-A(\tau)\Delta r_t$—whereas physical (Macaulay) duration of a zero-coupon bond is simply τ; the distinction between the two is central for the duration-rebalancing channel.

Idiosyncratic component as a multiplicative martingale To introduce bond-specific risk, each maturity carries an idiosyncratic multiplicative martingale $M_t(\tau)$ driven by a maturity-specific Brownian motion $Z_t(\tau)$,

$$\frac{dM_t(\tau)}{M_t(\tau)} = \sigma_\varepsilon(\tau)\,dZ_t(\tau), \qquad M_0(\tau) = 1, \tag{42}$$

with deterministic idiosyncratic volatility $\sigma_\varepsilon(\tau) \geq 0$. The traded price is the product of the systematic component and the martingale,

$$P(t,\tau) = P^{\text{sys}}(t,\tau)\, M_t(\tau), \tag{43}$$

so the log price is

$$p(t,\tau) \equiv \log P(t,\tau) = -A(\tau)(r_t - \bar{r}) - C(\tau) + m_t(\tau), \qquad m_t(\tau) \equiv \log M_t(\tau). \tag{44}$$

Following Vayanos and Vila (2021) and Greenwood, Hanson, and Vayanos (2024), we define $dP(t,\tau) \equiv P(t+dt, \tau - dt) - P(t,\tau)$, so along the life of a bond $d\tau = -dt$ and $dP(t,\tau)/P(t,\tau)$ is the instantaneous holding return.

Lemma 1 (Bond return dynamics)**.** *Under the short-rate dynamics* (40)*, the price decomposition* (43)*, and the roll-down convention $d\tau = -dt$,*

$$\frac{dP(t,\tau)}{P(t,\tau)} = \mu(t,\tau)\, dt + \sigma^{sys}(t,\tau)\, dB_t + \sigma^{id}(\tau)\, dZ_t(\tau), \tag{45}$$

where

$$\mu(t,\tau) \equiv A'(\tau)(r_t - \bar{r}) + C'(\tau) + \kappa_r A(\tau)(r_t - \bar{r}) + \frac{1}{2}A(\tau)^2\sigma_r^2, \tag{46}$$

$$\sigma^{sys}(t,\tau) \equiv -A(\tau)\sigma_r, \tag{47}$$

$$\sigma^{id}(\tau) \equiv \sigma_\varepsilon(\tau). \tag{48}$$

The instantaneous variance of the τ-maturity bond return is

$$\mathbb{V}ar_t\left[\frac{dP(t,\tau)}{P(t,\tau)}\right] = \left(A(\tau)^2\sigma_r^2 + \sigma_\varepsilon(\tau)^2\right) dt, \tag{49}$$

and for two maturities τ, τ' the instantaneous return covariance is

$$\mathbb{C}ov_t\left[\frac{dP(t,\tau)}{P(t,\tau)}, \frac{dP(t,\tau')}{P(t,\tau')}\right] = A(\tau)A(\tau')\sigma_r^2\, dt + \mathbf{1}_{\{\tau=\tau'\}}\sigma_\varepsilon(\tau)^2\, dt. \tag{50}$$

Proof. Two applications of Itô's lemma, both standard. Applied to (44) under $d\tau = -dt$, the log price has drift $A'(\tau)(r_t - \bar{r}) + C'(\tau) + \kappa_r A(\tau)(r_t - \bar{r}) - \frac{1}{2}\sigma_\varepsilon(\tau)^2$ and diffusions $-A(\tau)\sigma_r\, dB_t + \sigma_\varepsilon(\tau)\, dZ_t(\tau)$. Applied to $P = e^p$, the correction $\frac{1}{2}(dp)^2 = \frac{1}{2}\left(A(\tau)^2\sigma_r^2 + \sigma_\varepsilon(\tau)^2\right)dt$ cancels the Jensen drift of $m_t(\tau)$ exactly, leaving (46)–(48); the drift is therefore identical to the no-idiosyncratic-risk case. The variance and covariance follow from independence of B_t and the family $\{Z_t(\tau)\}$ across maturities (Vayanos and Vila, 2021). □

Remark 1 (Role of idiosyncratic risk)**.** Idiosyncratic volatilities add a diagonal $\sigma_\varepsilon(\tau)^2 dt$ term to the rank-one systematic covariance, making the covariance operator positive definite: even factor-neutral portfolios bear strictly positive variance. Portfolio-variance integrals $\int q(\tau)^2 \sigma^{\mathrm{id}}(\tau)^2 d\tau \cdot dt$ use the Dirac convention $\mathbb{E}[dZ_t(\tau) dZ_t(\tau')] = \delta(\tau - \tau') dt$, standard in continuum-of-bonds models (Vayanos and Vila, 2021).

J.2 Agents and Preferences

Two sectors trade the bonds: **preferred-habitat investors (PH)**, CARA optimizers over instantaneous dollar returns subject to a quadratic penalty on deviations of portfolio Macaulay duration from a target, and **arbitrageurs**, CARA mean–variance optimizers who provide risk-bearing capacity. Assumption 1 (asymmetric idiosyncratic risk bearing) applies throughout: PH investors bear idiosyncratic risk, $\sigma^{\mathrm{id}}(\tau) > 0$, while arbitrageurs run diversified books whose objective penalises only factor variance; Lemma 2 below shows this asymmetry is what gives the mandate a price effect at all.

Assumption 4 (Stable sector scale)**.** *The PH sector's reference portfolio scale is held constant at $\bar{W}$, with implicit inflows and outflows offsetting return-induced wealth fluctuations; arbitrageurs optimize directly over dollar positions and need no analogous parameter.*

Under CARA preferences, optimal dollar positions are wealth-independent, so Assumption 4 is a notational simplification: the FOCs and pricing implications coincide with a fully wealth-dynamic specification. Throughout, $(\mu(t,\tau), \sigma^{\mathrm{sys}}(t,\tau), \sigma^{\mathrm{id}}(\tau))$ denote the drift, systematic, and idiosyncratic volatilities of (46)–(48).

J.2.1 Portfolios and instantaneous risk

Let $x_t(\tau)$ and $y_t(\tau)$ denote the dollar position densities of Arbitrageurs and PH investors at maturity τ ($x_t(\tau)\, d\tau$ is the dollar amount invested in maturities $[\tau, \tau + d\tau]$), with effective-duration exposures

$$V_{A,t} \equiv \int_0^T x_t(u)\, A(u)\, du, \tag{51}$$

$$V_{H,t} \equiv \int_0^T y_t(u)\, A(u)\, du. \tag{52}$$

A sector holding the profile $\{q_t(\tau)\}$ has systematic dollar loading $-\sigma_r V_t$ on B_t, with $V_t \equiv \int_0^T q_t(\tau) A(\tau) d\tau$, and idiosyncratic loading $q_t(\tau)\sigma^{\mathrm{id}}(\tau)$ at each maturity; by

independence, the variance of its dollar returns is

$$\Sigma_t[q] \;\equiv\; \mathrm{Var}_t\!\left[\int_0^T q_t(\tau)\,\frac{dP(t,\tau)}{P(t,\tau)}\right]\Big/dt \;=\; \sigma_r^2\,V_t^2 \;+\; \int_0^T q_t(\tau)^2\,\sigma^{\mathrm{id}}(\tau)^2\,d\tau. \tag{53}$$

For PH investors ($q = y$), the full expression applies: $\Sigma_{H,t} = \sigma_r^2 V_{H,t}^2 + \int_0^T y_t(\tau)^2 \sigma^{\mathrm{id}}(\tau)^2 d\tau$. For Arbitrageurs, Assumption 1 retains only the factor component, $\Sigma_{A,t} = \sigma_r^2 V_{A,t}^2$; the symmetric expression $\Sigma_{A,t}^{\mathrm{sym}} = \sigma_r^2 V_{A,t}^2 + \int_0^T x_t(\tau)^2 \sigma^{\mathrm{id}}(\tau)^2 d\tau$ is used only in the variant of Appendix J.9.

J.2.2 Arbitrageurs

Arbitrageurs choose dollar positions $\{x_t(\tau)\}$ to maximise instantaneous CARA mean–variance utility over the dollar returns from their risky portfolio:

$$\max_{\{x_t(\tau)\}} \;\mathcal{U}_A \;=\; \int_0^T x_t(\tau)\big(\mu(t,\tau) - r_t\big)\,d\tau \;-\; \frac{\gamma_A}{2}\,\Sigma_{A,t}, \tag{54}$$

where $\gamma_A > 0$ is absolute risk aversion. Under Assumption 4, the risk-free return on any unused arbitrageur capital is a constant in $\{x_t(\tau)\}$ and does not affect the optimum, so it is omitted.

Under Assumption 1, $\Sigma_{A,t} = \sigma_r^2 V_{A,t}^2$, and the first-order condition with respect to $x_t(\tau)$, using $\partial V_{A,t}/\partial x_t(\tau) = A(\tau)$, is

$$\boxed{\mu(t,\tau) - r_t = \gamma_A \sigma_r^2 A(\tau)\,V_{A,t}} \tag{55}$$

for every $\tau \in (0, T]$: compensation for systematic short-rate risk only, proportional to effective duration $A(\tau)$ and the aggregate factor exposure $V_{A,t}$. Because (55) pins down only the scalar $V_{A,t}$, the arbitrageur is indifferent across factor-neutral reshufflings of the book; the profile $x_t(\tau)$ is determined residually by pointwise market clearing. In the *symmetric* variant in which the arbitrageur also bears idiosyncratic variance ($\Sigma_{A,t}^{\mathrm{sym}}$; Appendix J.9), the demand condition instead reads

$$\mu(t,\tau) - r_t = \gamma_A \sigma_r^2 A(\tau)\,V_{A,t} + \gamma_A \sigma^{\mathrm{id}}(\tau)^2\,x_t(\tau), \tag{56}$$

with the additional term compensating the idiosyncratic risk of the local position $x_t(\tau)$.

J.2.3 Preferred-habitat investors

The PH sector is a continuum of investors $i \in I = [0, 1]$ (mass-1 normalisation) with heterogeneous types $\{(D_i^*, \delta_i)\}$—a personal target duration and a personal penalty intensity—and common $(\gamma_H, \bar{W}, \sigma^{\mathrm{id}}(\tau))$. Heterogeneity is introduced at the setup

because the empirical strategy relies on it: it generates the cross-fund dispersion in β_f^{reb}, while the aggregate pricing equation keeps its affine form since individual FOCs are linear and aggregate to a single δ-weighted scalar (Section J.4.1). The mass-1 normalisation makes integrals over i both averages and aggregates; $|I| = N$ would rescale $(\gamma_H, \delta_i, \bar{W})$ in the standard CARA sense (Wilson, 1968) without changing any pricing object.

Investor i's Macaulay portfolio duration, normalised by the reference scale, is

$$D^i_{\text{port},t} \;=\; \frac{1}{\bar{W}} \int_0^T y^i_t(\tau)\, \tau \, d\tau; \tag{57}$$

with the residual $\bar{W} - \int y^i_t \, du$ implicitly in zero-duration cash, this is the Macaulay duration of the total portfolio. She chooses $\{y^i_t(\tau)\}$ to maximise CARA mean–variance utility augmented by a quadratic penalty on deviations from her target:

$$\max_{\{y^i_t(\tau)\}} \; \mathcal{U}^i_H \;=\; \int_0^T y^i_t(\tau)\big(\mu(t,\tau) - r_t\big)\, d\tau \;-\; \frac{\gamma_H}{2}\, \Sigma^i_{H,t} \;-\; \frac{\delta_i}{2}\big(D^i_{\text{port},t} - D^*_i\big)^2\, \bar{W}, \tag{58}$$

with $\Sigma^i_{H,t} = \sigma_r^2 (V^i_{H,t})^2 + \int_0^T y^i_t(\tau)^2 \sigma^{\text{id}}(\tau)^2 d\tau$, $V^i_{H,t} \equiv \int_0^T y^i_t(u) A(u) du$; the factor $\bar{W}$ keeps the penalty in dollar-utility units. As in (54), the risk-free return on unused capital is a constant in the choice variables and is omitted.

First-order condition Differentiating (58) with respect to $y^i_t(\tau)$, using $\partial V^i_{H,t}/\partial y^i_t(\tau) = A(\tau)$, $\partial D^i_{\text{port},t}/\partial y^i_t(\tau) = \tau/\bar{W}$, and $\mathcal{D}^i_{\text{dev},t} \equiv D^i_{\text{port},t} - D^*_i$, yields four contributions: the mean $\mu(t,\tau) - r_t$; the systematic variance $-\gamma_H \sigma_r^2 A(\tau) V^i_{H,t}$; the idiosyncratic variance $-\gamma_H \sigma^{\text{id}}(\tau)^2 y^i_t(\tau)$; and the duration penalty $-\delta_i \bar{W}\, \mathcal{D}^i_{\text{dev},t} \cdot (\tau/\bar{W}) = -\delta_i \tau\, \mathcal{D}^i_{\text{dev},t}$, where $\bar{W}$ cancels because the penalty scales with $\bar{W}$ while duration is normalised by it. Setting the sum to zero gives the per-investor PH demand condition:

$$\boxed{\mu(t,\tau) - r_t = \gamma_H\, \sigma_r^2\, A(\tau)\, V^i_{H,t} + \gamma_H\, \sigma^{\text{id}}(\tau)^2\, y^i_t(\tau) + \delta_i\, \tau\, \mathcal{D}^i_{\text{dev},t}} \tag{59}$$

for every $\tau \in (0, T]$ and every $i \in I$: mean–variance compensation for systematic and idiosyncratic risk, plus the compensation investor i requires for bonds that move her Macaulay duration away from target. The LHS does not depend on i and the RHS is linear in $(V^i_{H,t}, y^i_t(\tau), \delta_i \mathcal{D}^i_{\text{dev},t})$, which makes within-PH aggregation exact below.

J.3 Supply and Market Clearing

J.3.1 Fixed face-value supply

For each maturity $\tau \in [0, T]$, let $\bar{N}(\tau)$ denote the exogenous, fixed face value of outstanding bonds maturing at $t + \tau$. With $\bar{S}(\tau) \equiv \bar{N}(\tau) e^{-C(\tau)}$ the steady-state

market value, the market value of supply is

$$S_t(\tau) = \bar{N}(\tau)P(t,\tau) = \bar{S}(\tau)\exp\big(-A(\tau)(r_t-\bar{r})\big) \;\approx\; \bar{S}(\tau)\Big(1-A(\tau)(r_t-\bar{r})\Big), \quad (60)$$

the last step to first order in $(r_t - \bar{r})$: when rates rise, the market value of a fixed face-value supply contracts in proportion to its effective duration.

Assumption 5 (Local linearisation around the steady state)**.** *We work locally around $\bar{r}$ and use the first-order approximation in* (60)*; all aggregate-supply expressions are understood in this log-linear sense.*

Aggregate effective-duration supply is $\Theta_t \equiv \int_0^T S_t(u)A(u)\,du$; substituting (60),

$$\Theta_t \approx \bar{\Theta} - (r_t - \bar{r})\,\Xi, \qquad \bar{\Theta} \equiv \int_0^T \bar{S}(u)A(u)\,du, \qquad \Xi \equiv \int_0^T \bar{S}(u)A(u)^2\,du, \quad (61)$$

so duration supply is affine in $(r_t - \bar{r})$, contracting at speed $\Xi > 0$ as rates rise.

Market clearing requires, for each maturity,

$$x_t(\tau) + y_t(\tau) = S_t(\tau). \quad (62)$$

Weighting (62) by $A(\tau)$ and integrating over $\tau \in [0,T]$ gives the duration-space market-clearing condition

$$V_{A,t} + V_{H,t} = \Theta_t, \quad (63)$$

with $V_{A,t}, V_{H,t}$ as in (51)–(52) and Θ_t as in (61): the effective duration supplied by outstanding bonds must be absorbed by the two sectors' effective-duration demands. Idiosyncratic risk does not enter (63); it affects the variance of holdings, not the market value to be absorbed.

J.4 Aggregating Risk Premia and the Rebalancing Wedge

J.4.1 Within-PH aggregation: from per-investor FOC to sector FOC

The per-investor FOC (59) has an i-independent LHS and an RHS linear in $(V^i_{H,t}, y^i_t(\tau), \delta_i \mathcal{D}^i_{\mathrm{dev},t})$, so integrating over I yields a sector FOC of identical form. With $y_t(\tau) \equiv \int_I y^i_t\,di$, $V_{H,t} \equiv \int_I V^i_{H,t}\,di$, $\bar{\delta} \equiv \int_I \delta_i\,di$, and the δ-weighted sector deviation $\widetilde{\mathcal{D}}_{\mathrm{dev},t} \equiv \bar{\delta}^{-1}\int_I \delta_i \mathcal{D}^i_{\mathrm{dev},t}\,di$,

$$\mu(t,\tau) - r_t = \gamma_H \sigma_r^2 A(\tau)\, V_{H,t} + \gamma_H \sigma^{\mathrm{id}}(\tau)^2\, y_t(\tau) + \bar{\delta}\,\tau\,\widetilde{\mathcal{D}}_{\mathrm{dev},t}.$$

Heterogeneity in (D^*_i, δ_i) thus enters the aggregate pricing equation only through the single scalar $\widetilde{\mathcal{D}}_{\mathrm{dev},t}$; it drives the cross-section of individual demands and rebalancing intensities in Section J.8, not the aggregate price.

J.4.2 A symmetric benchmark: the aggregate three-term relation

This subsubsection drops Assumption 1 and lets both sectors bear idiosyncratic risk, so the arbitrageur condition is the symmetric form (56). The resulting relation underlies the priced-idiosyncratic variant of Appendix J.9 and the general drift of Section J.8; the main-specification pricing below does not use it. Multiplying the two FOCs by $1/\gamma_A$ and $1/\gamma_H$ respectively, summing, and using duration-space market clearing $V_{A,t}+V_{H,t} = \Theta_t$ together with pointwise market clearing $x_t(\tau)+y_t(\tau) = S_t(\tau)$:

$$\tau_{\text{agg}}\big(\mu(t,\tau) - r_t\big) = \sigma_r^2 A(\tau)\Theta_t + \sigma^{\text{id}}(\tau)^2 S_t(\tau) + \frac{\bar{\delta}}{\gamma_H}\tau\,\widetilde{\mathcal{D}}_{\text{dev},t}, \tag{64}$$

where $\tau_{\text{agg}} \equiv 1/\gamma_A + 1/\gamma_H$ is the aggregate risk tolerance: a factor term in $A(\tau)$, an idiosyncratic-supply term, and a duration-penalty term linear in τ.

J.4.3 Equilibrium pricing under asymmetric risk bearing

Under Assumption 1 the equilibrium drift follows *exactly* from the arbitrageur condition (55): expected excess returns load only on $A(\tau)$, with no approximation and no dropped terms. The duration penalty, whose marginal value is linear in τ, does not enter the price at any maturity; consistently, an unspanned component would admit asymptotic arbitrage as the maturity span grows.

Proposition 5 (No unspanned wedge)**.** *Let $W(\tau)$ be a deterministic component of $\mu(t,\tau) - r_t$, and for maturities $\tau_1, \ldots, \tau_N$ write $W = cA + W^\perp$ for its decomposition along and orthogonal to $A = (A(\tau_i))_i$.*

(i) If idiosyncratic risk is unpriced (fully diversified away), any $W^\perp \neq 0$ supports a factor-neutral portfolio with strictly positive expected excess return and vanishing risk—a riskless arbitrage. Equilibrium then requires $W(\tau) \propto A(\tau)$ exactly.

(ii) If idiosyncratic risk is priced, with variances bounded by $\bar{\sigma}^2_{\text{id}}$, the factor-neutral portfolio $w = W^\perp$ attains a Sharpe ratio of at least $\|W^\perp\|/\bar{\sigma}_{\text{id}}$. For $W(\tau) = \eta\tau$ with $\eta \neq 0$ on an unbounded maturity domain, $\|W^\perp\| \to \infty$ (since A is bounded while τ is not) and the attainable Sharpe ratio diverges; a finite maximal Sharpe ratio rules out the τ-linear wedge as the maturity span grows. On a bounded range, a bounded wedge is consistent with priced idiosyncratic risk (Appendix J.9).

Proof. Impose factor-neutrality $\sum_i w_i A(\tau_i) = 0$, eliminating exposure to dB_t; the portfolio's expected excess return is $\sum_i w_i W(\tau_i)$ and its variance is at most $\bar{\sigma}^2_{\text{id}} \sum_i w_i^2$

(zero when idiosyncratic risk is diversified away). Taking $w = W^{\perp}$: factor-neutrality holds by construction, the expected excess return equals $\|W^{\perp}\|^2$, and the standard deviation is at most $\bar{\sigma}_{\mathrm{id}}\|W^{\perp}\|$. (i) With unpriced idiosyncratic risk the portfolio has positive drift and no risk unless $W^{\perp} = 0$. (ii) With priced idiosyncratic risk the construction attains Sharpe at least $\|W^{\perp}\|/\bar{\sigma}_{\mathrm{id}}$. For $W(\tau) = \eta\tau$, the angle between $(\tau_i)_i$ and $(A(\tau_i))_i$ stays bounded away from zero as the span grows (A saturates at $1/\kappa_\eta$ while τ grows linearly), so $\|W^{\perp}\|^2$ grows in proportion to $\eta^2 \sum_i \tau_i^2$ and the Sharpe ratio diverges. □

This is the cross-sectional restriction of Ross (1976) and Huberman (1982): expected excess returns are spanned by the factor loadings up to diversifiable idiosyncratic terms. The duration mandate therefore reaches prices only through the sector's *factor duration* $V_{H,t}$ of (52). Substituting duration-space clearing $V_{A,t} = \Theta_t - V_{H,t}$ from (63) into (55):

$$\boxed{\mu(t,\tau) - r_t \;=\; \lambda_t\, A(\tau), \qquad \lambda_t \;=\; \gamma_A\, \sigma_r^2\big(\Theta_t - V_{H,t}\big).} \tag{65}$$

This is the canonical Vayanos–Vila pricing relation (their Lemma 1): expected excess returns load *only* on the effective-duration loading $A(\tau)$, with the price of factor risk set by the arbitrageurs' risk aversion times the duration they absorb. The market price of risk carries γ_A rather than a pooled sector risk tolerance because arbitrageurs alone bear the residual factor risk. (As $\delta_i \to 0$ the habitat sector holds its mean–variance share and $\gamma_A\sigma_r^2(\bar{\Theta} - \bar{V}_H) = \sigma_r^2\bar{\Theta}/\tau_{\mathrm{agg}}$ recovers the usual combined-tolerance term premium; with active mandates both the level and, more importantly, the state-sensitivity of λ_t shift.)

Remark 2 (Well-posedness)**.** Because PH investors bear idiosyncratic risk ($\sigma^{\mathrm{id}} > 0$), each investor's problem is strictly concave in the density $y_t^i(\cdot)$ pointwise, so the habitat allocation—and with it $V_{H,t}$ and the sector's Macaulay duration—is uniquely determined (the 2×2 characterisation is in Section J.8). Only the arbitrageur's profile $x_t(\tau)$ is indifferent across factor-neutral reshufflings, and it is pinned residually by $x_t(\tau) = S_t(\tau) - y_t(\tau)$.

Lemma 2 (Mandate neutrality under fully unpriced idiosyncratic risk)**.** *Suppose, contrary to Assumption 1, idiosyncratic risk is absent from* both *sectors' objectives ($\sigma^{id} \equiv 0$). Then in equilibrium $\mu(t,\tau) - r_t = \lambda_t A(\tau)$ with $V_{H,t}^i = \lambda_t/(\gamma_H\sigma_r^2)$ and $D_{port,t}^i = D_i^*$ for every PH investor. Hence $V_{H,t} = \lambda_t/(\gamma_H\sigma_r^2)$, $\lambda_t = \sigma_r^2\Theta_t/\tau_{agg}$, $\Omega = \Omega_0$, and $\eta_r = 0$: the duration mandate has no effect on equilibrium prices.*

Proof. With $\sigma^{\mathrm{id}} \equiv 0$ the arbitrageur FOC forces $\mu - r = \lambda_t A(\tau)$ exactly. Investor i's objective then depends on $y^i(\cdot)$ only through the scalars $V^i = \int_0^T y^i A\,du$ and $D^i = (1/\bar{W}) \int_0^T y^i u\,du$:

$$\mathcal{U}_H^i = \lambda_t V^i - \tfrac{\gamma_H}{2}\sigma_r^2 (V^i)^2 - \tfrac{\delta_i}{2}(D^i - D_i^*)^2 \bar{W}.$$

Since $A(\cdot)$ and $u \mapsto u$ are linearly independent on $(0, T]$, the pair (V^i, D^i) ranges over all of $\mathbb{R}^2$ as y^i varies, so the optimum separates: $V^i = \lambda_t/(\gamma_H \sigma_r^2)$ and $D^i = D_i^*$, each attained at zero marginal cost to the other (the achieving density is indeterminate). Aggregating and combining with $\lambda_t = \gamma_A \sigma_r^2 (\Theta_t - V_{H,t})$ gives $\lambda_t(1 + \gamma_A/\gamma_H) = \gamma_A \sigma_r^2 \Theta_t$, i.e. $\lambda_t = \sigma_r^2 \Theta_t / \tau_{\mathrm{agg}}$, whose state-sensitivity is the mean–variance value: $\Omega = \Xi/(\gamma_A \tau_{\mathrm{agg}}) = \Omega_0$ and $\eta_r = 0$. □

The logic is simple: when idiosyncratic risk is costless to hold, the sector meets its Macaulay target purely through the *shape* of its holdings, at no cost in factor space, so the mandate never touches the priced quantity. With $\sigma^{\mathrm{id}} > 0$ shape adjustments are costly, the two durations are coupled, and Proposition 7 below shows the coupling delivers $\eta_r > 0$, vanishing as $\sigma^{\mathrm{id}} \to 0$. The non-standard alternative, in which per-bond idiosyncratic risk is *priced* by both sectors and a τ-linear wedge survives on a bounded maturity range, is recorded in Appendix J.9.

J.5 Fundamental Pricing ODE

We now derive the ODE satisfied by the effective-duration loading $A(\tau)$ from the equilibrium pricing relation (65),

$$\mu(t, \tau) - r_t = \lambda_t\, A(\tau), \qquad \lambda_t \;=\; \gamma_A\, \sigma_r^2 \big(\Theta_t - V_{H,t}\big). \tag{66}$$

Since $A(\tau)$ is by construction the coefficient on $(r_t - \bar{r})$ in the log price, deriving its ODE requires matching coefficients of $(r_t - \bar{r})$ on both sides; the only relevant input is the *state-sensitivity* of λ_t, equivalently of the residual factor duration $\Theta_t - V_{H,t}$ that arbitrageurs bear, not its snapshot level.

The benchmark is the textbook Vasicek loading: with risk-neutral mean reversion κ^*, A_{BM} solves $A'_{BM} = 1 - \kappa^* A_{BM}$ with $A_{BM}(0) = 0$, i.e. $A_{BM}(\tau) = (1 - e^{-\kappa^* \tau})/\kappa^*$. In a one-factor Gaussian model all systematic risk premia are summarised by an effective risk-neutral mean reversion; we treat (κ^*, η_r) as reduced-form parameters that already incorporate them. The only state-dependent object in (66) is the residual factor duration $\Theta_t - V_{H,t}$, which we linearise around $\bar{r}$ before matching coefficients.

Assumption 6 (Local linearisation around the steady state)**.** *For small deviations of r_t from $\bar{r}$,*

$$\Theta_t - V_{H,t} \;\approx\; \big(\bar{\Theta} - \bar{V}_H\big) \;-\; \Omega\,(r_t - \bar{r}), \qquad \Omega \equiv -\frac{\partial(\Theta_t - V_{H,t})}{\partial r_t}\bigg|_{r_t=\bar{r}}, \tag{67}$$

the state-sensitivity of the arbitrageurs' residual factor duration. Ω combines the supply elasticity Ξ of (61) *with the habitat sector's rebalancing of $V_{H,t}$; in a fully self-consistent equilibrium it depends on $A(\cdot)$ and the joint distribution of (D_i^*, δ_i), a fixed point.*

J.5.1 Coefficient matching: deriving the ODE for $A(\tau)$

Substituting the Itô drift formula (46) into the equilibrium pricing relation (66), and writing $r_t = (r_t - \bar{r}) + \bar{r}$:

$$\big[A'(\tau) + \kappa_r A(\tau) - 1\big](r_t - \bar{r}) + \big[C'(\tau) + \tfrac{1}{2}A(\tau)^2\sigma_r^2 - \bar{r}\big] = \lambda_t A(\tau), \tag{68}$$

with $\lambda_t A(\tau) = \gamma_A \sigma_r^2\big[(\bar{\Theta} - \bar{V}_H) - \Omega(r_t - \bar{r})\big]A(\tau)$ from (67). Matching the coefficient of $(r_t - \bar{r})$ on both sides of (68) gives

$$A'(\tau) + \big(\kappa_r + \gamma_A \sigma_r^2 \Omega\big)A(\tau) = 1. \tag{69}$$

Write $\kappa_\eta \equiv \kappa_r + \gamma_A\sigma_r^2\Omega$. Absent the duration mandate the habitat holds its mean–variance optimum, the residual sensitivity takes a benchmark value Ω_0, and κ_η reduces to the standard term-premium mean reversion

$$\kappa^* \;\equiv\; \kappa_r + \gamma_A\sigma_r^2\Omega_0 \;=\; \kappa_r + \frac{\sigma_r^2\Xi}{\tau_{\text{agg}}}, \tag{70}$$

the usual combined-tolerance absorption of the supply elasticity (at the mean–variance share $\gamma_A(\bar{\Theta} - \bar{V}_H^{\text{MV}}) = \bar{\Theta}/\tau_{\text{agg}}$).[26] The duration mandate anchors the sector's factor demand—after a rate cut it expands factor duration by less, and after a hike it contracts it by less, than mean–variance sharing would dictate, selling into rallies and buying into sell-offs relative to the benchmark—making the residual $\Theta_t - V_{H,t}$ *more* state-sensitive, i.e. raising Ω above Ω_0 and contributing

$$\eta_r \;\equiv\; \gamma_A\sigma_r^2(\Omega - \Omega_0) \;>\; 0, \qquad \text{so that} \qquad \kappa_\eta = \kappa^* + \eta_r. \tag{71}$$

Assumption 7 (Uniform idiosyncratic supply-volatility profile)**.** *The supply-weighted idiosyncratic variance is approximately constant across maturities, $\sigma^{id}(\tau)^2\,\bar{S}(\tau) \approx$*

[26] The identity is exact as $\sigma^{\text{id}} \to 0$; the $O(\sigma^{\text{id}\,2})$ correction to Ω_0 is absorbed in the reduced-form κ^*. In the symmetric variant of Appendix J.9, Assumption 7 adds a further small additive term $\bar{\xi} \equiv \overline{\sigma^{\text{id}\,2}\bar{S}}/\tau_{\text{agg}}$ to κ^*; under Assumption 1 that term is absent.

$\overline{\sigma^{id\,2}\bar{S}}$ *for* $\tau \in (0, T]$*, of similar order of magnitude across the corporate and Treasury benchmark curves. In the symmetric variant of Appendix J.9, where idiosyncratic supply is priced, this contributes the* τ*-independent additive scalar* $\bar{\xi} \equiv \overline{\sigma^{id\,2}\bar{S}}/\tau_{agg}$ *that folds into* κ^* *via* $\kappa^* \to \kappa^* + \bar{\xi}$ *and keeps the closed-form Vasicek-style solution. Under Assumption 1 it is not needed.*

Proposition 6 (Fundamental pricing ODE)**.** *Under Assumptions 1 and 6, the structural loading* $A(\tau)$ *in the affine pricing formula* (41) *satisfies*

$$\boxed{A'(\tau) + \kappa_\eta\, A(\tau) = 1, \qquad A(0) = 0, \qquad \kappa_\eta = \kappa^* + \eta_r.} \tag{72}$$

The rebalancing channel enters only through the effective mean reversion: it raises the benchmark κ^* *defined in* (70) *to* κ_η *by the increment* $\eta_r > 0$ *defined in* (71)*. When* $\eta_r = 0$*, equation* (72) *collapses to the benchmark Vasicek ODE with loading* $A_{BM}(\tau) = (1 - e^{-\kappa^*\tau})/\kappa^*$*.*

The increment η_r is a state-sensitivity, not a snapshot level: the level $(\bar{\Theta} - \bar{V}_H)$ feeds only the steady-state term premium and the ODE for $C(\tau)$, while Ω alone enters the ODE for $A(\tau)$. Given the loading $A(\cdot)$, the linearised equilibrium delivers Ω in closed form.

Proposition 7 (Closed form for Ω and the unconditional sign of η_r)**.** *Let Assumptions 1 and 6 hold. Define the* σ^{id}*-weighted moments*

$$\tilde{I}_{AA} \equiv \int_0^T \frac{A(u)^2}{\sigma^{id}(u)^2} du, \quad \tilde{I}_{A\tau} \equiv \int_0^T \frac{A(u)\,u}{\sigma^{id}(u)^2} du, \quad \tilde{I}_{\tau\tau} \equiv \int_0^T \frac{u^2}{\sigma^{id}(u)^2} du, \quad \tilde{\Phi} \equiv \tilde{I}_{\tau\tau} - \frac{\tilde{I}_{A\tau}^2}{\tilde{I}_{AA} + 1/\sigma_r^2},$$

and each investor's λ-elasticity *of factor demand* $\beta_i \equiv \partial V_H^i / \partial \lambda$*. Then:*

(i) $\beta_i = \dfrac{\gamma_H \bar{W}\, \tilde{I}_{AA} + \delta_i\big(\tilde{I}_{AA}\tilde{I}_{\tau\tau} - \tilde{I}_{A\tau}^2\big)}{\gamma_H\big(1 + \sigma_r^2 \tilde{I}_{AA}\big)\big(\gamma_H \bar{W} + \delta_i \tilde{\Phi}\big)} > 0$*, independent of* D_i^* *and* strictly decreasing *in* δ_i*: the mandate anchors factor demand, lowering the elasticity to the price of factor risk.*

(ii) $\Omega = \dfrac{\Xi}{1 + \gamma_A \sigma_r^2 B}, \qquad B \equiv \displaystyle\int_I \beta_i\, di$*, strictly decreasing in* B*. Hence* $\eta_r = \gamma_A \sigma_r^2 (\Omega - \Omega_0) > 0$ *whenever* $\delta_i > 0$ *on a set of positive measure—with no restriction on the targets* $\{D_i^*\}$*, the holdings, or the maturity profile of* σ^{id}*—and* η_r *is strictly increasing in each* δ_i*, saturating as* $\delta_i \to \infty$*.*

(iii) As $\sigma^{id} \to 0$*,* $\beta_i \to 1/(\gamma_H \sigma_r^2)$ *for every* δ_i*, so* $\eta_r \to 0$ *(Lemma 2); for constant idiosyncratic volatility,* $\eta_r = O(\sigma^{id2})$*.*

Proof. Fix the state. Under (65) each investor's demand $y^i(\tau) = [(\lambda - \gamma_H \sigma_r^2 V^i) A(\tau) - \delta_i \tau \mathcal{D}^i_{\text{dev}}]/(\gamma_H \sigma^{\text{id}}(\tau)^2)$ must satisfy the self-consistency conditions $V^i = \int y^i A \, du$ and $\bar{W}(\mathcal{D}^i_{\text{dev}} + D^*_i) = \int y^i u \, du$, a 2×2 linear system in $(V^i, \mathcal{D}^i_{\text{dev}})$:

$$\gamma_H(1+\sigma_r^2 \tilde{I}_{AA}) V^i + \delta_i \tilde{I}_{A\tau} \mathcal{D}^i_{\text{dev}} = \lambda \tilde{I}_{AA}, \qquad \gamma_H \sigma_r^2 \tilde{I}_{A\tau} V^i + (\gamma_H \bar{W} + \delta_i \tilde{I}_{\tau\tau}) \mathcal{D}^i_{\text{dev}} = \lambda \tilde{I}_{A\tau} - \gamma_H \bar{W} D^*_i,$$

with determinant $\gamma_H(1+\sigma_r^2 \tilde{I}_{AA})(\gamma_H \bar{W} + \delta_i \tilde{\Phi})$. Cramer's rule gives V^i affine in λ with slope β_i as in (i); the numerator is positive by Cauchy–Schwarz, D^*_i enters only the intercept, and monotonicity follows because β_i is a ratio of two functions affine in δ_i whose cross-difference reduces, using $\tilde{I}_{AA}\tilde{\Phi} - (\tilde{I}_{AA}\tilde{I}_{\tau\tau} - \tilde{I}^2_{A\tau}) = \tilde{I}^2_{A\tau}\sigma_r^{-2}/(\tilde{I}_{AA} + 1/\sigma_r^2) > 0$, to a strictly negative constant. For (ii), differentiate the equilibrium with respect to r_t: $\partial V^i_H / \partial r = \beta_i \, \partial\lambda/\partial r$, and from $\lambda = \gamma_A \sigma_r^2 (\Theta_t - \int_I V^i_H \, di)$ with $\partial \Theta_t / \partial r = -\Xi$,

$$\frac{\partial \lambda}{\partial r} = -\frac{\gamma_A \sigma_r^2 \, \Xi}{1 + \gamma_A \sigma_r^2 B}, \qquad \Omega = -\frac{1}{\gamma_A \sigma_r^2} \frac{\partial \lambda}{\partial r} = \frac{\Xi}{1 + \gamma_A \sigma_r^2 B}.$$

Ω_0 is the same expression at $B_0 = \int_I \beta_i|_{\delta_i = 0} \, di > B$, so $\Omega > \Omega_0$. For (iii), write $\tilde{I}_{xy} = I_{xy}/\sigma^{\text{id}\,2}$: as $\sigma^{\text{id}} \to 0$ the terms quadratic in the moments dominate numerator and denominator of β_i, whose ratio tends to $1/(\gamma_H \sigma_r^2)$ for every δ_i. □

Remark 3 (Economic reading)**.** A rate cut lengthens the passive Macaulay duration of held-fixed portfolios (by $-(r_t - \bar{r}) \, \text{Cov}^{\text{old}}(A, \tau) > 0$, since A is increasing; Section J.8), and the mandate pulls holdings back toward target—countercyclical trading that, in factor space, is the dampened λ-elasticity of Proposition 7(i). The covariance plays no role in the *sign* of η_r, which is unconditional; it governs the sign of the fund-level passive drift used in Section J.8.

J.6 Analytical Solution for $A(\tau)$

The ODE (72) rests on the local linearisation (Assumption 6) and on treating (κ^*, η_r) as reduced-form parameters: their closed forms given $A(\cdot)$ are in Proposition 7, and the joint fixed point in $A(\cdot)$ is standardly separated from the cross-sectional pricing problem. The ODE is a pure Vasicek equation with mean reversion $\kappa_\eta = \kappa^* + \eta_r$; integrating against $e^{\kappa_\eta \tau}$ with $A(0) = 0$ gives

$$\boxed{A_\eta(\tau) = \frac{1 - e^{-\kappa_\eta \tau}}{\kappa_\eta}, \qquad \kappa_\eta = \kappa^* + \eta_r,} \tag{73}$$

with benchmark loading $A_{BM}(\tau) \equiv (1 - e^{-\kappa^* \tau})/\kappa^*$ obtained at $\eta_r = 0$. Being a pure Vasicek loading, A_η is bounded, positive, and increasing, and never turns negative: no positivity bound on η_r is required. The mandate is defined and measured in Macaulay

units, but no-arbitrage maps the rebalancing flow into yields through its *factor* content: the two coincide at short maturities ($A(\tau) \approx \tau$) and diverge at the long end. The structure is the canonical Vayanos–Vila (2021) Vasicek one, with κ_η playing the role of their risk-adjusted mean reversion and the shift sourced from duration-target rebalancing; the priced-idiosyncratic alternative with its τ-linear wedge is in Appendix J.9.

J.7 Implications: Price Sensitivity and Long-End Yield-Response Dampening

The continuously-compounded systematic yield is $y(t,\tau) = \frac{1}{\tau}[A(\tau)(r_t - \bar{r}) + C(\tau)]$, with short-rate sensitivity $A(\tau)/\tau$. The dampening implications of (73) follow.

Proposition 8 (Dampened prices and a hump-shaped yield-response dampening)**.** *Fix $\kappa^* > 0$ and $\eta_r > 0$. Then $A_\eta(\tau) < A_{BM}(\tau)$ for all $\tau > 0$, so both price and yield responses are strictly below benchmark, and the implied dampening of the yield response per unit short-rate shock,*

$$D(\tau) \;\equiv\; \frac{A_{BM}(\tau) - A_\eta(\tau)}{\tau} = \frac{1}{\tau}\left[\frac{1 - e^{-\kappa^*\tau}}{\kappa^*} - \frac{1 - e^{-\kappa_\eta\tau}}{\kappa_\eta}\right], \tag{74}$$

is hump-shaped*: it vanishes as $\tau \to 0$ and as $\tau \to \infty$, and attains a single interior peak τ_p, with $\tau_p > \tau_f \equiv \ln(\kappa_\eta/\kappa^*)/(\kappa_\eta - \kappa^*)$, the peak of the forward-rate dampening. To first order in η_r, the peak is at $\kappa^*\tau^\dagger = x^\dagger$, the unique positive root of $x^2 + x + 1 = e^x$ ($x^\dagger \approx 1.793$); i.e. $\tau^\dagger \approx 1.79/\kappa^*$. The exact peak drifts below $\tau^\dagger$ only at $O(\eta_r/\kappa^*)$ (at $\kappa^* = 0.1445$, $\eta_r = 0.03$: $\tau_p \approx 11.3y$ against $\tau^\dagger \approx 12.4y$); η_r controls the height.*

Proof. The map $\kappa \mapsto (1 - e^{-\kappa\tau})/\kappa$ is strictly decreasing in κ for each $\tau > 0$: its κ-derivative has numerator $e^{-\kappa\tau}(\kappa\tau + 1) - 1$, which is 0 at $\kappa\tau = 0$ with derivative $-\kappa\tau e^{-\kappa\tau} < 0$, hence negative. Since $\kappa_\eta > \kappa^*$, $A_\eta(\tau) < A_{BM}(\tau)$, and dividing by $\tau > 0$ gives the pointwise inequality. As $\tau \to 0$, $A_{BM} - A_\eta \sim \frac{1}{2}(\kappa_\eta - \kappa^*)\tau^2$, so $D(\tau) \sim \frac{1}{2}\eta_r\tau \to 0$; as $\tau \to \infty$, $A_{BM} - A_\eta \to 1/\kappa^* - 1/\kappa_\eta$, a positive constant, so $D(\tau) \to 0$. For single-peakedness, let $N \equiv A_{BM} - A_\eta$ and $h(\tau) \equiv N'(\tau)\tau - N(\tau)$, so that $D' = h/\tau^2$. Then $h(0) = 0$ and $h'(\tau) = N''(\tau)\,\tau$, with $N''(\tau) = \kappa_\eta e^{-\kappa_\eta\tau} - \kappa^* e^{-\kappa^*\tau}$. Since $(\kappa_\eta/\kappa^*)e^{-(\kappa_\eta - \kappa^*)\tau}$ is strictly decreasing, N'' changes sign exactly once, from positive to negative, at τ_f. Hence h is strictly increasing on $(0, \tau_f)$ and strictly decreasing thereafter; $h(0) = 0$ implies $h > 0$ on $(0, \tau_f]$, while $h(\infty) = -(1/\kappa^* - 1/\kappa_\eta) < 0$. So h has exactly one zero $\tau_p \in (\tau_f, \infty)$, with D strictly increasing before τ_p and strictly decreasing after: D is single-peaked. For the first-order location, $\partial_\kappa[(1 - e^{-\kappa\tau})/\kappa] = [\kappa\tau e^{-\kappa\tau} - (1 - e^{-\kappa\tau})]/\kappa^2$ gives $D(\tau) = \eta_r[(1 - e^{-\kappa^*\tau}) - \kappa^*\tau e^{-\kappa^*\tau}]/[(\kappa^*)^2\tau] + O(\eta_r^2)$;

writing $x = \kappa^*\tau$ and maximising $[1-(1+x)e^{-x}]/x$, the first-order condition reduces to $e^{-x}(x^2+x+1) = 1$, i.e. $x^2+x+1 = e^x$, with unique positive root $x^\dagger \approx 1.793$. $\square$

The instantaneous-forward dampening is also hump-shaped: $A'_{BM}(\tau) - A'_\eta(\tau) = e^{-\kappa^*\tau} - e^{-\kappa_\eta\tau}$ is strictly positive, vanishes at both ends, and peaks at $\ln(\kappa_\eta/\kappa^*)/(\kappa_\eta - \kappa^*) \to 1/\kappa^*$ as $\eta_r \to 0$—the like-for-like benchmark for excess-sensitivity findings on instantaneous forwards (Hanson and Stein, 2015). The curve pivots: a rate cut moves short yields nearly one-for-one, dampens intermediate-to-long yields below the Vasicek benchmark, and leaves the very long end where both loadings saturate. To first order in η_r the peak is fixed by the separately-calibrated κ^*, so the coincidence of the model's peak with the regression's peak is an over-identifying test.

J.8 Cross-Sectional Implications: Heterogeneous PH Investors Around a Shock

The aggregate FOC depends on the joint distribution of (D_i^*, δ_i) only through one δ-weighted scalar (Section J.4.1). We now derive, for each investor i, the steady-state duration $D_{\text{port}}^{i,\text{old}}$, the no-trade passive duration $D_{\text{port}}^{i,\text{passive}}$ after a shock, and the re-optimised duration $D_{\text{port}}^{i,\text{new}}$. The differences $(D_{\text{port}}^{i,\text{passive}} - D_{\text{port}}^{i,\text{old}},\ D_{\text{port}}^{i,\text{new}} - D_{\text{port}}^{i,\text{passive}})$ are the model counterparts of the empirical *passive duration deviation* and *active rebalancing* signals.

Investor i's problem and first-order condition are (58) and (59). Since PH investors bear idiosyncratic risk, the per-bond allocation is uniquely determined (Remark 2). We derive the cross-section under the general drift $\mu(t,\tau) - r_t = \lambda_t A(\tau) + \eta_t\tau + \chi_t(\tau)$, which nests both specifications: under Assumption 1, $\eta_t = \chi_t = 0$ and $\lambda_t = \gamma_A\sigma_r^2(\Theta_t - V_{H,t})$; in the symmetric economy of Appendix J.9, $\eta_t, \chi_t \neq 0$ and $\lambda_t = \sigma_r^2\Theta_t/\tau_{\text{agg}}$. The objects that carry the empirical content—$\omega(\delta_i)$, the convex combination, and $\beta_i^{\text{reb}} = (1-\omega_i)(\beta^{\text{MV}}+1) - 1$—do not depend on the forcing values and are identical across the two; only the level of the MV-baseline shift K below differs, and K is identified through β^{MV} rather than computed from primitives.

J.8.1 Decomposition of investor-i demand

Substituting the general drift into (59) and solving for $y_t^i(\tau)$ yields the demand decomposition

$$y_t^i(\tau) = \frac{1}{\gamma_H \sigma^{\text{id}}(\tau)^2}\Big[\big(\lambda_t - \gamma_H\sigma_r^2 V_{H,t}^i\big)A(\tau) + \big(\eta_t - \delta_i\,\mathcal{D}_{\text{dev},t}^i\big)\tau + \chi_t(\tau)\Big], \tag{75}$$

into an $A(\tau)$-shape component (investor-specific via $V^i_{H,t}$) and a τ-shape component (investor-specific via $\delta_i \cdot \mathcal{D}^i_{\text{dev},t}$, which depends on (D^*_i, δ_i)). At $\delta_i = 0$ the τ-shape vanishes for that investor; if all $\delta_i = 0$, the demand collapses to a single $y^{\text{MV}}_t(\tau)$ irrespective of D^*_i, the mean-variance baseline demand. With $\delta_i > 0$, the τ-shape acquires an investor-specific habitat tilt through $\delta_i \mathcal{D}^i_{\text{dev},t}$.

J.8.2 Steady-state portfolio duration

At the reference state $r_t = \bar{r}$, the aggregate equilibrium drift simplifies (the loadings on $(r_t - \bar{r})$ vanish at the steady state and the constant component is absorbed into $C(\tau)$). Substituting (75) into the definitions $V^i_{H,t} = \int y^i_t A \, d\tau$ and $D^i_{\text{port},t} = (1/\bar{W}) \int y^i_t \tau \, d\tau$ yields a 2×2 linear system in $(V^i_{H,t}, \mathcal{D}^i_{\text{dev},t})$, parameterised by δ_i, in which D^*_i enters only as a forcing term on the right-hand side. Because the system is linear in D^*_i at fixed δ_i, $D^i_{\text{port},t} = \mathcal{D}^i_{\text{dev},t} + D^*_i$ is also linear in D^*_i. By the consistency requirement that an investor with $D^*_i = D^{\text{MV}}_{\text{port}}$ must end up at $D^{\text{MV}}_{\text{port}}$ herself (since this D^*_i coincides with the common MV outcome), the linear function takes the convex-combination form:

$$D^{i,\text{old}}_{\text{port}} \;=\; D^{\text{MV}}_{\text{port}} \;+\; \omega(\delta_i)\left(D^*_i - D^{\text{MV}}_{\text{port}}\right), \qquad \omega(\delta_i) \in [0,1], \tag{76}$$

where $D^{\text{MV}}_{\text{port}}$ is the duration of the (common) mean–variance equilibrium when all $\delta_i = 0$, and $\omega(\delta_i)$ is the *habitat pull strength*, an increasing function of δ_i with $\omega(0) = 0$ (no penalty $\Rightarrow$ that investor at the common MV duration) and $\omega(\delta_i) \to 1$ as $\delta_i \to \infty$ (hard constraint $\Rightarrow$ that investor exactly at her target D^*_i).

Explicit functional form. For $\sigma^{\text{id}}(\tau) = \sigma^{\text{id}}_0$, solving the 2×2 system gives

$$\omega(\delta_i) \;=\; \frac{\delta_i \, \Phi}{\gamma_H {\sigma^{\text{id}}_0}^2 \bar{W} + \delta_i \, \Phi}, \qquad \Phi \equiv M_{\tau\tau} - \frac{M^2_{A\tau}}{M_{AA} + {\sigma^{\text{id}}_0}^2/\sigma^2_r} > 0, \tag{77}$$

with M_{xy} the *unweighted* maturity moments $M_{AA} = \int_0^T A(u)^2 du$, $M_{A\tau} = \int_0^T A(u)\, u \, du$, $M_{\tau\tau} = \int_0^T u^2 du$ (positivity of Φ from Cauchy–Schwarz); supply weights enter the forcing terms of the 2×2 system, not these moments, and $\Phi = {\sigma^{\text{id}}_0}^2 \tilde{\Phi}$ with $\tilde{\Phi}$ as in Proposition 7. Φ is the regularised Schur complement of the FOC system in the duration direction. Under homogeneity in $(\gamma_H, \bar{W}, \sigma^{\text{id}}_0)$, cross-sectional dispersion in β^{reb}_i enters entirely through δ_i.

J.8.3 Response to a short-rate shock

Following an unexpected shock $r_t = \bar{r} + \Delta r$, three quantities characterize investor i's response. The old (steady-state) duration $D^{i,\text{old}}_{\text{port}}$ is given by (76). The passive

duration, holding face-value positions fixed and letting prices respond by $-A(\tau)\Delta r$ to first order, satisfies

$$D^{i,\text{passive}}_{\text{port}} \;-\; D^{i,\text{old}}_{\text{port}} \;\approx\; -\Delta r \,\cdot\, \text{Cov}^{i,\text{old}}\big(A(\tau),\,\tau\big) \;+\; O(\Delta r^2), \tag{78}$$

where $\text{Cov}^{i,\text{old}}(A,\tau)$ is the centered maturity-covariance under investor i's pre-shock market-value weights $w^{i,\text{old}}(u) \equiv y^{i,\text{old}}(u)P(\bar{r},u)/\int y^{i,\text{old}}(v)P(\bar{r},v)\,dv$ (normalised to integrate to one):

$$\text{Cov}^{i,\text{old}}(A,\tau) \;=\; \int w^{i,\text{old}}(u)\,A(u)\,u\,du \;-\; \Big(\int w^{i,\text{old}}(u)\,A(u)\,du\Big)\cdot D^{i,\text{old}}_{\text{port}},$$

with the demeaning by $D^{i,\text{old}}_{\text{port}}\cdot\int wA\,du$ arising automatically from the shifted market-value denominator. Under the factor-consistent loading $A_\eta(\tau) = (1-e^{-\kappa_\eta\tau})/\kappa_\eta$, which is strictly increasing in τ, the symmetrisation identity $2\,\text{Cov}^{i,\text{old}}(A,\tau) = \mathbb{E}[(A(\tau_1)-A(\tau_2))(\tau_1-\tau_2)]$ has a sign-definite integrand, so $\text{Cov}^{i,\text{old}}(A,\tau) > 0$ for *any* non-degenerate portfolio, unconditionally and with no restriction on its maturity support. Applying the convex-combination (76) at both states gives the duration shift

$$D^{i,\text{new}}_{\text{port}} - D^{i,\text{old}}_{\text{port}} = (1-\omega(\delta_i))\,\big(D^{\text{MV,new}}_{\text{port}} - D^{\text{MV}}_{\text{port}}\big),$$

i.e., the MV-baseline shift scaled by the habitat-dilution factor $(1-\omega(\delta_i))$. The four-way ordering

$$D^{i,\text{old}}_{\text{port}} \;\le\; D^{i,\text{new}}_{\text{port}} \;\le\; D^{i,\text{new,MV}}_{\text{port}} \;\le\; D^{i,\text{passive}}_{\text{port}} \qquad \text{(rate cut)}, \tag{79}$$

reverses for $\Delta r > 0$. The first inequality is unconditional; the last holds iff $\text{Cov}^{i,\text{old}}(A,\tau) \ge K$, the *stabilisation condition* that the price-revaluation effect dominates the equilibrium return shift (K is the MV-baseline shift per unit shock, defined below), with the weaker condition $\text{Cov}^{i,\text{old}}(A,\tau) \ge (1-\omega_i)K$ sufficient for $D^{i,\text{new}} \le D^{i,\text{passive}}$; the middle inequality holds whenever $D^*_i \le D^{\text{MV,new}}_{\text{port}}$. As $\delta_i \to \infty$, $D^{i,\text{new}}_{\text{port}} \to D^*_i = D^{i,\text{old}}_{\text{port}}$ (perfect rebalancing).

J.8.4 Passive deviation, active rebalancing, and their identification

Define the two model-implied quantities:

$$\mathcal{D}ev^i_t \;\equiv\; D^{i,\text{passive}}_{\text{port}} - D^{i,\text{old}}_{\text{port}}, \tag{80}$$

$$\mathcal{R}eb^i_t \;\equiv\; D^{i,\text{new}}_{\text{port}} - D^{i,\text{passive}}_{\text{port}}. \tag{81}$$

$\mathcal{D}ev^i_t$ measures the passive duration drift induced mechanically by the rate shock; $\mathcal{R}eb^i_t$ measures the active duration adjustment from the optimum at the new state back toward the steady-state target.

By the ordering (79), $\mathcal{D}ev_t^i$ and $\mathcal{R}eb_t^i$ have opposite signs in the parameter region we focus on: when prices push portfolio duration up, the investor rebalances down, and vice versa. The ratio

$$\beta_i^{\text{reb}} \;\equiv\; \frac{\mathcal{R}eb_t^i}{\mathcal{D}ev_t^i}, \tag{82}$$

quantifies the fraction of passive drift that is offset by active rebalancing within the period.

Closed form, additive decomposition, and inversion. Under cross-sectional uniformity of $\text{Cov}^{i,\text{old}}(A,\tau) \approx \overline{\text{Cov}}^{\text{old}} > 0$ (exact under identical maturity-weight profiles, and holding to first order when fund holdings have similar maturity shape; $\overline{\text{Cov}}^{\text{old}} > 0$ unconditionally since A is monotone), the convex-combination structure delivers

$$\beta_i^{\text{reb}}(\delta_i) \;=\; (1-\omega(\delta_i))(\beta^{\text{MV}}+1)-1 \;=\; \beta^{\text{MV}} \;+\; \underbrace{\big(-\omega(\delta_i)(\beta^{\text{MV}}+1)\big)}_{\beta_i^{\text{Habitat}}\le 0}, \tag{83}$$

with boundary cases exact: $\omega = 0 \Rightarrow \beta_i^{\text{reb}} = \beta^{\text{MV}} = K/\overline{\text{Cov}}^{\text{old}} - 1 \in (-1, 0]$ under the stabilisation condition $K \le \overline{\text{Cov}}^{\text{old}}$, and $\omega \to 1 \Rightarrow \beta_i^{\text{reb}} \to -1$ (rigid rebalancing). Here $K > 0$ is the MV-baseline shift per unit shock, $D_{\text{port}}^{\text{MV,new}} - D_{\text{port}}^{\text{MV}} = -K\,\Delta r$; under Assumption 1 it takes the closed form $K = \gamma_A \sigma_r^2 \Omega\, \tilde{I}_{A\tau} / [\gamma_H \bar{W}(1+\sigma_r^2 \tilde{I}_{AA})]$ (differentiate the $\delta_i = 0$ solution of the 2×2 system in Proposition 7 with respect to r_t), and in the symmetric variant it takes a different value; all results below use only $K > 0$, and K is identified empirically through β^{MV}. Monotonicity: $d\beta_i^{\text{reb}}/d\delta_i = -(\beta^{\text{MV}}+1)\omega'(\delta_i) < 0$. The empirical inversion is

$$\widehat{\beta}_f^{\text{MV}} \;=\; \frac{\widehat{\beta}_f^{\text{reb}} + \widehat{\omega}_f}{1-\widehat{\omega}_f}, \qquad \widehat{\beta}_f^{\text{Habitat}} \;=\; \widehat{\beta}_f^{\text{reb}} - \widehat{\beta}_f^{\text{MV}},$$

which holds fund-by-fund as an algebraic identity under fund-specific $(\gamma_H^f, \bar{W}_f, \delta_f, D_f^*, \mathcal{U}_f, w^{f,\text{old}})$; under a uniform investment universe the closed form is $\beta_f^{\text{MV}} = K_f/\text{Cov}_{w^{f,\text{old}}} - 1$ with K_f the fund-specific analogue of K.

Mapping to the empirical signal. $(\mathcal{D}ev_t^i, \mathcal{R}eb_t^i)$ correspond to $\text{Dur}_{f,t}^{\text{dev}} \equiv \text{Dur}_{f,t}^{\text{pass}} - \text{Dur}_{f,t-1}$ and $\text{Dur}_{f,t}^{\text{reb}} \equiv \text{Dur}_{f,t} - \text{Dur}_{f,t}^{\text{pass}}$, where $\text{Dur}_{f,t}^{\text{pass}}$ is the no-trade counterfactual at post-shock prices.[27] The fund-level regression $\text{Dur}_{f,t}^{\text{reb}} = \beta_f^{\text{reb}}\,\text{Dur}_{f,t}^{\text{dev}} + \varepsilon_{f,t}$ yields

[27] The empirical $\text{Dur}_{f,t}^{\text{pass}}$ uses post-shock total bond value as wealth normaliser; the model's $D_{\text{port}}^{i,\text{passive}}$ uses pre-shock $\bar{W}$. The two agree at steady state and to first order in Δr.

$\beta_f^{\text{reb}} \in (-1, \beta_f^{\text{MV}}]$ with magnitude rising in δ_f; this slope underpins the ownership-weighted bond-level signal, with (83) splitting it into a common MV baseline plus an investor-specific habitat contribution.

Comparative statics in η_r are immediate: since $\kappa \mapsto (1 - e^{-\kappa\tau})/\kappa$ is strictly decreasing (Proposition 8), a tighter mandate—higher δ, hence higher η_r by Proposition 7—lowers $A_\eta(\tau)$ at every maturity and raises the dampening hump, moving its peak only at $O(\eta_r/\kappa^*)$.

J.9 Priced idiosyncratic risk: the linear-in-τ wedge

For completeness we record the non-standard alternative in which per-bond idiosyncratic risk is *priced by both sectors*—Assumption 1 replaced by symmetric risk bearing, so the arbitrageur condition is (56) and each bond's $\sigma^{\text{id}}(\tau)^2$ enters the equilibrium price, a departure from the canonical Vayanos–Vila baseline. The aggregated first-order condition (64) then retains all three terms,

$$\mu(t,\tau) - r_t = \underbrace{\frac{\sigma_r^2 \Theta_t}{\tau_{\text{agg}}} A(\tau)}_{\lambda_t A(\tau)} + \underbrace{\frac{\bar{\delta}}{\gamma_H \tau_{\text{agg}}} \widetilde{\mathcal{D}}_{\text{dev},t}\, \tau}_{\eta_t \tau} + \underbrace{\frac{\sigma^{\text{id}}(\tau)^2}{\tau_{\text{agg}}} S_t(\tau)}_{\chi_t(\tau)} : \tag{84}$$

a factor premium, a rebalancing wedge linear in τ, and an idiosyncratic supply wedge; these are the local objects η_t, χ_t used to select the per-bond habitat allocation in Section J.8. The preferred-habitat duration penalty now leaks into the per-bond price through the second term. Coefficient matching (as in Section J.5, with $\eta \equiv -\bar{\delta}\,\widetilde{\mathcal{D}}'/(\gamma_H \tau_{\text{agg}}) > 0$ the state-sensitivity of the wedge slope and the idiosyncratic term folded into κ^*) gives the constant-η ODE and loading

$$A'(\tau) + \kappa^* A(\tau) = 1 - \eta\,\tau, \qquad A(\tau) = A_{BM}(\tau) - \frac{\eta}{\kappa^*}\big[\tau - A_{BM}(\tau)\big]. \tag{85}$$

This loading is internally consistent but $A(\tau) \to -\infty$ as $\tau \to \infty$ (it turns negative once $\eta\tau$ is large), so it is well-posed only on a bounded maturity range. Setting $A(\tau) = 0$ at the longest maturity $T_{\max}$ and using $\kappa^* T_{\max} \gg 1$ gives the positivity bound

$$\eta \;<\; \frac{1}{T_{\max} - 1/\kappa^*}, \tag{86}$$

recorded here only for completeness; the calibration in Appendix I does not use this bound. By Proposition 5 the τ-linear wedge admits asymptotic arbitrage as the maturity span grows, since the arbitrageur's idiosyncratic risk-bearing $\int x(\tau)^2\, d\tau \propto \int \tau^2\, d\tau$ diverges. The main text therefore adopts the canonical diversifiable-idiosyncratic-

risk pricing, for which the wedge loads on $A(\tau)$ and the loading is the globally well-posed Vasicek form (73).