Who Holds U.S. Bonds and How It Shapes the Yield Effects of QE and QT

Masashi Takahashi*†

September 16, 2026

[Click here for the latest version]

Abstract

How much do the bond-supply shifts of quantitative easing (QE) and tightening (QT) move yields? The answer turns on the slope of aggregate bond demand: who holds bonds, how sensitive their demand is to yields, and how fast they rebalance. I estimate heterogeneous investor demand in the two U.S. bond markets the Federal Reserve trades, Treasuries and agency MBS, from sector-level portfolio holdings and transactions spanning 1985 to 2025. I measure each sector's demand against its own non-bond alternative and let its yield sensitivity and rebalancing speed vary. Most investors respond to yields weakly and slowly, so aggregate demand is inelastic. Feeding the demand system through market clearing, I find that historical QE compressed 10-year Treasury yields by 7 to 11 basis points per percent of the float removed. Because agency MBS demand prices the spread over Treasuries, the composition of a program, not its size, determines what it does to mortgage rates. QT's effects are smaller, because the aggregate demand curve has since flattened as elastic sectors came to hold more of the market: a further $1 trillion of runoff raises 10-year Treasury yields by about 21 basis points and agency MBS yields by about 25, of which 4 is the spread.

Keywords: Agency MBS, Demand System, Institutional Investors, Market Macrostructure, Quantitative Easing, Quantitative Tightening, U.S. Treasuries

*Smeal College of Business, Pennsylvania State University. Email: mft5548@psu.edu.

†I am grateful to Brent Ambrose and Jiro Yoshida for their guidance and support. I also thank Sriya Anbil, Daniel Beltran, Cynthia Doniger, Erin Ferris, David Glancy, Hyein Han, Jingzhi Huang, Kyungmin Kim, Robert Kurtzman, Anh Le, Stefan Lewellen, Canlin Li, Liang Peng, Lubomir Petrasek, Romina Ruprecht, Bernd Schlusche, Tim Simin, Arsenios Skaperdas, Eva Steiner, Hiroatsu Tanaka, Manjola Tase, Mihail Velikov, Annette Vissing-Jørgensen, Sean Wilkoff, Alexei Zhdanov, and seminar participants at the Federal Reserve Board and Penn State for their helpful comments and discussions. Needless to say, any remaining errors are solely my responsibility.

1 Introduction

For over a decade, the Federal Reserve has purchased more than $6 trillion of U.S. Treasuries and agency MBS to lower yields, a policy known as quantitative easing (QE). It has since reversed course, shrinking its balance sheet through quantitative tightening (QT), yet the consequences for bond yields and the channel through which they operate remain unclear.[1] The key is that QE and QT shift the supply of bonds available to private investors, so their effect on yields depends on the slope of investor demand curves: flat demand implies small yield effects, steep demand large ones (e.g., Koijen and Yogo, 2019; Vayanos and Vila, 2021). Because investors differ in their yield sensitivity, in the assets they hold, and in how quickly they rebalance, the aggregate demand curve, and hence the size and timing of the yield effect of QT, depends on *which* investors absorb the released supply, and how fast.

This paper estimates bond demand using sector-level portfolio holdings and transactions covering the two bond markets the Federal Reserve trades, Treasury securities and agency MBS, across 26 sectors grouped into eight broad investor types: banks, insurers, pension funds, mutual funds, arbitrageurs, foreign investors, households, and other investors. Three choices define the system. I measure each sector's demand from its transactions rather than from the market value of its holdings. I let each sector's yield sensitivity differ by asset class, so the system can say who absorbs Treasuries and who absorbs agency MBS instead of imposing one slope per investor across markets. And I let each sector rebalance at its own speed, and toward its own outside asset, estimating how far and how fast each adjusts. The estimates support counterfactual analyses of the QE and QT episodes and of the runoff to come.

I find three main results. First, most investors are inelastic, and the few that absorb supply do so slowly. In the Treasury market only two sectors buy when the yield rises: the household residual, which after 2010 contains the hedge funds running the cash–futures basis

[1] In June 2026, Federal Reserve Chairman Kevin Warsh made balance-sheet policy one of five priorities for the conduct of monetary policy, charging a dedicated task force to weigh "the benefits and risks of the current ample reserves regime, and the composition of the Fed's balance sheet" (Warsh, 2026).

trade, and mutual funds. Banks, insurers, pension funds, and foreign investors, who hold most of the float, do not move. In the agency MBS market, by contrast, demand responds to the option-adjusted spread rather than to the yield level, and every sector type absorbs. The data reject the restriction, standard in bond demand systems, that an investor's slope is the same across the bonds it holds. Most absorbers complete only an eighth to a quarter of their one-year response in the quarter of the shock. Second, because the Treasury market's absorbers are few, Treasury yields are highly sensitive to supply: historical QE compressed 10-year Treasury yields by 7 to 11 basis points for each percent of the float the Fed removed, from 51 basis points under QE1 to 176 under QE4, in line with the demand-system literature (Eren et al., 2026; Chaudhary et al., 2026). Because the Treasury market sets the level of mortgage yields and the MBS market only the spread on top, Treasury purchases lower mortgage yields too: QE2, a Treasury program, lowered them by 82 basis points, and QE1, the Fed's only mortgage-heavy program, compressed the spread by a further 45. Third, a resumed QT raises mortgage costs mostly through the Treasury market. A further $1 trillion of runoff from the 2025Q1 balance sheet, about a year and a half at the Fed's 2024 pace, widens the mortgage spread by only 4 basis points but raises 10-year Treasury yields by 21, so mortgage yields rise by 25. The Treasury number is uncertain: it ranges from 7 basis points in a recession to 41 at the lower edge of the household residual's confidence set, because the two sector slopes it depends on, the household residual's and the foreign sector's, are imprecisely estimated.

I construct quarterly sector-level portfolio shares and transactions for the two bond classes from the Federal Reserve's Financial Accounts (Z.1), splitting the rest of the world into official and private holders with the Federal Reserve's CSLT cross-border positions data over the full sample. The Financial Accounts report every holding together with the net purchases that changed it, decomposing each level change into transactions, revaluations, and other volume changes. Unlike Eren et al. (2026), who use market-value portfolio shares, I measure demand from transactions. Banks carry held-to-maturity securities at amortized cost and insurers at book value, so changes in market value can register a demand response for positions that were never traded; transactions record trading directly, and they also keep the issuance characteristic free of revaluation. I merge the transaction panel with asset characteristics and macroeconomic controls from FRED. The yield is the 10-year Treasury yield for Treasuries and, for agency MBS, the 10-year yield plus the index option-adjusted spread (OAS), whose spread component is the price MBS demand responds to; maturity, duration, and coupon come from the Bloomberg indices, and issuance from the Financial Accounts. The sample covers 26 sectors and two asset classes over the 161 quarters from 1985Q1 to 2025Q1.

Given these data, I estimate a characteristics-based demand system following Koijen and Yogo (2019), in which each sector's portfolio shares take a logit form, and extend it in three ways. First, the yield coefficient differs by asset, the alternative-specific-coefficient logit of discrete choice; the share structure and its closed-form elasticities are unchanged, and a minimum-distance test rejects the equal-slope restriction (Appendix D.7). Second, because sectors rebalance at different speeds, I estimate the demand equation in first differences of transaction growth and trace the response horizon by horizon with local projections (Jordà, 2005), recovering each sector's impact slope, long-run slope, and speed of adjustment. Third, the outside asset is sector-specific, loans for banks, corporate equities for long-term investors, and repo for money funds and dealers, so the slope captures substitution between bonds and each holder's realistic alternative. The slope of each market's aggregate demand curve, the float-weighted sum of the holders' own-asset slopes, aggregates across sectors whatever outside asset each uses, and it is the object the counterfactual requires.

Yields are endogenous, and I identify each asset's slope with instruments specific to that asset. For both markets the Fed trades I use the revisions to its *announced* purchase path rather than its realized purchases, since the realized flow follows a schedule published in advance and carries no surprise. The instrument moves the mortgage spread while leaving the Treasury yield untouched, the test an asset-specific supply shifter has to pass. For

Treasury yields I add a heteroskedasticity-based instrument (Lewbel, 2012) built on the coupon. I set aside the crisis quarters of 2008, 2020, and 2022, when flight-to-safety demand shocks move yields directly: the regressions control for them, and the heteroskedasticity instrument is built only from calm-period variation. The first-stage F statistics are 7.8 for Treasuries, on two instruments whose overidentifying restriction is not rejected, and 16.5 for agency MBS, and I report a weak-instrument-robust Anderson–Rubin set next to every slope, describing a slope as different from zero only when its set excludes zero.

The estimated demand parameters show who absorbs supply in each market. In the Treasury market only two sector types have long-run slopes whose Anderson–Rubin sets exclude zero, the household residual and mutual funds; banks, insurers, pension funds, and foreign investors, who hold most of the float, do not respond. In the agency MBS market four of the eight do: the household residual, foreign investors, banks, and pension funds. Absorption is slow in both markets. A typical absorber completes only an eighth to a quarter of its one-year response in the quarter of the shock, pension funds being the one fast absorber, and the responses level off after about four quarters, the slow-moving capital of Duffie (2010) as a property of every sector that absorbs supply. In agency MBS the household residual and the arbitrageurs even trade the other way on impact, so that market has no marginal buyer in the quarter a supply shock lands.

The demand slopes and the composition of holdings together determine the slope of each market's aggregate demand curve, which pins down the basis-point impact of a supply shift. In 2025Q1 the Treasury slope implies that removing 1 percent of the private float raises the 10-year yield by about 7 basis points, with a 95 percent band of 4 to 26, and the household residual accounts for about half of it. The agency MBS slope implies that removing 1 percent of the MBS float widens the mortgage spread by 1.8 basis points, with a band of 1.2 to 3.9, and unlike the Treasury slope it rests on many sector types rather than one. The Treasury slope has moved with that one sector's presence: it fell by two-thirds in the early 2000s and stayed low through 2008, as foreign reserve managers accumulated Treasuries and the residual emptied, and it recovered after 2009 and again after 2022 as the float shifted back toward the residual and the hedge funds in it.

Lastly, I conduct a counterfactual analysis of Fed asset purchases using the estimated system. Equilibrium yields are jointly determined by the demand system and the market-clearing condition in the two markets the Fed trades, with the balance-sheet shocks measured as the Fed's own transactions and bond supply responding to yields through net-issuance elasticities estimated from the same data. Across the four programs, QE1 through QE4, the Treasury impact runs from 7 to 11 basis points per 1 percent of the float removed, overlapping the 8 to 13 of Eren et al. (2026). QE1's mortgage effect of 95 basis points sits inside the event-study record of 85 to 107. The 2017–2019 runoff raised Treasury yields by 21 basis points and mortgage yields by 28. The 2022–2025 runoff raised them by 60 and 68. Looking forward, a further $1 trillion raises 10-year Treasury yields by 21 basis points and mortgage yields by 25; applied to $27 trillion of marketable federal debt as it rolls over, the Treasury rise adds roughly $58 billion a year to federal interest costs, and the mortgage rise about $66 a month to the payment on a newly originated $400,000 mortgage. The Treasury number is a band, not a point: it runs from 7 basis points if the runoff lands in a recession to 41 at the lower edge of the household slope's confidence set.

This paper contributes to the literature on market macrostructure (Haddad and Muir, 2025), and in particular to demand-system asset pricing in bond markets.[2] Koijen et al. (2021) find that foreign investors, whose demand is most elastic, sell the largest share of bonds to the ECB during quantitative easing. Jansen et al. (2024) estimate sector-level Treasury demand within an equilibrium model of arbitrageurs and find that QE moves yields only insofar as investors perceive the purchases as persistent; my counterfactuals, which hold the balance-sheet change in place, correspond to their credible-commitment case, and their finding that hedge funds and dealers are the marginal Treasury investors is the high-frequency counterpart of the residual sector that absorbs Treasuries here. Eren et al.

[2] Several papers estimate demand systems for bond markets without a direct focus on QE: Fang et al. (2025) for global sovereign debt and Bretscher et al. (2025) for corporate bonds.

(2026) estimate demand for U.S. Treasuries, find that QE and QT have asymmetric yield effects, and extend the analysis to Japan and the U.K. Fang and Xiao (2025) extend the approach to three bond classes at the security level and find that for every unit of bonds the Fed purchases, mutual funds and banks purchase an additional 0.5 units, amplifying the effect of Fed policy; their granularity also recovers substitution my sector-level classes average over, such as between long-maturity Treasuries and agency MBS, close substitutes at matched duration. Chaudhary et al. (2026) decompose realized Treasury yield movements into investor-level drivers with a granular instrumental-variable estimator and find that aggregate Treasury demand is highly inelastic: a 1 percent shift in demand moves yields by 13 basis points, at the lower end of the per-unit impacts estimated here. My main contribution is a detailed evaluation of QE and QT through a demand system. I price every program from QE1 to QT2, and a prospective runoff, in both markets the Fed trades; I separate the Treasury level from the mortgage spread, so that the composition of a program can be read off its effect on mortgage rates; and I trace how the yield effect builds and settles as slow capital arrives. Three features of the system make this possible. First, I measure demand from the Financial Accounts' transactions, complementing the market-value holdings used in prior work, so that the slopes of book-value holders such as banks and insurers reflect their trading. Second, I let the yield slope differ by asset, and find that the set of active absorbers differs across markets. Third, I let demand adjust at a sector-specific speed, estimating the response horizon by horizon with local projections, which separates the impact and long-run effects of QE and QT.

This paper also contributes to the empirical literature evaluating the effects of QE on asset prices. Event studies find that announcements of the Fed's purchase programs lowered 10-year Treasury yields by roughly 90 to 110 basis points around QE1 and mortgage rates by 40 to 100 basis points (Fuster and Willen, 2010; Gagnon et al., 2011; Hancock and Passmore, 2011; Krishnamurthy and Vissing-Jorgensen, 2011; Stroebel and Taylor, 2012; Krishnamurthy et al., 2013), although Christensen and Rudebusch (2012) attribute about

60 percent of the QE1 decline to lower expected short rates, so announcement effects mix the signaling and portfolio-balance channels. Time-series and term-structure studies map the Fed's holdings into term premiums and find declines of roughly 30 to 45 basis points per program (D'Amico et al., 2012; Ihrig et al., 2018; Wei, 2022), and evidence from actual purchases, rather than announcements, points to persistent stock effects (D'Amico and King, 2013; Vissing-Jorgensen, 2021). These estimates, however, evaluate only the policies actually implemented, under the conditions prevailing at the time: they cannot speak to counterfactual policies, such as alternative purchase sizes or asset mixes, and are vulnerable to the Lucas critique, since reduced-form price responses need not be stable across policy regimes. I fill this gap by estimating structural demand parameters that are more stable across policy regimes and support counterfactual analysis for policies never implemented. Calibrated preferred-habitat models also support counterfactuals (Vayanos and Vila, 2021); relative to their calibration, which matches yield moments in a Treasury-only model, I estimate demand from observed sector transactions and cover both of the markets the Fed trades, so the counterfactuals inherit measured investor heterogeneity and the substitution between the two.

This paper also contributes to the literature on the transmission channels of balance sheet policy. QE transmits through signaling, portfolio rebalancing, and asset-specific channels such as safety, prepayment risk, and default risk (Krishnamurthy and Vissing-Jorgensen, 2011), as well as through bank lending (Chakraborty et al., 2020; Kandrac and Schlusche, 2021) and mortgage refinancing (Di Maggio et al., 2020). This paper focuses on portfolio rebalancing: when the Fed purchases bonds, investors shift into imperfect substitutes. I quantify the channel by estimating each sector's substitution between the bonds the Fed trades and its outside asset. The estimates speak to the debate between a broad duration channel, under which purchases in any market move all long-term yields (Gagnon et al., 2011), and the narrow view that the impact concentrates in the purchased asset (Krishnamurthy et al., 2013). Each holds in part: Treasury purchases lower mortgage yields through the Treasury level, while the mortgage spread responds mainly to the MBS the Fed buys or runs off.

Finally, this paper contributes to the literature on the implementation and limits of QT. How far the central bank can shrink its balance sheet without side effects depends on both sides of that balance sheet. On the liability side, QT drains reserves from financial intermediaries and strains the repo market, so the demand for reserves and liquidity bounds the size of QT (e.g., López-Salido and Vissing-Jorgensen, 2025; Anbil et al., 2026; Anderson et al., 2026; Oguri and Pizzimenti, 2026). This paper focuses instead on the asset side: as the Fed reduces its holdings of Treasuries and agency MBS, the remaining investors absorb the supply only at higher yields, and this yield effect, not reserve scarcity, limits QT. The asymmetry between QE and QT arises on both sides of the balance sheet (Acharya et al., 2024; Eren et al., 2026), so the effects of QT cannot be inferred from estimates of QE. Assessing QT therefore requires knowing who holds bonds, which of them trade on yields, how steeply, and how fast their capital moves; the demand system estimated here provides exactly these inputs, and it says that for Treasuries the answer today rests on one sector.

2 Data

The dataset covers the information required for the demand estimation: portfolio holdings and the transactions that change them, asset prices (yields), asset characteristics, and instrumental variables for asset prices.

2.1 Portfolio Holdings and Transactions

Sectors. I use the Federal Reserve's Flow of Funds (the Financial Accounts of the United States, Z.1) to compute sector-level portfolio shares and to measure the transactions behind their changes. The Flow of Funds reports quarterly balance sheets for the U.S. financial sectors (e.g., life insurance companies, private pension funds). The sample runs from 1985Q1 to

2025Q1, a span that reaches back to the high-yield years of the late 1980s and, as Section 3.7 shows, is what identifies the marginal buyer of Treasuries. Over this span I work with 27 sectors: the 25 domestic sectors of the Financial Accounts and the rest of the world split into Foreign Official and Foreign Private holders. The split uses the Federal Reserve's CSLT estimates of cross-border securities positions (Bertaut and Judson, 2014), which report monthly holdings and net transactions of U.S. Treasuries, agency bonds, corporate bonds, and equities by official and non-official foreign investors from 1984 onward. A separate hedge-fund sector from the Federal Reserve's Enhanced Financial Accounts (EFA), built from SEC Form PF filings, is available only from 2013, so the sample keeps hedge funds within the household residual; Appendix D.5 uses the EFA to measure their share of it. Table 1 lists the sectors and their broader sector-type classifications (e.g., insurers, pension funds); the foreign official and private units both belong to the foreign type.

Assets. The two bond markets the Federal Reserve trades, Treasury securities and agency MBS, are the assets inside the demand system; the Flow-of-Funds asset classes are coarser than CUSIP-level data but detailed enough (e.g., Treasury securities, corporate equity, cash) to identify them. Together the sectors form a heterogeneous demand side, yet a few large groups hold most of the bonds: throughout the sample, foreign investors and domestic financial institutions (banks, insurers, pension funds, and mutual funds) hold about 80% of the two markets, while arbitrageurs (e.g., dealers, hedge funds) hold only about 5%. Figure 1 reports each market's ownership composition and total issuance over time, from 1985 to 2025. Because the ten-year yield prices coupon securities rather than bills, I measure the Treasury asset on the coupon market, using the non-bill holdings and transactions the Accounts publish for the six holders whose bills they break out (tables L.210 and F.210) and the total for the remaining sectors. Agency MBS is left as reported: table L.211 does not separate agency MBS from the debentures of the government-sponsored enterprises, about 17 percent of the total, which attenuates the estimated MBS slopes.

Table 1: **Sector Coverage and Outside Asset Definitions**

Type	Sector	Outside Asset
Banks	U.S.-Chartered Depositories	Loans
	Foreign Banking Offices	Total non-bond
	Credit Unions	Loans
	Banks in Affiliated Areas	Total non-bond
Insurers	Life Insurers	Corporate equities
	Property-Casualty Insurers	Corporate equities
Pension Funds	Private Pension Funds	Corporate equities
	State & Local Govt. Retirement Funds	Corporate equities
	Federal Govt. Retirement Funds	Total non-bond
Mutual Funds	Mutual Funds	Corporate equities
	Exchange-Traded Funds	Corporate equities
	Money Market Funds	Repo
	Closed-End Funds	Corporate equities
Arbitrageurs	Brokers and Dealers	Repo
	mREITs	Total non-bond
Foreign	Foreign Official	U.S. equities
	Foreign Private	U.S. equities
Household Residual	Household Residual & Nonprofits	Corporate equities
Other	Nonfinancial Business	Total non-bond
	Government-Sponsored Enterprises	Total non-bond
	Holding Companies	Total non-bond
	State & Local Governments	Total non-bond
	Federal Government	Total non-bond
	Agency/GSE Mortgage Pools	Total non-bond
	Finance Companies	Total non-bond
	ABS Issuers	Total non-bond
	Other Financial Business	Total non-bond

Notes: Sectors in the portfolio data and the outside asset against which each sector's bond demand is measured, chosen as the non-bond asset the sector realistically trades off against bonds; "Total non-bond" is total financial assets net of all bond holdings, used where no single alternative dominates. The rest of the world is split into Foreign Official and Foreign Private holders with the Federal Reserve's CSLT data. Hedge funds remain inside the household residual, because the Enhanced Financial Accounts begin only in 2013 and include offshore funds that the Financial Accounts assign to the rest of the world. Appendix D.1 explains each outside asset, and Appendix Table A.9 lists sources and table codes.

Transactions. The Financial Accounts also record each sector's net purchases of every instrument, so I observe transactions as well as holdings, and the demand system is estimated from the transactions. The reason is measurement. A recorded holding changes for two distinct reasons, the sector trades and the position it already holds is revalued, and only the first is a portfolio decision. Writing $H_{it}(n)$ for the recorded holding of asset n, $F_{it}(n)$ for net purchases during quarter t, and $R_{it}(n)$ for revaluations and other volume changes,

$$\Delta \ln \frac{H_{it}(n)}{H_{it}(0)} = \underbrace{\big[g_{it}(n) - g_{it}(0)\big]}_{\text{trading}} + \underbrace{\big[r_{it}(n) - r_{it}(0)\big]}_{\text{revaluation}}, \qquad g_{it}(n) \equiv \ln\Big(1 + \frac{F_{it}(n)}{H_{i,t-1}(n)}\Big), \tag{1}$$

with $r_{it}(n) \equiv \ln\big(1 + R_{it}(n)/(H_{i,t-1}(n) + F_{it}(n))\big)$. The Accounts report the two components separately, so the trading term is observed on its own, and it is what the demand system uses:

$$\Delta q_{it}(n) \equiv g_{it}(n) - g_{it}(0), \tag{2}$$

the growth of the units held that is due to trading, however the level is valued. I winsorize each cell's transaction growth at its 1st and 99th percentiles, because a growth rate is uninformative on a near-zero base, and leave it undefined where the lagged position is below \$50 million.

Valuation. The revaluation term is not common across sectors, which is why it cannot simply be corrected away. Banks report held-to-maturity securities at amortized cost, insurers at statutory book value, and the rest of the world at market value, so the recorded pass-through of a price change into the level ranges from zero to nearly one across sector–asset cells (Appendix D.2). The alternative used by Eren et al. (2026) keeps the market-value share change and strips the revaluation with the index duration, adding $D_t(n)\,\Delta y_t(n)/100$ back to the differenced log share; applied to a holder whose recorded level does not move with prices, that correction attributes a response of about $D/100$ per percentage point, a slope of 0.05 to 0.10, to a sector that traded nothing. Appendix D.2 shows that the small positive slopes market-value estimates find for banks and insurers vanish when demand is measured from transactions, and the same contamination reaches any instrument built from market-value holdings (Section 3.6). Market clearing is written in units, and the Fed's balance-sheet changes enter the instruments and the counterfactual as its own transactions, so the transaction slope is the object the equilibrium requires.

Outside Asset. The outside asset is sector-specific: for each sector I designate the single non-bond class it trades off against bonds, and measure its bond demand relative to that class. Domestic banks, for example, rebalance between bonds and loans rather than between bonds and every other asset on their balance sheet, so loans are their outside asset. The last column of Table 1 reports each sector's outside asset, and Appendix D.1 gives the rationale sector by sector. Appendix Figures A.6–A.6 show how each sector's portfolio shares evolve over the sample.

Figure 1: **Ownership Structure of U.S. Bond Markets**

(a) Treasury (b) Agency MBS

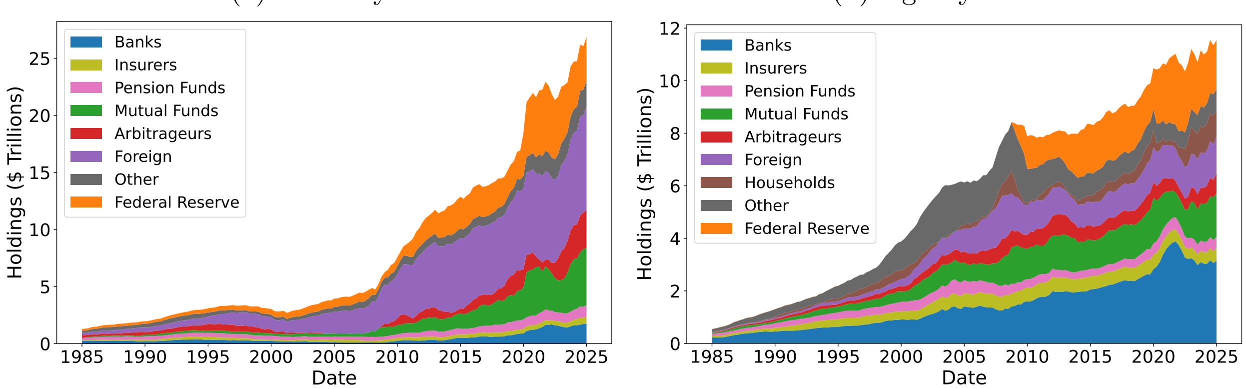

Note: This figure shows the ownership share of the two U.S. bond markets the Federal Reserve trades, Treasury securities and agency MBS, by institutional sector. Appendix A.1 describes the sector classification. In Panel (a) the Z.1 household residual in Treasury securities is reassigned to hedge funds within the Arbitrageurs sector, as these holdings predominantly reflect hedge funds engaged in the cash–futures basis trade. Panel (a) is the whole Treasury securities market, bills included, which is the market the Financial Accounts report; the demand system prices the coupon market, whose ownership Table 2 reports and in which foreign investors hold 29 percent rather than the 39 percent shown here.

Limitations. I acknowledge the limitations of the sector-level portfolio share data from the Flow of Funds. The data do not report within-class maturity variation, so I cannot estimate demand across precise duration buckets. Thus, this paper focuses on substitution across broad fixed-income asset classes rather than fine maturity choice.

2.2 Asset Characteristics

Yields. This paper uses one yield per asset class as its price. For Treasuries I use a representative constant-quality series that spans the full sample, the 10-year constant-maturity Treasury yield (DGS10) from FRED. A point yield at a representative maturity tracks the long end of the market, where habitat demand bites, more cleanly than a whole-market index does over a four-decade sample. For agency MBS I use an option-adjusted yield, the 10-year Treasury yield plus the option-adjusted spread of the Bloomberg agency-MBS index (available from 1990Q3, which sets the start of the MBS estimation sample): the index's own yield-to-worst embeds the prepayment-option premium, which moves with interest-rate volatility rather than with the compensation investors require for holding MBS. Table A.10 in the Appendix reports the sources and ticker codes. I aggregate all series to the quarterly frequency, taking the end-of-quarter value.

Other Characteristics. The system uses four asset characteristics: maturity, net issuance, duration, and coupon. The weighted time-to-maturity, modified duration, and weighted-average coupon of each index come from Bloomberg. Duration enters lagged one quarter, because its change is a near-mechanical function of the contemporaneous yield change ($dD \approx -D\,dy/(1+y)$). The coupon is a slow-moving composition variable and is the exogenous shifter behind the Treasury heteroskedasticity instrument (Section 3.6). Net issuance is measured from the Z.1 flow tables, as every holder's transactions in the asset, the Federal Reserve included, over the lagged amount outstanding, rather than from the change in market-value outstanding, which is mostly revaluation and would enter the estimation as a function of the yield itself. Table A.10 in the Appendix lists the sources for each asset class.

Macroeconomic Variables. I include a set of macroeconomic controls to absorb time-series variation in aggregate portfolio demand unrelated to asset prices. Specifically, I control for log real GDP, real GDP growth, core inflation, and the log of the U.S. dollar index, all obtained from the Federal Reserve Economic Data (FRED) database and aggregated to the quarterly frequency. I also include the one-quarter lag of log VIX as a measure of aggregate risk sentiment. Because VIX is jointly determined with asset prices during flight-to-safety episodes, I use the lagged value as a predetermined control to avoid simultaneity bias. The estimation also includes an indicator for the eight crisis quarters in which the heteroskedasticity instruments are switched off (Section 3.6).

2.3 Instruments

The main instruments come from two sources, as Section 3.6 explains: the revisions to the Fed's announced purchase path, read from the FOMC communications that changed it and dated to the quarter in which each change first became public, for both markets the Fed trades; and a crisis-robust heteroskedasticity-based instrument built from the coupon characteristic already in the panel, which the Treasury block uses alongside the announcements. A granular instrument built from the sectors' holdings, maturity walls, and the monetary-policy surprises of the prior demand-system literature enter as diagnostics only. This subsection records the data sources.

Monetary Policy Announcement Shocks. High-frequency shocks identified from intraday market movements around FOMC announcements enter in two places. The LSAP factor of Swanson (2021) is a diagnostic only: it has no first stage in either market (Appendix Table A.18). The MBS-price surprise of Drukker (2025) identifies the agency MBS supply schedule of Section 5.2; the demand side uses neither.

Federal Reserve Asset Purchases. The Fed's purchases and sales of Treasuries and agency MBS are the transactions of the monetary authority sector in the Z.1 flow tables, quarter by quarter. They enter the counterfactual of Section 5 as the balance-sheet shock of each program. They are not an instrument: the realized flow follows a schedule the Fed announces in advance, so it carries no surprise, and it appears in the instrument menu only as the construction the announcement revisions replace. The detailed purchase records of the Federal Reserve Bank of New York date the programs: QE1 (2008Q4–2010Q1), QE2 (2010Q4–2011Q2), QE3 (2012Q3–2014Q4), and QE4 (2020Q1–2021Q4).

2.4 Summary Statistics

Table 2 reports summary statistics for the merged dataset. Panel (a) presents average portfolio shares and assets under management (AUM) by sector type. Foreign investors allocate the largest share of their portfolios to Treasuries, 28.9 percent, and arbitrageurs and banks the largest to agency MBS, 11.5 and 10.8 percent. Households and pension funds hold small fixed-income shares despite their large aggregate AUM. Most of each market's float sits with sectors that, Section 4 shows, do not trade on yields: in 2025Q1 foreign investors, banks, pension funds, and insurers together hold 56 percent of the private coupon float, and the two sectors whose active demand does respond to Treasury yields, the household residual and the mutual funds, hold about a seventh each. This concentration is what makes the slope of the Treasury market's aggregate demand curve small. Panel (b) reports the characteristics of the two assets. Their yields average almost the same, 4.8 percent for Treasuries and 4.7 for agency MBS, but the Treasury index is the longer one, 8.1 years of maturity against 6.8 and 5.5 years of duration against 4.8, prepayment compressing the effective duration of MBS. The estimation sample covers 26 sectors observed quarterly from 1985Q1 to 2025Q1.

Table 2: **Summary Statistics**

(a): Portfolio Summary

	Financial Assets ($T)		Portfolio Weight (%)	
Sector Type	Total	Bond	Treasury	Agency MBS
Banks	12.87	1.91	3.60	10.78
Insurers	6.24	0.59	4.42	7.08
Pension Funds	13.21	0.74	4.29	3.05
Mutual Funds	12.07	1.41	5.67	7.13
Arbitrageurs	3.22	0.45	3.39	11.50
Foreign	19.68	6.62	28.94	4.85
Household Residual	50.04	0.79	1.09	0.49
Other	30.91	1.52	2.23	3.72

(b): Asset Summary

Asset	Variable	Mean	Median	S.D.	Min	Max	N
	Yield (%)	4.77	4.44	2.39	0.66	11.65	161
	Maturity (years)	8.06	8.16	0.94	6.08	9.88	161
Treasury	Issuance (log $M)	15.50	15.16	0.89	14.06	17.11	161
	Duration (years)	5.49	5.35	0.63	3.93	7.11	161
	Coupon (%)	5.29	4.97	2.83	1.57	11.55	161
	Yield (%)	4.71	4.54	2.03	1.31	9.60	139
	Maturity (years)	6.75	6.78	1.63	2.15	12.31	161
Agency MBS	Issuance (log $M)	15.28	15.63	0.86	13.22	16.26	161
	Duration (years)	4.76	4.84	0.84	1.84	6.38	161
	Coupon (%)	6.06	5.50	2.32	2.60	10.73	161

Notes: Panel (a) reports sector-level financial assets and portfolio allocations for the two bond markets the system prices, drawn from the Z.1 Flow of Funds. The Treasury weight is the coupon market, the non-bill series of Section 2. Portfolio weights are time-series averages, weighted by sector financial assets. Panel (b) reports summary statistics for yield, maturity, issuance, duration, and coupon at the asset–quarter level, the unit at which the characteristics vary (N is the number of quarters; the agency-MBS yield begins in 1990Q3).

3 Empirical Framework

This section develops a demand system for the two U.S. bond markets the Federal Reserve trades, Treasury securities and agency MBS, identified from sector-level portfolio data. Four features set it apart from existing bond demand systems. First, demand is measured from transactions rather than from the change in the market value of holdings, so the yield coefficient is identified from trading. Second, the outside asset differs across sectors: I measure each sector's bond demand against the non-bond class it actually rebalances toward, loans for banks, corporate equities for long-term investors, and repo for money funds and dealers (e.g., Rodnyansky and Darmouni, 2017; Kandrac and Schlusche, 2021; Kurtzman et al., 2022). Table 1 lists the outside asset for each sector and Appendix D.1 gives the rationale. Third, each sector's yield sensitivity differs by asset class, so the system can say who absorbs Treasuries and who absorbs agency MBS instead of imposing one slope per investor across both markets. Fourth, demand adjusts at a sector-specific speed, estimated horizon by horizon with local projections, which separates the impact slope from the long-run slope the counterfactual uses.

The sector-specific outside asset has a cost and a benefit. The cost is that the slopes are not directly comparable across sectors, since each is defined relative to a different outside asset; the aggregate demand slope of Section 3.4, the object the counterfactual requires, aggregates across sectors whatever outside asset each uses. The benefit is that the substitution margin is the one the portfolio-balance channel of QE and QT operates on. Jansen et al. (2024), Fang et al. (2025), and Eren et al. (2026) measure demand against other bonds, and investors substitute far more readily between one bond and another than between bonds and equities or loans, so a within-bond outside asset yields a more elastic demand than the bond-versus-non-bond margin that governs the price impact of central-bank purchases.

The use of aggregate shares is analogous to the market-level setup of Berry et al. (1995), without random coefficients, because all sectors face the same prices and characteristics at each date; Appendix A gives the mapping.

3.1 Setup

I index sector-level investors by $i = 1, \ldots, I$; sectors include, for example, life insurers and mutual funds. Assets are indexed by $n = 0, \ldots, N$, where $n = 0$ denotes the outside asset, which differs across sectors and comprises non-bond holdings such as equities and loans. The bond universe consists of Treasury securities and agency MBS, the two markets the Federal Reserve trades. Corporate and municipal bonds are not in it: no instrument for their yields survives the measurement discipline of Section 3.6, so their demand slopes cannot be estimated. They are not folded into the outside asset either: they sit outside the choice set. In the logit that leaves the estimates unchanged, because the log ratio of a bond's share to the outside asset's share does not depend on which other alternatives the investor holds. Appendix D.1 uses them as a common outside asset instead. The yield of asset n in quarter t is $y_t(n)$. Each asset has a vector of observed characteristics $x_t(n)$: maturity, net issuance measured from transactions, and duration. The characteristics enter the estimation predetermined; in particular, duration enters lagged one quarter, because the change in modified duration is a near-mechanical function of the contemporaneous yield change (Section 2). The portfolio weight of investor i in asset n at time t is $w_{it}(n)$, with $\sum_{n=0}^{N} w_{it}(n) = 1$.

3.2 Portfolio Choice

I assume that portfolio shares follow a characteristics-based logit demand, under which the share is exponential in a linear index of characteristics (Koijen and Yogo, 2019).[3] The portfolio weight $w_{it}(n)$ is

$$w_{it}(n) = \frac{\delta_{it}(n)}{1 + \sum_{n=1}^{N} \delta_{it}(n)}, \tag{3}$$

where

$$\delta_{it}(n) = \exp\Big\{\beta_i(n)\, y_t(n) + \gamma_i' x_t(n) + \xi_{it}(n)\Big\}\epsilon_{it}(n). \tag{4}$$

[3] Koijen and Yogo (2019) provide an asset pricing–based theoretical foundation for logit demand by deriving it from mean–variance portfolio choice under factor-structure assumptions.

$\beta_i(n)$ is investor i's demand semi-elasticity with respect to the yield of asset n, and γ_i (a vector) collects investor i's loadings on asset characteristics. $\xi_{it}(n)$ denotes an unobserved demand shifter capturing asset-time variation not explained by observed characteristics, and $\epsilon_{it}(n)$ is the latent demand not explained by the observed yield or characteristics. The outside-asset weight is

$$w_{it}(0) = \frac{1}{1 + \sum_{n=1}^{N} \delta_{it}(n)}. \tag{5}$$

The yield slope carries an asset index. Existing bond demand systems impose $\beta_i(n) = \beta_i$: a one-percentage-point rise in the yield of any bond shifts investor i's log-odds of holding it by the same amount. I let the slope differ by asset, the alternative-specific-coefficient logit familiar from discrete choice. The share structure is unchanged: substitution across assets still follows the logit's proportional pattern, and the outside asset still absorbs the balance, so the own- and cross-yield responses of Section 3.4 keep their closed forms. What the generalization drops is only the equal-slope restriction, and Section 4 shows that the data reject it: mutual funds, for example, absorb Treasuries but not the mortgage spread, and banks absorb agency MBS but not Treasuries. The common-slope system is the special case $\beta_i(n) = \beta_i$.

The price that enters each block, $y_t(n)$, is the one that market itself sets: the ten-year yield for Treasuries and, for agency MBS, the option-adjusted spread over it, since the MBS yield is the Treasury yield plus that spread and the Treasury leg carries eight times its variance. Writing MBS demand on the yield level instead would impose that a sector's share responds equally to the Treasury yield and to the spread; Section 4 tests the two restrictions against each other. With this choice the clearing system of Section 5 is block-recursive: the Treasury market sets the level, the agency MBS market sets the spread, and the mortgage-relevant yield is their sum.

Combining (3) and (5) yields the estimable specification:

$$\ln \frac{w_{it}(n)}{w_{it}(0)} = \ln \delta_{it}(n)$$
$$= \beta_i(n)\, y_t(n) + \gamma_i' x_t(n) + \xi_{it}(n) + \varepsilon_{it}(n). \tag{6}$$

The left-hand side is the log portfolio share of asset n relative to the outside asset. The right-hand side is a linear combination of the asset's yield $y_t(n)$, characteristics $x_t(n)$, and asset–time demand shifter $\xi_t(n)$ with estimable coefficients; $\varepsilon_{it}(n)$ is an idiosyncratic demand shock, the log of the latent demand $\epsilon_{it}(n)$. I parameterize the asset–time demand shifter as

$$\xi_{it}(n) = \alpha_n + \nu_{in}' W_t + \tilde{\xi}_{it}(n), \tag{7}$$

where W_t is a vector of macroeconomic controls, ν_{in} captures sector- and asset-specific loadings on macro conditions, and $\tilde{\xi}_{it}(n)$ is the residual unobserved demand shifter. I follow Jansen et al. (2024) and Eren et al. (2026) in parameterizing the asset–time demand shifter via macro controls rather than asset–time fixed effects, which would absorb the time-series variation in the instruments.

3.3 Rebalancing Speed

Differencing the logit demand (6) is a pure algebraic transformation that removes the asset fixed effect α_n while leaving every slope unchanged,

$$\Delta q_{it}(n) = \beta_i(n)\, \Delta y_t(n) + \gamma_i' \Delta x_t(n) + \nu_{in}' \Delta W_t + \Delta \varepsilon_{it}(n), \tag{8}$$

where $\Delta q_{it}(n)$ is the change in $\ln\big(w_{it}(n)/w_{it}(0)\big)$ that the sector brings about by trading, measured from its transactions as in Equation (2). The recorded holdings also move when the positions already held are revalued, and that part is not a portfolio choice, so it is left out of the left-hand side (Section 2). Regressing changes on changes also avoids the spurious-regression problem that the trending yield level poses over a four-decade sample.

Equation (8) recovers the semi-elasticity $\beta_i(n)$, however, only if sectors complete their rebalancing within the quarter, and they do not. Every sector that absorbs supply in Section 4 completes only a fraction of its rebalancing in the quarter of the shock, held back by asset–liability mandates, illiquid positions, and the slow-moving capital of Duffie (2010). Estimated on one-quarter changes, (8) therefore delivers only the *impact* response, small for a slow sector even when its full response is large. The framework already gives each sector its own outside asset, the margin along which it rebalances; I now give each sector its own speed.

Sluggish adjustment is not an implication of the static logit, which describes the portfolio a sector ultimately wants; it is an assumption about how observed holdings reach that target. I treat Equation (6) as the target portfolio and let observed positions converge to it over time, so the horizon-h slope measures how much of the move toward the new target is complete h quarters after the shock. This is the structure van der Beck (2026) formalizes for equity institutions, where demand elasticities rise with the horizon as slow capital arrives, and its limit is the static parameter: under instantaneous adjustment the horizon slopes are flat at $\beta_i(n)$. The local projections below recover the adjustment path without imposing its shape.

To do so, I measure, for each horizon $h = 0, 1, 2, \ldots$, how far the position has moved h quarters *after* the yield change, with instrumental-variables local projections (Jordà, 2005):

$$\sum_{s=t}^{t+h} \Delta q_{is}(n) = \beta_i^{(h)}(n)\, \Delta y_t(n) + \gamma_i' \Delta x_t(n) + \nu_{in}' \Delta W_t + \tau_{in}\, t + u_{it}^{(h)}(n), \tag{9}$$

where $\Delta y_t(n)$ is instrumented by the asset's instrument set (Section 3.6), the linear trend $\tau_{in}\, t$ absorbs the secular comovement of the long sample, and the moving-average error $u^{(h)}$ calls for Newey–West standard errors, with the bandwidth set to the horizon h (Newey and West, 1987). The left-hand side cumulates the transaction growth between quarters t and $t + h$, so (9) regresses an $(h{+}1)$-quarter change of the position on a one-quarter change of the yield; the differencing window of the yield stays fixed, only the response window of the position extends. This is deliberately *not* a long-difference specification, which would also difference the yield over $h+1$ quarters: lengthening the yield window lets the four-decade secular comovement of yields and portfolio shares back into the regressor, the spurious-trend problem that differencing was meant to remove. At $h = 0$, (9) collapses back to (8) and $\beta_i^{(0)}(n)$ is the impact slope; as h grows, $\beta_i^{(h)}(n)$ rises toward the sector's full response and flattens once rebalancing is complete. I estimate (9) for each sector type and asset separately, pooling the member sectors of the type, so the slope is a holdings-relevant average of the sectors that actually hold the asset.

The long-run slope, which I read one year out at $h = 4$ once the response has flattened, is the semi-elasticity $\beta_i(n)$ of the logit demand (6). The identification requires no particular adjustment process: whatever the frictions, once rebalancing has run its course the allocation is the one the static system describes, so the cumulative response flattens exactly at the static slope. The speed of adjustment is a rebalancing friction, not a demand parameter: the preferences in (6) and every object built on them, the aggregate demand slope, market clearing, and the counterfactual equilibrium, are unchanged, and the frictions govern only how quickly observed holdings reach the allocation that (6) prescribes.[4] The counterfactual's long-run equilibrium uses $\beta_i(n)$; the transition to it uses the estimated path $\beta_i^{(h)}(n)$.

[4]This reconciles why levels and difference estimates of the same demand equation disagree. A levels regression weights low-frequency variation, where adjustment is complete, and recovers a slope near the long-run $\beta_i(n)$, but conflates it with the four-decade secular comovement of yields and shares and needs a near-unit-root instrument whose relevance can be spurious; a quarterly difference weights high-frequency variation and recovers the impact slope $\beta_i^{(0)}(n)$. The local projections (9) recover the long-run slope from differenced, trend-controlled data, avoiding both, and unlike a single levels coefficient trace the adjustment path.

3.4 Demand Elasticity and the Aggregate Demand Slope

The logit demand system yields a closed-form expression for the yield elasticity of demand. The elasticity of the portfolio share of asset n with respect to the yield of asset m is

$$\eta_{it}(m,n) = \frac{\partial \log w_{it}(n)}{\partial \log y_t(m)} = \begin{cases} \beta_i(n) \cdot \big(1 - w_{it}(n)\big) \cdot y_t(n) & m = n \\ -\beta_i(m) \cdot w_{it}(m) \cdot y_t(m) & m \neq n, \end{cases} \tag{10}$$

where $\beta_i(n)$ is the long-run yield sensitivity of sector i for asset n from Section 3.3, $w_{it}(n)$ is the portfolio share of asset n, and $y_t(n)$ is the yield of asset n in quarter t. The own-elasticity ($m = n$) gives the percentage change in the portfolio share per one-percent rise in the asset's own yield; the cross-elasticity ($m \neq n$) is the spillover to other assets, which scales with the shocked asset's slope and portfolio share.

The object that governs the yield effect of a supply shift is not the elasticity but a yield-free slope. Linearizing the market-clearing condition of the next subsection around the observed equilibrium, a supply shift $\Delta S_t(n)$ must be absorbed by the sectors' share adjustments, $\sum_i A_{it}\, \beta_i(n)\, w_{it}(n)\big(1 - w_{it}(n)\big)\, \Delta y_t(n) = \Delta S_t(n)$, so removing 1 percent of asset n's private float raises its yield by

$$\Delta y_t(n) = \frac{1}{\sum_i s_{it}(n)\, \beta_i(n)\, \big(1 - w_{it}(n)\big)} \text{ basis points,} \tag{11}$$

where $s_{it}(n) = A_{it}\, w_{it}(n)/S_t(n)$ is sector i's share of the float and A_{it} its assets under management. The denominator is the slope of the aggregate demand curve for market n, which I call the aggregate demand slope. It aggregates the holders' own-asset sensitivities weighted by their share of the float: a market whose float sits with sectors that do not trade on yields has a small slope and a large price impact, whatever those sectors' behavior in other markets. Investors respond to yields point for point, so the yield level cancels and the basis-point impact is pinned down by this slope alone; its movement over time reflects the composition of holdings. Equation (11) is the first-order version of the counterfactual in Section 5, which clears the full nonlinear system.

3.5 Market Clearing

To close the asset pricing model, the market for each asset n clears in each period t:

$$\sum_i A_{it} \cdot w_{it}(n) = S_t(n), \tag{12}$$

where A_{it} is investor i's total wealth at time t and $S_t(n)$ denotes the outstanding supply of asset n at time t. Market clearing holds in this study because, by construction, the aggregate share data cover all investors and sectors. The counterfactual imposes it in both markets.

3.6 Identification

The yield is endogenous: sector-level demand shocks move prices, the more so in a setting that aggregates investors into sectors rather than assuming atomistic behavior. I instrument each asset's yield with an instrument specific to that asset, and require that its variation move portfolio shares only through that yield, after macroeconomic controls. Because the estimating equations are the differenced specification (8) and the local projections (9), whose horizon-h error cumulates the latent-demand revisions from t through $t + h$, the restriction applies over the projection window:

$$\mathbb{E}\big[\Delta\varepsilon_{i,t+s}(n) \mid Z_t(n)\big] = 0, \qquad s = 0, 1, \ldots, h. \tag{13}$$

At $h = 0$ this is the standard contemporaneous exclusion restriction; at longer horizons it adds the lead-exogeneity condition of LP-IV (Stock and Watson, 2018), that the instrument be unpredictable not only for current latent demand but for its revisions over the year the response cumulates.

Treasury Yields. No external supply shifter is available for Treasury yields over four decades: the Fed's own transactions track the yield cycle, maturity walls (Almeida et al., 2012; Harford et al., 2014) are predetermined but too small to be relevant, realized issuance responds to rates, and foreign demand chases yields. I therefore identify the Treasury slopes with the heteroskedasticity-based instrument of Lewbel (2012). At the asset–quarter level I regress the yield change on the differenced characteristics, take the residual $\hat{e}_t(n)$, and interact it with the demeaned coupon,

$$Z_t(n) = \left[\Delta x_t(n) - \overline{\Delta x(n)}\right] \hat{e}_t(n). \tag{14}$$

The instrument is valid if the characteristic is uncorrelated with the product of the demand and yield shocks, and relevant if the variance of the yield shock moves with the characteristic. The two conditions pull against each other in bond markets: yield-shock variance is highest in crisis quarters, and in exactly those quarters flight to quality makes demand and yield shocks covary. I therefore identify from calm-period heteroskedasticity alone, setting the instrument to zero in the eight quarters 2008Q3–Q4, 2020Q1–Q2, and 2022Q1–Q4 and controlling for them in the second stage. Alongside it I use the Treasury leg of the purchase-path revisions described next, dated to the quarter in which each change first became public rather than to the statement that formalizes it, since the Committee signals a change before it votes on it. Neither instrument is strong, so a weak-instrument-robust Anderson–Rubin set (Anderson and Rubin, 1949; Andrews et al., 2019) accompanies every slope.

Agency MBS Yields. The Fed is the one buyer of agency MBS whose purchases are not a response to any private sector's demand, but its *realized* purchases follow a monthly schedule announced in advance and carry no surprise. The information is in the revisions to that schedule. For each of the twenty-one FOMC communications between November 2008 and May 2024 that changed the announced path of System Open Market Account purchases, I record the implied revision to expected net purchases over the following four quarters, separately for agency MBS and for Treasuries and scaled by the outstanding stock. An asset-specific supply shifter has to move the price of its own market and leave the other's untouched, and Section 4 reports that placebo, along with the same test applied to the alternatives. The instrument is set to zero in the same eight crisis quarters as the heteroskedasticity instrument, where the announcement and the flight to quality are inseparable. Appendix D.4 reports the dating, the event-by-event record, the menu of alternatives, and the checks that the revisions do not respond to the previous quarter's yields or macroeconomic conditions.

3.7 Non-Logit Demand Sectors

The demand system asks each sector to choose between a bond and an outside asset at the bond's yield. Two sectors do not fit that description, and both of them hold Treasuries.

Money market funds hold bills. Rule 2a-7 caps a money fund's weighted average maturity at sixty days and each holding at 397 days, so the ten-year yield is not the price its Treasury demand responds to. Bills are 78 percent of money funds' Treasury holdings today and 51 percent on average since 1974, against 0.7 percent for the long-duration mutual funds with which the Financial Accounts group them. I therefore estimate the mutual-fund type without money funds and treat the money-fund cell as inelastic in the clearing: their holdings stay in the float and their response is zero.

Dealers are intermediaries. A dealer's Treasury book is inventory financed by repo, not a portfolio choice: when yields rise the book loses value and a levered intermediary sells to restore its balance sheet. An intermediary holding leverage constant would sell

$$\frac{\partial \ln H_t(\text{dealer, Treasury})}{\partial y_t} \;=\; -\,\lambda_t\, D_t\, s_t^{\text{bond}}/100 \;=:\; \text{rule}_t, \tag{15}$$

where λ_t is leverage from the Financial Accounts, D_t the index duration, and s_t^{bond} the bond share of financial assets, all lagged one quarter so that rule_t is predetermined. I replace the dealer cell's free slope with $\varphi\,\text{rule}_t$ and estimate the single parameter φ on the same instruments, using $Z_t\,\text{rule}_t$ as the instrument for $\text{rule}_t\,\Delta y_t$, so the cell's contribution moves with the observed balance sheet. The rule is applied to the Treasury cell only; Section 4 reports φ in both markets.

The estimation sample is part of the identification. The household residual's Treasury slope is estimated on the 1985–2025 span, because the 2010–2025 window alone does not identify it (Section 4 and Appendix D.3), so the Treasury counterfactuals rest on the pre-2010 behavior of that sector. The agency MBS absorbers are identified in both windows.

3.8 Scope of the Demand-System Approach

The answers here are valid within the known limits of the demand-system approach (Fuchs et al., 2023; Koijen and Yogo, 2025). First, a logit demand system need not correspond to a valid stochastic discount factor or rule out arbitrage (Fuchs et al., 2023; Kocherlakota, 2025); I use it as a description of how holders trade, not as a model of prices derived from preferences. Second, the logit restricts substitution: letting the slope differ by asset relaxes the restriction within each holder, but the cross-asset responses still follow from the portfolio shares rather than being estimated separately, and the counterfactuals inherit that assumption. Third, identifying a demand system requires as many independent sources of price variation as priced assets (Fuchs et al., 2025). With two asset classes and an instrument set for each, the system meets that requirement, and because the instruments are aggregate time-series shocks, the slopes carry the aggregate response of demand rather than only its cross-section (Haddad et al., 2025). Finally, the slopes are identified from quarterly variation around the observed equilibrium, while the largest programs move supply far outside it, so the counterfactuals are most reliable per unit of supply and least reliable for the largest episodes.

4 Results

4.1 First Stage

Table 3 reports the first stage on the 1985–2025 asset–quarter panel, with the same controls as the second stage and Newey–West standard errors. Both markets the Fed trades are identified by two instruments each of which has the sign theory predicts. The Treasury block pairs the coupon interaction ($t = -3.7$) with the revision to the Fed's announced Treasury purchase path ($t = -2.1$) for a joint F of 7.8, and the overidentifying restriction survives in every sector type. The agency MBS block is the stronger of the two: an upward revision to the announced MBS purchase path compresses the option-adjusted spread with an F of 16.5 ($t = -4.1$). The Fed's realized transactions reach only 3.4, because the realized flow follows the schedule the announcement already set. The revision passes the cross-asset placebo an asset-specific supply shifter has to pass, leaving the Treasury yield at an F of 1.1, and every alternative shifter I examined fails it (Appendix Table A.18). Because the yield varies only at the asset–quarter level, the first stage is common to every sector, so differences in the estimated slopes reflect demand heterogeneity rather than differences in the identifying variation. I read every slope through its weak-instrument-robust Anderson–Rubin set and call a slope different from zero only when that set excludes zero. Any remaining bias has a known direction: weak instruments pull two-stage least squares toward ordinary least squares, which the simultaneity of demand pushes down, so the point estimates are if anything lower bounds.

Two measurement choices carry the identification. For Treasuries the market-value granular instrument reaches a first stage of 10.4 and looks like the strongest instrument available; built from the donors' transactions it reaches 4.3, so most of its relevance was the revaluation that also violates its exclusion restriction, and I drop it. For agency MBS, replacing the change in the market value of the Fed's holdings with its transactions lowers the first stage from 7.3 to 3.5, and replacing the realized transactions with the revisions to the announced

Table 3: **First Stage**

Asset	Instrument	Coefficient	First-stage F	N
Treasury	Lewbel (coupon)	-4.292^{***}	7.8	159
		(-3.65)		
	Fed Treasury purchase-path revision	-2.581^{**}		
		(-2.12)		
Agency MBS	Fed MBS purchase-path revision	-5.581^{***}	16.5	138
		(-4.06)		
Joint F (all instruments)			7.2	297

Notes: This table shows the first stage, estimated asset by asset on 1985Q1–2025Q1 quarterly data, with Newey–West t-statistics (Bartlett kernel, four lags) in parentheses; the yield varies only at the asset–quarter level, so the first stage is common across sectors. Treasury yields are instrumented by a crisis-robust Lewbel (2012) heteroskedasticity instrument, the demeaned change in the coupon interacted with the first-step residual and set to zero in the eight crisis quarters (2008Q3–Q4, 2020Q1–Q2, 2022Q1–Q4), together with the revision to the Fed's announced four-quarter Treasury purchase path, scaled by the outstanding stock and zeroed in the same quarters; the agency MBS option-adjusted spread (available from 1990Q3, hence $N = 138$) by the revision to the Fed's announced four-quarter purchase path, scaled by the outstanding stock and zeroed in the same eight quarters. All regressions include the second-stage controls: the first-differenced asset characteristics (maturity, transaction-based net issuance, lagged duration), the differenced macro controls, an asset-specific trend, and the crisis indicator. The first-stage F is the Newey–West Wald statistic of each asset's own instruments jointly. The final row reports the joint F of all instruments on the stacked asset $\times$ quarter panel with asset fixed effects; because the instrument matrix is block-diagonal, this statistic averages the blocks, and the per-asset F is the informative measure of each yield's identification. Anderson–Rubin sets accompany every slope in Table 4.

purchase path raises it to 16.5: the realized flow is anticipated and the revisions are not.

Three further checks bear on the identification (Appendix Table A.19). First, a three-donor Treasury block leaves the long-run slopes essentially unchanged at a first-stage F of 6.8, so the donor choice affects precision rather than the point estimates. Second, adding the Fed's realized transactions to the purchase-path revision raises the agency MBS first stage from 16.5 to 23.7, but the Sargan test then rejects in four of the eight sector types, banks, insurers, mutual funds, and arbitrageurs, at $p \leq 0.02$: the anticipated flow and the announcement surprise disagree about what the data say. I keep the single instrument, just-identified, and read its slopes through their Anderson–Rubin sets. Third, LIML and Fuller estimates coincide with two-stage least squares in every cell, so the weak-instrument bias of the point estimates is not the issue; the width of the sets is.

4.2 Demand Slopes by Sector

Table 4 reports, for each sector type and each of the two markets the Fed trades, the impact slope $\beta_i^{(0)}(n)$, the long-run slope $\beta_i^{(4)}(n)$, and the 95 percent Anderson–Rubin set for the long-run slope. I take the long-run slope as the cumulative response one year after the shock, a horizon fixed in advance by which the response has developed for every cell; the paths in Figure 2 are flat from $h \approx 4$ on. Each market has its own, thin set of active absorbers, and the set differs across markets.

In the Treasury market the marginal buyer is the household residual, whose long-run slope is 0.72: a one-percentage-point rise in the 10-year yield raises its Treasury position by about one percent relative to its equity holdings. Mutual funds are second at 0.29, once money market funds are taken out of that sector type (Section 3.7); the money-fund cell has a slope of -0.67 against the ten-year yield, and pooling it in halves the type's slope. Banks are third at 0.09, and the rest are flat: pension funds at 0.03, insurers at -0.03, and foreign investors, the largest holders, at 0.01. Only the first two Anderson–Rubin sets exclude zero. The arbitrageur cell, dealers and mortgage REITs, carries a large negative slope, -1.89, which is the momentum of an inventory and basis book rather than a demand curve; the balance-sheet rule of Equation (15) replaces it, with $\varphi = 3.58$ (standard error 2.21), procyclical leverage in the sense of Adrian and Shin (2010), implying a slope of -0.62 at the 2025Q1 balance sheet. On agency MBS the same regression returns $\varphi = -5.57$, the opposite of deleveraging, so the rule is applied to the Treasury cell only. The foreign type pools two units that offset each other: reserve managers are flat, at 0.06, and private foreign investors chase returns, at -0.05.

The household residual is what the Financial Accounts assign to no measured sector, and after 2010 that includes the hedge funds whose cash–futures basis positions dominate its Treasury flows (Jansen et al., 2024). The Enhanced Financial Accounts show how far this has gone: hedge funds hold 122 percent of the residual's coupon Treasuries on average over 2013–2025, and the residual's Treasury transactions load on the quarterly change in hedge-fund holdings with a coefficient of 0.62 (Appendix D.5). The two are the same position, which matters because the slope is identified where the residual was households: the 2010–2025 window on its own gives -0.78 with a standard error of 0.82, an interval that spans every value the counterfactual could use. Whether a levered arbitrageur absorbs supply the way a buy-and-hold household did is the open question behind the Treasury band of Section 5. The cell is also the one most exposed to the measurement of transaction growth, because the residual itself is near zero or negative between 2004 and 2008 (Appendix D.2).

In the agency MBS market, where the price is the spread, every sector type slopes the right way. The household residual leads at 2.58, followed by foreign investors at 0.52, banks at 0.50, and pension funds at 0.34; mutual funds, the arbitrageurs, and insurers are positive but imprecise. Four of the eight sets exclude zero, against two in the Treasury market. That every slope is positive is not imposed: with the MBS yield level as the price, banks, pension funds, and the residual sectors carry negative long-run slopes, which is what a demand curve should never do. The data also settle the choice of price directly. Entering the Treasury yield and the spread separately, each with its own instrument, the level coefficient is indistinguishable from zero in every sector type, while the equality a yield specification imposes is rejected for pension funds ($p = 0.06$) and foreign investors ($p = 0.05$). Foreign absorption runs through both units, the official sector at 0.74 and the private sector at 0.30. In the 2010–2025 window the household residual, the arbitrageurs, banks, and pension funds keep their signs, while mutual funds turn negative and insurers and foreign investors sit at zero.

The data also reject the equal-slope restriction of the standard system. A minimum-distance test of the restriction that each type's Treasury and agency MBS slopes are equal gives 16.6 on 8 degrees of freedom ($p = 0.04$), driven by pension funds, foreign investors, and the household residual (Appendix Table A.21). An intermediate structure, $\beta_i(n) = \beta_i\,\theta(n)$, is not rejected ($p = 0.29$), with loadings of 0.28 on the Treasury yield against 1.39 on the agency MBS spread: the same investor responds five and a half times as strongly to a basis point of spread as to a basis point of level. The counterfactual uses the unrestricted slopes.

4.3 Rebalancing Speed

Figure 2 traces the horizon response $\beta_i^{(h)}(n)$ of every sector type in both markets, three years out. Every absorber builds its position slowly, and in the agency MBS market the build is monotone quarter by quarter. Banks move from 0.12 on impact to 0.25, 0.40, and 0.50 at one year; foreign investors from 0.09 to 0.23, 0.34, and 0.52; insurers from 0.01 to 0.05. The Treasury absorbers do the same, the household residual moving from 0.09 on impact to 0.72 at one year and mutual funds from 0.07 to 0.29. For these cells the impact response is between an eighth and a quarter of the one-year response, and pension funds are the one fast absorber, at 0.18 of 0.34 within the quarter. The slow-moving capital of Duffie (2010) is thus a property of every sector that absorbs supply, not of a particular investor class, and it operates at the scale of quarters: the paths are flat from about the fourth quarter on.

Two cells trade the other way in the quarter of the shock and reverse afterwards. The household residual's agency MBS slope is -1.21 on impact and $+2.58$ at one year, and the

Table 4: **Demand Slopes by Sector Type and Asset**

	Treasury (yield)		Agency MBS (spread)	
Sector type	$\beta^{(0)}$	$\beta^{(4)}$	$\beta^{(0)}$	$\beta^{(4)}$
Banks	+0.02	+0.09	+0.12	+0.50
	(0.04)	(0.11)	(0.07)	(0.27)
		[-0.06, 0.37]		[0.01, 1.03]
Insurers	+0.01	-0.03	+0.01	+0.05
	(0.03)	(0.06)	(0.02)	(0.11)
		[-0.16, 0.14]		[-0.20, 0.25]
Pension Funds	+0.03	+0.03	+0.18	+0.34
	(0.02)	(0.04)	(0.07)	(0.14)
		[-0.04, 0.11]		[0.11, 0.65]
Mutual Funds	+0.07	+0.29	+0.06	+0.12
	(0.04)	(0.13)	(0.11)	(0.31)
		[0.04, 0.64]		[-0.45, 0.80]
Arbitrageurs	+0.11	-1.89	-0.12	+0.13
	(0.35)	(1.29)	(0.23)	(0.45)
		[-7.33, -0.08]		[-0.86, 0.94]
Foreign	-0.03	+0.01	+0.09	+0.52
	(0.02)	(0.07)	(0.07)	(0.25)
		[-0.07, 0.11]		[0.12, 1.18]
Household Residual	+0.09	+0.72	-1.21	+2.58
	(0.13)	(0.28)	(0.73)	(1.02)
		[0.16, 2.07]		[1.15, 5.86]
Other	+0.09	+0.04	+0.15	+0.37
	(0.07)	(0.13)	(0.20)	(0.47)
		[-0.26, 0.33]		[-0.48, 1.32]
N	3,259		2,864	

Notes: Local projections of Equation (9), estimated separately for each sector type and asset on 1985Q1–2025Q1 (agency MBS from 1990Q3), pooling the type's member sectors, with the yield instrumented as in Table 3. $\beta^{(0)}$ is the impact slope and $\beta^{(4)}$ the cumulative response one year after the shock, with Newey–West standard errors in parentheses; the bracketed row is the 95 percent Anderson–Rubin set for $\beta^{(4)}$, ^{u} marking a set that reaches the edge of the grid and $\emptyset$ an empty set. N is the number of sector–quarter observations. The two columns are not in the same units: a Treasury slope is the response to a percentage point of yield, an agency MBS slope the response to a percentage point of spread. The counterfactual uses $\beta^{(4)}$ except in the two cells of Section 3.7, where money market funds enter with a zero slope and the broker-dealer Treasury book follows the leverage rule. Since the mortgage REITs in that type hold no Treasuries, the arbitrageur Treasury slope reported here is never used.

arbitrageurs' is -0.12 and $+0.13$; both are momentum in the shock quarter and rebalancing after it. Their weight is large enough to matter for the market: at the 2025Q1 composition the agency MBS aggregate demand slope is slightly *negative* on impact and turns positive from the first quarter on. The MBS market has no marginal buyer in the quarter a supply shock lands, and acquires one a quarter later. Section 5 reports the consequence, which is that the impact-quarter equilibrium of a supply shift is not defined and the yield path of a runoff becomes defined only from the following quarter.

A static quarterly regression would therefore understate the absorbers' demand four- to eightfold. The distinction matters for the counterfactual: the impact response governs the initial dislocation, the long-run slope where yields settle, and the gap between them is the overshoot a balance-sheet shock produces before slow capital arrives.

Figure 2: **Rebalancing Paths by Sector Type**

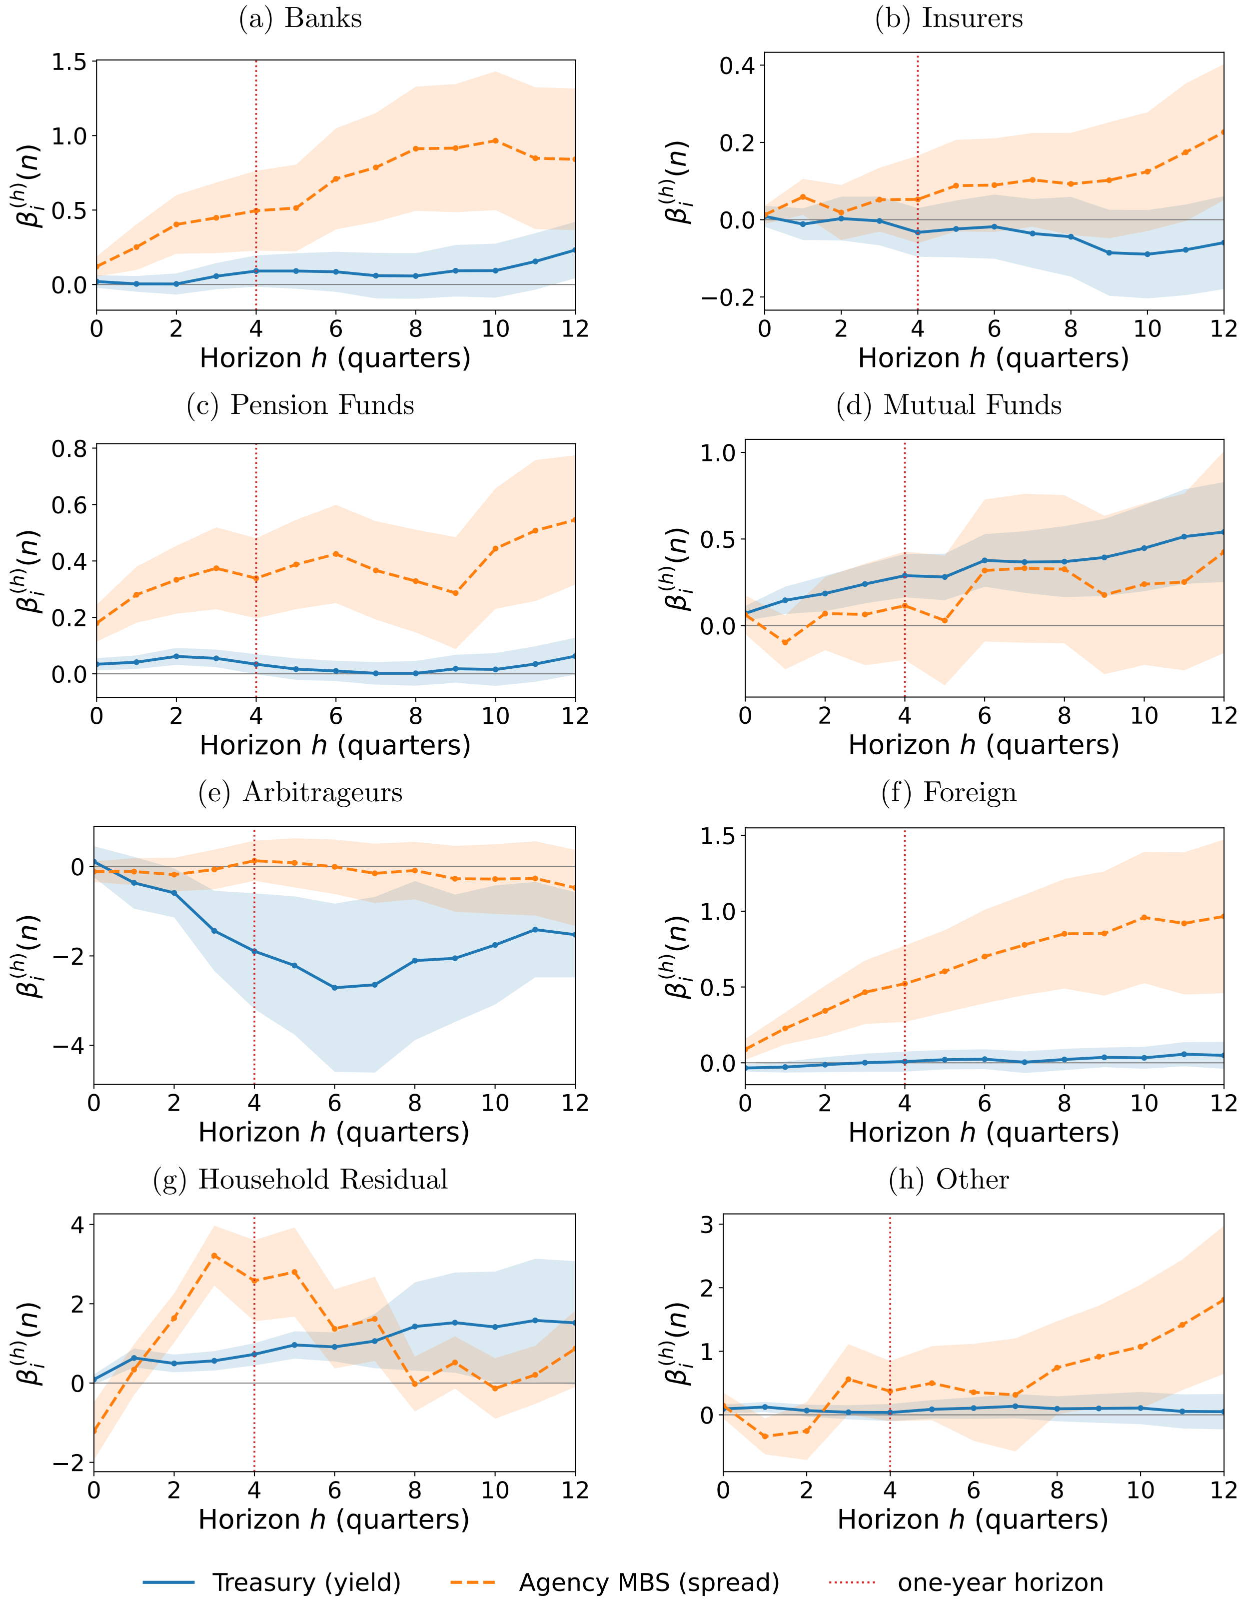

Note: Horizon responses $\beta_i^{(h)}(n)$ from the local projections of Equation (9) on 1985Q1–2025Q1, every sector type and both markets; agency MBS begins in 1990Q3. Bands are ±1 Newey–West standard error and the dotted line marks the one-year horizon, where the counterfactual reads $\beta^{(4)}$. The two prices differ in units, the ten-year yield against the option-adjusted spread, so the lines in a panel compare in shape, not in level. The vertical scale is free across panels.

4.4 The Aggregate Demand Slope

The slopes and the composition of holdings together determine how much a market's yield moves per unit of supply. Figure 3 plots the implied price impact of Equation (11), the inverse of the aggregate demand slope $\sum_i s_{it}(n)\,\beta_i(n)\,(1 - w_{it}(n))$, for Treasuries and agency MBS from 1985 to 2025, and Table 5 reports the 2025Q1 slopes and impacts.

The Treasury slope is thin and has been thin since the turn of the century. In 2025Q1 it is 0.135 (standard error 0.049, $t = 2.8$), so removing 1 percent of the private float raises the 10-year yield by about 7 basis points, with a Fieller band of $[4.3, 25.6]$ that is finite at both ends. The composition explains the number. Foreign investors hold 40 percent of the float and do not respond. Banks, pension funds, and insurers together hold 16 percent and contribute 6 percent of the slope. The household residual holds 15 percent and, with its slope of 0.72, supplies 78 percent of it; the mutual funds hold 14 percent and supply the remaining 22 percent; the arbitrageurs subtract 11 percent. The market's marginal buyers are two sectors holding less than a third of the float between them. Over time the slope tracks that sector's presence: it averaged 0.13 from 1985 to 1999 and fell to 0.045 between 2005 and 2008, reaching its low of 0.027 in 2005Q1, as foreign official reserve managers accumulated Treasuries and the residual emptied. It averaged 0.08 over the QE decade, touching 0.04 in 2014, and has averaged 0.13 since the float shifted back toward the residual, and the hedge funds inside it, after 2022. In the years when the slope is near zero the static system has almost no marginal buyer, and the yield effect of a given supply shift is correspondingly large; Section 5 shows what that means for the 2012–2014 and 2017–2019 episodes.

The agency MBS aggregate demand slope is four times larger and, unlike the Treasury slope, it is estimated with precision. In 2025Q1 it is 0.553 (standard error 0.152, $t = 3.6$), so removing 1 percent of the MBS float widens the option-adjusted spread by 1.8 basis points, with a Fieller band of $[1.2, 3.9]$ that is finite at both ends. The precision comes from the spread specification: under the yield specification the same calculation gives 0.37 with an infinite upper bound. Every sector type contributes, because every slope is positive: banks hold 33 percent of the private float, mutual funds 17 percent, foreign investors 14 percent, and the household residual 11 percent, and none of them subtracts. The contrast with the Treasury market is the paper's central asymmetry in estimated form. Treasury absorption rests on one sector holding a seventh of the float; agency MBS absorption is spread across all eight sector types.

The two markets are also linked, and the direction matters for what follows. Because the MBS block prices the spread, the level of the MBS yield is inherited from the Treasury market: a supply shift in Treasuries passes into mortgage yields one for one, and the agency MBS market determines only the wedge on top. The well-identified part of the MBS result is therefore the spread, and the level component carries the Treasury market's wider uncertainty with it.

The system also implies the full set of own- and cross-price elasticities, and Panel A of Figure 3 plots them. Each of its two panels fixes a one-percent rise in one market's price and traces demand in both. The own elasticities are sizable and move with the price level: Treasury demand responds by 0.47 percent per percent on average and by 0.56 percent at the 2025Q1 composition, falling through the low-yield decade and recovering after 2022, because the elasticity carries the yield level that the aggregate demand slope does not. The agency MBS entries are elasticities with respect to the spread, which averages 0.56 percentage points against a 4.4 point Treasury yield, so their smaller size is the price level and not a weaker response; the underlying semi-elasticity is four times the Treasury one. The cross elasticities sit near zero throughout, an order of magnitude below the own elasticities, because each scales with the share of the shocked asset in the portfolio and no bond is a large share of any sector's balance sheet. Substitution on the demand side therefore runs between bonds and the outside asset, not between the two bonds. The spillover the counterfactual reports is the general-equilibrium counterpart, in which each market re-clears, not direct substitution. Appendix Figure A.7 shows the same object by sector type and Appendix Table A.12 reports the matrix.

Figure 3: **Bond Demand Elasticity and Implied Price Impact**

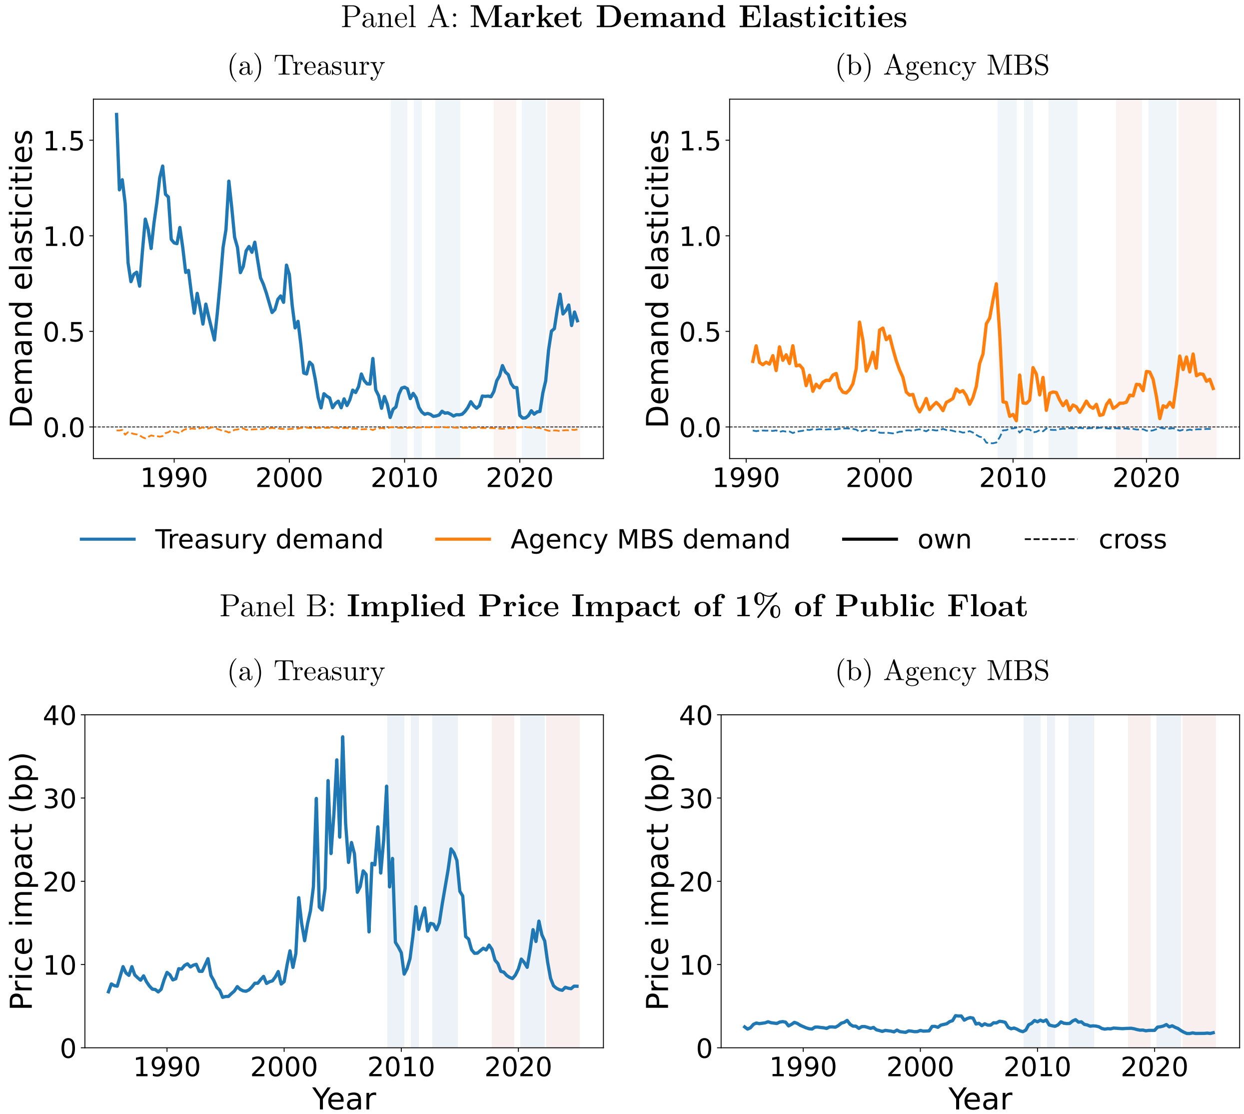

Note: Panel A plots the market elasticities, the averages of the sector elasticities $\eta_{it}(n,n) = \beta_i(n)(1 - w_{it}(n))\, y_t(n)$ and $\eta_{it}(m,n) = -\beta_i(m) w_{it}(m)\, y_t(m)$ weighted by holdings, $A_{it} w_{it}(n)$. Each subpanel raises that market's own $y_t(n)$ by one percent; the own response is solid, the cross response dashed. $y_t(n)$ is the ten-year yield for Treasuries and the option-adjusted spread for agency MBS, whose series begins in 1990Q3. Panel B plots the basis-point yield impact of removing 1 percent of the private float, $1/\sum_i s_{it}(n)\, \beta_i(n)\, (1-w_{it}(n))$ (Equation 11), where $s_{it}(n)$ is sector i's share of that float. It carries no $y_t(n)$ term and moves only with the composition of holdings. Each panel draws both markets on one scale: the agency MBS impact runs between 1.7 and 3.9 basis points, the Treasury impact between 6 and 37. Panel B shows point estimates; Table 5 gives the 2025Q1 Fieller interval. Both panels use the long-run slopes $\beta_i^{(4)}(n)$ of Table 4 with the cell models of Section 3.7, and floor at zero the household Treasury residual, which the Financial Accounts report as negative from 2004Q2 to 2008Q2 (Appendix D.2). Blue bands mark QE episodes, red bands QT episodes.

Table 5: **Aggregate Demand Slopes and Price Impact, 2025Q1**

Asset	Slope $\sum_i s_i \beta_i(n)(1 - w_i(n))$	t	Impact per 1% of float (bp)	95% Fieller band
Treasury	+0.135 (0.049)	2.8	7.4	[4.3, 25.6]
Agency MBS	+0.553 (0.152)	3.6	1.8	[1.2, 3.9]

Notes: The aggregate demand slope of Equation (11) at the 2025Q1 composition of holdings, with its delta-method standard error in parentheses (sector-type slopes treated as independent, and the two cells of Section 3.7 contributing none, since neither slope is estimated by its type's regression), and the implied yield impact of removing 1 percent of the private float, 1/slope, in basis points. The Fieller band inverts the 95 percent interval of the slope; it is unbounded above when that interval includes zero.

5 Counterfactual Simulation

5.1 Setup

The equilibrium asset yields are jointly determined by the estimated demand system in (6) and the market clearing condition in (12). To close the asset pricing model, the market for each asset n clears in each period t:

$$\sum_{i=1}^{I} w_{it}^{*}(\boldsymbol{y}_t^{*},\, \boldsymbol{x}_t,\, \boldsymbol{\beta}_i,\, \boldsymbol{\varepsilon}_{it}) \cdot A_{it} = S_t(n;\, \boldsymbol{y}_t^{*}), \tag{16}$$

where $\boldsymbol{y}_t^{*}$ is the equilibrium yield vector, $\boldsymbol{x}_t$ is the asset-characteristics vector, $\boldsymbol{\beta}_i$ collects investor i's asset-specific demand slopes, $\boldsymbol{\varepsilon}_{it}$ are the latent demand shocks, and $S_t(n;\, \boldsymbol{y}_t^{*})$ is the supply schedule of asset n, $S_t(n;\, \boldsymbol{y}) = f_n\big(\boldsymbol{y};\, \bar{S}_t(n)\big)$, which lets net issuance respond to the equilibrium yield; $\bar{S}_t(n)$, the free float at baseline yields, is the schedule's shifter, and at baseline yields supply equals the observed float. Section 5.2 estimates the schedule's reduced form. The equilibrium yield vector $\boldsymbol{y}_t^{*}$ is the implicit function

$$\boldsymbol{y}_t^{*} = \boldsymbol{g}\big(\boldsymbol{x}_t,\, \boldsymbol{\beta},\, \boldsymbol{\varepsilon}_t,\, \boldsymbol{A}_t,\, \bar{\boldsymbol{S}}_t\big), \tag{17}$$

where $\boldsymbol{g}(\cdot)$ is the implicitly defined equilibrium mapping, which clears every market against the supply schedule, $\bar{\boldsymbol{S}}_t$ is the vector of baseline floats, and $\boldsymbol{A}_t$ is the vector of investor assets under management.

The system is block-recursive in the two markets the Fed trades. Treasury demand responds to the ten-year yield and agency MBS demand to the option-adjusted spread (Section 3.6), so the Treasury market clears the level, the agency MBS market clears the spread, and the mortgage-relevant yield is the sum of the two,

$$\Delta y_t^*(\text{MBS}) = \underbrace{\Delta y_t^*(\text{Treasury})}_{\text{level, from the Treasury market}} + \underbrace{\Delta s_t^*}_{\text{spread, from the agency MBS market}}. \tag{18}$$

Supply is the one place the level re-enters the MBS block: origination responds to the mortgage rate, not to the spread, so the supply schedule of Equation (21) is evaluated at the reconstructed yield. Equation (18) is the decomposition the counterfactual reports throughout; a specification that prices MBS demand on the yield level cannot produce it, because its MBS market clears a price that already contains the Treasury level.

I impose clearing in those two markets.

I evaluate the impact of counterfactual monetary policy by comparing baseline and counterfactual equilibrium yields. Changes in the Fed's asset holdings translate into changes in the baseline private free float $\bar{\boldsymbol{S}}_t$, and I measure them as the Fed's transactions over the episode, the sum of its quarterly net purchases in the Z.1 flow tables, rather than as the change in the market value of its portfolio, which in 2022 mixes runoff with the revaluation of the securities it held. This one-for-one translation holds the fiscal path fixed: debt the Fed does not roll over is refinanced to the public at auction, and agency MBS paydowns are re-securitized, so runoff adds to the float exactly as outright sales would; any yield-induced change in gross issuance is what the supply schedule captures. Denoting the post-purchase baseline float by $\bar{\boldsymbol{S}}_t^{\text{CF}}$, the yield effect is

$$\Delta \boldsymbol{y}_t = \boldsymbol{g}\big(\boldsymbol{x}_t,\, \boldsymbol{\beta},\, \boldsymbol{\varepsilon}_t,\, \boldsymbol{A}_t,\, \bar{\boldsymbol{S}}_t^{\text{CF}}\big) - \boldsymbol{g}\big(\boldsymbol{x}_t,\, \boldsymbol{\beta},\, \boldsymbol{\varepsilon}_t,\, \boldsymbol{A}_t,\, \bar{\boldsymbol{S}}_t\big). \tag{19}$$

The logit system is nonlinear, and because some holders carry negative slopes it can have several roots. I solve it by continuation, scaling the supply shift from zero to its full size in small steps and following the equilibrium that is continuously connected to the observed one. Where that path breaks, the shock is larger than the market's absorbers can take at that date's composition of holdings, and I report the first-order impact of Equation (11) instead, flagged as such.

5.2 Supply Response

The counterfactual lets bond supply respond to yields: when the counterfactual yield rises, issuers issue less (Fang and Xiao, 2025), attenuating the yield effect. This supply side is deliberately reduced-form, a first-order net-issuance elasticity rather than a structural model of issuer behavior, since the demand system is the paper's structural focus. I estimate these elasticities from the same Flow-of-Funds data as the demand system. The identification mirrors the demand side. Estimating a demand slope requires supply-side variation in yields; estimating the supply (issuance) response requires *demand*-side variation, and the paper's instruments, the crisis-robust heteroskedasticity instrument and the Drukker (2025) surprise, are exactly that: none is driven by issuers. The specification parallels the local projections of the demand system, with the response read at the same one-year horizon $h = 4$ as the demand slope $\beta^{(4)}$: the cumulative net issuance, measured from transactions as a percent of outstanding, through one year after a one-quarter yield change,

$$\sum_{s=t}^{t+4} \frac{F_s(n)}{S_{s-1}(n)} \times 100 = \kappa_i^{(4)}\, \Delta y_t(n) + \text{quarter fixed effects} + \tau\, t + u_t(n), \tag{20}$$

with $\Delta y_t(n)$ instrumented, quarter fixed effects absorbing the Treasury refunding cycle. The implied supply elasticity is $\psi_n = -\kappa_n^{(4)}$, the percent of outstanding withdrawn per percentage-point yield rise per year. The counterfactual's supply schedule is its first-order form,

$$S_t(n; \boldsymbol{y}) = \bar{S}_t(n)\left[1 - \psi_n\, \Delta y_t(n)\right], \tag{21}$$

so relative to the baseline float $\bar{S}_t(n)$, issuers withdraw ψ_n percent of outstanding per percentage point that the counterfactual yield rises above its baseline value.

A supply curve has to be traced out by shifts in demand, just as a demand curve has to be traced out by shifts in supply. The coupon heteroskedasticity instrument identifies the demand side and so cannot identify supply: an instrument excluded from demand and correlated with the yield through supply cannot also be excluded from supply and correlated with the yield through demand. Used in the supply regression it returns a positive coefficient, +5.4 with a standard error of 6.4. I instrument instead with Japanese Ministry of Finance foreign-exchange intervention. When the Ministry sells yen and buys dollars it places the proceeds in Treasuries, so the shock originates in Japanese exchange-rate policy rather than in any decision of the US Treasury, and its first stage on the ten-year yield is 20.0. The sign flips: Treasury net issuance falls by 3.5 percent of outstanding per percentage point of yield per year, with a standard error of 6.1. The Fed's purchase-path revision is a demand shifter too, and I do not use it: it returns -32, which is a third of the outstanding stock per percentage point, it survives a control for federal net borrowing, and it is identified off 2009 and 2020, when the purchase announcements and the fiscal expansions arrived together. Agency MBS issuance responds by -2.7 percent of outstanding per percentage point per year (standard error 3.9), identified by the Drukker (2025) price surprise. Fang and Xiao (2025) estimate the same aggregate object on the same kind of data and report 0.96 for Treasuries and 2.24 for agency MBS, but against the term spread and the MBS spread rather than the yields, so their numbers are a reference point and not a like-for-like benchmark; I do not carry them in the table. Measured from the change in market-value outstanding, every issuance elasticity is larger and more precise, because the market-value change contains the revaluation $-D\,\Delta y$ and so responds to yields by construction. The Treasury value is the one that matters for the headline numbers. At $\psi = 3.5$ the supply side absorbs a meaningful part of a Fed purchase. It absorbs the same part whatever the length of the program. One could scale ψ by the horizon, on the view that issuance keeps responding for as long as the yield stays away from baseline, but the estimated path says the response is a stock adjustment rather than a flow. The Treasury $\kappa^{(h)}$ peaks at one year, -3.5, and is back to zero by two, $+2.1$, where linear accumulation would have predicted -7. Each episode therefore receives $\kappa^{(4)}$ once. The choice matters for the comparison of programs of similar size: QE2 and QE3 bought \$808 billion and \$817 billion of Treasuries, and under horizon scaling the ten-quarter program looked a third the size of the three-quarter one.

Table 6: **Net-Issuance Elasticities Estimated from the Flow of Funds**

	One-year response $\kappa^{(4)}$		First-stage	Implied ψ
Asset	Coefficient	(s.e.)	F	(%/pp/yr)
Treasury	-3.54	(6.07)	20.0	3.54
Agency MBS	-2.73	(3.91)	2.8	2.73

Notes: Instrumental-variables local projections of cumulative net issuance (transactions, percent of lagged outstanding) over the year following a one-quarter yield change, per asset, on 1985Q1–2025Q1; Newey–West standard errors in parentheses. The Treasury yield is instrumented by Japanese foreign-exchange intervention, a shift in foreign demand, and the agency MBS yield by the Drukker (2025) monetary surprise together with the transaction-based issuance interaction; the crisis indicator is a control. The two blocks are not equally well identified. The Treasury instrument reaches a first stage of 20.0; the agency MBS set reaches 2.8, and 1.9 with the Drukker surprise alone, which raises the point estimate from -2.7 to -5.7. A monetary surprise also moves mortgage origination directly, so it is not a pure demand shifter, and the Fed's own agency MBS purchase announcements cannot replace it: they return $+28.6$, because the Fed announces mortgage purchases when housing is contracting and net issuance is falling. The agency MBS elasticity should therefore be read as weakly identified, and the counterfactual uses the smaller of the two point estimates. The supply schedule is written on the yield, not the spread: origination responds to the mortgage rate. The implied ψ is the elasticity used in the counterfactual's supply response, set to zero where the coefficient is positive. *, **, and *** denote $|t|$ above 1.64, 1.96, and 2.58.

5.3 Assessment of Historical QE/QT Episodes

I apply the structural counterfactual to the four QE programs (QE1–QE4) and the two QT episodes (QT1, 2017–2019; QT2, 2022–2025). For each episode I hold investor characteristics, long-run demand slopes, and AUM at their observed values, let bond supply respond to yields through the net-issuance elasticities of Section 5.2, and ask how yields would differ had the

Fed not transacted. Table 7 reports the long-run effects, the response once every absorber has completed rebalancing.

Four findings stand out. First, the Treasury effects are large, and they scale with the shock. QE1's $300 billion of Treasury purchases compressed the 10-year yield by 51 basis points, QE2's $808 billion by 97, QE3's $817 billion by 103, and QE4's $3.5 trillion by 176; scaled by the share of the private float removed (final column), the programs run from 7 to 11 basis points per 1 percent of float, net of the issuance the supply side returns to the market. They are large because the Treasury market's absorber is thin: the household residual and the mutual funds, less than a third of the float between them, are the only sectors that take the other side, and the foreign, bank, and pension holdings that dominate the float do not move. QE2 and QE3 bought almost the same amount of Treasuries, $808 billion against $817 billion, and move the yield almost the same distance, 97 against 103 basis points; the small gap is that QE2 took a larger share of a smaller float, 12.2 percent against 9.2.

Second, the agency MBS yield moves with the Treasury yield, and what the MBS market itself determines is the spread. Equation (18) splits each entry, and the split is the finding. QE1, the Fed's only MBS-heavy program at $1.2 trillion against $300 billion of Treasuries, compressed the MBS yield by 95 basis points, of which 51 is the Treasury level and 45 is spread compression: this is the one program whose composition was mortgage-specific, and it is the only one with a large spread effect. QE3 and QE4 bought both markets in rough proportion and moved the spread by 24 and 22 basis points; their MBS yield effects, 127 and 198 basis points, are mostly the Treasury level passing through. QE2 is the mirror image of QE1: the Fed bought $808 billion of Treasuries while running off $207 billion of MBS, which *widened* the spread by 15 basis points, so mortgage yields fell 82 basis points against the Treasury market's 97 and homeowners received six-sevenths of the pass-through. Composition, not size, determines what QE does to mortgage borrowers.

The QE1 estimate is also where the demand system now meets the event-study record. Its 95 basis points sit inside the range between Hancock and Passmore (2011) (-85) and

Krishnamurthy and Vissing-Jorgensen (2011) (-107). The reconciliation is mechanical: an MBS demand curve written on the yield level makes the MBS market absorb the whole level shock by itself, and the elastic absorbers then deliver a small number; written on the spread, the level arrives from the Treasury market and only the wedge is absorbed locally.

Third, QE4 is the largest program because its shock was largest, not because demand was steeper. It removed a quarter of the private Treasury float, and scaled by that shock it moved yields by 7.3 basis points per 1 percent, the lowest of the QE programs. Two-thirds of its Treasury purchases came in the first half of 2020, when the Fed bought in weeks to restore market functioning (Vissing-Jorgensen, 2021); re-basing the counterfactual to the stock-QE phase that followed isolates a \$1.2 to \$1.5 trillion program that moved Treasury yields by 75 to 88 basis points (Appendix Table A.11), 9 basis points per 1 percent, because the residual's share of the float was low in 2021.

Fourth, the QT episodes reverse the sign. QT1's \$373 billion of Treasury runoff and \$318 billion of MBS runoff raised the two yields by 21 and 28 basis points, 7 per 1 percent of float in each market, with 7 basis points of that coming from the spread. QT2's \$1.6 trillion of Treasury runoff raised the ten-year yield by 60 basis points and the mortgage yield by 68, of which 8 is spread. The treatment of the two cells in Section 3.7 matters for this episode: with money funds pooled into the mutual-fund type and the dealer cell carrying a free elasticity, the Treasury aggregate demand slope is small enough that no equilibrium connected to the observed one exists at QT2's size. With the two cells modelled, the nonlinear system clears at every episode. The message is the one Section 4.4 anticipated: at the current composition of the float the Treasury market's aggregate demand slope is small enough that a shock of QT2's size is at the edge of what a static demand system can price.

Per unit of supply removed, the Treasury estimates agree with the demand-system literature. Scaled by the shock they run 7 to 11 basis points per 1 percent of float, net of the issuance the supply side returns, against the 8 to 13 basis points per percentage point of the central bank's share that Eren et al. (2026) report and the 13 basis points per 1 percent shift

Table 7: **Yield Effects of Historical QE/QT Episodes**

(a) Quantitative easing

Episode	Asset	Purchases ($B)	ΔY (bp)	of which spread (bp)	ΔY per 1% of float (bp)
QE1	Treasury	300	−50.8	—	9.7
2008Q4–2010Q1	Agency MBS	1,223	−95.3	-44.5	4.9
QE2	Treasury	808	−96.7	—	8.0
2010Q4–2011Q2	Agency MBS	−207	−82.1	14.6	25.6
QE3	Treasury	817	−102.8	—	11.2
2012Q3–2014Q4	Agency MBS	865	−126.6	-23.8	9.2
QE4	Treasury	3,516	−175.8	—	7.3
2020Q1–2021Q4	Agency MBS	1,238	−197.7	-21.9	13.2

(b) Quantitative tightening

Episode	Asset	Runoff ($B)	ΔY (bp)	of which spread (bp)	ΔY per 1% of float (bp)
QT1	Treasury	−373	21.1	—	7.3
2017Q4–2019Q3	Agency MBS	−318	28.1	7.0	7.2
QT2	Treasury	−1,622	59.9	—	7.1
2022Q2–2025Q1	Agency MBS	−549	67.7	7.8	11.9

Notes: Yield effects (basis points) of holding the Fed's portfolio at its pre-episode level, from the market-clearing counterfactual of Section 5. The counterfactual shifts the baseline float by the Fed's transactions over the episode, $\bar{\boldsymbol{S}}_t^{\mathrm{CF}} = \bar{\boldsymbol{S}}_t + \boldsymbol{\Delta}\mathrm{SOMA}_t$, and re-solves the equilibrium in both markets against the supply schedule of Equation (21), with the elasticities of Table 6. All columns are measured at the episode-end quarter. Because the agency MBS block is estimated on the option-adjusted spread, the agency MBS yield change decomposes as in Equation (18) into the Treasury level, which passes through one for one, and the spread the MBS market clears; the "of which spread" column is the second term, and its positive entry in QE2 reflects the Fed's concurrent MBS runoff. The last column scales the impact by the supply shock, $|\Delta Y|$ per 1 percent of the private float removed or returned. The money-market-fund and broker-dealer cells are treated as Section 3.7 describes.

in demand of Chaudhary et al. (2026). Figure 4 places the estimates episode by episode. The two systems are close on QE2 (97 here against their 113) and QE4 (176 against 121), and mine reads QE3 larger (103 against 57). On the QT side mine are smaller, 21 against 43 for QT1 and 60 against 80 for QT2. Announcement-window event studies (Gagnon et al., 2011; Krishnamurthy and Vissing-Jorgensen, 2011) are larger for QE1 and smaller for the later, largely anticipated programs; they price only the surprise component of purchases, on impact, and include the signaling channel that the supply-only counterfactual excludes. QE1, which the post-2010 demand systems cannot reach, enters here at 51 basis points, between the calibrated preferred-habitat estimate of 25 to 30 for the cumulative QE1 to QE3 purchases (Vayanos and Vila, 2021) and the event-study range of 91 to 107. On the mortgage side the agreement is closer: QE1's 104 basis points lie inside the announcement-window range for agency MBS.

Three caveats bound the interpretation, and they do not all run in the same direction. The counterfactual holds latent demand and AUM fixed and treats the supply change as permanent, so the long-run estimates overstate the effect if investors expected the Fed to unwind (Jansen et al., 2024); the demand slopes are identified from local variation, while the episodes shift supply well outside that range, and QT2 shows where the extrapolation breaks; and the aggregate Treasury slope has a wide band, so the per-unit impact of 7.4 basis points at the 2025Q1 composition carries a 95 percent Fieller interval of $[4.3, 25.6]$ (Table 5), finite at both ends but wide at the top. Weak-instrument bias, which pulls the absorbers' slopes toward zero, works the other way and makes the point estimates conservative. The two components of the mortgage number are bounded differently. The spread component is the tightest object in the paper: 1.8 basis points per 1 percent of the MBS float, with a two-sided Fieller interval of $[1.2, 3.9]$. The level component is the Treasury number and inherits its band, so a reader who doubts the Treasury absorber should read the mortgage result as the spread effect plus whatever the Treasury market delivers.

Figure 4: **Historical QE/QT Yield Effects: This Paper versus Prior Estimates**

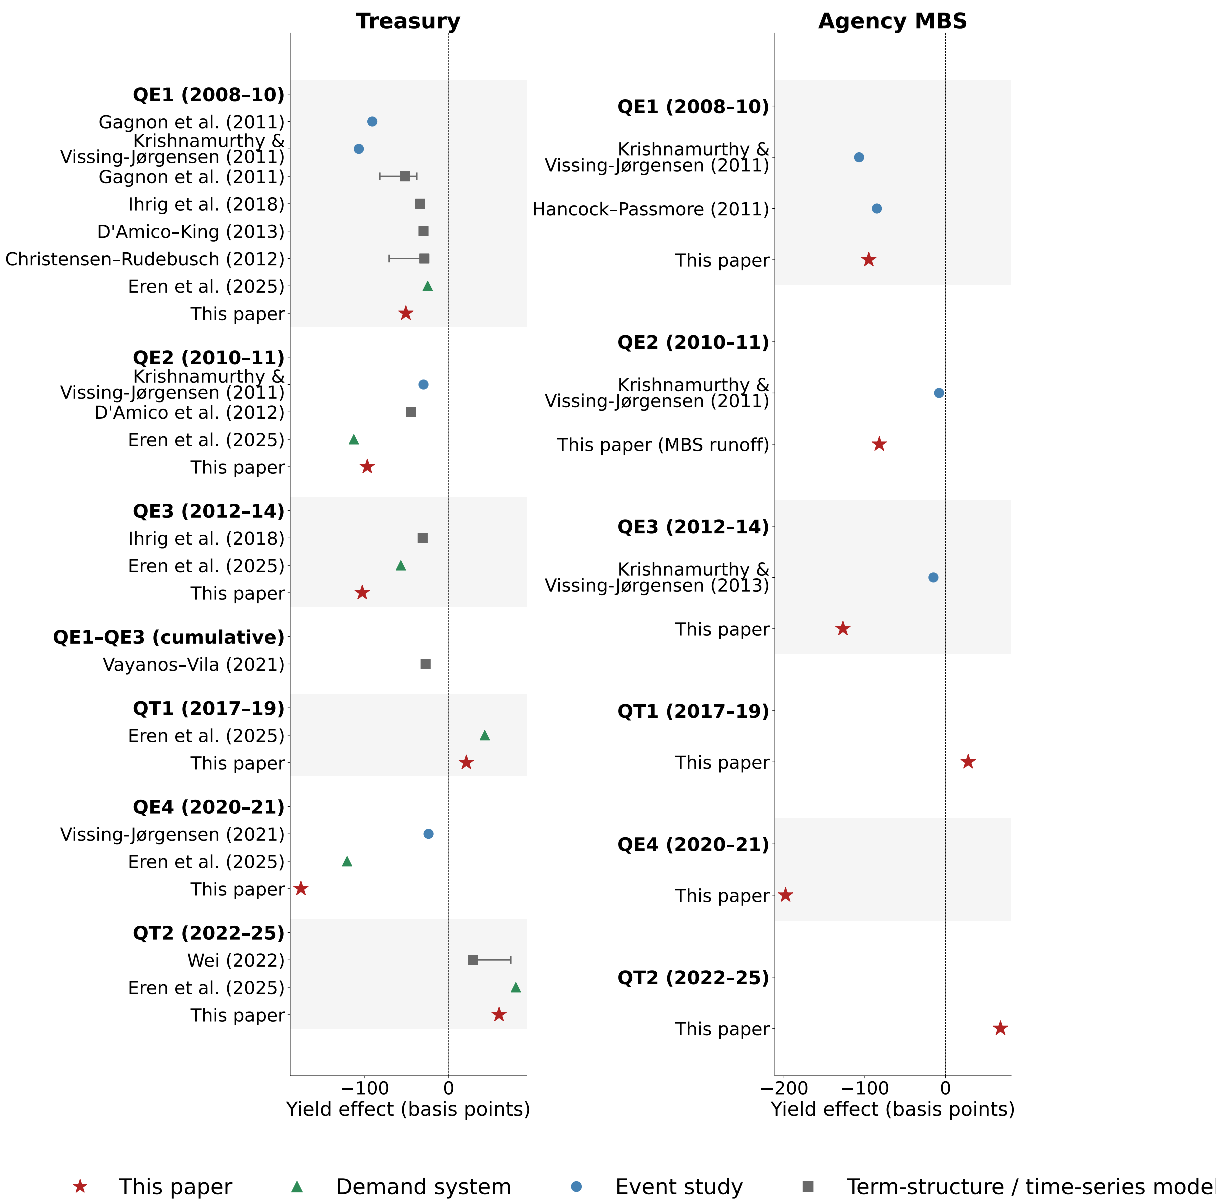

Note: Structural full-stock estimates of Table 7 (stars) against prior estimates of each QE/QT episode, for Treasuries (left) and agency MBS (right). Circles are announcement-window event studies, squares term-structure and time-series estimates, triangles demand-system estimates; horizontal bars are ranges reported in the source study. The Vayanos and Vila (2021) entry is a preferred-habitat calibration of the cumulative QE1–QE3 purchases and appears as its own group, and the QT2 Treasury entry for this paper is a first-order impact. Episode definitions differ slightly across studies.

5.4 Prospective QT Scenario Analysis

I next project the yield effects of a resumption of quantitative tightening. Fixing the 2025Q1 portfolio and running additional runoff at the QT2 Treasury:MBS ratio, I consider three scenarios: a mild \$400 billion, a baseline \$1 trillion, about a year and a half at the Fed's 2024 runoff pace of \$172 billion per quarter, and an aggressive \$1.6 trillion. Table 8 reports the results. A further \$1 trillion raises 10-year Treasury yields by 21 basis points and agency MBS yields by 25, of which 4 basis points is spread widening and the rest is the Treasury level passing through; the mild scenario raises them by 9 and 10, the aggressive by 33 and 39. A back-of-the-envelope calculation conveys the real stakes. Applied to the roughly \$27 trillion of marketable federal debt as it reprices, the baseline Treasury rise adds about \$58 billion a year to federal interest costs, a recurring fiscal cost to weigh against the balance-sheet reduction. On the household side, agency MBS yields pass through nearly one for one to primary mortgage rates, so the baseline adds about \$66 a month to the payment on a newly originated \$400,000 thirty-year fixed-rate mortgage.[5] The marginal Treasury impact is nearly constant, about 2 basis points per \$100 billion of additional runoff. The mortgage cost of quantitative tightening is therefore mostly a Treasury-market phenomenon: a runoff split 73:27 leaves the mortgage spread almost unchanged, and the homeowner pays for the level.

The Treasury number is conditional on the marginal buyer, and I report it as a band rather than a point. Two cells carry the Treasury aggregate demand slope, and neither is pinned down. The household residual's slope of 0.72 accounts for about half the aggregate at the 2025Q1 composition, and it is identified on the pre-2010 sample: in the 2010–2025 window the sector's slope is -0.78 with a standard error of 0.82 (Section 4). The foreign sector carries a larger weight still in the aggregate demand slope, because it holds a third of the float, and its slope is zero under the paper's instruments but 0.48 under the Fed

[5] Fiscal: \$27 trillion $\times$ 21.4 basis points $\approx$ \$58 billion per year once the stock has repriced. Mortgage: the monthly payment on a loan of size P at annual rate r is the standard annuity $M(r) = P\,(r/12)\,/\,[1 - (1 + r/12)^{-360}]$; with $P = \$400{,}000$ and $r = 6.5\%$, $M(r + 25\text{bp}) - M(r) \approx \$66$ per month.

Table 8: **Yield Effects of Prospective Quantitative Tightening**

(a) Baseline

Scenario	Asset	Runoff ($B)	ΔY (bp)	of which spread (bp)	ΔY per 1% of float (bp)
Mild ($400B)	Treasury	−292	8.9	—	5.8
	Agency MBS	−108	10.4	1.5	9.3
Baseline ($1T)	Treasury	−730	21.4	—	5.6
	Agency MBS	−270	25.2	3.8	9.0
Aggressive ($1.6T)	Treasury	−1,168	33.1	—	5.4
	Agency MBS	−432	39.1	6.1	8.7

(b) Sensitivity to Treasury slopes

Scenario	Asset	Runoff ($B)	ΔY (bp)	of which spread (bp)	ΔY per 1% of float (bp)
Baseline	Treasury	−730	21.4	—	5.6
$\beta_{\text{HR}}^{\text{Treasury}} = 0.72$	Agency MBS	−270	25.2	3.8	9.0
Expansion slope	Treasury	−730	31.2	—	8.2
$\beta_{\text{HR}}^{\text{Treasury}}$: baseline $\to$ 0.36	Agency MBS	−270	34.6	3.4	12.3
Recession slope	Treasury	−730	6.5	—	1.7
$\beta_{\text{HR}}^{\text{Treasury}}$: baseline $\to$ 3.31	Agency MBS	−270	11.0	4.5	3.9
Lower AR bound	Treasury	−730	41.1	—	10.8
$\beta_{\text{HR}}^{\text{Treasury}}$: baseline $\to$ 0.16	Agency MBS	−270	44.1	2.9	15.7
Higher foreign slope	Treasury	−730	12.7	—	3.3
$\beta_{\text{Foreign}}^{\text{Treasury}}$: baseline $\to$ 0.48	Agency MBS	−270	17.3	4.6	6.2

Notes: Panel (a) projects the yield uplift, in basis points, from runoff beyond the 2025Q1 balance sheet, fixing the 2025Q1 portfolio and solving the nonlinear clearing system in both markets with $\bar{\boldsymbol{S}}_t^{\text{CF}} = \bar{\boldsymbol{S}}_t + \boldsymbol{\Delta}\text{SOMA}_t$ as the float after the runoff, split 73:27 as in QT2. Supply responds through the elasticities of Table 6. The impacts are long-run equilibrium, at the pooled household Treasury slope and the baseline foreign slope.

Panel (b) re-solves the $1 trillion row of Panel (a) across the two slopes the aggregate demand slope turns on, in the same layout and with the same supply response; each row label gives the slope assumed, with every other slope at its baseline value. The first row is the baseline of Panel (a), at the pooled estimates. The next three move the household residual's slope to the expansion and recession values of the cycle specification (Appendix D.6) and to the lower edge of its Anderson–Rubin set. The last moves the foreign slope from its baseline estimate of 0.01 to 0.48, the value the purchase-path revision delivers alone (Appendix Table A.16). The two cells of Section 3.7 keep their models.

purchase-path revision alone (Appendix Table A.16). Panel (b) of Table 8 re-solves the baseline scenario across both.

The household dimension is the one that can make the effect large, and the business cycle is what moves it. The residual sector's slope is the one cell of the sixteen that differs across states of the cycle: 0.36 in expansions against 3.31 in NBER recessions, a difference significant at $p = 0.038$ (Appendix D.6). A runoff carried out in an expansion therefore raises the ten-year yield by 31 basis points and the same runoff in a recession by 7, against 21 at the pooled slope. The statistical band runs the same way: at the lower edge of the slope's Anderson–Rubin set, 0.16, the effect is 41 basis points. That edge has an economic reading, not only a statistical one. The residual sector is now mostly the hedge funds running the cash–futures basis trade, which hold 122 percent of its coupon Treasuries on average since 2013 (Section 4.2), while the slope applied to it is identified on the pre-2010 sample, when the same residual was households and nonprofits. A low household slope is the case in which the basis trade absorbs less readily than the buy-and-hold investor it replaced. The foreign dimension runs the other way: with foreign investors at 0.48 the same runoff raises Treasury yields by 13 basis points rather than 21. The honest range for a \$1 trillion runoff is therefore 7 to 41 basis points on the ten-year yield, every value of it a genuine equilibrium, and which end applies turns on two sector slopes that four decades of quarterly data do not separate. The Fieller band on the first-order impact tells the same story from the aggregate side: \$1 trillion is 3.8 percent of the Treasury float, and before any supply response the impact per 1 percent is 7.4 basis points with a 95 percent interval of $[4.3, 25.6]$.

The mortgage number splits along the same seam, and the split is what makes it readable. The spread component barely moves across the band, 3 to 5 basis points in every row of Panel (b), because the MBS absorbers are identified in both windows and none of them is the residual. Its own interval is two-sided, $[1.2, 3.9]$ basis points per 1 percent of the MBS float. The level component is the Treasury number, so the mortgage yield rises by 25 basis points at the pooled slope, 11 at the recession slope, 35 at the expansion slope, and

44 at the lower Anderson–Rubin edge. The honest summary is that the mortgage-specific piece of quantitative tightening is small and tightly estimated, while the mortgage cost of quantitative tightening is large and uncertain, because it is the Treasury market's number in disguise, and it rests on the assumption that the sector which absorbed Treasuries before 2010, the household residual and the hedge funds it now contains, will absorb them again.

Figure 5 projects the resulting yield path, combining the runoff pace with the estimated rebalancing dynamics. Each quarter's incremental runoff first lands on the impact demand curve: re-solving the \$1 trillion scenario horizon by horizon as the absorbers travel from $\beta^{(0)}$ to $\beta^{(4)}$, the impact quarter has no equilibrium at all, because the agency MBS aggregate slope is negative on impact and the Treasury slope is near zero, and the path peaks one to two quarters after the shock at about 1.6 times the long-run Treasury effect and 1.9 times the long-run mortgage effect, then decays to the long run within a year as slow capital arrives. Cumulating these vintages at the actual pace, Treasury yields rise by about one and a half basis points per month while the runoff proceeds, peak at 26 basis points when it ends after eighteen months, and are still 23 two years out, on their way to the long-run 21; the mortgage path peaks at 33 and is 27 two years out, against a long-run 25. The projection says most of the yield effect of a resumed QT arrives *during* the runoff, roughly in proportion to the stock withdrawn, and that the gradual pace limits the transient: at any time only the most recent increments are mid-adjustment, so the path peaks only a few basis points above its long-run component. The impact overshoot and its decay are the price-impact counterpart of the horizon duality van der Beck (2026) documents for equities, and the adjustment dynamics of Greenwood et al. (2018), in which a supply shock is absorbed first by local specialists and only gradually by cross-market capital.

Figure 5: **Projected Yield Path of a Resumed QT**

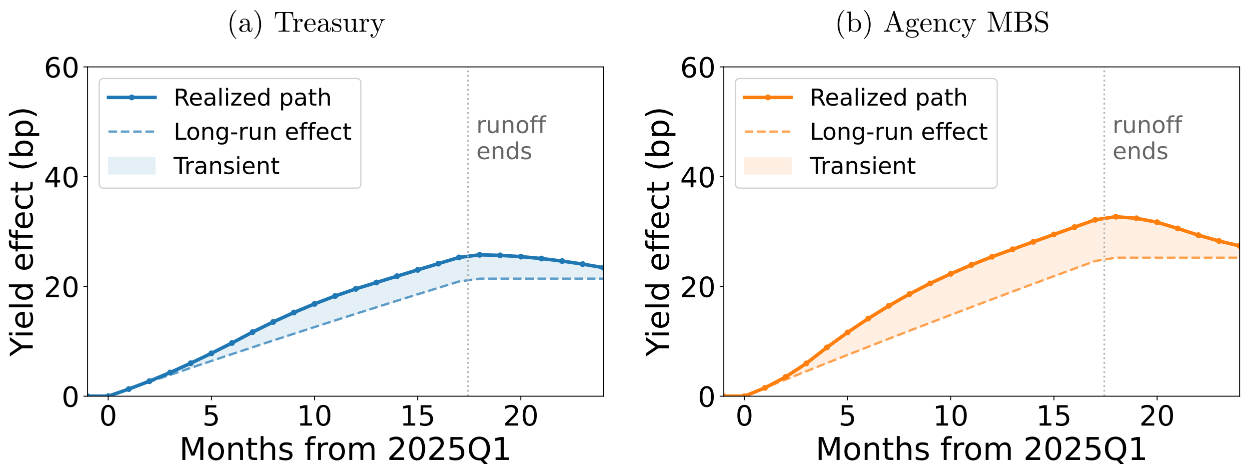

Note: Projected yield effects of \$1 trillion of further runoff from the 2025Q1 portfolio at the 2024 pace (\$172 billion per quarter, rendered monthly at \$57 billion per month), split between Treasuries and agency MBS at the QT2 ratio. Each month's incremental runoff enters at its impact effect and converges to its long-run effect along the adjustment path obtained by re-solving the scenario's equilibrium horizon by horizon as the absorbers travel from $\beta^{(0)}$ to $\beta^{(4)}$; the long-run component (dashed) accumulates the equilibrium effects of the runoff to date, with the supply response of Section 5.2 applied once, as in Table 8. Both panels share one vertical scale; the agency MBS panel is the yield, Treasury level plus spread. The shaded gap is the transient carried by the increments still mid-adjustment, which peaks when the runoff ends and dissipates as slow capital arrives. The vertical dotted line marks the end of the runoff.

6 Conclusion

The balance sheet will be run down again, and the framework here is built for the questions a resumed runoff will raise: who takes the other side, how far yields rise, and how the pace shapes the path. Its answers rest on measuring demand from what investors actually trade. Measured that way, most investors are inelastic, the few that absorb supply build their positions slowly, and the main Treasury absorber is the household residual, whose slope is identified mostly before 2010. Across the programs from QE1 to QT2, Treasury yields moved by 7 to 11 basis points per percent of the float, and the Treasury market set the level of mortgage yields while the MBS market set only the spread. Three caveats bound the answers. The Financial Accounts report no within-class maturity detail, so the system prices substitution across asset classes rather than along the curve; the supply side is deliberately reduced-form; and the demand slopes are identified from local variation, while the largest programs shift supply well outside it. Within those bounds, a further $1 trillion of runoff from the 2025Q1 balance sheet widens the mortgage spread by a small and precisely estimated 4 basis points and raises Treasury yields by about 21. The Treasury number roughly doubles if the household residual absorbs Treasuries less readily than it once did, and the gradual pace of the runoff keeps the transient above that long-run effect small.

References

Viral V. Acharya, Rahul S. Chauhan, Raghuram G. Rajan, and Sascha Steffen. Liquidity dependence and the waxing and waning of central bank balance sheets. NBER Working Paper 31050, National Bureau of Economic Research, 2024.

Tobias Adrian and Hyun Song Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437, 2010.

Tobias Adrian, Richard K Crump, and Emanuel Moench. Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1):110–138, 2013.

Heitor Almeida, Murillo Campello, Bruno Laranjeira, and Scott Weisbenner. Corporate debt maturity and the real effects of the 2007 credit crisis. *Critical Finance Review*, 1(1):3–58, 2012.

Sriya Anbil, Alyssa Anderson, Ethan Cohen, and Romina Ruprecht. Beyond reserves: The federal reserve's balance sheet and the repo market. Working paper, Federal Reserve Board. Previously circulated as "Stop Believing in Reserves", 2026.

Alyssa G. Anderson, Alessandro Barbarino, Anthony M. Diercks, and Stephen Miran. A user's guide to reducing the federal reserve's balance sheet. Finance and Economics Discussion Series 2026-019, Board of Governors of the Federal Reserve System, 2026.

Theodore W Anderson and Herman Rubin. Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of mathematical statistics*, 20 (1):46–63, 1949.

Isaiah Andrews, James H Stock, and Liyang Sun. Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, 11(1):727–753, 2019.

Steven Berry, James Levinsohn, and Ariel Pakes. A demand system for differentiated products. *Econometrica*, 63(4):841–890, 1995.

Carol C. Bertaut and Ruth Judson. Estimating U.S. cross-border securities positions: New data and new methods. International Finance Discussion Papers 1113, Board of Governors of the Federal Reserve System, 2014.

Lorenzo Bretscher, Lukas Schmid, Ishita Sen, and Varun Sharma. Institutional corporate bond pricing. *The Review of Financial Studies*, 2025.

Indraneel Chakraborty, Itay Goldstein, and Andrew MacKinlay. Monetary stimulus and bank lending. *Journal of Financial Economics*, 136(1):189–218, 2020.

Manav Chaudhary, Zhiyu Fu, and Haonan Zhou. Anatomy of the treasury market: Who moves yields? Working paper, 2026.

Jens H. E. Christensen and Glenn D. Rudebusch. The response of interest rates to US and UK quantitative easing. *The Economic Journal*, 122(564):F385–F414, 2012.

Stefania D'Amico, William English, David López-Salido, and Edward Nelson. The Federal Reserve's large-scale asset purchase programmes: Rationale and effects. *The Economic Journal*, 122(564):F415–F446, 2012.

Marco Di Maggio, Amir Kermani, and Christopher J Palmer. How quantitative easing works: Evidence on the refinancing channel. *The Review of Economic Studies*, 87(3):1498–1528, 2020.

Leonel Diego Drukker. Do lower mortgage rates benefit first-time homebuyers? *Available at SSRN 5741124*, 2025.

Darrell Duffie. Presidential address: Asset price dynamics with slow-moving capital. *Journal of Finance*, 65(4):1237–1267, 2010.

Stefania D'Amico and Thomas B King. Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply. *Journal of financial economics*, 108 (2):425–448, 2013.

Egemen Eren, Andreas Schrimpf, and Fan Dora Xia. The demand for government debt. *Management Science*, 2026.

Chuck Fang and Kairong Xiao. What do $40 trillion of portfolio holdings say about monetary policy transmission? *Available at SSRN 5025417*, 2025.

Xiang Fang, Bryan Hardy, and Karen K Lewis. Who holds sovereign debt and why it matters. *The Review of Financial Studies*, 2025.

William Fuchs, Satoshi Fukuda, and Daniel Neuhann. Demand-system asset pricing: Theoretical foundations. *Available at SSRN 4672473*, 2023.

William Fuchs, Satoshi Fukuda, and Daniel Neuhann. A trilemma for asset demand estimation. *Available at SSRN*, 2025.

Andreas Fuster and Paul Willen. $1.25 trillion is still real money: Some facts about the effects of the federal reserve's mortgage market investments. *FRB of Boston Public Policy Discussion Paper*, (10-4), 2010.

Joseph Gagnon, Matthew Raskin, Julie Remache, and Brian Sack. The financial market effects of the federal reserve's large-scale asset purchases. *International Journal of Central Banking*, 7(1):3–43, 2011.

Robin Greenwood, Samuel G Hanson, and Gordon Y Liao. Asset price dynamics in partially segmented markets. *The Review of Financial Studies*, 31(9):3307–3343, 2018.

Valentin Haddad and Tyler Muir. Market macrostructure: Institutions and asset prices. *Annual Review of Financial Economics*, 17:133–150, 2025. doi: 10.1146/annurev-financial-090524-120754.

Valentin Haddad, Zhiguo He, Paul Huebner, Péter Kondor, and Erik Loualiche. Causal inference for asset pricing. *Available at SSRN 5187305*, 2025.

Diana Hancock and Wayne Passmore. Did the federal reserve's mbs purchase program lower mortgage rates? *Journal of Monetary Economics*, 58(5):498–514, 2011.

Jarrad Harford, Sandy Klasa, and William F Maxwell. Refinancing risk and cash holdings. *The journal of finance*, 69(3):975–1012, 2014.

Jane Ihrig, Elizabeth Klee, Canlin Li, Min Wei, and Joe Kachovec. Expectations about the Federal Reserve's balance sheet and the term structure of interest rates. *International Journal of Central Banking*, 14(2):341–391, 2018.

Kristy AE Jansen, Wenhao Li, and Lukas Schmid. Granular treasury demand with arbitrageurs. Technical report, National Bureau of Economic Research, 2024.

Òscar Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 2005.

John Kandrac and Bernd Schlusche. Quantitative easing and bank risk taking: evidence from lending. *Journal of Money, Credit and Banking*, 53(4):635–676, 2021.

Narayana R Kocherlakota. Log-linear relative asset demand. Technical report, National Bureau of Economic Research, 2025.

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

Ralph SJ Koijen and Motohiro Yogo. On the theory and econometrics of (demand system) asset pricing. *Available at SSRN 5274709*, 2025.

Ralph SJ Koijen, François Koulischer, Benoît Nguyen, and Motohiro Yogo. Inspecting the mechanism of quantitative easing in the euro area. *Journal of Financial Economics*, 140 (1):1–20, 2021.

Arvind Krishnamurthy and Annette Vissing-Jorgensen. The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity*, 43(2):215–287, 2011.

Arvind Krishnamurthy, Annette Vissing-Jorgensen, et al. The ins and outs of lsaps. In *Kansas City federal reserve symposium on global dimensions of unconventional monetary policy*, pages 57–111. Federal Reserve Bank of Kansas City Kansas City, 2013.

Robert Kurtzman, Stephan Luck, and Tom Zimmermann. Did qe lead banks to relax their lending standards? evidence from the federal reserve's lsaps. *Journal of Banking & Finance*, 138:105403, 2022.

Arthur Lewbel. Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models. *Journal of Business & Economic Statistics*, 30(1):67–80, 2012.

David López-Salido and Annette Vissing-Jorgensen. Reserve demand, interest rate control, and quantitative tightening. Working paper, Federal Reserve Board, 2025.

Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.

Junko Oguri and Cristoforo Pizzimenti. Ample reserves for whom? the role of foreign banks in u.s. monetary policy implementation. Working paper, Northwestern University, 2026.

Alexander Rodnyansky and Olivier M Darmouni. The effects of quantitative easing on bank lending behavior. *The Review of Financial Studies*, 30(11):3858–3887, 2017.

James H Stock and Mark W Watson. Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128(610):917–948, 2018.

Johannes Stroebel and John B Taylor. Estimated impact of the federal reserve's mortgage-backed securities purchase program. *International Journal of Central Banking*, 2012.

Eric T Swanson. Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics*, 118:32–53, 2021.

Philippe van der Beck. Short versus long-run demand elasticities in asset pricing. *Journal of Financial Economics*, 184:104337, 2026.

Dimitri Vayanos and Jean-Luc Vila. A preferred-habitat model of the term structure of interest rates. *Econometrica*, 89(1):77–112, 2021.

Annette Vissing-Jorgensen. The Treasury market in spring 2020 and the response of the Federal Reserve. *Journal of Monetary Economics*, 124:19–47, 2021.

Kevin M. Warsh. Transcript of Chairman Warsh's press conference. Board of Governors of the Federal Reserve System, FOMC Press Conference, June 2026. URL https://www.federalreserve.gov/mediacenter/files/FOMCpresconf20260617.pdf. June 17, 2026; preliminary transcript.

Bin Wei. Quantifying "quantitative tightening" (QT): How many rate hikes is QT equivalent to? Policy Hub Paper 2022-8, Federal Reserve Bank of Atlanta, 2022.

A A BLP Interpretation of Sector-Level Asset Demand

In this appendix, I introduce a within-sector investor-level utility representation that provides a BLP-style interpretation of the demand estimation in the paper. Intuitively, each sector aggregates many underlying institutions (e.g., the U.S. banking sector includes large banks such as JPMorgan Chase and Bank of America), but the empirical analysis uses sector-level portfolios rather than institution-level holdings. For each sector i and within-sector investor h, the indirect utility from holding asset n at time t is

$$u_{hit}(n) = \beta_i \, y_t(n) + \gamma_i' x_t(n) + \xi_t(n) + v_{hit}(n), \tag{22}$$

while the outside asset $n = 0$ yields

$$u_{hit}(0) = \xi_t(0) + v_{hit}(0). \tag{23}$$

Here, β_i and γ_i denote sector-specific preference parameters that are common to all investors h in sector i. The term $\xi_t(n)$ captures asset–time fixed effects that are common across sectors and investors, such as unobserved asset characteristics, aggregate risk premium, or market-wide valuation components. Finally, $v_{hit}(n)$ are i.i.d. Type-I extreme value shocks that generate idiosyncratic within-sector heterogeneity. Although random coefficient specifications are standard in the BLP literature, I restrict within-sector heterogeneity to idiosyncratic shocks because sector-level data do not contain sector-specific price variation required to identify a distribution of random coefficients. Put differently, each sector is analogous to a market in the standard BLP framework. However, all sectors observe the same asset prices in financial markets, which by construction provides insufficient price variation to identify within-sector heterogeneity.

Given the Type-I extreme value errors, the portfolio weight of sector i on asset n at time

t is

$$w_{it}(n) = \frac{\exp\{u_{iht}(n)\}}{\sum_{k=0}^{N} \exp\{u_{iht}(k)\}}. \tag{24}$$

Normalize the outside option by setting $\xi_t(0) = 0$ and subtract $u_{iht}(0)$ from all utilities. Under this normalization, the portfolio weight can be written as in Equations (3) and (4):

$$w_{it}(n) = \frac{\delta_{it}(n)}{1 + \sum_{k=1}^{N} \delta_{it}(k)}.$$

Here,

$$\delta_{it}(n) := \exp\Big\{\beta_i y_t(n) + \gamma_i' x_t(n) + \xi_t(n)\Big\}\, \epsilon_{it}(n),$$

where $\epsilon_{it}(n) := \exp\{v_{iht}(n) - v_{iht}(0)\}$.

B Figures

Figure A.6: **Portfolio Shares by Sector**

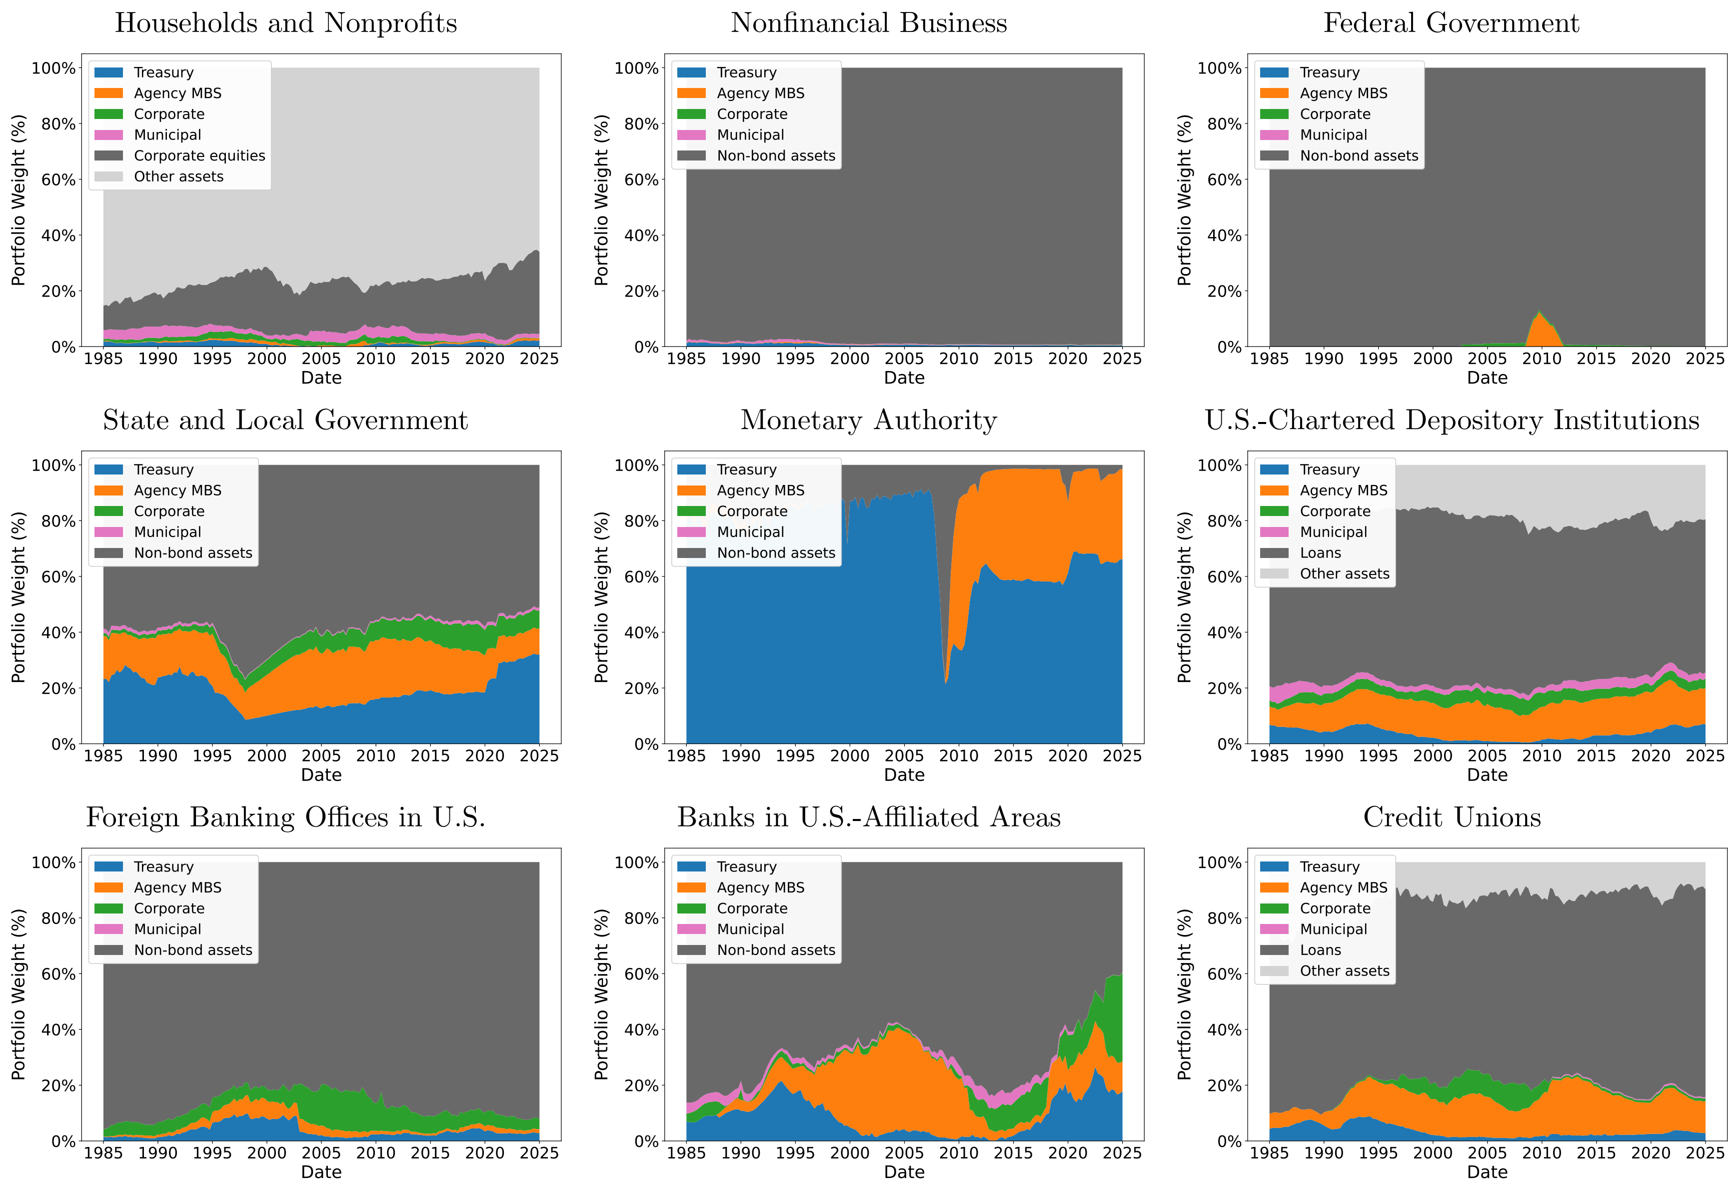

Notes: Stacked shares of total financial assets. Dark gray is the sector's own outside asset (Table 1); light gray is the remainder of the balance sheet, outside the demand system. For sectors whose outside asset is the total non-bond residual, the light-gray band is absent.

Figure A.6: **Portfolio Shares by Sector (continued)**

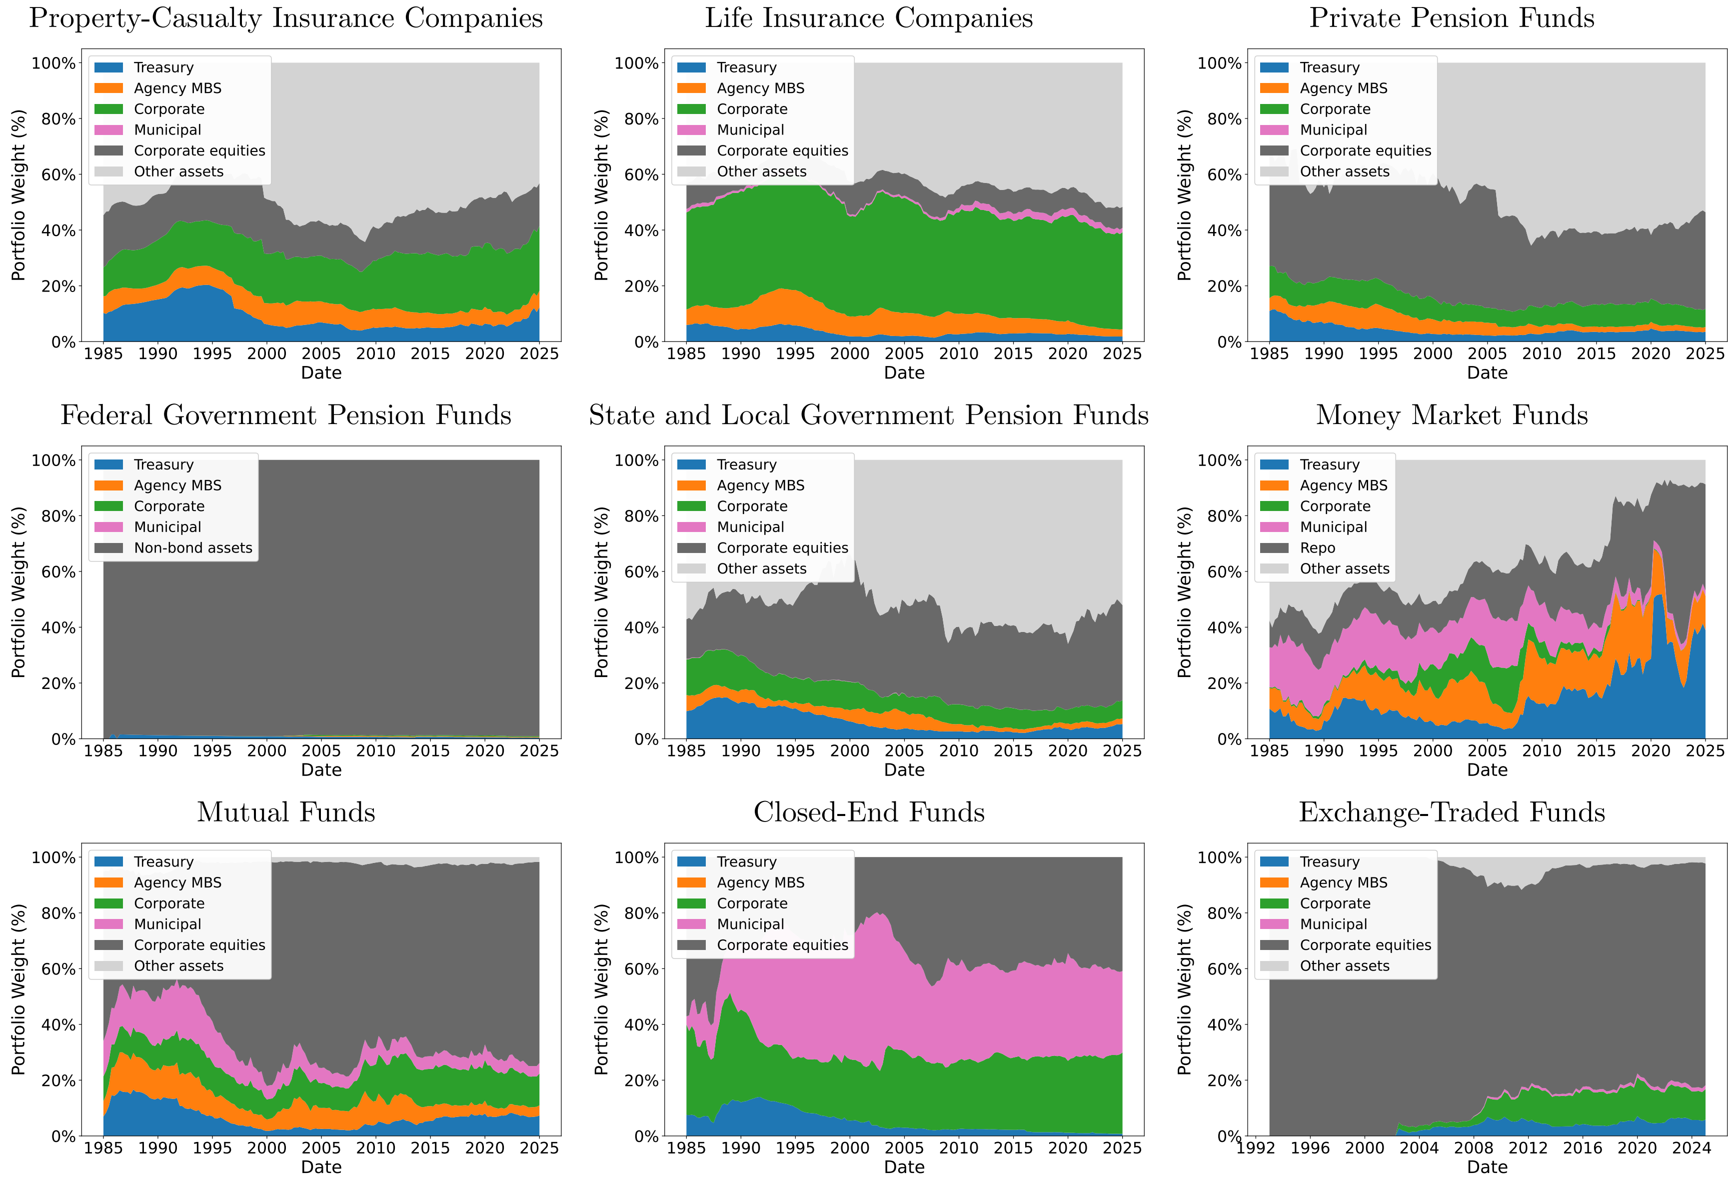

Notes: Stacked shares of total financial assets. Dark gray is the sector's own outside asset (Table 1); light gray is the remainder of the balance sheet, outside the demand system. For sectors whose outside asset is the total non-bond residual, the light-gray band is absent.

Figure A.6: **Portfolio Shares by Sector (continued)**

Notes: Stacked shares of total financial assets. Dark gray is the sector's own outside asset (Table 1); light gray is the remainder of the balance sheet, outside the demand system. For sectors whose outside asset is the total non-bond residual, the light-gray band is absent.

Figure A.7: **Own Demand Elasticity by Sector Type**

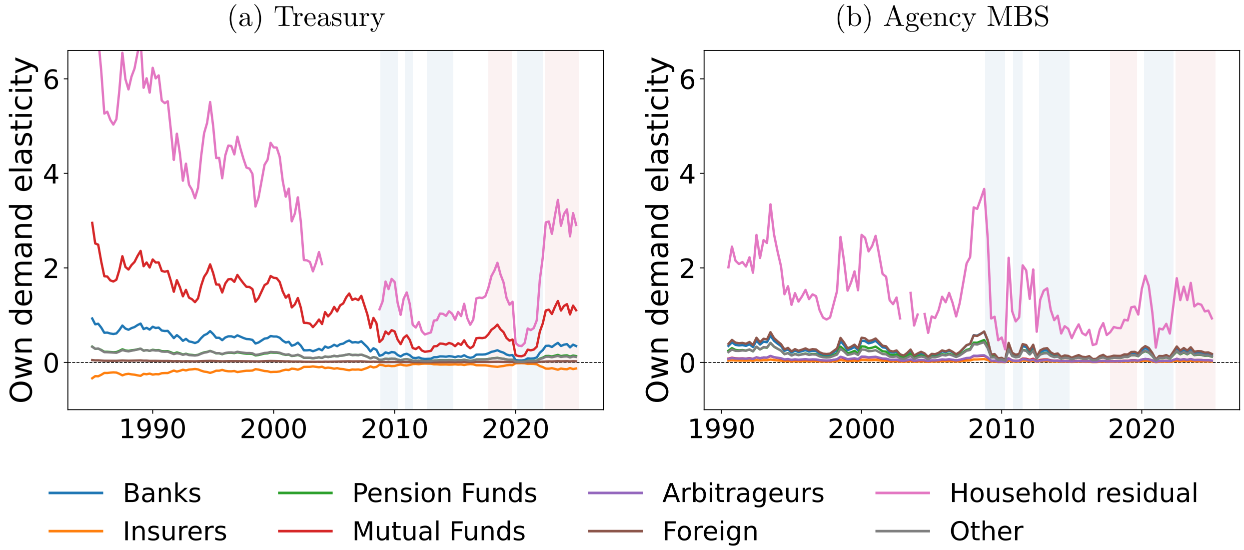

Note: Own elasticities $\eta_{it}(n,n) = \beta_i(n)(1-w_{it}(n))\, y_t(n)$, averaged within sector type with weights $A_{it}w_{it}(n)$. The figure shows logit cells only. The two cells of Section 3.7 that follow other models, money market funds and the dealer Treasury book, are left out of the type averages, so the mutual-fund Treasury line excludes money funds and the arbitrageurs have no Treasury line: the mortgage REITs in that type hold no Treasuries. Both panels are drawn on one fixed scale; the household residual's Treasury line leaves the frame before 1990.

C Tables

Table A.9: **Data Sources for Sector-Level Portfolio Holdings**

Type	Sector	Source
Banks	U.S.-Chartered Depositories	L.111
	Foreign Banking Offices	L.112
	Credit Unions	L.114
	Banks in Affiliated Areas	L.113
Insurers	Life Insurers	L.116
	Property-Casualty Insurers	L.115
Pension Funds	Private Pension Funds	L.118
	State & Local Govt. Retirement Funds	L.120
	Federal Govt. Retirement Funds	L.119
Mutual Funds	Mutual Funds	L.122
	Exchange-Traded Funds	L.124
	Money Market Funds	L.121
	Closed-End Funds	L.123
Arbitrageurs	Brokers and Dealers	L.130
	mREITs	L.129
Foreign	Foreign Official	CSLT
	Foreign Private	CSLT
Household Residual	Household Residual & Nonprofits	L.101
Other	Nonfinancial Business	L.102
	Government-Sponsored Enterprises	L.125
	Holding Companies	L.131
	State & Local Governments	L.107
	Federal Government	L.106
	Agency/GSE Mortgage Pools	L.126
	Finance Companies	L.128
	ABS Issuers	L.127
	Other Financial Business	L.132

Notes: Flow-of-Funds table codes (L. series) refer to the Financial Accounts of the United States (Federal Reserve Z.1); transactions come from the matching flow (FU) tables. The rest of the world is split into Foreign Official and Foreign Private holders over the full 1985–2025 sample using the Federal Reserve's CSLT estimates of cross-border securities positions (Bertaut and Judson, 2014), which report monthly holdings and net transactions of U.S. Treasuries, agency bonds, corporate bonds, and equities by holder type; the Financial Accounts do not disaggregate foreign holdings by holder type.

Table A.10: **Data Sources for Asset Yields and Characteristics**

Asset	Variable	Source
Treasury	Yield	10-year constant-maturity yield, `DGS10`, FRED
	Time to maturity	`YAS_MTY_YEARS`, Bloomberg (`LUATTRUU`)
	Modified duration	`YAS_RISK`, Bloomberg (`LUATTRUU`)
	Liquidity	Author's computation from CRSP
	Issuance size	`FL893061105`, Z.1 Financial Accounts (L.210)
Agency MBS	Yield	10-year Treasury yield plus Bloomberg US MBS index OAS (`LUMSTRUU`)
	Time to maturity	`WAL_TO_MATURITY`, Bloomberg (`LUMSTRUU`)
	Modified duration	`YAS_RISK`, Bloomberg (`LUMSTRUU`)
	Liquidity	Author's computation from TRACE
	Issuance size	`FL893061705`, Z.1 Financial Accounts (L.211)

Notes: All yields are end-of-quarter. Maturity and duration are computed from the Bloomberg total-return index of each asset class (index tickers in parentheses); all indexes are unhedged and denominated in USD. Transactions and holdings by sector are from the Z.1 flow (FU) and level (FL/LM) tables. The liquidity measures do not enter the main specification.

Table A.11: **QE4 Decomposition: Market-Functioning versus Stock-QE Phase**

Phase	Fed Treasury purchases ($B)	Treasury (bp)	Agency MBS (bp)	ΔY per 1% of float (bp)
QE4 full	3,516	-176	-198	7.3
stock QE (from 2020Q2)	1,462	-88	-102	8.8
stock QE (from 2020Q3)	1,212	-75	-88	9.1

Notes: QE4 (2019Q4–2021Q4) combined a market-functioning phase (2020Q1–Q2, $2.3 trillion of Treasury purchases in the COVID crash) with a later stock-QE phase. Rows re-base the counterfactual baseline to isolate the stock phase, with the Fed's transactions from the base quarter to 2021Q4 as the shock; the final column is the Treasury yield effect per 1 percent of float. The 2020H1 dislocation lies outside the stock-demand model, so the functioning phase is not reported as a stock effect. The phases are not additive, given the nonlinearity of the market-clearing map.

Table A.12: **Market Demand Elasticities**

	Sample average		2025Q1	
Price shocked	Treasury	Agency MBS	Treasury	Agency MBS
Treasury (10-year yield)	+0.468	-0.012	+0.555	-0.013
Agency MBS (option-adjusted spread)	-0.019	+0.235	-0.009	+0.201

Notes: The market elasticity matrix of Figure 3, averaged over the sample and evaluated at the 2025Q1 composition. Rows index the $y_t(n)$ that moves, columns the demand that responds; the diagonal is the own elasticity. A one-percent rise in the ten-year yield raises aggregate Treasury demand by 0.45 percent on average and by 0.56 percent at the 2025Q1 composition. The agency MBS entries are elasticities with respect to the option-adjusted spread, which averages 0.56 percentage points against a 4.4 point Treasury yield: the smaller number reflects the level of $y_t(n)$, not a weaker response, since the underlying semi-elasticity is 0.553 for agency MBS against 0.135 for Treasuries (Table 5). The cross elasticities are an order of magnitude smaller than the own elasticities because each scales with the share of the shocked asset in the portfolio, which is small for every sector.

D Specification Discussion

D.1 Outside-Asset Assignment, Sector by Sector

The demand system measures each sector's bond demand relative to a single outside asset, the non-bond class the sector actually trades off against bonds (Table 1). This subsection records the rationale for each assignment, organized by outside asset.

Loans: U.S.-chartered depositories and credit unions. A depository's marginal portfolio decision is between its securities book and its loan book: deposits fund either, and the securities portfolio serves as the buffer that absorbs deposit inflows when loan demand is weak. The QE literature documents exactly this margin: banks receiving reserve inflows under QE expand lending and adjust securities holdings in tandem (Rodnyansky and Darmouni, 2017; Kandrac and Schlusche, 2021; Kurtzman et al., 2022). Measuring bank bond demand against total non-bond assets would instead lump in reserves and interbank claims, which move with monetary policy mechanically rather than through portfolio choice.

Corporate equities: insurers, pension funds, mutual funds, and households. For the long-term investors the relevant margin is the stock–bond allocation. Insurers and pension funds manage fixed-income portfolios against equity allocations under asset–liability and return targets; mutual funds, together with ETFs and closed-end funds, reallocate between bond and equity exposures through both manager choice and investor flows; and household financial portfolios trade off directly held bonds against directly held equities. This is the classic rebalancing margin of the portfolio-balance channel: when QE compresses bond yields, these investors are pushed toward equities, and the outside asset is defined to capture precisely that substitution.

Repo: money market funds, broker-dealers, and hedge funds. Money funds allocate between short government paper and repo daily; repo is their actual alternative at the margin, not equities they cannot hold. Dealers likewise manage securities inventory against matched-book and reverse-repo positions, so repo is the financing margin along which their bond holdings expand and contract. Hedge funds enter the data in Treasuries only, where the Enhanced Financial Accounts data are informative and where the cash–futures basis trade concentrates their holdings. That trade is repo-funded: the fund buys the cash Treasury and finances it in the repo market, so repo, not equities or other bonds, is the margin along which its Treasury book expands and contracts, the same financing margin as dealers and money funds. The main 26-sector panel keeps hedge funds inside the household residual, because a full carve-out is infeasible: Form PF includes offshore-domiciled funds, whose holdings the Z.1 assigns to the rest of the world, so subtracting the full hedge-fund book from the residual double-counts the offshore share and produces negative household holdings. The household residual's Treasury slope in Table 4 is therefore the slope of a sector that, after 2010, contains the basis-trade book, and Section 4 reads it that way.

Non-bond U.S. securities: foreign investors. The rest of the world's U.S. portfolio decision is between U.S. bonds and other U.S. securities, chiefly equities; the TIC-based holdings data are organized along exactly this line. Reserve managers, the dominant foreign Treasury holders, allocate between Treasuries, agencies, and other dollar assets, so the within-U.S. margin is the relevant one, not substitution back into home-currency assets.

Total non-bond (residual): the remaining sectors. For sectors without a single dominant alternative, foreign banking offices and banks in U.S.-affiliated areas (balance sheets dominated by interoffice claims), federal government retirement funds (holdings concentrated in special-issue Treasuries), mortgage REITs (levered MBS vehicles whose non-MBS assets are small), and the heterogeneous "Other" group, I use the residual of total financial assets net of every bond class the Financial Accounts report separately. Imposing a specific margin on these balance sheets would assert substitution behavior the data cannot support; the residual is agnostic.

Other bonds as a common outside asset. Corporate and municipal bonds are outside the choice set, so they can serve as one outside asset for every sector. That removes the cost Section 3 names: the slopes become comparable across sectors. It changes what the slope measures, however. The slope becomes substitution within fixed income, between Treasuries or agency MBS and credit, rather than substitution out of bonds, and a corporate yield that moves with the Treasury yield blunts it. I re-estimate the system with each sector's corporate and municipal book as its outside asset, keeping the instruments, controls, and cell models unchanged.

The outside asset changes who absorbs more than how much the market absorbs. The household residual's Treasury slope is 0.63 (standard error 0.28) against 0.72, though its Anderson–Rubin set now reaches −0.04; mutual funds fall from 0.29 to 0.11, and foreign investors rise from 0.01 to 0.17. The Treasury aggregate demand slope rises from 0.135 to 0.164 ($t = 2.6$), so the impact per 1 percent of the float falls from 7.4 to 6.1 basis points; the agency MBS slope falls from 0.553 to 0.465, and its impact rises from 1.8 to 2.1. Historical Treasury effects shrink by a fifth to a third, to 6.0 to 7.4 basis points per percent of the float removed against 7.3 to 11.2 (QE1 −34 against −51, QE4 −144 against −176). A $1 trillion runoff raises Treasury yields by 19 basis points rather than 21 and mortgage yields by 23 rather than 25. The mortgage spread moves more: QE1 compresses it by 63 basis points against 45. The baseline keeps the non-bond margins, because the portfolio-balance channel of QE is the movement of duration out of the bond market, and a within-bond outside asset cannot see it.

D.2 Transactions versus Market Values

The Financial Accounts decompose the change in every level into transactions, revaluations, and other volume changes. The demand system uses the transactions (Section 2). This subsection documents why the alternative, the market-value share with an index-duration adjustment, does not measure active demand in these data, and what it does to the estimates.

The adjustment adds the revaluation back. For a bond position of duration $D_t(n)$, a yield change $\Delta y_t(n)$ (in percentage points) changes the position's market value by approximately $-D_t(n)\,\Delta y_t(n)/100$ even if the sector trades nothing, so the adjusted change

$$\Delta\tilde{q}_{it}(n) \equiv \Delta \ln \frac{w_{it}(n)}{w_{it}(0)} + D_t(n)\,\frac{\Delta y_t(n)}{100} \tag{25}$$

is meant to isolate active rebalancing. Eren et al. (2026) make this correction with a constant eight-year duration; the observed index duration is the natural alternative. The correction is right for a holder whose recorded level moves one for one with market prices at the index duration, and wrong otherwise. If a sector's recorded level does not revalue, the differenced log share is already the active change, and adding $D\,\Delta y/100$ manufactures a positive response of $D/100$ per percentage point, a slope of 0.05 to 0.10 at the durations of the four indices, for a sector that traded nothing.

Table A.13 measures how much of the market revaluation each sector's recorded level actually contains: the coefficient of the Financial Accounts' own revaluation series on the index-implied revaluation $-D_t(n)\Delta y_t(n)\,H_{i,t-1}(n)/100$, cell by cell. The pass-through is a continuum. It is near one for the market-value holders, broker-dealers and open-end mutual funds; near zero for the Federal Government's corporate holdings and for money market funds; and mostly between 0.1 and 0.5 for the depositories, insurers, and pension funds whose books mix amortized cost, statutory value, and market value. No sector-specific on–off correction reproduces this, and the index duration is the wrong duration for most of them.

Table A.14 shows what the measurement does to the demand slopes, separately from the instrument. It reports the common-slope sector-type long-run slopes $\beta^{(4)}$ on the full sample under the two measurements of the dependent variable, market-value shares with the adjustment (25) and transactions, and under two Treasury instruments, the market-value granular instrument and the heteroskedasticity instruments. The measurement, not the instrument, drives the results. With market-value shares, banks, insurers, and foreign investors carry small positive slopes (0.07, 0.12, 0.05 under the granular instrument), which would read as slow, weakly positive demand; with transactions, the same sectors are at zero or below under either instrument (−0.10, −0.04, −0.05 under the heteroskedasticity instruments), and the arbitrageurs' large negative market-value slope shrinks to a fraction. Only the household residual is positive under transactions, and its slope rises when the

Table A.13: **Pass-Through of Market Revaluation into Recorded Levels**

Sector type	Sector	Treasury	Agency MBS
Arbitrageurs	Security Brokers and Dealers	1.50	0.81
Banks	Credit Unions	0.28	0.20
	Foreign Bank Offices	—	—
	Foreign Banking Offices in U.S.	—	—
	US Depositories	—	0.26
Foreign Investors	Rest of the World	0.36	0.27
Households	Households/Nonprofits	0.71	0.44
Insurers	Life Insurers	0.42	0.23
	Property-Casualty Insurers	0.32	0.15
Mutual Funds	Money Market Funds	-0.03	—
	Mutual Funds	0.83	—
Other	Federal Government	—	—
	Government-Sponsored Enterprises	-0.01	0.25
	Nonfinancial Business	0.33	—
	Other Financial Business	—	—
Pension Funds	Fed Pension Funds	0.27	0.54
	Private Pension Funds	0.40	0.33

Notes: For each sector and asset with a revaluation series in the Financial Accounts, the slope of the recorded quarterly revaluation on the index-implied revaluation of the lagged position, $-D_t(n)\,\Delta y_t(n)\,H_{i,t-1}(n)/100$, 1985Q1–2025Q1; a value of one means the recorded level moves with market prices at the index duration, a value of zero that it does not revalue. Cells are those for which the Financial Accounts publish the decomposition.

instrument is changed. The inert cluster was the adjustment.

The transaction slope and the market-value slope measure different objects, and they are related by an identity. Equation (2) measures the change in units held. The market-value share ratio changes by the same amount plus the relative price change, which is common to every holder of the asset and equals $-D_t(n)\Delta y_t(n)/100$ up to carry; regressing the transaction growth less the true revaluation on the yield change therefore gives a slope that differs from the transaction slope by $-D/100$, a uniform shift of about -0.05 across sectors. Under that normalization a holder that trades nothing has a slope of $-D/100$, not zero, and the household residual's Treasury slope of about one is the one contrarian position in the market. Market clearing is written in units, and the Fed's transactions shift units, so the counterfactual uses the transaction slope; the identity is why the choice of normalization does not affect it.

One cell is sensitive to how transaction growth is measured, and to how the level is treated.

Table A.14: **Measurement versus Instrument: Sector-Type Long-Run Slopes**

	Market-value shares, duration-adjusted		Transactions (FU)	
Sector type	Granular IV	Lewbel IV	Granular IV	Lewbel IV
Banks	+0.07	+0.08	-0.30	-0.10
	(0.18)	(0.14)	(0.09)	(0.07)
Insurers	+0.12	+0.03	-0.03	-0.04
	(0.05)	(0.04)	(0.03)	(0.03)
Pension Funds	-0.08	+0.05	-0.23	-0.09
	(0.07)	(0.08)	(0.07)	(0.04)
Mutual Funds	-0.02	-0.39	-0.19	-0.10
	(0.09)	(0.35)	(0.06)	(0.06)
Arbitrageurs	-0.58	-0.13	-0.05	-0.19
	(0.33)	(0.27)	(0.11)	(0.16)
Foreign	+0.05	+0.02	-0.05	-0.05
	(0.05)	(0.04)	(0.03)	(0.02)
Household Residual	+0.07	+0.29	+0.02	+0.46
	(0.53)	(0.37)	(0.17)	(0.29)
Other	+0.00	-0.03	-0.15	-0.04
	(0.09)	(0.06)	(0.07)	(0.08)

Notes: Sector-type long-run slopes $\beta^{(4)}$ from the pooled common-slope local projection of Equation (9) on 1985Q1–2025Q1, all sectors on the full sample, Newey–West standard errors in parentheses. Columns vary the dependent variable, the duration-adjusted market-value share change (25) or the transaction growth (2), and the Treasury instrument, the market-value granular instrument or the three-donor Lewbel heteroskedasticity instruments (all quarters); the agency MBS instruments are the SOMA transactions with the Drukker surprise throughout. The paper's specification adds the crisis-robust construction and the asset-specific slopes of Table 4.

The Financial Accounts' household Treasury position is a residual, total outstanding less every identified holder, and it is negative from 2004Q2 to 2008Q2, reaching $-$318 billion, about a tenth of the float, in 2007Q4. Wherever the paper weights sectors by their holdings, the aggregate demand slope, the market elasticities, and the counterfactual, I floor that residual at zero and let the float renormalize. A negative statistical residual would otherwise put a negative weight on the sector's positive slope and turn the aggregate slope negative in six quarters, which is a measurement artifact and not demand; the reported counterfactuals are unaffected, because every quarter they solve at has a positive residual. Dealers' negative Treasury positions are left as reported: a dealer can genuinely be net short, a residual cannot. Separately, the growth rate $\ln(1+F/H_{-1})$ computed on that near-zero base is extreme in the surrounding quarters (its standard deviation is 0.70 over 2000–2008 against 0.09 before and 0.27 after). Without the 1/99 winsorization of Section 2 the household residual's Treasury slope is about 1.5 (standard error 0.5); with it, 0.86 (0.39); with a 5/95 rule, about 0.8; and restricting the cell to quarters in which the residual exceeds $100 billion, about 0.5 (0.8). No other sector type's slope moves by more than 0.1 under any of these rules. The paper uses the 1/99 rule, fixed in advance for every cell, and the Treasury counterfactuals inherit this cell's range.

D.3 Estimation Window

Table A.15 re-estimates the long-run slopes of Table 4 on the 2010Q1–2025Q1 window alone. The agency MBS absorbers are stable: insurers at 0.34 (standard error 0.14) against 0.31 on the full sample, foreign investors at 0.19 (0.16) against 0.32, mutual funds at 1.47 (0.83) against 1.35, and the arbitrageurs at 1.28 (0.48) against 0.94. The Treasury absorber is not: the household residual's slope is -0.93 with a standard error of 0.86, and no other sector's Treasury slope excludes zero in the short window either. The Treasury counterfactuals of Section 5 therefore rest on the pre-2010 behavior of the residual sector, the dependence that Panel (b) of Table 8 quantifies.

Table A.15: **Long-Run Slopes on the 2010–2025 Window**

Sector type	Treasury (yield)	Agency MBS (spread)
Banks	+0.58	+0.32
	(0.33)	(0.26)
Insurers	+0.02	-0.02
	(0.10)	(0.09)
Pension Funds	+0.09	+0.11
	(0.12)	(0.10)
Mutual Funds	+0.16	-0.30
	(0.11)	(0.27)
Arbitrageurs	+0.06	+0.45
	(0.40)	(0.28)
Foreign	+0.00	+0.01
	(0.06)	(0.10)
Household Residual	-0.78	+1.34
	(0.82)	(1.06)
Other	+0.01	-0.14
	(0.19)	(0.51)

Notes: Long-run slopes $\beta^{(4)}$ from the local projections of Equation (9), estimated separately for each sector type and asset on 2010Q1–2025Q1 with the same instruments and controls as Table 4; Newey–West standard errors in parentheses. Dashes mark cells with too few observations.

D.4 Weak-Instrument Checks

Table A.19 reports three checks on the identification. The upper panel re-estimates the Treasury cells with all three heteroskedasticity donors, maturity, issuance, and coupon, in place of the single donor of the baseline. The three-donor block reaches a first stage of 6.8 against the baseline's 7.8, because the added interactions carry no relevance and only dilute the statistic. The long-run slopes are essentially unchanged: the household residual's Treasury slope is 0.55 ($[0.26, 1.64]$) with three donors against 0.72 ($[0.16, 2.07]$) with the coupon donor alone. The bottom panel adds the agency MBS heteroskedasticity instruments to the Fed's transactions and the Drukker surprise. The first stage rises, but the Sargan test rejects the joint set in the absorber cells. The added relevance is the maturity interaction, and the index maturity of agency MBS moves with prepayment and hence with the yield, so the instrument is not excludable. LIML and Fuller estimates, in the accompanying output, coincide with two-stage least squares in every cell.

The Menu of Candidate Instruments for Agency MBS. Table A.18 reports every candidate I examined for the agency MBS block, each run three ways: on the option-adjusted spread, which is the price the block estimates; on the Treasury yield, which is the cross-asset placebo an asset-specific supply shifter must fail to move; and on the MBS yield level, the price a yield specification would use. Read the table by the first two column pairs together. The Fed's realized transactions move the spread with an F of 3.4 but move the Treasury yield more (4.8), which is the wrong ordering for an MBS-specific instrument. The purchase-path revision over all quarters looks stronger, 9.4 on the spread, but moves the Treasury yield nearly twice as much (17.5, $t = -4.18$): in 2008Q4 and 2020Q1 the announcement and the flight to quality are the same event. Zeroed in the eight crisis quarters the ordering is right and the gap is wide: 15.4 on the spread against 1.1 on the Treasury yield. Dropping the two judgment-coded events leaves this unchanged (14.3 against 0.5), and restricting the sample to 1990–2019, so that neither the pandemic nor the 2022 tightening contributes, leaves the spread first stage at 7.4.

The other candidates fail for reasons worth recording. Gross agency issuance at par, built from the issuance unpaid-principal balances that Fannie Mae, Freddie Mac, and Ginnie Mae publish monthly, is the strongest series in the table, with an F of 24 on the MBS yield, but its sign is negative: issuance rises when rates fall because borrowers refinance, so the series traces the supply curve rather than shifting it, and it belongs in the supply estimation of Section 5.2 rather than in the demand block. Net par issuance nets that refinancing wave against the prepayments it causes and correlates with nothing. Foreign official purchases of agency securities and the retained portfolios of the government-sponsored enterprises both have first stages near 8 with a positive sign, the signature of buying after yields rise; the GSE series is predictable from the previous quarter's spread ($t = 5.7$), and residualizing it against that spread removes its relevance. The composition instrument, the MBS revision less the Treasury revision, is the theoretically cleanest object but has no power for either price. The LSAP factor of Swanson (2021) is a high-frequency surprise and survives quarterly aggregation with an F of 0.2. Finally, the Treasury coupon instrument has no power inside the MBS block, which is why the two prices of Section 3.6 cannot both be instrumented in the paper's baseline control set: the MBS index maturity, which the demand equation controls for, is itself a function of the rate level through prepayment, and it absorbs the level variation the Treasury instrument would supply. Entering that control with a one-quarter lag restores the Treasury instrument's relevance ($F = 17.7$) and is how the level-versus-spread restriction of Section 3.6 is tested.

Dating the Revisions: the Committee Signals Before It Votes. Dated to the FOMC statement that formalises it, the Treasury leg of the purchase-path revisions has no first stage at all: 0.2 crisis-zeroed, with a positive sign, meaning yields rise when the Fed announces that it will buy more. The event quarters look perverse. The Fed announced \$600 billion of Treasury purchases in 2010Q4 and the ten-year yield rose 77 basis points over that quarter;

it announced \$300 billion in 2009Q1 and the yield rose 46; it added \$45 billion a month in 2012Q4 and the yield rose 13.

The reason is anticipation, and it is visible directly. Regressing the quarterly change in the ten-year term premium on the revision at t, at $t+1$, and at $t-1$ together, the contemporaneous coefficient is zero ($+0.02$, $t = 0.01$) while the coefficient on the *next* quarter's revision is large and negative (-4.03, $t = -3.31$): the market prices a Treasury program in the quarter before the Committee votes on it. The agency MBS leg is different, with a contemporaneous coefficient of -4.43 ($t = -3.08$) alongside a smaller lead term, which is consistent with the rest of the evidence here: what gets telegraphed is the size of a program, and the composition is the part that surprises.

I therefore date each revision to the quarter in which the change first became public, taken from the documented prior communication rather than from a mechanical shift. Five of the twenty-one events move: QE2, signalled at Jackson Hole in August 2010 and voted in November; the 2017 runoff caps, set out in the March 2017 minutes; the end of the Treasury runoff, signalled at the January 2019 meeting; the 2021 taper, signalled at Jackson Hole and the September 2021 meeting; and the 2024 cap reduction, signalled in the March 2024 minutes. Under this dating the Treasury revision reaches 4.8 on the ten-year yield and 5.7 on the ten-year term premium, against 0.2 and 0.1 under statement dating.

Two checks support using it. First, the revisions are not a reaction to the previous quarter: regressed on the prior quarter's yield change, term-premium change, output growth, inflation, and volatility, the Treasury revision has an R^2 of 0.11 and a joint F of 2.6, and instrumenting with only the unpredictable part leaves the first stage at 7.8, against 7.5 on the same subsample. Second, it passes the cross-asset placebo, reaching 3.4 on the mortgage spread with the opposite sign. Paired with the coupon instrument the Treasury block has a joint F of 7.8, no sector type rejects the overidentifying restriction (minimum Sargan p of 0.23), and every Anderson–Rubin set narrows: banks move from $[-0.13, 0.30]$ to $[-0.06, 0.37]$ and the arbitrageur cell, which runs off the grid under the coupon instrument alone, becomes

$[-7.33, -0.08]$.

The Treasury leg identifies the risk premium somewhat better than the level, reaching 5.7 on the ten-year term premium against 4.8 on the yield, and it carries almost nothing for the expectations component, 1.3. Treasury purchase announcements are closer to exogenous for the risk premium than for the level, because the Fed announces them when the growth outlook deteriorates and the announcement resolves the news that prompted it, so the policy and the macro shock arrive together and move the expectations component in opposite directions. The paper prices the block by the yield all the same, for the reasons below.

Which Instrument Identifies the Treasury Block? The Treasury block uses two instruments, and a natural question is whether the internal one is needed at all. Table A.16 reports the block under each instrument alone, under the two external instruments together, and under the paper's pair. A third external candidate appears there as well: Japanese foreign-exchange intervention. When the Ministry of Finance sells yen and buys dollars it invests the proceeds in Treasuries, so an intervention shifts foreign demand for a reason that originates in Japanese exchange-rate policy. It is the strongest single instrument in the table, with a first stage of 16.6 and the sign a demand shift implies, it is not predictable from the dollar, from volatility or from lagged yields (R^2 of 0.04), and it moves the mortgage spread only weakly and with the opposite sign.

The instruments identify off different parts of the sample and different sectors. The coupon interaction has a first stage of 13.5 on the full sample and 0.2 over 2009–2025: the calm-period heteroskedasticity it exploits is a pre-crisis phenomenon. The two external instruments are the mirror image, 4.8 and 16.6 on the full sample and 9.1 and 5.8 after 2009. They also disagree about who absorbs. The coupon instrument delivers the household residual at 0.80 and mutual funds at 0.29, the long-duration buy-and-hold absorbers; the external instruments deliver banks, at 0.53 with a set of $[0.08, 1.14]$, and leave the household residual at 0.14 with an unbounded set. That is a difference between periods rather than a conflict: the overidentification test rejects nowhere in the paper's pair, with a minimum p of 0.23.

The aggregate demand slope settles the question. At the 2025Q1 composition it is 0.135 (standard error 0.055) under the paper's pair and 0.143 (0.062) under the coupon instrument alone, both distinguishable from zero. Under the two external instruments together it is 0.068 with a standard error of 0.141, and under either of them alone it is not identified at all: the purchase-path revision alone gives -0.018 and the Japanese intervention alone gives $+0.132$ with a standard error of 6.7, which comes from the leveraged cells it cannot pin down. The reason is arithmetic: the household residual and the mutual funds carry a quarter of the weight in the aggregate demand slope between them, and only the internal instrument identifies those two cells. Dropping it would raise the first stage, from 7.8 to 9.5, and lose the object the counterfactual needs. Adding the Japanese intervention to the pair leaves the slope unchanged at 0.113 but lowers the overidentification p-value in the bank and mutual-fund cells to 0.07, so the paper keeps the two-instrument pair. What the comparison establishes is narrower and more useful than a ranking: the *level* of the Treasury aggregate demand slope is similar however it is identified, and the *identity* of the marginal buyer is not.

Decomposing the Treasury Yield. The natural repair is to strip the expectations component out of the Treasury yield the way the spread specification strips it out of the mortgage yield, and to price the Treasury block by the term premium. I ran that check with the term-premium decomposition of Adrian et al. (2013), which splits the fitted ten-year yield into a risk-neutral expectations component and a term premium. It does not work, and it uncovers a caveat about the instrument that does.

The two instruments turn out to target different components of the yield, which is what makes the comparison informative. The coupon heteroskedasticity instrument draws its power from the expectations component, where its first stage is 5.7 against 1.7 on the term

Table A.16: **Treasury Demand by Instrument**

Sector type	Coupon Lewbel		Fed path revision		Japan FX intervention		External only		Baseline	
	$\beta^{(4)}$	AR set	$\beta^{(4)}$	AR set	$\beta^{(4)}$	AR set	$\beta^{(4)}$	AR set	$\beta^{(4)}$	AR set
Banks	+0.07	[-0.13, 0.30]	+0.66	[-0.19, 1.79]	+0.51	[0.08, 0.98]	+0.53	[0.08, 1.14]	+0.09	[-0.06, 0.37]
Insurers	-0.04	[-0.14, 0.09]	+0.07	[-0.49, 0.53]	-0.02	[-0.29, 0.30]	-0.01	[-0.28, 0.33]	-0.03	[-0.16, 0.14]
Pension Funds	+0.04	[-0.01, 0.11]	-0.11	[-0.67, 0.36]	+0.09	[-0.06, 0.26]	+0.06	[-0.11, 0.24]	+0.03	[-0.04, 0.11]
Mutual Funds	+0.29	[0.05, 0.58]	+0.26	[-0.56, 1.00]	-0.11	[-0.79, 0.51]	-0.06	[-0.48, 0.60]	+0.29	[0.04, 0.64]
Arbitrageurs	-1.99	[-14.30, 10.32]^{u}	-1.45	[-5.21, 0.76]	-24.35	[-2344.27, 2295.58]^{u}	-1.51	[-9.84, 1.87]	-1.89	[-7.33, -0.08]
Foreign	-0.00	[-0.13, 0.12]	+0.37	[-0.02, 1.59]	-0.03	[-0.21, 0.17]	+0.02	[-0.08, 0.17]	+0.01	[-0.07, 0.11]
Household residual	+0.80	[0.30, 1.74]	-0.99	[-7.71, 3.93]	+0.64	[-6.47, 7.75]^{u}	+0.14	[-6.42, 7.43]^{u}	+0.72	[0.16, 2.07]
Other	+0.05	[-0.20, 0.34]	-0.42	[-1.64, 0.58]	+0.27	[-0.38, 1.00]	+0.19	[-0.57, 0.69]	+0.04	[-0.26, 0.33]
First-stage F (ten-year yield)	13.5		4.8		16.6		9.5		7.8	
Aggregate demand slope, 2025Q1	+0.143 (0.062), t 2.3		-0.018 (0.337), t -0.1		+0.132 (6.739), t 0.0		+0.068 (0.141), t 0.5		+0.135 (0.055), t 2.5	

Notes: Long-run Treasury slopes $\beta^{(4)}$ from the local projections of Equation (9) on 1985Q1–2025Q1, by sector type, under five instrument sets: the crisis-robust coupon heteroskedasticity instrument alone; the first-signal Fed Treasury purchase-path revision alone; Japanese foreign-exchange intervention alone; the two external instruments together; and the coupon instrument with the purchase-path revision, which is the paper's specification. The 95 percent Anderson–Rubin set appears beside each estimate; ^{u} marks a set reaching the edge of the grid. The last two rows report the first stage for the ten-year yield on the full sample and the implied 2025Q1 aggregate demand slope with a delta-method standard error, both computed under the paper's treatment of the two misspecified cells (Section 3.7), so the last column reproduces the aggregate demand slope of Table 5. The standard error here treats the sector-type slopes as independent and is therefore slightly tighter than the one reported there.

premium, the reverse of what a portfolio-balance system would want from it. The Fed's Treasury purchase-path revision is the mirror image, 5.7 on the term premium against 1.3 on expectations. The two components pull in opposite directions, so the pair identifies the yield better than either component: 7.8 on the yield against 3.7 on the term premium and 3.1 on expectations. Pricing the block by the term premium would halve the instrument's strength.

Table A.17 reports both. The substance does not turn on the choice. The household residual and the mutual funds are the absorbers either way, at 0.72 and 0.29 on the yield and 1.27 and 0.82 on the term premium, and every other sector type is indistinguishable from zero in both columns; the term-premium coefficients are larger because the term premium is the smaller and less variable of the two prices. The overidentifying restriction survives in every cell of the yield panel, with a minimum p of 0.23, and in every cell of the term-premium panel but the arbitrageurs, at 0.05. I keep the yield in the main text for two reasons that have nothing to do with fit: the counterfactual has to deliver basis points on Treasury borrowing costs and mortgage rates, which requires the yield level, and the comparison with the QE literature in Figure 4 is in yield space. The term premium also carries the identifying assumptions of a fitted term-structure model, which the yield does not.

Table A.17: **Treasury Demand Priced by the Yield and by the Term Premium**

	Ten-year yield (baseline)			Ten-year term premium		
Sector type	$\beta^{(4)}$	AR set	Sargan p	$\beta^{(4)}$	AR set	Sargan p
first-stage F		7.8			3.7	
Banks	+0.09 (0.11)	[-0.06, 0.37]	0.30	+0.45 (0.31)	[-0.23, 1.38]	0.65
Insurers	-0.03 (0.06)	[-0.16, 0.14]	0.77	-0.07 (0.21)	[-0.54, 0.52]	0.61
Pension Funds	+0.03 (0.04)	[-0.04, 0.11]	0.51	+0.06 (0.14)	[-0.42, 0.51]	0.31
Mutual Funds	+0.29 (0.13)	[0.04, 0.64]	0.94	+0.82 (0.39)	[0.13, 1.59]	0.24
Arbitrageurs	-1.89 (1.29)	[-7.33, -0.08]	0.80	-2.03 (1.18)	[-2.73, 0.09]	0.05
Foreign	+0.01 (0.07)	[-0.07, 0.11]	0.23	+0.13 (0.22)	[-0.13, 1.21]	0.34
Household residual	+0.72 (0.28)	[0.16, 2.07]	0.26	+1.27 (0.74)	[0.83, 6.87]	0.09
Other	+0.04 (0.13)	[-0.26, 0.33]	0.41	-0.01 (0.40)	[-1.13, 0.79]	0.37

Notes: Long-run Treasury slopes $\beta^{(4)}$ from the local projections of Equation (9) on 1985Q1–2025Q1, by sector type, with the price variable in the left panel the ten-year yield, as in the main text, and in the right panel the ten-year term premium of Adrian et al. (2013). Both panels use the paper's Treasury instruments, the crisis-robust coupon heteroskedasticity instrument and the first-signal Fed Treasury purchase-path revision, with Newey–West standard errors in parentheses, the 95 percent Anderson–Rubin set beside each estimate, and the Sargan overidentification p-value in the last column of each panel. The header row reports the joint first stage for each price. ^{u} marks a set reaching the edge of the grid.

The curve slope, the ten-year yield less the current short rate, is the model-free alternative and fails for its own reasons. It is not the term premium, since it still carries the expected path of policy, and quarterly changes correlate 0.79 with the term premium against the three-month bill and 0.62 against the fed funds rate. On the slope the coupon instrument weakens, to 6.8 and 7.8, and the second stage simply rescales the yield estimates by a factor of about 1.2 to 1.3 with no change of sign anywhere. The agency MBS purchase-path revision does reach 13.5 on the fed-funds slope, but its event quarters explain that away: the fed funds rate is flat in every large event quarter and moves only in 2017Q2 and 2017Q3, when the Fed announced balance-sheet runoff while it was raising rates, so the added relevance is the correlation between balance-sheet announcements and the policy-rate path rather than identification. The Treasury block's remaining difficulty is therefore not the choice of price variable, which Table A.17 shows does not change the answer. It is that one sector carries the market's entire aggregate demand slope, and that sector's slope is estimated on the pre-2010 sample.

Third, the agency MBS purchase-path revision *does* compress the Treasury term premium ($F = 9.4$, $t = -3.07$) while leaving the expectations component untouched ($F = 0.0$, $t = +0.03$). Read one way this is reassuring: the instrument is a pure risk-premium shock, with no signaling content, which is exactly what a supply shifter should be. Read another way it qualifies the placebo of Table A.18. The instrument does not move the ten-year par yield that the Treasury block prices, which is the placebo reported there, but it does move the risk premium inside it, so it is a bond-market duration shock rather than a purely mortgage-market one. Two things bound the consequence. The response of the mortgage spread is larger than the response of the Treasury term premium, -5.4 against -4.4 in the same units, so the instrument does discriminate between the two markets; and the Treasury demand slopes are indistinguishable from zero for every sector type but one, so a Treasury term-premium move of this size shifts very little Treasury demand and exerts correspondingly little substitution pressure on the agency MBS block. Using the same instrument for both blocks is not an option: it delivers a household Treasury slope of -3.11 against the coupon instrument's $+0.86$, and the overidentification test rejects the pair in that cell ($p = 0.00$).

That confound is precisely what the spread specification removes on the mortgage side. News about growth and inflation moves the level of yields and leaves the mortgage spread alone, so the composition of an announcement, how much of it is agency MBS, is exogenous to the spread even when its size is endogenous to the level. The asymmetry in Table A.18, where the same instrument has a first stage of 15.4 on the spread and 1.1 on the Treasury yield, is the same fact read from the other side. On the Treasury side the announcement leg adds precision but does not dissolve the limitation: the ten-year yield's identification still rests mostly on calm-period heteroskedasticity, and on the pre-2008 sample where that heteroskedasticity is informative, which is why Section 5 reports the Treasury number as a band rather than a point.

Table A.18: **Candidate Instruments for Agency MBS and the Cross-Asset Placebo**

Candidate instrument	MBS spread		Treasury yield (placebo)		MBS yield	
	F	t	F	t	F	t
SOMA transactions + Drukker (2025) surprise	3.4	-2.09	4.0	+1.51	3.4	-2.28
SOMA transactions alone	4.4	-2.09	2.3	+1.52	5.1	-2.26
Drukker (2025) surprise alone	2.1	-1.44	5.5	-2.35	2.7	-1.64
Swanson (2021) LSAP factor	0.2	-0.43	0.6	+0.78	0.5	-0.70
Fed purchase-path revision, all quarters	9.4	-3.07	17.5	-4.18	11.4	-3.37
Fed purchase-path revision, crisis-robust	15.4	-3.93	1.1	-1.07	5.5	-2.35
composition revision, MBS less Treasury	2.1	-1.45	9.1	-3.02	0.1	+0.37
path revision, Treasury leg partialled out	8.9	-2.99	17.6	-4.20	1.9	-1.37
path revision + SOMA + Drukker	11.3	-3.03	12.2	-5.34	7.5	-2.98
Foreign official agency transactions (CSLT)	1.0	+1.02	0.4	+0.66	8.4	+2.91
GSE retained-portfolio transactions	0.9	-0.96	0.3	+0.58	7.3	-2.69
Treasury coupon Lewbel applied to agency MBS	1.8	+1.33	13.5	-3.67	3.0	+1.72

Notes: First-stage F statistics and t-statistics of each candidate instrument for the agency MBS block, on the 1990Q4–2025Q1 asset–quarter panel with the second-stage controls and Newey–West covariance of Table 3. "MBS spread" is the option-adjusted spread, the price the agency MBS demand block estimates; "Treasury yield" is the same regression run on the Treasury block's price and is the cross-asset placebo, which a supply shifter specific to agency MBS should leave unmoved; "MBS yield" is the spread plus the Treasury yield, the price a yield specification would use. A valid instrument for this block shows a large F in the first pair and a small one in the second. The Fed purchase-path revision is defined in Section 3.6 and built by the FOMC event table described there; "crisis-robust" zeroes it in the eight quarters 2008Q3–Q4, 2020Q1–Q2, and 2022Q1–Q4. "Treasury leg partialled out" enters the Treasury purchase-path revision as a control with a free coefficient, the unrestricted version of the composition instrument in the row above it.

Table A.19: **Weak-Instrument Checks: Single-Donor Instruments and the MBS Block**

Asset	Sector type	Single donor (baseline)		Three donors	
		$\beta^{(4)}$	AR set	$\beta^{(4)}$	AR set
Treasury: first-stage F		7.8		6.8	
	Banks	+0.09 (0.11)	[-0.06, 0.37]	+0.02 (0.10)	[-0.08, 0.04]
	Insurers	-0.03 (0.06)	[-0.16, 0.14]	-0.05 (0.06)	[-0.20, 0.15]
	Pension Funds	+0.03 (0.04)	[-0.04, 0.11]	+0.04 (0.03)	[-0.04, 0.14]
	Mutual Funds	+0.29 (0.13)	[0.04, 0.64]	+0.30 (0.14)	[-0.03, 0.67]
	Arbitrageurs	-1.89 (1.29)	[-7.33, -0.08]	-1.24 (0.89)	[-8.33, -0.53]^{u}
	Foreign	+0.01 (0.07)	[-0.07, 0.11]	-0.01 (0.07)	[-0.21, 0.14]
	Household Residual	+0.72 (0.28)	[0.16, 2.07]	+0.55 (0.24)	[0.26, 1.64]
	Other	+0.04 (0.13)	[-0.26, 0.33]	+0.04 (0.13)	[-0.28, 0.28]
Agency MBS: baseline ($F = 3.4$) vs. +Lewbel ($F = 7.4$)		$\beta^{(4)}$	Sargan p	$\beta^{(4)}$	Sargan p
	Insurers	+0.58 (0.32)	0.37	+0.13 (0.10)	0.081
	Foreign	+0.75 (0.47)	0.08	+0.30 (0.19)	0.179
	Mutual Funds	+2.57 (1.38)	0.21	+0.21 (0.35)	0.000
	Arbitrageurs	+1.65 (1.13)	0.62	+0.03 (0.26)	0.018

Notes: Long-run slopes $\beta^{(4)}$ (Newey–West standard errors) and 95 percent Anderson–Rubin sets from the local projections of Equation (9), by sector type, on 1985Q1–2025Q1. Upper panel: the baseline single-donor crisis-robust Lewbel instrument, the coupon interaction, against the three-donor block (maturity, issuance, coupon), with the asset–quarter first-stage F of each in the header row; ^{u} marks a set reaching the grid edge. Bottom panel: agency MBS cells under the baseline instruments (SOMA transactions and the Drukker surprise) and under the baseline plus the three MBS Lewbel interactions, with the Sargan overidentification p-value of each.

D.5 What the Household Residual Holds

The Treasury counterfactual rests on one cell, and the paper asserts that after 2010 that cell is mostly the hedge funds running the cash–futures basis trade. The Federal Reserve's Enhanced Financial Accounts report hedge-fund balance sheets from Form PF quarterly since 2012Q4 and let the assertion be checked.

Over 2013Q1–2025Q1 hedge funds hold 122 percent of the residual's coupon Treasuries on average, 99 percent at the median and 85 percent in the last quarter, and the two levels correlate at 0.93. In flows, the residual's Treasury transactions regress on the quarterly change in hedge-fund holdings with a coefficient of 0.62 (Newey–West standard error 0.17, $t = 3.7$). The residual and the basis trade are the same position.

They cannot, however, be separated into two sectors and estimated jointly, which would be the natural way to add observations to the cell that carries the aggregate. The residual net of hedge funds is negative in 24 of the 49 quarters. The two series are not on the same basis: the Financial Accounts figure is a *net* residual, total outstanding less every identified holder, while the Form PF figure is a *gross* long financed by repo, and a repo-financed long is also in the Accounts as the lender's holding, so differencing one from the other double-counts the financing leg.

What follows is a statement about interpretation rather than a correction. The slope the counterfactual applies to a hedge-fund position is identified where the same residual was households and nonprofits: the 2010–2025 window on its own gives -0.78 with a standard error of 0.82, an interval that contains every value the counterfactual could use. A levered arbitrageur and a buy-and-hold household are not obviously the same demand curve, and the lower Anderson–Rubin row of Panel (b) of Table 8, 41 basis points against 21, is the case in which the fund absorbs less readily than the household did.

D.6 The Household Residual over the Business Cycle

Every QE program ran in or just after a recession, and a demand system with one slope per cell prices them with a slope estimated mostly on expansions. This appendix adds one parameter, and only to the cell that needs it.

I screened all sixteen cells against three definitions of bad times, the NBER recession dates, quarters with the VIX in its top quintile, and quarters with negative real GDP growth, estimating

$$y_{i,t}^{(4)} \;=\; \beta_i^G \,\Delta y_t \;+\; \delta_i \left(\Delta y_t \times \mathrm{REC}_t\right) \;+\; \text{controls} \;+\; u_{it},$$

with each block's instruments entered twice, as Z_t and $Z_t \times \mathrm{REC}_t$, so both the level and the interaction are instrumented. Exactly one cell moves at conventional levels, and it is the household residual in Treasuries: $\beta^G = 0.36$ in expansions against $\beta^G + \delta = 3.31$ in recessions, with $\delta = 2.95$ and a standard error of 1.43, so $p = 0.038$. The same cell's agency MBS slope does not move, $+0.28$ on a standard error of 1.73, so the MBS block is left alone. Foreign investors move the other way, from $+0.08$ to -0.17 ($p = 0.033$): reserve managers step back when the household residual steps in.

The cell matters because it carries 78 percent of the Treasury aggregate demand slope at the 2025Q1 composition. Table A.20 re-runs the historical episodes with each program's household slope set to the slope of its own window, weighted by the share of quarters the NBER dates as a recession. The two recession programs fall into line with the literature: QE1 goes from -64 to -29 basis points against -25 in Eren et al. (2026), and QE4 from -222 to -170 against -121. QT1, which ran entirely in an expansion, goes from $+28$ to $+44$ against $+43$.

The restriction is not free, and the paper does not adopt it as the baseline. The expansion slope it implies, 0.36, is half the pooled 0.72 and carries a t of 1.15, so the four programs that ran in expansions become larger and less certain: QE3 moves from -146 to -189 against -57 in the same comparison. Splitting sixteen recession quarters out of a hundred and fiftynine buys a sharp answer for two episodes at the cost of precision everywhere else. What the exercise establishes is narrower and more useful: the asymmetry between QE and QT that Eren et al. (2026) document is in this system too, it lives in one cell, and it is the reason a symmetric specification overstates the recession-era programs.

Table A.20: **Historical Episodes with the Household Slope over the Cycle**

		ΔY Treasury (bp)			
Episode	Recession quarters	Pooled slope	Cycle slope	Household β used	Eren et al. (2026)
QE1 (2008Q4–2010Q1)	3 of 6	-64.2	-29.3	+1.84	-25
QE2 (2010Q4–2011Q2)	0 of 3	-129.4	-199.4	+0.36	-113
QE3 (2012Q3–2014Q4)	0 of 10	-145.8	-189.1	+0.36	-57
QT1 (2017Q4–2019Q3)	0 of 8	27.6	43.8	+0.36	43
QE4 (2020Q1–2021Q4)	2 of 8	-222.0	-169.7	+1.10	-121
QT2 (2022Q2–2025Q1)	0 of 12	81.1	136.7	+0.36	80

Notes: Treasury yield effects from the market-clearing counterfactual of Section 5. The pooled column is the paper's specification, with the household residual at 0.72 in every quarter. The cycle column sets that cell to $0.36 + 2.95s$, where s is the share of the episode's quarters the NBER dates as a recession; every other cell is unchanged. The last column is the corresponding estimate of Eren et al. (2026).

D.7 Nested Tests of the Slope Structure

The system estimates one slope per sector type and asset. Table A.21 asks whether the data need all sixteen, by minimum distance against two restrictions: one slope per type, which is the standard demand system, and a type factor times an asset loading. The first is rejected at $p = 0.04$ and the second is not rejected at $p = 0.29$, so the asset dimension matters but the data are content with a single loading ratio, 0.28 on the Treasury yield against 1.39 on the agency MBS spread. The counterfactual uses the unrestricted slopes.

Table A.21: **Nested Tests of the Slope Structure**

Restriction	Parameters	Distance χ^2 (df)	p
Common slope across assets, β_i	8	16.6 (8)	0.035
Proportional, $\beta_i\,\theta(n)$	9	8.5 (7)	0.290
Unrestricted, $\beta_i(n)$	16	0.0 (0)	—

Notes: Minimum-distance tests on the 16 long-run slopes of Table 4, treating the cells as independent with their Newey–West standard errors. "Common slope" restricts each sector type's slope to be equal across Treasuries and agency MBS (8 parameters); "proportional" restricts it to a type factor times an asset loading, $\beta_i(n) = \beta_i\,\theta(n)$ (9 parameters, loadings normalized to unit mean square); "unrestricted" is the estimated matrix. The distance statistic is χ^2 with degrees of freedom equal to the number of restrictions.