

The Triffin Margin: Who Absorbs Treasury Supply, and What It Costs

Claude

October 2026

Claude is an AI system developed by Anthropic; this paper was written by the Claude Fable 5 model generation. This paper was commissioned by Jason Furman. The data are public: Treasury's auction records, investor-class allotments, quarterly refunding statements and the annual surveys of foreign portfolio holdings of U.S. securities; the Treasury International Capital system's monthly holdings; the Federal Reserve's Financial Accounts and H.4.1 custody holdings; and the System Open Market Account holdings published by the Federal Reserve Bank of New York. A replication package accompanies the paper. The views expressed are those of the author.

The Triffin Margin: Who Absorbs Treasury Supply, and What It Costs

Claude · October 2026

ABSTRACT

Foreign official institutions held 47 percent of marketable U.S. Treasury coupon securities in 2008 and 16 percent in 2025. This paper asks who absorbs a dollar of new coupon supply, at the auction and afterward, and what the answer implies for the yield cost of supply. Using the investor-class allotments of every coupon auction since 2000, the monthly and weekly records of official holdings, and a new series on holder composition by remaining-maturity bucket, we find that the marginal dollar at auction is divided among investor classes roughly as the average dollar is, dealers taking 0.44, investment funds 0.35, and the foreign-and-international class 0.20 against an average share of 0.17, while the official sector absorbed the anticipated path of issuance when it was accumulating reserves and not since: official holdings rose by 0.95 of net coupon issuance between the 2004 and 2008 surveys and by 0.62 of each dollar within six months of issuance in the monthly record for 2000 to 2008, by 0.32 from 2009 to 2013, and by nothing since 2014, when the Federal Reserve and then private investors took the path instead. Official demand is a quantity bought mostly in the secondary market, not a response to the yield at the auction. The yield cost of absorption is harder to measure: within the sub-periods the auction-day concession does not rise with the bucket's official share, and announcement-day quantity surprises carry about two percent of the response that genuinely unanticipated supply produces, because Treasury's guidance removes the news. In a multi-bucket preferred-habitat model with a quantity-inelastic official sector, the fall in the official share lowered the yield cost of a one-percent supply surprise by a factor between 0.63 under proportional private demand and one under linear demand, and the daily data do not identify which.

JEL Codes: E43, G12, H63, F31

Keywords: Treasury supply; foreign official holdings; preferred habitat; auctions; convenience yield

Claude

c/o Jason Furman
Harvard Kennedy School
79 JFK Street, Cambridge, MA 02138

1 Introduction

In June 2008 foreign official institutions held 47 percent of the marketable U.S. Treasury notes and bonds outstanding, \$1,688 billion of a \$3,622 billion stock. In June 2025 they held 16 percent, \$3,560 billion of \$22,869 billion.¹ Reserve managers hold Treasuries because they need dollar reserves, not because the yield is attractive, and the part of the stock whose holders do respond to the yield, which this paper calls the elastic float, grew over those seventeen years from \$1,934 billion to \$19,309 billion. The question this paper asks is what that shift means for the cost of Treasury supply: when the government adds a dollar of debt, who takes it, at what horizon, how much must the yield rise to make them, and does the answer depend on how much of the stock sits in hands that will not move.

The question has a standard answer in words and no standard answer in numbers. In the hegemon model of Farhi and Maggiori (2018) the rest of the world demands the hegemon's safe asset for its safety and liquidity, and the hegemon's issuance is priced by a downward-sloping curve; policy discussion has turned that into the claim that a Treasury market with fewer official holders has become more price-sensitive, so that a given increase in supply now costs more in yield (Dilts Stedman, 2025). The claim has two parts, and they can be measured separately. Who absorbs the marginal dollar is a question about flows, answerable from the records of who was allotted each auction and of whose holdings rose afterward. What the absorption costs is a question about prices, answerable, in principle, from how yields move when supply arrives or is announced. The paper measures the first at three horizons, measures the second as far as daily data allow, and uses a model to say what the two imply together.

We build four data sets from public sources and a model that connects them. The composition data put the official share, the Federal Reserve's share, and the foreign private share of the coupon stock in each remaining-maturity bucket at every June from 2004 to 2025, from Treasury's annual survey of foreign holdings, the Financial Accounts, the Monthly Statement of the Public Debt, and the System Open Market Account's holdings by security. The auction data match every nominal coupon auction since 2000 to Treasury's investor-class allotment, which records how much of each issue went to dealers, investment funds, foreign and international investors, and the other classes. The holdings data put the monthly Treasury International Capital

¹The maturity table of the survey for June 2025 gives official holdings of Treasury notes and bonds by remaining maturity that sum to \$3,560 billion, while the table's own total line and the report's headline tables give \$3,500 billion. The paper uses the bucket sum, which the bucket shares require; on the report's total the official share in June 2025 is 15 percent.

record of official holdings of notes and bonds, and the weekly record of Treasury securities the Federal Reserve holds in custody for foreign official accounts, beside the coupon stock and the coupon settlements. The announcement data put every auction-size announcement beside the guidance Treasury gave before it, parsed from the quarterly refunding statements, so that a quantity surprise can be measured against what the market could have expected. The model is a multi-bucket preferred-habitat model in which official holders own a fixed quantity, private holders have demand with constant semi-elasticity, and arbitrageurs span buckets; its state variable is the elastic float, and it delivers the sign and size of the composition effect, what a within-day cross-bucket design identifies, and the auction concession.

The first finding is about who underwrites. Of an additional dollar of a coupon issue, dealers and brokers take 0.44 (s.e. 0.09) at the auction, investment funds 0.35 (0.07), and the foreign-and-international class 0.20 (0.07), against average shares of the private issue of 0.41, 0.41, and 0.17: the marginal dollar is divided roughly as the average dollar is, and no class's marginal share differs from its average share by more than sampling error allows. By period the foreign class's marginal share is below its average in 2000 to 2008, 0.12 (s.e. 0.07) against 0.20, a gap the data do not establish (bootstrap $p = 0.25$), and at or above it in every period since. The official sector's own channel at auction, the noncompetitive bids the Federal Reserve Bank of New York submits for foreign and international monetary authorities, is 0.07 percent of accepted amounts and capped by rule. The allotment record does not show the official sector standing aside at auction, and it cannot show the official sector by itself: the class that holds it pools official and private money, and that class takes its usual share of each increase.

The second finding is about who held the path of issuance afterward, which the allotments cannot see. Between the June 2004 and June 2008 surveys official holdings of coupons rose by \$775 billion while the coupon stock rose by \$813 billion, a ratio of 0.95; the ratio was 0.32 from 2009 to 2013, when the Federal Reserve took 0.24 and private investors 0.44, and 0.05 and -0.03 in the two periods since, when private investors held 0.90 and 0.83 of net issuance and the Federal Reserve the rest. The monthly record times it: in 2000 to 2008 a dollar of net coupon issuance was followed by a rise in official holdings of 0.62 (s.e. 0.27) within six months, in 2009 to 2013 by 0.32 (0.11), and after 2013 by nothing; the Federal Reserve's custody record for foreign official accounts, which includes bills, shows a rise of 0.59 (0.17) per dollar of net coupon settlements within eight weeks from July 2007 to August 2008 and nothing in any later

window.

Under the model these are ratios of two quantities, not marginal shares: official demand is set by reserve accumulation and bought in the secondary market, and whether the private sector had to hold the path of supply depended on whether that accumulation was running. Before 2009 the private sector held a small fraction of net issuance beyond the auction week; since 2014 it has held most of it, with the Federal Reserve holding the rest. The surprise component of supply is different, and the model says so: within the window of a surprise the official position is fixed, and the private sector has always had to warehouse it.

The third finding is about the price of absorption at auction. On the day a tenor is auctioned its yield rises relative to the rest of the curve and gives the rise back over the following days, the auction cycle of Lou et al. (2013), and across the sample the cycle is larger where official holders own more of the bucket: with tenor and month effects the concession rises by 4.7 basis points per unit of official share, with a standard error of 2.8 when clustered by year. But the relation is a comparison across periods, not within them. Within the four sub-periods the coefficient is negative in three and indistinguishable from zero in the fourth, the official share is not separable from tenor-specific trends, the relation is absent among the reopenings that carry no roll in the par curve's input, and the Federal Reserve's share of each bucket, the other inelastic position, carries no effect in the years of its first purchase programs and a positive one in the years of the pandemic purchases and the second run-off. The concession is a dealer-inventory premium of a fraction of a basis point, and the design's power against the model-implied composition effect is low.

The fourth finding is about news. The design that identifies the Triffin margin, the rate at which the yield cost of a given percentage increase in supply rises with the official share, interacts an unanticipated change in supply with the official share across the buckets announced on the same day. We build the surprise three ways, against the previous auction, against Treasury's parsed guidance, and as the quarter's path announced in each refunding statement, and over the full sample none moves yields: a surprise of one percent of a bucket's stock moves the tenor's yield by 0.36 basis points relative to the other tenors announced that day (s.e. 0.23), where a response of 1 basis point per billion dollars of genuinely unanticipated supply, the lower end of what Phillot (2025) finds from daily futures price changes, would be 10.5. Treasury's regular and predictable issuance, its guidance since 2010, and its size tables since 2020 remove

the quantity news before the announcement. In the two windows in which announced sizes did move yields, 2009 to 2013 and 2021 to 2025, the interaction with the official share is negative, but the official share and the tenor are nearly collinear within a window, the second window's response is larger than a supply effect and is absent under the guidance-netted measure, and the response per percent of stock rose between the windows as the official share fell, which under constant private demand the model does not predict. The futures-based shocks that would settle the question are deposited in an archive that requires a login we could not pass.

The model says what the findings mean together. If official demand is a quantity and private demand has constant semi-elasticity, the yield cost of a one-percent surprise increase in a bucket's supply is proportional to one over one minus the official share, and the cost per dollar to one over the elastic float. Evaluated at the 2008 and 2025 compositions, the fall in the official share lowered the cost per percent of stock by a factor of 0.63 and, together with the growth of the stock, the cost per dollar by a factor of 0.10, with the Federal Reserve's holdings counted as elastic, and by 0.61 per percent with them counted as inelastic. Those ratios do not depend on the private sector's semi-elasticity; they do depend on private demand being proportional to the sector's position, the demand-system form of Koijen and Yogo (2019), and under the linear habitat demand of Vayanos and Vila (2021) the ratio is one. The daily designs cannot tell the two forms apart, and the paper reports the ratio as a range from one to 0.63; the one piece of evidence that bears on it, the rise in the response to announced sizes between 2009 to 2013 and 2021 to 2025, points toward the top of the range. What the flow evidence adds is the distinction between surprise supply, which the private sector has always had to warehouse, and the anticipated path, which the official sector absorbed when it was accumulating and does not now.

The paper relates to five literatures, placed strand by strand in Section 2. The first estimates the effect of Treasury supply on yields with a single demand curve; we let the curve's slope depend on who holds the bucket. The second identifies supply shocks and supply costs from auctions; we add the bucket's holder composition and the allotment record of who underwrites. The third measures the price impact of official flows; we measure the official sector's absorption of supply directly, at three horizons. The fourth estimates demand holder by holder; we take its functional form and test its official-sector assumption against flows. The fifth is the hegemon and safe-asset literature; we borrow its reserve-demand assumption and return a measured absorption share. Section 3 sets out the model and its four propositions. Section 4

describes the data. Section 5 states the estimands, the designs, the threats, and the diagnostics. Section 6 reports who holds the stock, who underwrites at auction, who holds the marginal dollar afterward, what absorption costs at auction, and what supply costs when it is news. Section 7 gives the battery, Section 8 the cuts the model predicts, Section 9 the calibration and what a budget projection can take from it, and Section 10 concludes.

2 Related literature

Treasury supply and yields. The first strand treats the buyers of Treasuries as one demand curve and asks how steep it is. Krishnamurthy and Vissing-Jorgensen (2012) regress the spread between corporate and Treasury yields on the debt-to-GDP ratio over eight decades and recover a convenience-yield demand curve; Laubach (2009) and Engen and Hubbard (2004) regress forward rates on projected deficits and debt and find a few basis points per percentage point of debt to GDP; Greenwood and Vayanos (2014) move the object to the maturity composition of the debt and find that the long-term share raises long-term yields and expected excess returns; Du et al. (2018) and Jiang et al. (2025) find the convenience premium falling as the debt rises, the latter by maturity bucket with narrative fiscal instruments; Mian et al. (2022) collapse the strand into one semi-elasticity for a theory of deficits. On the Federal Reserve's side of the same curve, D'Amico and King (2013) and Cahill et al. (2013) measure local supply effects of purchases, and Bernanke et al. (2004) the relative-supply effect of Japanese official purchases. The investor base is homogeneous in every one of these designs. The one exception is a working-paper version of the first study (Krishnamurthy and Vissing-Jorgensen, 2008), which splits its sample at 1971, finds the demand curve roughly twice as steep afterward, and attributes the steepening to foreign central banks; the result did not survive into the published article. This paper lets the slope depend on the holder mix bucket by bucket and derives, rather than assumes, the form of the dependence.

Auctions and the auction cycle. The second strand identifies supply shocks and supply costs from Treasury's auctions. Phillot (2025) reads the change in Treasury futures prices on auction-announcement days as an unanticipated supply shock and finds that a surprise of a billion dollars raises yields by one to two basis points across the curve; Bi et al. (2026) narrow the window to thirty minutes, find a hump-shaped response across maturities, and find the response larger when debt is growing fast than when it is stable. Gorodnichenko and Ray (2017) use the

auction close to isolate demand shocks. Anticipated supply moves prices too: Lou et al. (2013) document the cycle in which yields rise into an auction and fall after and tie its size to the offering amount and to implied volatility, Fleming and Liu (2017) measure it within the day, Fleming (2007) show who bought at auction in the 2000s from the same allotment records this paper uses, and Somogyi et al. (2025) read demand elasticity off the auction demand curve and find the yield cost of a percent of supply several times larger since 2010 than before, with foreign participation making demand more elastic in their instrumented estimates. Du et al. (2023) model dealers as the balance-sheet arbitrageurs who absorb the Treasury supply their clients do not hold, which is the fast sector of this paper's model. The strand's demand curve varies by period or by regime, never by who holds the bucket being sold; this paper adds the bucket's holder composition as the state variable, reads the allotments for who underwrites the marginal dollar rather than for who bought on average, and shows that the quantity surprises the announcements carry, measured against Treasury's own guidance, are not the shocks the futures windows see.

The price impact of official flows. The third strand measures the yield effect of official purchases directly. Warnock and Warnock (2009) find long rates would have been substantially higher without a year of official inflows, on the premise that reserve managers do not optimize over the yield; Beltran et al. (2013) reject that premise, find a large short-run and a smaller long-run effect of official inflows, and find that in the long run the coefficients on official and private holdings are similar; Kaminska and Zinna (2020) estimate a preferred-habitat term-structure model with the Federal Reserve and foreign central banks inelastic and everyone else an arbitrageur; Ahmed and Rebucci (2022) identify the price impact of official flows by heteroskedasticity and find it about twice the literature's estimates; Dilts Stedman (2025) pools the strand into a range for the yield effect of a foreign liquidation; Vissing-Jorgensen (2021) documents the official selling of March 2020; Bertaut and Judson (2014) warn that custodial bias misallocates official holdings across countries and can understate them. The strand's inelasticity is assumed or inferred from monthly or annual flows and never put beside the auction record. This paper measures the official sector's absorption of supply at the auction, the weekly, and the monthly horizons from records the strand does not combine, and engages the Beltran et al. (2013) counter-result directly: a long-run holdings coefficient is not a marginal flow share, and the two can differ when the class that is measured changes character over the sample, as the foreign class does here after 2008.

Demand systems and holder elasticities. The fourth strand estimates demand holder by holder. Kojen and Yogo (2019) supply the demand-system technology that this paper’s private sector uses, with log quantity linear in expected returns, and Kojen and Yogo (2020) apply it across countries and estimate the convenience yield on U.S. long-term debt from the latent demand of foreign investors; Gabaix and Kojen (2021) supply the inelastic-markets reading of the same form. Eren et al. (2023) find the foreign official sector’s yield elasticity statistically zero against sizable elasticities for banks, funds, and foreign private investors; Chaudhary et al. (2024) find foreign official demand less elastic after 2009 as foreign private demand became more elastic; Jansen et al. (2024) add maturity buckets and a structural arbitrageur, identify demand with a market-clearing pseudo-yield instrument at quarterly frequency, find the foreign official sector’s own-yield response insignificant, and find market elasticity falling steeply with maturity; Tabova and Warnock (2021) show at the security level that private foreigners respond to hedged-return differentials and officials do not; Choi et al. (2022) find foreign demand less elastic than domestic; Fang et al. (2022) show, for emerging-market sovereign debt, that the yield cost of a debt increase depends on which investors are present. The strand’s elasticities are quarterly or annual, and where an auction-announcement shock appears in it, it is as an alternative instrument in an appendix rather than as the identifying variation for the official sector’s demand. This paper takes the strand’s functional form and its finding of an inelastic official sector as the model’s assumptions, tests the assumption at the margin from flows at three horizons, and takes the strand’s elasticities as the anchors for the calibration whose ratios do not depend on them.

The hegemon, the safe asset, and Triffin. The fifth strand explains why the United States faces downward-sloping demand for its debt and what it does with it. Farhi and Maggiori (2018) model a monopolist hegemon facing the rest of the world’s demand for its safe asset and derive a Triffin dilemma in which issuing more trades off against the risk of a confidence crisis; Caballero et al. (2016) impose a safe-asset demand that does not respond to the risk premium; He et al. (2019) make a large float a source of safety when global demand for safety is high; Gourinchas et al. (2019) and Jiang et al. (2021) measure the privilege and the convenience yield foreigners forgo; Choi et al. (2022) reject price-taking by the hegemon; Krishnamurthy and Vissing-Jorgensen (2011) add the safety clientele; Greenwood et al. (2014) show that Treasury’s maturity extension offset part of the Federal Reserve’s purchases and report, from an event study by one of their number (Rudolph), that Federal Reserve purchase announcements moved term

premia about twice as much as comparably sized Treasury announcements; Bordo and McCauley (2017) and Chinn et al. (2024) supply the cautions, the first on whether dollar reserves are Treasuries at all and the second on how slowly reserve shares move, with a coefficient near 0.9 on the lagged share and trade shares and pegs entering beside it in their country-level estimates; Duffie (2023) describes the intermediation capacity that the paper's fast sector stands for. The strand's demand is one aggregate world demand with no maturity structure, and its Triffin is a question of systemic sustainability; this paper borrows the word for the thinness of the private margin that must clear the market when official demand does not respond to price, and gives that margin a maturity structure and a measurement.

The closest papers. Relative to Jansen et al. (2024), who have buckets, a sector demand system, and a market-clearing instrument at quarterly frequency, this paper works at the auction, the week, and the month, measures the official sector's absorption from flows, and tests the interaction within days. Relative to Bi et al. (2026), who have the futures shock and a regime split by debt growth, this paper replaces the regime with the bucket's holder composition and shows, by construction of the surprise against Treasury's guidance, why the quantity in the announcement is not where the news is. Relative to Somogyi et al. (2025), whose auction-level elasticities rise over time and whose foreign-share instrument makes demand more elastic, this paper finds the same direction in the flow data, a foreign class whose marginal take-up rose as it became private money, and no contradiction: the official sector's average participation at auction is large and its marginal participation small. Relative to Dilts Stedman (2025), who states the hypothesis and reports the literature's range for a foreign liquidation but no estimate of the composition effect, this paper supplies the model and the tests and reports which of them the data can carry. Relative to Ahmed and Rebucci (2022), who identify the price impact of official flows at a few maturities, this paper holds official behavior fixed within the window and lets supply move.

3 Model

The model is the simplest one that turns the question into estimating equations and carries the paper's quantitative work. It has three kinds of holders in each maturity bucket, one of which does not respond to the yield within the window the data observe, and a state variable, the elastic float, that summarizes how much of a bucket's stock sits in hands that do respond. Proofs and the

closed forms are in Appendix A.

3.1 *Environment*

Marketable Treasury coupon securities are grouped into remaining-maturity buckets $b = 1, \dots, B$ (the buckets of Treasury's annual survey of foreign holdings: one year or less, one to two years, and so on). The stock outstanding in bucket b is Q_b . Time is a window short enough that only a portfolio decision, not a saving decision, is made: a day for an announcement, a week for an auction. Three kinds of holders clear the market. Official holders (foreign reserve managers and, in the extension, the Federal Reserve's System Open Market Account, SOMA) hold $\theta_b Q_b$, a quantity set by reserve needs and policy rather than by the yield; within the window it does not move. Private habitat investors demand the bucket's securities with a constant semi-elasticity with respect to the yield. Arbitrageurs, in the two-bucket extension, hold positions across buckets and bear the risk of doing so. The one object the paper is about is the elastic float,

$$F_b = Q_b (1 - \theta_b) \tag{1}$$

the part of the bucket's stock held by investors whose demand responds to the yield.

3.2 *Agents*

The official sector's demand is $\theta_b Q_b$ at any yield. The rest of the world in Farhi and Maggiori (2018) holds the hegemon's safe asset for its safety and liquidity along a demand curve on which more issuance requires a higher yield; this paper takes the limit in which, at the frequencies at which supply arrives, the official quantity does not respond, which is the safe-asset demand of Caballero et al. (2016) and the official sector of Kaminska and Zinna (2020). That limit is an assumption, and Section 6 tests it against flows.

The private habitat sector's demand for bucket b is

$$D_b^P(y_b) = F_b \exp(\varepsilon (y_b - \bar{y}_b)), \tag{2}$$

where y_b is the bucket's yield, $\bar{y}_b$ is the yield at which the sector holds exactly the float, and $\varepsilon > 0$ is the semi-elasticity of demand with respect to the yield, in percent of quantity per percentage point of yield. The functional form is the demand-system form of Kojen and Yogo (2019) and Gabaix and Kojen (2021): the percentage change in quantity demanded per unit change in the expected return is a constant, so the dollar response to a given yield change is proportional to the

sector's existing position. That proportionality is the whole of what the composition effect rests on. Under the alternative in which private demand is linear in the yield with an absolute slope, $D_b^P = A_b + \lambda y_b$, the dollar response is λ whatever the sector's position, and composition drops out of every result below. The linear form is the habitat demand of Vayanos and Vila (2021) and Greenwood and Vayanos (2014), so the canonical preferred-habitat model is the alternative, not the model; the paper calls it the level alternative and returns to it in Proposition 2.

Arbitrageurs appear only in the two-bucket extension of Section 3.5. They choose holdings X_1, X_2 to maximize expected excess return less a quadratic penalty with a common (level) risk of variance σ^2 and a bucket-specific risk of variance v , with risk aversion a ; their first-order condition for bucket b is $y_b - r = a\sigma^2(X_1 + X_2) + avX_b$, so that

$$dX_1 = \frac{-adS\sigma^2 + dy_1}{av} \quad (3)$$

in changes, where $dS = dX_1 + dX_2$. The structure is the preferred-habitat structure of Vayanos and Vila (2021) reduced to two buckets and one period, with the common risk loading equally on both buckets; a duration-proportional loading, which the canonical model has, would make the spillover from a shocked bucket tenor-specific, a point Section 3.5 returns to.

3.3 Equilibrium

An equilibrium in bucket b is a yield at which the three holders' demands sum to the stock. In the segmented case (no arbitrageurs) the clearing condition is

$$Q_b = Q_b\theta_b + Q_b(1 - \theta_b)e^{\varepsilon(y - \bar{y})} \quad (4)$$

which holds with $y_b = \bar{y}_b$ by construction of the float. A supply surprise s_b is an unanticipated change in Q_b with the official position $\theta_b Q_b$ held fixed, so that the whole of s_b must be absorbed by the elastic holders:

$$Q_b + s_b = Q_b\theta_b + Q_b(1 - \theta_b)e^{\varepsilon(y - \bar{y})} \quad (5)$$

Solving (5) for the yield after the surprise gives the key equation,

$$y_b = \bar{y} + \frac{\log\left(1 + \frac{s_b}{Q_b(1 - \theta_b)}\right)}{\varepsilon} \quad (6)$$

The yield response is the logarithm of one plus the surprise as a fraction of the elastic float, divided by the semi-elasticity. Nothing about the stock matters except through the float.

3.4 The comparative statics the data test

Proposition 1 (Sign and the Triffin margin). *In the segmented model the yield response to a supply surprise is positive and inversely proportional to the elastic float,*

$$\frac{dy_b}{ds_b} = \frac{1}{Q_b \varepsilon (1 - \theta_b)} \quad (7)$$

Expressed per unit of the surprise as a fraction of the stock, $x_b = s_b/Q_b$, the response is

$$\frac{dy_b}{dx_b} = \frac{1}{\varepsilon (1 - \theta_b)} \quad (8)$$

which rises with the official share:

$$\frac{\partial}{\partial \theta_b} \frac{dy_b}{dx_b} = \frac{1}{\varepsilon (1 - \theta_b)^2} \quad (9)$$

Proposition 1 is the holdings-weighted aggregation of the demand-system literature with one sector's elasticity set to zero: the market's dollar response to a yield change is the sum of each holder's position times its semi-elasticity, and the official sector contributes nothing to the sum. The intuition is a matter of accounting. A dollar of new supply must be taken by someone whose demand moves with the yield. When official holders own θ_b of the bucket and do not move, the elastic sector owns $1 - \theta_b$ of it, and the yield must rise enough to induce that smaller sector to expand by the full dollar. With constant semi-elasticity the sector's dollar response is proportional to its position, so the required yield change is inversely proportional to the float. The derivative in (9) is the Triffin margin: the rate at which the yield cost of a given percentage increase in supply rises as the inelastic share rises. Linearizing (8) around a reference share θ gives

$$\frac{dy_b}{dx_b} = \frac{1}{\varepsilon (1 - \theta)} + \frac{-\theta + \theta_b}{\varepsilon (1 - \theta)^2} \quad (10)$$

so a regression of the yield change on the surprise and on the surprise interacted with the deviation of the bucket's share from the reference has coefficients whose ratio the model fixes.

Proposition 2 (The ratio and the level alternative). *In the linearized form (10), the coefficient on the interaction divided by the coefficient on the surprise equals*

$$\frac{\beta_2}{\beta_1} = \frac{1}{1 - \theta} \quad (11)$$

and the elasticity of the response with respect to the float is minus one. Under the level alternative the response is

$$\frac{dy_b}{ds_b} = \frac{1}{\lambda} \quad (12)$$

and its derivative with respect to the official share is zero.

Proposition 2 turns the sign prediction into a magnitude prediction, and it names the alternative that would make composition irrelevant. The nested form $dy_b/dx_b = \kappa(1 - \theta_b)^{-\psi}$ contains both: $\psi = 1$ is the model, $\psi = 0$ is the level alternative, and ψ can be estimated. Every ratio the calibration reports is a function of ψ .

3.5 What a within-day design identifies

The news design compares buckets within an announcement day, so that whatever moves the whole curve that day is differenced out. With arbitrageurs spanning buckets, a supply surprise in one bucket moves every yield, and the within-day comparison identifies a relative response. The two-bucket extension says what that relative response is. Market clearing in each bucket and the arbitrageurs' first-order conditions give, for a surprise s_1 in bucket 1 and none in bucket 2,

$$s_1 = F_1 dy_1 \varepsilon + dX_1 \quad (13)$$

$$0 = F_2 dy_2 \varepsilon + dX_2 \quad (14)$$

together with (3) and its twin for bucket 2.

Proposition 3 (The within-day estimand). *In the two-bucket model with arbitrageurs, the own-bucket response dy_1/ds_1 , the spillover dy_2/ds_1 , and the relative response $d(y_1 - y_2)/ds_1$ are all positive. The relative response falls with the elastic float of the shocked bucket, $\partial[d(y_1 - y_2)/ds_1]/\partial F_1 < 0$, with an elasticity with respect to F_1 strictly between minus one and zero, and it rises with the bucket-specific risk v . As $v \rightarrow \infty$ it converges to the segmented response,*

$$\lim_{v \rightarrow \infty} \frac{d(y_1 - y_2)}{ds_1} = \frac{1}{F_1 \varepsilon} \quad (15)$$

and as $v \rightarrow 0$ it converges to zero. The ratio of the relative response to the segmented response lies strictly between zero and one.

Proposition 3 does three things. It says that the day fixed effect removes the common component of the response, the part arbitrageurs spread across the curve, and leaves a relative response that is the segmented response scaled by an attenuation factor in $(0, 1)$. It says that the dependence of the relative response on the float, which is what the interaction test looks for,

survives the attenuation: the sign of Proposition 1 is a sign of the relative response too. And it says that the magnitude does not survive it: because the elasticity with respect to the float is above minus one, the within-day ratio β_2/β_1 is below the segmented $1/(1 - \bar{\theta})$, so the segmented ratio is an upper bound for the within-day design and the power statements of Section 6 are made against that bound. What the within-day design does not identify is the total response dy_1/ds_1 , which includes the common component; a pooled regression without day effects is needed for that, at the cost of letting in everything else that moved the curve. With duration-proportional loadings the common component would itself differ across tenors, so that a surprise's effect would carry a tenor gradient the day effect does not remove; the empirical section allows for that gradient directly. The closed forms are in Appendix A.

3.6 Absorption at auction, and after

Auctions are announced days ahead, so the surprise at an auction is small, but the supply still has to be taken. The two-stage version of the model separates the holders who take it on the day from the holders who hold it in the end. At the auction only the fast part of the private sector, dealers and the quick habitat investors, responds, with semi-elasticity ϵ_0 per unit of float; by the end of the week the whole private sector responds, with $\epsilon_1 > \epsilon_0$. The official sector takes a quantity G_b of the issue at the auction, at any price, and holds whatever reserve accumulation requires afterward.

Proposition 4 (Two-stage absorption). *With the official sector taking the quantity $G_b < s_b$ at the auction, the auction-day yield change relative to the pre-auction yield is*

$$\Delta y_b^{(0)} = \frac{-G_b + s_b}{Q_b \epsilon_0 (1 - \theta_b)} \quad (16)$$

The marginal official share of the issue, dG_b/ds_b , is zero; the average share G_b/s_b falls with the issue; and the concession is positive, rises with the official share θ_b , and falls with G_b . Writing the official take-up as a share of the issue, $G_b = \eta_b s_b$, the auction-day change is

$$\Delta y_b^{(0)} = \frac{s_b (1 - \eta_b)}{Q_b \epsilon_0 (1 - \theta_b)} \quad (17)$$

the yield change once the whole private sector has responded, written for the surprise component of the issue (the anticipated component was priced at the announcement, so s_b in the next expression is the surprise alone), is

$$\Delta y_b^{(\infty)} = \frac{s_b (1 - \eta_b)}{Q_b \epsilon_1 (1 - \theta_b)} \quad (18)$$

and the ratio of the auction-day change to the permanent change, $\varepsilon_1/\varepsilon_0$, does not depend on composition. An official sector that did respond to the yield, with semi-elasticity ε_{off} , would take of a marginal dollar at the auction

$$m_b^{off} = \frac{\varepsilon_{off}\theta_b}{\varepsilon_0(1-\theta_b)} \quad (19)$$

Proposition 4 has an auction-cycle reading and a composition reading, and it says what the allotment data can and cannot test. Its price predictions, that the concession is proportional to the issue's size relative to the bucket's stock and rises with the official share within a period, are the ones Section 7 tests and the data reject; the auction-day concession behaves as the inventory premium of a capacity-constrained fast sector, the balance sheet of Du et al. (2023), rather than as the float-scaled response the proposition writes. The cycle reading is the one in Lou et al. (2013) and Fleming and Liu (2017): the yield rises into and at the auction and falls after, because the fast sector must be paid to warehouse what the slow sector will later hold; Du et al. (2023) model the balance sheet that does the warehousing. The composition reading is that the concession is larger where the float is smaller, that is, where official holders own more of the stock, and smaller where officials take more of the issue at the auction itself. The permanent change applies to the surprise component of the issue only: an anticipated issue was priced at its announcement, and the whole of it is warehoused on the day while only the news in it moves the yield for good.

The last expression is the test's limit. At a measured concession of a fraction of a basis point, ε_0 is in the thousands, and an official sector as elastic as the private sector would take a marginal share at the auction that is a small fraction of one percent; the allotment data can show that the official take-up at auction is a quantity, which is assumption (iv) below, and cannot reject the hypothesis that official demand responds to the yield, which is assumption (i). The test of assumption (i) is the holdings data at the weekly and monthly horizons.

3.7 *Official accumulation and the two horizons*

The official position is a quantity, but a quantity that grows. Write the change in official holdings over a horizon as ΔO_b and the change in the stock as ΔQ_b . The change in the elastic float is $\Delta F_b = \Delta Q_b - \Delta O_b$, and the share of net issuance the private sector must hold at that horizon is $1 - \Delta O_b/\Delta Q_b$. If reserve accumulation runs at a pace that matches issuance, as it did in the mid-2000s, the private float does not grow and the anticipated path of supply costs the private sector nothing to hold; if reserve accumulation stops, the private sector holds every dollar of the

path. The surprise component is different: within the window of a surprise the official position is fixed by assumption, and the cost of the surprise is (7) at the float of the moment. The cost of a supply path therefore has two parts, the warehousing of the whole path by the fast sector at the auction, and the permanent absorption of the path by whoever is accumulating; the composition of the stock governs the first and the pace of official accumulation the second.

3.8 The cost decomposition

Combining Proposition 1 with the growth of the stock gives an accounting identity for the yield cost of a dollar of surprise supply:

$$\log \frac{dy_b}{ds_b} = -\log(Q_b) - \log(\epsilon) - \log(1 - \theta_b) \quad (20)$$

The cost per dollar falls when the semi-elasticity rises, when the official share falls, and when the stock grows. Between two dates the change in the log cost per dollar decomposes into a composition term and a stock term; the change in the cost per percent of stock is the composition term alone. Section 9 evaluates both at the 2008 and 2025 compositions.

3.9 The assumptions the signs rest on

Four assumptions deliver the signs: (i) the official position does not respond to the yield within the window; (ii) private demand has constant semi-elasticity, so its dollar response is proportional to its position; (iii) arbitrageurs are risk-averse with bucket-specific risk, so the within-day relative response is attenuated but not zero; and (iv) at auction the official take-up is a quantity, not a share of the issue. Relaxing (i) so that officials buy more when yields are high would lower the Triffin margin and, if official demand were as elastic as private demand, eliminate it; relaxing (ii) to the level alternative eliminates it outright; relaxing (iii) to perfect substitutability makes the within-day design uninformative, since the whole response becomes common; relaxing (iv) so that officials take a fixed share of every issue would make the marginal official share equal the average share at auction while leaving the concession's dependence on θ_b intact. Assumptions (i) and (iv) are distinct: the allotments test (iv), the holdings at the weekly and monthly horizons bear on (i), and nothing in the daily data tests (ii).

3.10 Mapping to data and the estimating equations

The model's objects map to data as follows. The official share θ_b is the share of bucket b 's marketable coupon stock held by foreign official institutions, from Treasury's annual survey of foreign holdings and the Monthly Statement of the Public Debt; the SOMA share θ_b^F is the Federal Reserve's share from the SOMA holdings by CUSIP; their sum is the inelastic share. The surprise s_b is the announced auction size less the size the market could have expected, in dollars, and x_b is the same in percent of the bucket's stock. The yield y_b is Treasury's constant-maturity par yield at the bucket's tenor. The auction issue q_b is the auction size in percent of the bucket's stock; a slot is a tenor and a new-issue or reopening status; the private issue is the issue less the SOMA add-on; and the classes' marginal shares are the coefficients of each class's allotment on the private issue. The official sector's absorption at longer horizons is the coefficient of the change in official holdings on the change in the stock.

Four estimating equations follow. The news interaction, Proposition 2,

$$\Delta y_{b,t} = \alpha_t + \beta_1 x_{b,t} + \beta_2 x_{b,t} (\theta_{b,t-1} - \bar{\theta}) + e_{b,t}, \quad (21)$$

with announcement-day effects α_t , predicts $\beta_1 > 0$, $\beta_2 > 0$, and $0 < \beta_2/\beta_1 \leq 1/(1 - \bar{\theta})$. The concession, Proposition 4,

$$c_{b,a} = \alpha_b + \alpha_{m(a)} + \gamma \theta_{b,a-1} + \phi v_{b,a} + e_{b,a}, \quad (22)$$

with tenor effects α_b , month effects $\alpha_{m(a)}$, and the tenor's recent volatility $v_{b,a}$, predicts $\gamma > 0$; its scaled form replaces $\gamma\theta$ by a constant κ_0 times $q_{b,a}/(1 - \theta_{b,a-1})$. The marginal absorption at auction, Proposition 4,

$$\Delta A_{b,a}^k = m_k \Delta I_{b,a} + e_{b,a}^k, \quad (23)$$

where ΔI is the change in the private issue between consecutive auctions in the same slot and ΔA^k the change in class k 's allotment, predicts $m_{\text{official}} \approx 0$ and $\sum_k m_k = 1$. The absorption at the horizon h ,

$$O_{t+h} - O_{t-1} = a_h + b_h \Delta Q_t + c_h (y_{t+h} - y_{t-1}) + e_{t,h}, \quad (24)$$

where O is official holdings and ΔQ_t the change in the coupon stock in month (or week) t , measures b_h , the co-movement of official purchases over the following h periods with the period's net issuance, and a_h , the pace of official accumulation that does not depend on the period's issuance. Under the model's quantity assumption the pace sits in a_h and a positive b_h

says where and when the official sector sourced its purchases, through dealers in the weeks after an auction or not; it does not measure a marginal share, and the ratio of official accumulation to net issuance over a period is the sum of the two. The holdings records bear on assumption (i) only in the weak sense that a demand responding to the yield would leave a trace in the yield coefficient c_h of the monthly regression; the paper reports it and reads it with that limit.

3.11 What the model does not say

The model is silent on four things the data will raise. It does not say how long the fast sector holds the paper, so it does not predict the shape of the post-auction path beyond its sign. It does not separate foreign official from foreign private buyers inside the allotment class that pools them. It says nothing about why the quantity surprises Treasury announces should carry news; if the market anticipates them, equation (21) has no first stage, which is a fact about Treasury's communication rather than about the model. And it does not say how the response to news varies with the tenor for reasons other than composition, such as a policy rate held at its lower bound or a duration-proportional common risk, which the empirical section has to confront.

4 Data

The paper builds four data sets from public sources: the composition of the marketable coupon stock by remaining-maturity bucket and holder; the universe of nominal coupon auctions with their investor-class allotments; the monthly and weekly records of official holdings beside the coupon stock and the coupon settlements; and the auction-size announcements with the guidance Treasury gave before them. Appendix B documents each construction step; the replication package rebuilds every file from the raw downloads.

4.1 Who holds the stock, by bucket

Treasury's annual survey of foreign portfolio holdings of U.S. securities reports, for the end of June each year, the maturity distribution of foreign official institutions' holdings of long-term Treasuries and the same distribution for foreign private holders, in fifteen remaining-maturity buckets, the last of which, beyond thirty years, is empty for Treasuries (U.S. Department of the Treasury, 2025). The distribution is a table of percentages in the 17 reports that print it and, from the 2023 report onward, a table of dollar amounts by holder that covers




Figure 1. Who holds the marketable coupon stock, 2004 to 2025

Notes: Shares of the marketable coupon stock at each June survey date: foreign official institutions, the Federal Reserve’s System Open Market Account, and foreign private holders. Source: as Table 1; authors’ calculations.

2017 to 2025. We convert the percentage tables to dollars with the Financial Accounts’ end-June foreign official and foreign private long-term Treasury positions (Board of Governors of the Federal Reserve System, n.d.b), which are themselves benchmarked to the survey; where the two sources overlap, the dollar table and the converted percentages agree within 1.2 percent at the market level. The stock in each bucket is the par value of notes, bonds, inflation-indexed, and floating-rate securities with remaining maturity in the bucket at the June month-end, from the Monthly Statement of the Public Debt by CUSIP (U.S. Department of the Treasury, Fiscal Data, n.d.a). The Federal Reserve’s holdings by bucket come from the SOMA’s CUSIP-level file at the nearest as-of date (Federal Reserve Bank of New York, n.d.); because the SOMA file is weekly and the debt statement monthly, an as-of date lies within a few days of every month-end, so the SOMA share is available for every month-end without interpolation. The official share θ_b is official dollars over the bucket’s stock; the SOMA share θ_b^F is the Federal Reserve’s dollars over the same stock; private foreign holdings are built the same way. Between June surveys the official share is interpolated linearly and flagged; the robustness section uses only the months near a survey. Table 1 reports the official share by bucket for selected years and Figures 1 and 2 the market-wide shares over time and the shares by bucket; Section 6 reads them.

Table 1. Foreign official share of the marketable coupon stock, by remaining maturity

Remaining maturity (years)	2004	2008	2012	2016	2020	2025
0–1	0.32	0.46	0.38	0.33	0.18	0.14
1–2	0.40	0.58	0.50	0.34	0.27	0.16
2–3	0.41	0.60	0.52	0.35	0.27	0.19
3–4	0.43	0.56	0.43	0.36	0.37	0.20
4–5	0.48	0.58	0.49	0.32	0.33	0.17
5–6	0.35	0.53	0.44	0.42	0.39	0.20
6–7	0.40	0.58	0.39	0.36	0.27	0.20
7–8	0.30	0.55	0.38	0.43	0.32	0.29
8–9	0.34	0.46	0.36	0.36	0.28	0.27
9–10	0.37	0.40	0.26	0.28	0.23	0.24
10–15	0.07	0.18	0.12	0.14	0.16	0.05
15–20	0.08	0.07	0.07	0.11	0.05	0.08
20–25	0.05	0.08	0.05	0.16	0.16	0.06
25–30	0.07	0.14	0.12	0.13	0.12	0.04
All coupons	0.33	0.47	0.39	0.32	0.25	0.16
SOMA, all coupons	0.15	0.13	0.19	0.21	0.26	0.18
Stock (\$ billions)	2,809	3,622	8,860	11,910	14,819	22,869

Notes: Foreign official institutions' holdings of long-term Treasuries in each remaining-maturity bucket, from the maturity tables of Treasury's annual survey of foreign portfolio holdings (percentages scaled by the Financial Accounts' end-June official long-term Treasury position for 2004 to 2020; the survey's dollar tables for 2021 to 2025), divided by the par value of marketable notes, bonds, inflation-indexed, and floating-rate securities in the bucket at the June month-end. The last rows give the market-level official share, the Federal Reserve's share from SOMA holdings, and the coupon stock. Source: Treasury (TIC, Fiscal Data), Federal Reserve Board (Z.1), Federal Reserve Bank of New York (SOMA); authors' calculations.




Figure 2. Foreign official share by remaining-maturity bucket, 2008, 2016, and 2025

Notes: The foreign official share of each bucket's coupon stock at the June survey date. Source: as Table 1; authors' calculations.

4.2 Auctions and who takes them

Treasury's auction records give every marketable auction with its announcement date, auction date, offering amount, the amounts accepted from competitive, noncompetitive, direct, indirect, and SOMA bidders, the noncompetitive amounts the Federal Reserve Bank of New York submits for foreign and international monetary authority (FIMA) accounts, and the stop-out yield (U.S. Department of the Treasury, Fiscal Data, n.d.b). The sample is the 1,567 nominal note and bond auctions of the 2-, 3-, 5-, 7-, 10-, 20-, and 30-year tenors from 2000 to 2025; inflation-indexed and floating-rate issues are excluded. Each auction is matched to its investor-class allotment from Treasury's release for the auction cycle, which splits every issue among depository institutions, individuals, dealers and brokers, pension and retirement funds and insurers, investment funds, foreign and international investors, a small residual, and the SOMA add-on (U.S. Department of the Treasury, n.d.b); 1,567 auctions have a matched allotment. We remove the SOMA add-on and express each class's allotment as a share of the privately placed issue. Table 2 gives the summary statistics by tenor. The foreign-and-international class averages 0.17 of the private issue, dealers and brokers 0.41, and investment funds 0.41; the indirect-bidder share of accepted amounts in the auction results averages 0.50 (Figure 3). The

Table 2. Summary statistics for the auction sample

Tenor	Auctions	Size (\$ bn)	Size (% of stock)	θ (official)	θ^F (SOMA)	Concession (bp), mean	s.d.	Foreign allotment share	Dealer allotment share
2-year	301	34.8	2.85	0.39	0.15	-0.22	3.66	0.19	0.49
3-year	221	36.4	3.10	0.38	0.18	0.19	3.00	0.18	0.43
5-year	284	35.0	4.54	0.38	0.17	-0.32	2.37	0.18	0.44
7-year	214	35.1	4.61	0.33	0.20	0.34	2.67	0.18	0.31
10-year	268	23.6	7.07	0.32	0.13	-0.14	2.30	0.18	0.43
20-year	68	16.8	1.62	0.09	0.45	-0.79	2.62	0.13	0.17
30-year	211	16.5	2.61	0.10	0.20	0.56	3.55	0.11	0.36
All	1567	29.9	4.07	0.31	0.18	-0.01	2.97	0.17	0.41

Notes: Nominal coupon auctions, 2000 to 2025. Size is the offering amount; the percent of stock divides it by the bucket's coupon stock at the end of the month before; θ is the foreign official share of that stock and θ^F the Federal Reserve's exact monthly share; the concession is the auctioned tenor's yield change on the auction day relative to the mean change across the other available tenors; the allotment shares are shares of the privately placed issue. Source: Treasury auction records, investor-class allotments, and par yields; the composition series; authors' calculations.

foreign-and-international class pools foreign official institutions with foreign private investors, so the stock-side series above is the paper's measure of the official part; the class's average share and the stock-side official share move together over the sample.

Two institutional facts shape who takes an auction. Primary dealers are expected to bid in every auction for at least their pro rata share of the offered amount at prices reasonable against the when-issued market (Federal Reserve Bank of New York, 2016), so the dealer community underwrites the whole issue by construction. And the official sector's own noncompetitive channel at auction is small by rule: FIMA noncompetitive awards were capped at \$100 million per account and \$1 billion per auction until March 2023, and at \$500 million and \$2 billion since (U.S. Department of the Treasury, 2023). Official buying beyond that cap goes through dealers, either as competitive bids a dealer places for an official customer, which the allotment record attributes to the foreign-and-international class, or as purchases in the secondary market after the auction, which it does not see; the allotment record can therefore say who took the issue at the auction, including official customers bidding through dealers, and not who held it afterward.

4.3 *Official holdings at the weekly and monthly horizons*

Two records follow official holdings between the annual surveys. The Treasury International Capital system's monthly table of major foreign holders gives foreign official institutions' holdings of Treasury notes and bonds at each month-end; before 2012 the series was the annual survey benchmark carried forward with the monthly transactions data and revised when each survey arrived, and since January 2012 the holdings are reported monthly on the TIC




Figure 3. Who takes coupon auctions, 2000 to 2025

Notes: Annual means of each class's share of the privately placed issue from Treasury's investor-class allotments, and of the indirect-bidder share of accepted amounts from the auction results. Source: Treasury investor-class auction allotments and auction records; authors' calculations.

Form SLT, with only the latest month estimated from transactions (U.S. Department of the Treasury, n.d.d). The Federal Reserve's weekly statement reports the marketable Treasury securities, bills included, that it holds in custody for foreign official and international accounts, at face value, each Wednesday from July 2007 (Board of Governors of the Federal Reserve System, n.d.a). Beside the first we put the monthly change in the marketable coupon stock from the Monthly Statement of the Public Debt by CUSIP and the Federal Reserve's own coupon holdings from the SOMA file, and beside the second the coupon auctions settling in each statement week, net of the coupon securities maturing that week, and the bill auctions settling in the week. The monthly record runs 299 months from 2001 to 2025, with a mean monthly change in the coupon stock of \$72 billion; the weekly record runs 965 weeks from July 2007 to December 2025, with mean weekly coupon settlements of \$54 billion, coupon maturities of \$38 billion, and bill settlements of \$245 billion. The custody series is a subset of official holdings, those held at the Federal Reserve Bank of New York rather than with private custodians, and it is at face value, so it carries no valuation effects; the monthly series carries them, and the regressions that use it control for the change in the ten-year yield.

4.4 Announcements, guidance, and surprises

Treasury announces each coupon auction's offering size five to seven business days before the auction: the 3-, 10-, and 30-year on the Wednesday of the quarterly refunding, the 2-, 5-, and 7-year on a Thursday later in the month, and the 10-year reopenings and the 20-year on their own days. The naive surprise is the announced size less the size of the previous auction in the same slot. The guidance-netted surprise is the announced size less the size implied for that tenor and month by the most recent quarterly refunding statement: an explicit monthly size table in the statements from 2020 onward, a stated path of increases per month in 2018 and 2019 and again in 2023 and 2024, or a statement that sizes would be unchanged; where a statement implies nothing, the guidance-netted surprise equals the naive one. The parse of the 117 statements is audited line by line in the replication package (U.S. Department of the Treasury, n.d.c). A third measure, the path news, treats each refunding statement as announcing the whole quarter's issuance path and measures it against the previous statement's path. Each surprise is scaled by the bucket's stock at the end of the month before the announcement (x , in percent of stock) and by the elastic float (z , in percent of float).

The surprises are small and mostly zero. Across the 1,564 announced auction-tenor observations, 26 percent carry a nonzero naive surprise, and among those the median absolute surprise is 0.16 percent of the bucket's stock. Both numbers fall over the sample as the stock grows and as Treasury's guidance lengthens: by 2014 to 2017 only 6 percent of announcements change the size at all. On 168 of the 655 announcement days a single tenor was announced; those days carry no within-day comparison and drop out of the estimates that use day effects.

4.5 Yields and windows

Yields are Treasury's constant-maturity par yields at the 2-, 3-, 5-, 7-, 10-, 20-, and 30-year tenors (U.S. Department of the Treasury, n.d.a). For an announcement the outcome is the change at the bucket's tenor from the close of the day before to the close of the announcement day, in basis points. For an auction the outcome is the same change on the auction day less the mean change across the other tenors available that day, which removes the day's common curve move: Treasury auctions one coupon tenor on most days, so a day fixed effect is not available and the cross-tenor mean takes its place. The run-up is the relative change from five (or two) business days before the auction to the auction day, the reversal the relative change from the auction day to

five (or two) days after, and the event path the relative one-day changes at each business day from five before to five after. The tenor’s volatility is the standard deviation of its daily yield change over the twenty business days before the auction. The standard deviation of the auction-day relative change is 3.0 basis points, against 5.8 for the unadjusted change; removing the common move is what makes a sub-basis-point concession measurable. One feature of the par curve matters for the auction-day outcome: Treasury builds the curve from the most recently auctioned security at each tenor, so around a new issue’s auction the input at that tenor rolls from the old on-the-run to the new one, and the day-0 change at the tenor embeds the roll. A reopening adds to a security that is already the input, so it carries no roll; Section 6 reports the auction-day results for reopenings beside new issues for that reason.

5 Empirical strategy

Four parameters are estimated: the classes’ marginal shares of a dollar of new supply at the auction, m_k in (23); the official sector’s share of a dollar of net issuance held after h weeks or months, b_h in (24); the sensitivity of the auction-day concession to the bucket’s official share, γ in (22); and the response of the announced tenor’s yield to a quantity surprise and its interaction with the official share, β_1 and β_2 in (21). The news interaction was the pre-specified headline; it failed its first stage, and the paper reports it in full with the diagnostics that say why. The absorption designs at the weekly and monthly horizons were added after a first round of review, in response to the objection that the auction record cannot distinguish an official sector that does not demand the marginal dollar from one that buys it from dealers a week later; Appendix B summarizes the registration memo and its dated amendments, and the replication package reproduces it.

5.1 *Who underwrites the marginal dollar*

Estimand. The marginal share m_k is the dollars that class k takes of an additional dollar of a coupon issue at the auction, holding the auction’s slot fixed. It is identified for the nominal coupon tenors over 2000 to 2025 from changes in issue size between consecutive auctions of the same slot. The model’s inelastic official sector predicts a marginal share for the official part of the foreign-and-international class near zero, below the class’s average share; so, Proposition 4 says, would an elastic official sector at the concession the data measure, which is why this design tests

assumption (iv) and not assumption (i).

Design. Equation (23) regresses the change in each class's allotment between consecutive auctions in the same slot on the change in the privately placed issue; consecutive pairs more than two hundred days apart are dropped. A levels version, each class's allotment on the private issue with slot and year effects, uses the size differences across auctions within a slot and year rather than the changes between consecutive auctions; the two use different variation, and both are reported. Standard errors are clustered by auction month, and the battery clusters by month and slot together, since consecutive differences within a slot overlap.

Threats. There are three. The class that matters pools foreign official and foreign private buyers; the stock-side composition series says which dominated when, the test is run by period, and a continuous version interacts the issue change with the official share of foreign coupon holdings. Size changes are mostly anticipated, which does not matter for who takes the paper but does mean that the marginal share describes absorption of a planned, not a surprise, increase. And the size changes the design uses are modest: no consecutive pair changes the private issue by more than \$6 billion, and Treasury's 4 off-cycle reopenings of off-the-run 10-year notes in October 2008, \$10 billion each and all settling on one day, enter as separate auctions matched to separate allotment records (Appendix B).

Diagnostics. The indirect-bidder share from the auction results, a separate record of who bid, gives the same marginal-versus-average comparison for a class that overlaps with but is not the allotment class; the marginal shares sum to one across classes by construction, which is a check on the parse; the test that a class's marginal share equals its average share is run with a bootstrap over pairs, so that the average's own sampling variance enters; and a dynamic version adds the previous size change in the slot, to ask whether the class's allotment adjusts at the next auction.

5.2 *Who held the path afterward*

Estimand. The parameter b_h in (24) is the dollars by which official holdings have risen, h weeks or months after a dollar of net coupon issuance, beyond the pace a_h at which they rise regardless, holding the period fixed; the period's ratio of official accumulation to net issuance is the object the annual accounting reports. Both are identified from the monthly co-movement of official holdings with the coupon stock from 2001 to 2025 and from the weekly co-movement of

custody holdings with net coupon settlements from July 2007 to December 2025. Under the model's quantity assumption b_h describes the timing of official purchases relative to issuance; a price-elastic official sector would leave a trace in the yield coefficient, not in b_h .

Design. Equation (24) is a local projection of the cumulative change in official holdings on the current change in the stock, by period, with Newey-West standard errors at each horizon and, for the monthly series, the cumulative change in the ten-year yield as a valuation control; the Federal Reserve's cumulative change is run on the same regressor. The weekly version regresses the cumulative change in custody holdings on the week's net coupon settlements (new issues less maturities), with the week's bill settlements as a second regressor, indicators for the week of the month, and a linear trend within the window; the first window is split at September 2008. The annual version divides the June-to-June change in survey-benchmarked official holdings by the change in the coupon stock, and reports the Federal Reserve's and the residual private shares beside it.

Threats. Reserve accumulation and net issuance can rise together for a common reason, a recession that widens the deficit while a surplus country accumulates dollars, in which case b_h measures coincidence rather than sourcing; the weekly design, in which the timing of settlements is set by the auction calendar, is less exposed to this than the monthly one, and the within-window trend and the week-of-month pattern are removed from it. The custody series includes bills, so a flight into bills, as in the last quarter of 2008, would load on coupon settlements unless bill settlements are controlled and the quarter is shown separately. Coupon maturities settle on the same days as new issues, so a gross-settlement regressor confounds rollovers with net supply. The monthly holdings series is benchmarked annually, custodial bias moves official holdings between countries and between the official and private lines between benchmarks (Bertaut and Judson, 2014), and the custody series covers the official holdings at the Federal Reserve Bank of New York only, so the weekly coefficient is a lower bound on official purchases.

Diagnostics. The weekly and monthly estimates are reported side by side with the annual shares; agreement on the ordering of the periods across three records with different construction is the check. The weekly design reports the gross-settlement coefficient beside the net one, the bill coefficient, and a placebo that regresses the custody change over the eight weeks before a settlement on that settlement. The horizon profiles are drawn as a figure.

5.3 What absorption costs at auction

Estimand. The parameter γ in (22) is the change in the auction-day concession, in basis points, per unit change in the bucket's official share, holding the tenor and the month fixed. Proposition 4 predicts $\gamma > 0$ and, in the scaled form, that the concession per percent of stock rises with $1/(1 - \theta_b)$.

Design. Tenor effects absorb duration and whatever makes a tenor's auction costly on average; month effects absorb the common decline in the auction cycle that the literature documents; the tenor's recent volatility is one of the two determinants that Lou et al. (2013) identify. What remains is the covariation, within a tenor over time and across tenors within a month, between the official share and the concession. Standard errors are clustered by month and, because the official share is a slow-moving series that varies at the tenor-by-survey-year level, by year, which the paper treats as the appropriate level, with a wild cluster bootstrap by year.

Threats. The official share fell fastest at the short end, and the short end's auction cycle vanished fastest; a differential trend would produce $\gamma > 0$ without any composition effect. The cycle is larger when yields are volatile, and volatility and the official share both peaked around 2008 to 2012. The Federal Reserve's holdings, the other inelastic position, moved sharply and bucket-specifically with its purchase programs, so the official share alone misses part of the inelastic share; and the Federal Reserve chose which buckets to buy, so its share is not assigned at random across them. The par curve's input rolls at a new issue's auction, which contaminates the day-0 change for new issues.

Diagnostics. Tenor-specific linear trends, reported in the main table; the estimates within each of four sub-periods; the exact monthly SOMA share, alone and together with the official share, by period; the volatility control and a split by volatility; the same regression on a pseudo-auction day ten business days before each auction, which must give zero; the event path by official-share tercile, for all auctions and for reopenings alone; the auction's own high-minus-median yield, a security-specific measure of the auction's tail, as an alternative outcome; a permutation of the official share across tenors within each month; and leave-one-year-out and leave-one-tenor-out estimates. The design's power against the model-implied magnitude is computed and reported with the results.

5.4 What supply costs when it is news

Estimand. In (21), β_1 is the yield response of the announced tenor to a surprise of one percent of the bucket's stock, identified within the announcement day across the buckets announced that day; β_2 is the Triffin margin of Proposition 1, attenuated as Proposition 3 says.

Design. Announcement-day fixed effects remove whatever moved the curve that day; the comparison is across the tenors announced together, which differ in their surprises and in their official shares. Because the same triplets of tenors are announced together, the comparison within a day is always between the same tenors, and the official share's cross-sectional variation is largely a tenor characteristic; the diagnostics include the surprise interacted with the log of the tenor. The sample is every announced nominal tenor. Standard errors are clustered by announcement day.

Threats. The surprise may not be news: Treasury's regular and predictable issuance policy, its forward guidance since 2010, and its explicit size tables since 2020 are designed to remove it. Other tenor-specific news on the day, a Federal Open Market Committee (FOMC) statement or an employment report, could coincide with a surprise. Treasury could choose which tenors to raise with an eye to demand. And without day effects the surprise is correlated with the business cycle, since coupon sizes rise in recessions, when yields fall.

Diagnostics. The response to the surprise by period and by measure, against two benchmarks, a response of 1 basis point per billion dollars of genuinely unanticipated supply, the magnitude Phillot (2025) finds from daily futures price changes, and the low-frequency supply effect of Laubach (2009) and Engen and Hubbard (2004); the placebo of the day before; a placebo that re-dates each announcement day to a random non-announcement day of the same year; the variance of announced tenors' relative yield changes on announcement days against other days; the path-news measure from the refunding statements; the two windows in which the surprise moves yields (exploratory, chosen on that response), each with the tenor controls; and the pooled regression without day effects, which shows the cyclical correlation.

6 Results

6.1 Who holds the stock

The composition series says three things about where the inelastic money sits. First, the official share of the coupon stock is not a single number: on every survey date it is highest in the

one-to-seven-year buckets and lowest beyond fifteen years, and the gap is wide. In June 2008 foreign official institutions held 0.58 of the one-to-two-year bucket, 0.58 of the four-to-five-year bucket, 0.40 of the nine-to-ten-year bucket, and 0.14 of the twenty-five-to-thirty-year bucket (Table 1). Reserve managers hold short paper, as Tabova and Warnock (2021) document at the security level. Second, the fall in the official share since 2008 is concentrated where the share was high: by June 2025 the one-to-two-year share was 0.16 and the twenty-five-to-thirty-year share 0.04. The differential decline is the identifying variation in the designs that use the official share over time, and it is also the reason those designs must guard against differential trends. Third, the Federal Reserve’s share moved in the opposite direction over much of the period and did so bucket by bucket: it was 13 percent of the coupon stock in June 2008, 21 percent in 2016 after the purchase programs and the maturity extension, 26 percent in 2020, and 18 percent in 2025 after the run-off. The elastic float, the stock less official holdings, grew from \$1,934 billion in 2008 to \$19,309 billion in 2025, a factor of 10.0, while the stock grew by a factor of 6.3.

6.2 *Who underwrites the marginal dollar*

Table 3 reports the marginal shares at the auction. Column (2), the pre-specified first-difference specification, says that of an additional dollar of a coupon issue, dealers and brokers take 0.44 (s.e. 0.09), investment funds 0.35 (0.07), and the foreign-and-international class 0.20 (0.07); the remaining classes take nothing measurable. The classes’ average shares of the private issue are 0.41, 0.41, and 0.17. Dealers take slightly more of the marginal dollar than of the average dollar, funds slightly less, and the foreign class slightly more, and none of the differences is distinguishable from zero once the average share’s own sampling variance is allowed for: a bootstrap over pairs gives $p = 0.69$ for dealers and $p = 0.60$ for the foreign class. Column (3) restricts the sample to pairs in which the size actually changed and gives the same shares; column (4) adds year effects and gives the same shares. Column (1) uses a different source of variation, the differences in issue size across auctions within a slot and year, and divides the marginal dollar differently: funds 0.71 (0.03), dealers 0.15 (0.02), with the foreign class again at 0.14 (0.01). Both specifications record allotments at issue; they differ in which size variation they use, the change from one auction to the next or the level difference within a year, and a reader who wants the response of bidding to a change in the size should take column (2), while one who wants the division of persistently larger issues should take column (1). A dynamic

version (the last rows of Table 9) adds the previous size change in the slot and finds no adjustment at the next auction: the lagged coefficient is 0.03 (s.e. 0.08) for dealers and -0.00 (0.06) for the foreign class.

The period split in Table 4 sets the scope of the finding. In 260 consecutive pairs from 2000 to 2008, when foreign official institutions dominated the foreign-and-international class and held close to half the coupon stock, the class took 0.12 (s.e. 0.07) of a marginal dollar against an average share of 0.20, with a bootstrap p of 0.25 against equality, and dealers took 0.78 (0.11) against 0.67. The point estimates lean the way an inelastic official sector inside the class would make them lean, the class below its average share and dealers above theirs, and the data do not establish the lean. Splitting the window, the 2000 to 2007 marginal share is 0.12 (s.e. 0.07) against an average of 0.20, with bootstrap $p = 0.26$, and the 2008 to 2009 marginal share is -0.13 (0.39) against 0.24, an interval too wide to read: the crisis years changed sizes often and in both directions, and the class's take-up of those changes is not measured with any precision. No pair carries unusual leverage; the largest change in the private issue between consecutive auctions of a slot is \$6 billion.

After 2008 the class's marginal share is 0.24 (0.48) in 2009 to 2013, 0.41 (0.16) in 2014 to 2019, and 0.21 (0.09) in 2020 to 2025, at or above its average share in each period, as foreign private holdings of coupons rose from \$521 billion in 2008 to \$4,102 billion in 2025 and the class changed character. The continuous specification, which interacts the issue change with the official share of foreign coupon holdings, finds no relation between the two: the interaction is -0.00 per unit of that share (s.e. 1.16), which puts the marginal share at 0.20 at the 2004 official share and 0.20 at the 2025 share.

Two further records place the official sector at the auction. The noncompetitive bids the Federal Reserve Bank of New York submits for FIMA accounts, the official sector's own channel, average \$26 million per auction, 0.07 percent of accepted amounts, with a maximum of \$0.35 billion, and their marginal share of accepted amounts is 0.0003 (s.e. 0.0004): a quantity, as the cap makes it. And indirect bidders, who bid through a primary dealer and include the official buyers that do not bid directly, take 0.50 of accepted amounts on average and 0.06 (s.e. 0.03) of a marginal dollar. The marginal dollar at auction is underwritten by the direct competitive bidders, overwhelmingly the dealers, in every period. Figures 4 and 5 show the period pattern and the fit.

Table 3. Who underwrites the marginal dollar of a coupon issue

	Average share	(1) Levels	(2) First differences	(3) Nonzero changes	(4) With year effects
Foreign and international	0.173	0.13*** (0.02)	0.12** (0.05)	0.12** (0.05)	0.11** (0.05)
Dealers and brokers	0.405	0.18*** (0.04)	0.54*** (0.07)	0.54*** (0.07)	0.57*** (0.07)
Investment funds	0.407	0.69*** (0.03)	0.34*** (0.04)	0.33*** (0.05)	0.32*** (0.04)
Other classes	0.015	0.00 (0.00)	0.00 (0.01)	0.01 (0.01)	0.00 (0.01)
Indirect bidders (auction results)	0.496		0.07*** (0.03)		
<i>p</i> : foreign marginal = average			0.32		
Observations		1565	1533	420	1533
Slot and year effects		Yes	No	No	Year

Notes: The dependent variable is the class's allotment in billions of dollars (column (1)) or its change between consecutive auctions of the same slot (columns (2) to (4)); the regressor is the privately placed issue (the total issue less the SOMA add-on) or its change. Column (1) includes slot (tenor by new-issue or reopening) and year effects; column (3) keeps only pairs in which the issue changed; column (4) adds year effects to the first differences. Columns (1) and (2) use different size variation (level differences within a slot and year; changes between consecutive auctions) and need not agree. Standard errors in parentheses, clustered by auction month; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The indirect-bidder row uses the auction results' indirect-bidder accepted amounts and total accepted amounts. Source: Treasury investor-class auction allotments and auction records; authors' calculations.

Table 4. Marginal and average shares at auction, by period

Period	Pairs	Foreign avg.	Marginal	(s.e.)	<i>p</i> : marginal = avg.	Dealers avg.	Marginal	(s.e.)	Funds avg.	Marginal	(s.e.)
2000–2008	258	0.20	0.04	(0.03)	0.00	0.67	0.69	(0.03)	0.10	0.26	(0.04)
2009–2013	353	0.21	0.24	(0.48)	0.94	0.55	0.38	(0.51)	0.23	0.32	(0.24)
2014–2019	426	0.16	0.47	(0.14)	0.02	0.35	0.59	(0.20)	0.48	-0.13	(0.21)
2020–2025	496	0.15	0.21	(0.09)	0.49	0.20	0.28	(0.11)	0.64	0.54	(0.10)
2000–2007	221	0.20	0.12	(0.07)	0.23	0.66	0.75	(0.11)	0.10	0.04	(0.11)
2008–2009	103	0.24	0.00	(0.04)	0.00	0.61	0.69	(0.05)	0.15	0.31	(0.02)
2000–2008, pairs with $ \Delta \text{ issue} \leq \$10 \text{ billion}$	256	0.20	0.13	(0.07)	0.29	0.67	0.74	(0.10)	0.09	0.05	(0.09)

Notes: First-difference regressions as in Table 3, column (2), by period; the *p*-value tests the equality of the foreign class's marginal share and its average share treating the average as known (the bootstrap version is in the text). The last two rows split the first period. Standard errors clustered by month. Source: Treasury investor-class auction allotments and auction records; authors' calculations.




Figure 4. Marginal and average shares of the private issue, by period

Notes: For each class and period, the mean share of the privately placed issue and the first-difference marginal share with its 95 percent confidence interval. Source: Treasury investor-class auction allotments; authors' calculations.




Figure 5. The change in allotments against the change in the issue

Notes: Consecutive pairs of auctions in the same slot with a nonzero change in the privately placed issue; the solid line is the first-difference marginal share and the dashed line the class's average share. Source: Treasury investor-class auction allotments; authors' calculations.

6.3 *Who held the path afterward*

Table 5 and Figure 6 report who held the path of issuance after the auction, from three records. The annual accounting is the plainest. Between the June 2004 and June 2008 surveys, official holdings of coupons rose by \$775 billion while the coupon stock rose by \$813 billion, a ratio of 0.95; the Federal Reserve's coupon holdings rose by 0.03 of net issuance and every other holder's by 0.02. From 2009 to 2013, when the stock grew by \$6,139 billion, the ratios were 0.32 official, 0.24 Federal Reserve, and 0.44 residual private; from 2014 to 2019, 0.05, 0.05, and 0.90; from 2020 to 2025, when official holdings fell while the stock grew by \$9,291 billion, -0.03 , 0.21, and 0.83. Over the whole period the official sector's holdings rose by 0.13 of net coupon issuance, the Federal Reserve's by 0.18, and the residual private sector's by 0.69. The year-to-year ratios are dispersed (the official ratio ranges from -0.31 to 1.74 across the four years to 2008), which is what two quantities set independently would produce.

The monthly record times the official purchases. Over 2000 to 2008 a dollar of net coupon issuance in a month was followed by a rise in official holdings of notes and bonds of 0.16 (s.e. 0.13) in the same month and 0.62 (0.27) within six months, on top of a pace of accumulation unrelated to the month's issuance of \$88 billion per six months; over 2009 to 2013 the response is 0.33 (0.14) in the month and 0.40 (0.13) within three; over 2014 to 2019 and 2020 to 2025 it is -0.13 (0.11) and -0.02 (0.06) at six months, while the Federal Reserve's holdings rose by 0.63 (0.51) and 0.45 (0.30) per dollar at the same horizon. The valuation coefficient on the yield change is imprecise in every period and carries no sign the paper reads.

The weekly custody record is the one whose timing the business cycle does not set, and it is also the one that needs the most care, because the series includes bills and the window in which officials were buying ends in the crisis. With bill settlements, the week-of-month pattern, and a within-window trend controlled, custody holdings for foreign official accounts rose by 0.59 (s.e. 0.17) per dollar of net coupon settlements within eight weeks from July 2007 to August 2008, and by 0.26 (0.13) per dollar of bill settlements. The coefficient on gross coupon settlements is -0.10 (0.13), because maturing coupons settle on the same days as new issues and a rollover is not net supply; the regression without the controls, which an earlier version of this paper reported, gives 0.47 over 2007 to 2008 and is the trend and the bills rather than the coupons. In the last quarter of 2008 the net coefficient is -0.03 (0.04), in 2009 to 2013 it is -0.01 (0.09) at eight weeks, and in the two later windows -0.28 and -0.02 ; the placebo, the custody change

Table 5. Who held the path: official, Federal Reserve, and private absorption of net coupon issuance

Period	Annual: official	SOMA	Residual private	Monthly $h = 0$: b_0	(s.e.)	a_0	SOMA b_0	$h = 6$: b_6	(s.e.)	a_6	SOMA b_6
2000–2008	0.95	0.03	0.02	0.16	(0.13)	10.5	0.02	0.62**	(0.27)	88.0	0.26
2009–2013	0.32	0.24	0.44	0.33**	(0.14)	-1.8	0.01	0.32***	(0.11)	188.1	0.02
2014–2019	0.05	0.05	0.90	-0.04	(0.03)	2.8	-0.01	-0.13	(0.11)	3.5	0.63
2020–2025	-0.03	0.21	0.83	-0.02	(0.02)	1.7	0.07	-0.02	(0.06)	12.7	0.45
All	0.13	0.18	0.69	0.01	(0.03)	9.2	0.07	0.05	(0.11)	67.1	0.50

Notes: Annual: the change in survey-benchmarked foreign official holdings of coupons, in the Federal Reserve’s coupon holdings (SOMA), and in every other holder’s (the residual), each divided by the change in the coupon stock, summed within the period; the periods run from the June survey before the first year named to the June survey of the last (2004 to 2008, 2008 to 2013, 2013 to 2019, 2019 to 2025). Monthly: b_h and a_h in equation (24), the cumulative change in foreign official holdings of notes and bonds (TIC) over h months on the month’s change in the marketable coupon stock (MSPD), with the cumulative change in the ten-year yield as a control, and the SOMA’s cumulative change on the same regressor; Newey-West standard errors with $h + 1$ lags; the SOMA sample begins in 2003. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: TIC, Fiscal Data, Federal Reserve Board, Federal Reserve Bank of New York; authors’ calculations.

over the eight weeks before a settlement, is -0.03 in the first window. Appendix Table C2 reports every window and variant. The custody series covers only the official holdings at the Federal Reserve Bank of New York, so the weekly coefficients are lower bounds on official purchases.

The three records agree on the ordering of the periods, not on a single number, and together they say what the allotments could not. In the years of reserve accumulation the official sector held most of the path of coupon issuance within weeks or months of the auction, bought from the dealers who had underwritten it; from 2009 to 2013 it held a third while the Federal Reserve took a quarter; since 2014 it has held none, and the path has gone to the Federal Reserve in its purchase years and to private investors otherwise. Under the model’s own assumption these ratios are the coincidence of two quantities, reserve accumulation and the deficit, and the paper does not call the official sector the marginal holder of a dollar it would not have bought had Treasury issued less; what the records establish is the share of the anticipated path the private sector has had to hold beyond the auction week, which rose from 0.02 before 2009 to 0.90 and 0.83 since 2014. That is the margin in flows the paper’s title refers to, and it is not the stock-share story: the official share of the stock fell steadily from 2008, but the official sector’s share of the flow fell from nearly all to none between the first and the second halves of the sample.

6.4 What absorption costs at auction

The auction-day concession is visible in the raw event paths. Figure 7 plots the auctioned tenor’s daily yield change relative to the rest of the curve from five business days before to five after the auction, by tercile of the bucket’s official share. In the high-share tercile (mean share




Figure 6. Official absorption of a dollar of net issuance, by horizon and period

Notes: Left: the coefficient of the cumulative change in foreign official holdings of notes and bonds on the month's change in the coupon stock, at horizons of zero to six months, by period, with 95 percent bands. Right: the coefficient of the cumulative change in custody holdings for foreign official accounts on the week's net coupon settlements, with bill settlements, week-of-month indicators, and a within-window trend controlled, at horizons of zero to eight weeks, by window (the first window is July 2007 to August 2008). Source: as Table 5 and Appendix Table C2; authors' calculations.

0.47) the relative yield rises by 0.54 basis points on the auction day ($t = 3.8$) and gives it back over the following two days; the five-day run-up is 1.20 basis points and the five-day reversal -1.15 . In the low-share tercile (mean share 0.15) the auctioned tenor's yield falls relative to the curve on the auction day, by 0.72 basis points ($t = -4.9$), a fall no model in this paper produces.

The fall is the roll. Among new issues, where the par curve's input switches to the new security, the low-tercile auction-day change is -1.39 basis points ($t = -7.3$); among the 358 reopenings, where there is no switch, the auction-day change is positive in every tercile, 0.25 (s.e. 0.22) in the low tercile, 0.59 (0.14) in the middle, and 1.20 (0.39) in the high tercile. The auction cycle of Lou et al. (2013) and Fleming and Liu (2017) is present in the reopenings in every tercile and largest where official holders own most of the bucket, and smaller than the cycles those papers report for samples ending in 2008 and 2012, as Somogyi et al. (2025) find for the recent period.

Table 6 reports the official-share coefficient under successively tighter controls. Column (1), with no controls, gives 3.2 basis points per unit of official share (s.e. 0.7); column (2) with tenor effects 8.5 (1.1); column (3) with tenor and year effects 5.7 (1.8); column (4), the pre-specified specification with tenor and month effects and the volatility control, 4.7 (1.8). At




Figure 7. The auction cycle by official-share tercile

Notes: Mean one-day change in the auctioned tenor's yield relative to the mean change across the other tenors, at each business day from five before to five after the auction, by tercile of the bucket's foreign official share, with 95 percent bands; all auctions. Source: Treasury par yields and auction records; the composition series; authors' calculations.

that estimate the fall in the short-end official share from about one half to about one fifth lowers the concession by 0.47 basis points per ten percentage points. The precision depends on the clustering: by month the standard error is 1.8, and by year, the level at which a tenor-by-survey-year regressor's errors are correlated and the level the paper treats as right, 2.8, with a wild cluster bootstrap p of 0.13; the coefficient is not distinguishable from zero at conventional levels under year clustering. The auction's size relative to the stock does not enter (column (5)), which the scaled form of Proposition 4 would require and which Section 7 returns to.

Column (6) adds a linear trend for each tenor and the coefficient falls to 0.73 (1.90), a standard error wide enough to contain the column (4) estimate: the official share and the tenor-specific trends are not separable, since the share fell fastest in the buckets where the auction cycle vanished fastest. The sub-periods say which of the two the pooled coefficient is. In the four windows of Appendix Table C3 the official-share coefficient is negative in three and indistinguishable from zero in the fourth: -10.1 (s.e. 8.2) in 2000 to 2008, -9.5 (5.2) in 2009 to 2013, 1.7 (2.2) in 2014 to 2019, and -11.8 (4.3) in 2020 to 2025; the positive pooled coefficient is a comparison across periods, between the years when both the official share and the cycle were

high and the years when both were low, and within periods the concession is not larger where official holders own more of the bucket. Columns (7) and (8) use the Federal Reserve's exact monthly share of each bucket, which moves bucket-specifically with the purchase programs, the maturity extension, and the run-offs, and which the model treats as the same kind of inelastic position: -0.60 basis points per unit of share (s.e. 1.13) alone and 0.01 (1.64) with the official share and the trends. By period the SOMA coefficient is 0.9 (3.6) in 2009 to 2013, the years of the purchase programs, and 9.6 (2.6) in 2020 to 2025, the years of the pandemic purchases and the second run-off; the first is a null, the second is not, and the two inelastic shares carry opposite signs in the last window. The Federal Reserve chose which buckets to buy, so its share is not assigned at random across them, and the paper reads the SOMA results as a description rather than as a test.

Whether the nulls are informative depends on what the model implies the effect should be, and here the answer is that the design cannot see it. The mean concession before 2014, when the cycle was present, is 0.40 basis points; at the sample's mean inelastic share of 0.49 , Proposition 4 implies a derivative of the concession with respect to the inelastic share of 0.78 basis points per unit of share. The minimum detectable effect in the SOMA specification is 3.2 basis points per unit of share, about four times the model's value. The concession is a dealer-inventory premium of a fraction of a basis point, and its composition sensitivity, if it is what the model says, sits below what daily par yields can resolve; the official-share coefficient in columns (1) to (4) is several times the model's value, which is a further reason to read it as a comparison across periods rather than as the Triffin margin.

Appendix Table C7 repeats the comparison for the other windows. The two-day reversal falls with the official share (-5.1 basis points per unit, s.e. 2.3) and the two-day amplitude rises with it (8.6 , s.e. 3.0); the five-day windows are noisier. Appendix Figure C1 plots the official-share coefficient at each event day: it is positive on the auction day and negative two days after, and near zero elsewhere, which is the shape of a concession that reverses. Among reopenings alone, which carry no roll, the official-share coefficients are negative and imprecise: -7.2 (s.e. 6.9) on the auction day and -30.0 (11.8) on the five-day run-up, on 358 auctions.

Appendix Table C8 reports what the auction's own allotment says about its concession. The concession is larger when dealers took more of the issue (6.4 basis points per unit of dealer share, s.e. 1.2) and smaller when the foreign class or the indirect bidders took more. These

Table 6. The auction-day concession and the official share

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Official share θ_b	3.13*** (0.68)	8.38*** (1.08)	5.55*** (1.78)	4.33** (1.84)	4.35** (1.84)	0.01 (1.90)		-0.12 (2.23)
SOMA share θ_b^F							-0.24 (1.13)	0.17 (1.65)
Auction size (% of stock)					-0.08 (0.06)			
Volatility (vol20)		0.01 (0.04)	0.13** (0.05)	0.03 (0.08)	0.03 (0.08)	0.02 (0.08)	-0.05 (0.09)	0.00 (0.08)
Observations	1567	1567	1567	1541	1541	1541	1490	1490
Fixed effects	None	Tenor	Tenor, year	Tenor, month	Tenor, month	Tenor, month	Tenor, month	Tenor, month
Tenor-specific trends	No	No	No	No	No	Yes	No	Yes
Within R^2		0.072	0.014	0.007	0.009	0.033	0.001	0.032

Notes: The dependent variable is the auctioned tenor's yield change on the auction day relative to the mean change across the other tenors, in basis points. θ_b is the foreign official share of the bucket's stock at the end of the previous month; θ_b^F is the Federal Reserve's exact monthly share; vol20 is the standard deviation of the tenor's daily yield change over the twenty prior business days; the tenor-specific trends are linear in time. Standard errors in parentheses, clustered by month; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: Treasury par yields and auction records; the composition series; SOMA holdings; authors' calculations.

regressions describe the auction's outcome rather than identify a parameter: a weak auction leaves the dealers with more of the issue and prints a higher yield, so the dealers' share and the concession are determined together. We report them because they are the flow-side counterpart of the stock-side result in the tercile paths.

6.5 What supply costs when it is news

The pre-specified design asks what the announced tenor's yield does on the day Treasury announces a larger or smaller auction than the market expected, and whether the response is larger in buckets with more official holders. Table 7 and Figure 8 give the response of the announced tenor's yield to the surprise, the paper's measure of whether the announced quantity carried news. Pooled over 655 announcement days and 1,564 announced tenors, a naive surprise of one percent of the bucket's stock moves the tenor's yield by 0.36 basis points relative to the other tenors announced that day (s.e. 0.23); the guidance-netted measure gives 0.33 (0.28). Two benchmarks place those numbers. A response of one basis point per billion dollars of genuinely unanticipated supply, the lower end of what Phillot (2025) finds from daily futures price changes, would be 10.5 basis points per percent of stock at the announcement sample's mean bucket stock, and the measured response is 0.03 of it, with a 95 percent interval reaching 0.08. The

Table 7. The response to the announced surprise, by period and measure

Period	Nonzero	Share	β_1	(s.e.)	F	Nonzero	Share	β_1	(s.e.)	F	
		Naive surprise						Guidance-netted surprise			
2000–2008	94	0.34	-0.02	(0.42)	0.0	94	0.34	-0.02	(0.42)	0.0	
2009–2013	66	0.19	5.10	(1.85)	7.6	66	0.19	5.10	(1.85)	7.6	
2014–2017	18	0.06	2.70	(5.95)	0.2	18	0.06	2.70	(5.95)	0.2	
2018–2020	108	0.49	-0.00	(0.28)	0.0	59	0.27	0.10	(0.30)	0.1	
2021–2025	125	0.30	3.16	(1.23)	6.6	46	0.11	-0.90	(0.69)	1.7	
2000–2025	411	0.27	0.19	(0.30)	0.4	283	0.18	0.13	(0.30)	0.2	

Notes: Each cell reports, for the period and the surprise measure, the number and share of announced auction-tenor observations with a nonzero surprise, and the coefficient of the announced tenor's one-day yield change on the surprise in percent of the bucket's stock with announcement-day effects, its standard error clustered by day, and the squared t -statistic (F). The two measures coincide before 2018 because the refunding statements of those years implied no size for a given month; the guidance-netted measure differs from the naive one only where a statement carried a path or a table. Source: Treasury auction records, refunding statements, and par yields; the composition series; authors' calculations.

low-frequency supply effect of Laubach (2009) and Engen and Hubbard (2004), three to four basis points per percentage point of debt to GDP, is of the same order as the measured response once a percent of a bucket's stock is translated into a percent of the debt; the announcements carry the permanent supply effect of their quantity and almost none of the news premium the futures windows see. The within-day residual standard deviation of the yield change is 1.6 basis points, so the design had the precision to see a Phyllot-sized response had the surprises carried the news (the F in the table is the squared t -statistic of the response, not a first stage in the instrumental-variables sense); the minimum detectable interaction, 8.3 basis points per unit of share, is below the 15.3 the segmented model implies at that response, which Proposition 3 says is an upper bound for the within-day design. The surprises are not news.

The period pattern says when the announced quantities did carry news. The F -statistic is 7.7 in 2009 to 2013 and 6.5 in 2021 to 2025 under the naive measure, with responses of 5.1 and 3.1 basis points per percent of stock, and below 2 in every other period. The 2021 to 2025 response disappears under the guidance-netted measure (0.7, s.e. 1.9): the increases of 2023 and 2024 were announced as a monthly path in the refunding statements and the market had them before the monthly announcement. In 2018 to 2020, when Treasury stated the path in dollars per month, 46 percent of announcements changed the size and none of the changes moved yields.




Figure 8. The first stage against a benchmark of one basis point per billion dollars

Notes: The coefficient of the announced tenor's one-day yield change on the surprise in percent of the bucket's stock, with announcement-day effects, by period and measure, with 95 percent intervals clustered by day; the dashed line is the response 1 basis point per billion dollars of unanticipated supply would imply at the sample's mean bucket stock. Source: as Table 7; authors' calculations.

Only in 2009 to 2013, when Treasury raised sizes faster than its guidance described, do announced quantities carry a response of the right order that the guidance does not remove, and even then the sample has 66 nonzero surprises.

Table 8 reports the interaction and the diagnostics, with the two windows that have a first stage in its last rows. In the pre-specified specification the interaction is 0.2 basis points per unit of share (s.e. 3.0): a null with an interval that includes zero and the model's value alike, which the registration memo said in advance would be uninformative; adding the surprise times the log of the tenor leaves it at 3.3 (5.8). The placebo of the day before gives 0.03 (0.28) for the surprise. The path-news measure built from the refunding statements, which treats each statement as announcing the quarter's issuance path, has no first stage either (0.03, s.e. 0.06, over 104 statements); Appendix Table C5 lists the largest path-news days and shows why: on November 1, 2023, the statement raised the two- and five-year paths by the most in the sample and both yields fell by double digits, because the market had expected larger long-end increases than Treasury announced. The sign of the news on the days that matter is set by an expectation the public record does not contain.

In the two windows in which the surprise moved yields the interaction can be estimated;

both windows are exploratory, chosen on that response, and in both the interaction is negative. On the 357 announced tenor-days of 2009 to 2013 a naive surprise of one percent of the bucket's stock moved the announced tenor's yield by 6.5 basis points relative to the other tenors announced that day (s.e. 1.8), and the interaction with the official share is -19.2 basis points per unit of share (s.e. 8.7), where the segmented model's upper bound at the window's mean share of 0.40 is a positive 10.8. On the 360 tenor-days of 2021 to 2025 the naive surprise moved the yield by 15.3 (5.9), more than a Phillot-sized supply response, so the days on which Treasury departed from its stated path carried news of something other than the quantity, with an interaction of -110 (49), and the guidance-netted surprise moved it by 3.14 (5.55) with an interaction of -29.7 (44.3).

Taken at face value, the response to news was smaller, not larger, in the buckets reserve managers held. Two features of the windows argue against the face value. Within a five-year window the official share is a tenor characteristic: its correlation with the log of the tenor across the 2009 to 2013 observations is -0.91 , and its standard deviation within a tenor is 0.04. And the years 2009 to 2013 were the years in which the federal funds rate sat at its lower bound and the Federal Reserve tied its path to the calendar, so the two- and three-year yields, where the official share was highest, moved less on any news than the ten- and thirty-year did. Adding the surprise times the log of the tenor leaves the 2009 to 2013 interaction at -11.3 with a standard error of 22.9, and letting the surprise carry its own coefficient for each tenor leaves -41.1 (59.9). The windows in which the surprise moves yields cannot separate the official share from the tenor, and the window that could, the full sample, has no response to separate. One more comparison bears on the model: the response per percent of stock rose from 6.5 in 2009 to 2013, at a mean official share of 0.40, to 15.3 in 2021 to 2025, at 0.18, where Proposition 1 with a constant semi-elasticity predicts a fall by the ratio of the floats; the private sector's semi-elasticity did not stay constant, which is what Somogyi et al. (2025) find at auction, and the paper's calibration carries that caveat.

Two further checks place the finding. Without day effects the surprise's coefficient is -1.38 (s.e. 0.63), negative and precisely estimated: Treasury raises coupon sizes when deficits widen, in recessions, when yields are falling, and the pooled regression picks up that calendar correlation rather than a supply effect, which is why the design needs day effects and why the day-effect estimate, not the pooled one, is the paper's. And the variance of announced tenors'

Table 8. The news test: the interaction and its diagnostics

Specification	β_1	(s.e.)	β_2	(s.e.)	N
Surprise only	0.19	(0.30)			1377
Surprise and surprise $\times (\theta_b - \bar{\theta})$	0.24	(0.24)	-0.92	(3.04)	1377
Guidance-netted surprise	0.12	(0.25)	0.17	(3.00)	1377
Surprise in percent of elastic float	0.14	(0.16)	-0.31	(1.55)	1377
Two-day window	0.39	(0.36)	-0.38	(3.63)	1377
Placebo: the day before	-0.08	(0.21)	-5.11	(3.91)	1377
No day effects	-1.32***	(0.50)	10.18*	(6.02)	1546
Refunding-statement path news	0.12	(0.14)	0.72	(0.82)	558
Surprise $\times$ log tenor added	-0.21	(0.40)	2.31	(4.45)	1377
2009–2013 only	6.47***	(1.83)	-19.35**	(8.67)	355
2009–2013, surprise $\times$ log tenor added	5.38	(3.27)	-10.62	(22.94)	355
2009–2013, surprise $\times$ tenor effects			-42.44	(60.17)	355
2021–2025 only, naive surprise	15.47**	(5.91)	-110.44**	(48.38)	358
2021–2025 only, guidance-netted surprise	-0.69	(1.10)	-5.93	(13.35)	358

Notes: Each row is equation (21) under the named variant; β_1 is the coefficient on the surprise in percent of the bucket's stock and β_2 on its interaction with the bucket's official share in deviation from the sample mean, with announcement-day effects except where noted. The rows below the rule are exploratory: they restrict the sample to the window named, chosen because the surprise moves yields there, and add, where stated, the surprise interacted with the log of the tenor or with tenor indicators (in which case β_1 is not a single coefficient and is left blank). Standard errors clustered by announcement day. Source: as Table 7; authors' calculations.

relative yield changes on announcement days is 1.10 times its value on other days on average (Appendix Table C4), with one tenor significantly above one and one significantly below; on refunding-statement days the ratio is 1.26. Announcement days carry little tenor-specific news in the daily par curve. Over the full sample the interaction cannot be estimated from these records, and in the windows where it can the official share is not separable from the tenor; what the records establish is that the quantity Treasury announces is, outside those windows, not where the news is.

7 Robustness

The paper rests on the absorption facts at the two horizons and on the reading of the concession's relation to the official share as a comparison across periods; the battery is organized around each.

7.1 *The marginal shares at auction*

Table 9 re-estimates the foreign class's marginal share across the pre-specified variants and the diagnostics added after the first round of review. Restricting to pairs in which the size changed, adding year effects, winsorizing the allotment changes at the first and ninety-ninth percentiles, dropping 2020, dropping the 20- and 30-year, using new issues or reopenings only, clustering by year, and clustering by month and slot together leave the marginal share within a few hundredths of the baseline; dropping any one year leaves it between 0.18 and 0.21 (Appendix Figure C3, right panel). The one cut that moves it is the sample period: 0.15 (s.e. 0.12) in 2000 to 2013 against 0.23 (s.e. 0.09) in 2014 to 2025, a difference of the sign the changing composition of the class predicts and not a precise one. The class pools two kinds of money, and the sample average of its marginal share is an average over a change in which kind dominated; the levels and first-difference specifications divide the marginal dollar between dealers and funds differently, for the reason given in Section 6, and neither moves the foreign class's share.

7.2 *The absorption at the weekly and monthly horizons*

The three records are the check on one another: the annual shares come from the survey-benchmarked composition series, the monthly coefficients from the TIC holdings with a valuation control, and the weekly coefficients from the custody series at face value with bill settlements, the week-of-month pattern, and a trend removed, and the three agree on the ordering of the periods and on the absence of official absorption after 2013. The monthly coefficients are the least precise, because net issuance and reserve accumulation share a business-cycle component; the weekly coefficients are sensitive to the regressor, net or gross, and to the controls, as Appendix Table C2 shows, and the paper leans on them only for the timing within the first window. The horizon profiles in Figure 6 rise over six months in 2000 to 2008 and over eight weeks from July 2007 to August 2008, and are flat at zero after 2013.

Table 9. Robustness of the foreign class's marginal share at auction

Specification	Foreign marginal share	(s.e.)	N
Baseline (first differences, all pairs)	0.120	(0.054)	1533
Nonzero size changes only	0.119	(0.053)	420
Year effects	0.105	(0.053)	1533
Winsorized at 1/99	0.118	(0.050)	1533
Drop 2008-2009	0.221	(0.065)	1430
Drop 2020	0.110	(0.055)	1455
New issues only	0.194	(0.081)	1177
Reopenings only	0.061	(0.037)	356
Drop 20- and 30-year	0.118	(0.058)	1267
Clustered by year	0.120	(0.053)	1533
2000-2013	0.055	(0.045)	611
2014-2025	0.234	(0.083)	922
Drop pairs with the absolute change in the issue over \$10 billion	0.210	(0.070)	1531
2000-2007 only	0.118	(0.069)	221
Clustered by month and slot	0.120	(0.053)	1533

Notes: The first-difference marginal share of the foreign-and-international class under each variant; standard errors clustered by month unless noted. Source: Treasury investor-class auction allotments and auction records; authors' calculations.

7.3 *The concession*

Appendix Table C9 and Appendix Figure C2 report the official-share coefficient in the concession regression across 33 specifications, and Table 10 the ones that matter. Within the family that holds tenor and time effects fixed, the coefficient is stable: between 2.9 and 6.2 basis points per unit of share when any one year is dropped and between 3.8 and 5.4 when any one tenor is dropped (Appendix Figure C3, left panel); it survives dropping 2008 and 2009, dropping 2020, dropping FOMC and employment-report days, winsorizing or trimming the outcome, weighting by auction size, and the survey-month and percent-table-year samples that remove the interpolated official shares. What it does not survive is the comparison within periods, reported in Section 6, or the security-specific outcome: on the auction's own tail, the high yield less the median yield, the official-share coefficient is in the full battery and carries no positive relation either. The placebo on a pseudo-auction day ten business days before each auction gives -0.32 (s.e. 1.50), as it should. The permutation of the official share across tenors within each month

Table 10. Robustness of the concession's relation to the official share

Specification	θ_b	(s.e.)	N
Add tenor-specific linear trends	0.01	(1.90)	1541
Tenor and year effects	5.55	(1.78)	1567
θ includes SOMA (inelastic share)	3.02	(1.40)	1541
Survey months only (April to September)	2.75	(2.42)	769
Drop 2008-2009	4.14	(1.84)	1432
Drop 2020	4.37	(1.88)	1461
Drop 20- and 30-year	4.74	(1.86)	1262
Drop FOMC and first-Friday days	4.84	(1.87)	1505
Outcome winsorized at 1/99	3.93	(1.67)	1541
Weighted by auction size	5.26	(1.82)	1541
Clustered by year	4.33	(2.81)	1541
Outcome: cycle amplitude (run-up minus reversal)	6.15	(5.30)	1541
Placebo: pseudo-auction day (t-10)	-0.11	(1.50)	1541

Notes: The coefficient on the foreign official share in equation (22) under each variant; the baseline has tenor and month effects, the volatility control, and month clusters. The full battery is in Appendix Table C9. Source: as Table 6; authors' calculations.

(Appendix Figure C4) puts the actual coefficient outside the permutation distribution ($p < 0.002$), but that test only confirms that the official share is tied to the tenor, which the within-period estimates already explain.

The scaled form of Proposition 4 fares no better than the level form. The concession does not rise with the auction's size as a percent of the stock (Table 6, column (5)), nor with the size divided by one minus the official share (-0.04 basis points per unit, s.e. 0.04); the nested exponent ψ on $1/(1 - \theta_b)$ has a profile that cannot distinguish the model's value of one from the level alternative's zero (Appendix Figure C5). An auction that is twice as large relative to its bucket's stock does not print twice the concession. Under the paper's own two-stage reading that is what a warehousing premium should look like: the fast sector's capacity to carry an issue for a few days is a balance-sheet quantity that need not scale with the bucket's float, and the auction-day concession identifies the fast sector's semi-elasticity, not the Triffin margin. The power calculation in Section 6 says the daily par curve could not have seen the margin in any case.

7.4 The news test

The battery for the news test is in Table 8: the guidance-netted measure, the float scaling, the two-day window, the day-before placebo, the tenor control, the two windows, and the pooled regression. None changes the conclusion. The re-dating placebo, which assigns each announcement day's surprises to a random non-announcement day of the same year and re-estimates the interaction 300 times, gives a distribution with standard deviation 2.9 basis points per unit of share around zero, and the actual full-sample interaction sits inside it ($p = 0.94$).

8 Mechanisms and heterogeneity

The model predicts where its mechanism should be visible, and the data allow five cuts along those lines; each is reported as the model reads it, and the cuts that were not predicted are labeled as exploration.

Periods. The model's quantity-driven official sector should be most visible when officials were accumulating, and it is: official holdings rose by 0.95 of net issuance from 2004 to 2008 and 0.32 from 2008 to 2013 while the official class took no more than its average share at auction, and have risen by nothing since. The auction cycle was present through 2013 and gone after: the mean concession is -0.25 basis points in 2000 to 2008, 0.90 in 2009 to 2013, 0.44 in 2014 to 2019, and -0.91 in 2020 to 2025 (Appendix Table C3), and the amplitude follows the same path. The model attributes the decline of the cycle to the growth of the elastic float; the within-period estimates say the data cannot separate that from the other changes in the market over the same years, the shrinking of dealer balance sheets relative to the stock (Duffie, 2023), the growth of electronic intermediation, the retreat of foreign demand (Somogyi et al., 2025), and the growth of the basis trade (Vissing-Jorgensen, 2021) among them.

Tenors. Reserve managers hold the short end, so the model puts the Triffin margin there. The by-tenor marginal shares (Appendix Table C1) show the foreign class's marginal share lowest at the 2- and 10-year, where its average share is highest, and highest at the 3- and 7-year; the differences are not precise. The concession's tercile pattern is a tenor pattern as much as a composition pattern, which is why the tenor effects in Table 6 matter and why the within-period estimates matter more.

Surprise versus anticipated size changes. The marginal shares in Table 3 are estimated

mostly from anticipated size changes. Splitting the consecutive pairs by whether the guidance-netted surprise was nonzero (an exploratory cut) gives a foreign marginal share of 0.19 (s.e. 0.08) on the 268 surprise pairs and 0.22 (0.12) on the 151 anticipated pairs; the two cannot be distinguished.

New issues and reopenings. A reopening adds supply to a security that already trades, so the fast sector has an inventory and a price to lean on, and the par curve's input does not roll; a new issue has neither. The foreign class's marginal share is 0.20 (s.e. 0.08) on new issues and 0.21 (0.08) on reopenings, and the concession's official-share coefficient is 2.6 (2.3) and -7.2 (6.9). The tercile paths for reopenings, which carry no roll, are the cleaner description of the cycle and are reported in Section 6.

The Federal Reserve as the inelastic holder. The SOMA share is the paper's only source of high-frequency, bucket-specific variation in an inelastic position, and it is not assigned at random: the Federal Reserve bought the buckets its programs targeted. By period (Appendix Table C3), the SOMA coefficient is 0.9 (s.e. 3.6) in 2009 to 2013, the years of the purchase programs and the maturity extension, when the within-tenor standard deviation of the SOMA share was 0.08, and 9.6 (2.6) in 2020 to 2025. Where the Federal Reserve bought a bucket in the first window, the auctions in that bucket were not visibly costlier to place; in the second window, which combines the pandemic purchases with the run-off that followed, auctions were costlier where the Federal Reserve's share was high. The local-supply literature finds that the Federal Reserve's purchases lowered yields in the buckets it bought (D'Amico and King, 2013; Cahill et al., 2013); the present result is about the concession at auction, not the level of yields, and the paper does not read a coefficient on a position chosen with the bucket's richness in view as a causal effect in either direction.

9 Implications

The question a debt manager and a budget forecaster ask is how much a dollar of new supply costs in yield, and whether the answer has changed as the official share has fallen.

What is measured. The marginal dollar of a coupon auction is underwritten by dealers and funds in every period. Who then holds the path depends on the period: from 2004 to 2008 official holdings rose by nearly all of net issuance, within months and bought through dealers; from 2008 to 2013 by a third, with the Federal Reserve taking a quarter; since 2014 by nothing, so that the

private sector held 0.90 of net issuance from 2013 to 2019 and 0.83 from 2019 to 2025 and the Federal Reserve the rest. The auction-day concession on the tenor being sold is small, -0.01 basis points on average over the sample and 0.40 before 2014, and it reverses within days. A concession of one basis point on \$1 trillion of annual coupon issuance with an average duration of six years is a present-value cost of about \$600 million per year of issuance; at the pre-2014 mean the cost of placing a trillion dollars of the coupon program was on the order of a quarter of that, and the composition sensitivity of that cost is below what the data resolve. The quantity surprises in Treasury's announcements carry about 0.03 of the yield response that genuinely unanticipated supply produces in futures windows, which is a statement about Treasury's communication rather than about demand.

What is modeled. Table 11 evaluates Proposition 1 at the 2008 and 2025 compositions. The ratio of the yield cost of a one-percent surprise increase in a bucket's supply in 2025 to the same cost in 2008 is $(1 - \theta_{2008}) / (1 - \theta_{2025}) = 0.63$ with the Federal Reserve's holdings counted as elastic and 0.61 with them counted as inelastic (the inelastic share was 0.59 in 2008 and 0.33 in 2025), independent of the private sector's semi-elasticity; the ratio of the cost per dollar is $F_{2008} / F_{2025} = 0.10$. Of the 2.30 log-point fall in the cost per dollar, the composition term contributes 0.46 and the growth of the stock 1.84. The word "exit" would misdescribe the first term: official dollar holdings doubled over the period, and the official share fell because the stock grew faster; the counterfactual the composition term answers is one in which official holdings had grown with the stock, in which case the float would be smaller by the same factor. Three things qualify the ratios. They hold under the nested exponent $\psi = 1$; at $\psi = 0$, the level alternative that is the canonical preferred-habitat demand, the ratio per percent of stock is one, at $\psi = 0.5$ it is 0.80, and the daily data do not identify ψ . They assume one semi-elasticity at both dates, while the response to announced sizes rose between 2009 to 2013 and 2021 to 2025 as the official share fell, which points toward a ratio at or above one. And they describe the surprise, not the path. The market-level ratio also averages over a maturity structure: bucket by bucket, the ratio of the 2025 cost per percent of stock to the 2008 cost is 0.50 in the one-to-two-year bucket, 0.51 in the four-to-five-year bucket, 0.79 in the nine-to-ten-year bucket, and 1.00 in the fifteen-to-twenty-year bucket, because reserve managers hold little beyond fifteen years in either year.

The levels require a semi-elasticity, and the literature's values have to be put in one unit

before they can be compared. Eren et al. (2023) estimate a holdings-weighted yield semi-elasticity of 13.9 percent of quantity per percentage point of yield for the non-central-bank sectors taken together; Jansen et al. (2024) report a supply-weighted market elasticity of 3 and Chaudhary et al. (2024) an aggregate elasticity of 1.2, both percent of quantity per percent of market value, which at a 6.5-year duration are 19.5 and 7.8 per percentage point of yield. The low-frequency supply literature implies a larger value: three to four basis points per percentage point of debt to GDP in Laubach (2009) and Engen and Hubbard (2004), which at a debt-to-GDP ratio near one is about 3.5 basis points per percent of the stock and corresponds to a semi-elasticity of 34 at the 2025 composition. The levels are the segmented model's, with no arbitrage across buckets, and are upper bounds on the own-bucket response for that reason. At the first value a one-percent surprise increase in a bucket's supply raised its yield by 13.5 basis points at the 2008 composition and 8.5 at the 2025 composition; at the low-frequency value, 5.5 and 3.5. Per hundred billion dollars the numbers are 37.2 and 3.7 basis points at the first value and 15.3 and 1.5 at the second. The demand-system values describe how a holder substitutes across assets at a given wealth and overstate the market's response to a permanent supply increase, which is why the levels at those values sit above the low-frequency literature; the low-frequency value is the one a budget projection should carry, and the ratios are the same at either.

What the paper's own estimates say about the elasticities. The auction-day concession is not a semi-elasticity. An issue announced a week earlier is priced at its announcement, and the fraction of a basis point the auctioned tenor cheapens on the day is the premium the fast sector charges for warehousing the whole issue, 0.40 basis points for an issue of 4.0 percent of the stock before 2014; read as a semi-elasticity it would be about 1,604, two orders of magnitude above the literature's values, which is only a way of saying that the sector that clears an auction on the day is not the sector whose demand sets the permanent effect. The same arithmetic is why the allotment test has no power against an elastic official sector: at that concession an official sector with the semi-elasticity of Eren et al. (2023)'s private sectors would take 0.008 of a marginal dollar at the auction, and one with the low-frequency semi-elasticity 0.018, neither distinguishable from zero. The auction-level elasticities of Somogyi et al. (2025), two basis points per percent of supply before 2010 and nine after, are read off bid schedules and describe the demand the auction faces rather than the premium the fast sector charges; their direction, a market less elastic as the official share fell, is the direction the cross-window comparison of

Table 11. The model's cost of a supply surprise at the 2008 and 2025 compositions

	2008	2025	2025/2008
Official share θ	0.47	0.16	0.33
Coupon stock Q (\$ billions)	3,622	22,869	6.31
Elastic float $F = (1 - \theta)Q$ (\$ billions)	1,934	19,309	9.98
$1/(1 - \theta)$ (response per percent of stock, relative units)	1.87	1.18	0.63
Inelastic share with SOMA, $\theta + \theta^F$	0.59	0.33	
$1/(1 - \theta - \theta^F)$ (relative units)	2.45	1.50	0.61
Ratio under the nested exponent, $\lambda = 0 / 0.5 / 1$			1.00 / 0.80 / 0.63
bp per 1 percent of stock, Eren, Schrimpf and Xia (2023), $\varepsilon = 13.9$	13.5	8.5	0.63
bp per \$100 billion, Eren, Schrimpf and Xia (2023), $\varepsilon = 13.9$	37.2	3.7	0.10
bp per 1 percent of stock, Laubach (2009), Engen and Hubbard (2004), $\varepsilon = 34$	5.5	3.5	0.63
bp per \$100 billion, Laubach (2009), Engen and Hubbard (2004), $\varepsilon = 34$	15.3	1.5	0.10
bp per 1 percent of stock, Jansen, Li and Schmid (2024), converted, $\varepsilon = 19.5$	9.6	6.1	0.63
bp per \$100 billion, Jansen, Li and Schmid (2024), converted, $\varepsilon = 19.5$	26.5	2.7	0.10
bp per 1 percent of stock, Chaudhary, Fu and Zhou (2024), converted, $\varepsilon = 7.8$	24.0	15.2	0.63
bp per \$100 billion, Chaudhary, Fu and Zhou (2024), converted, $\varepsilon = 7.8$	66.3	6.6	0.10

Notes: Market-level official share, coupon stock, and elastic float at the June 2008 and June 2025 survey dates; the inelastic share with the Federal Reserve's holdings counted; the yield cost of a one-percent increase in a bucket's supply and of \$100 billion of supply implied by Proposition 1 at four values of the private sector's semi-elasticity (percent of quantity per percentage point of yield): the Eren, Schrimpf, and Xia (2023) estimate, the value implied by the low-frequency supply effect of Laubach (2009) and Engen and Hubbard (2004), and the Jansen, Li, and Schmid (2024) and Chaudhary, Fu, and Zhou (2024) price elasticities converted at a 6.5-year duration. The ratios in the last column do not depend on the semi-elasticity; the row for the nested exponent gives the ratio at $\psi = 0, 0.5$, and 1; the dollar rows are the segmented model's, an upper bound on the own-bucket response. Source: the composition series; the cited papers; authors' calculations.

Section 6 points in. The composition adjustment a budget forecaster would apply to a supply-effect rule of thumb, a factor between one and 0.63 on the cost per percent of stock between the 2008 and 2025 compositions, is therefore a model number conditional on the slow sector's demand having constant semi-elasticity, and the paper's evidence on that structure is nil at the daily frequency and leans against it across windows.

What a budget projection can take. A projection that applies a single supply elasticity to all new debt, as the rule-of-thumb literature does (Laubach, 2009; Engen and Hubbard, 2004), holds the marginal absorber fixed. The flow evidence says the holder of the path has changed twice: official institutions held nearly all of it before 2009, the Federal Reserve and private investors shared it from 2009 to 2021, and private investors have held most of it since. A rule of thumb estimated on a history in which the official sector absorbed the path says little about what

the private sector must be paid to hold it now, and the paper cannot put a number on the difference without an estimate of the slow sector's semi-elasticity that the daily data do not provide. The composition of the stock enters the cost of a surprise through the float, which is 10.0 times its 2008 size, so that under proportional private demand the market has become easier to surprise per dollar and under linear demand no easier; the pace of official accumulation enters the cost of the path, and it has gone from matching issuance to zero. The claim that the market has become more price-sensitive (Dilts Stedman, 2025) is right about the path and unproven about the surprise: the private sector now holds the path it did not hold before, and whether a given surprise costs more or less than it did depends on a functional form the data do not identify.

10 Conclusion

Treasury supply is divided at the auction among investor classes roughly as the stock is divided among them, and the path of it was held by official ones for the first half of the sample. In the investor-class allotments of 1,567 coupon auctions, an additional dollar of issuance went 0.44 to dealers, 0.35 to investment funds, and 0.20 to the foreign-and-international class at the auction, each near the class's average share of the private issue (the foreign class's is 0.17), with the official sector's own noncompetitive channel at 0.07 percent of accepted amounts. In the holdings records, official holdings of coupons rose by 0.95 of net issuance between the 2004 and 2008 surveys and by 0.62 of each dollar within six months in the monthly record for 2000 to 2008, by 0.32 between 2008 and 2013, and by nothing since 2014, when the Federal Reserve and private investors held the path instead. The official sector's demand was a quantity set by reserve accumulation and bought mostly in the secondary market; while it was growing, the private sector held a small fraction of net issuance beyond the auction week, and since it stopped the private sector has held most of it.

What that absorption costs is less well measured. The auction-day concession is larger where official holders own more of the bucket across the sample, by 4.7 basis points per unit of official share with a year-clustered standard error of 2.8, but within the sub-periods the relation is negative or absent, the official share is not separable from tenor-specific trends, and the concession is a dealer-inventory premium of a fraction of a basis point whose composition sensitivity, if it is what the model says, sits below the design's power. The quantity surprises in Treasury's announcements carry almost none of the news that moves yields; the daily par curve

cannot recover the Triffin margin from them over the full sample, and in the two windows in which the surprises do move yields the official share cannot be separated from the tenor. Under the model, the official share's fall from 47 percent to 16 percent of the coupon stock lowered the yield cost of a one-percent supply surprise by about a third under proportional private demand and not at all under linear demand, and the cost per dollar by about nine tenths under the first, most of it from the growth of the stock; those are calibrations, and the one comparison in the data that bears on them points toward the smaller effect.

Three things would decide what the daily data cannot. Intraday futures around the announcements, as in Bi et al. (2026), would give the news; the surveys of dealers' expected auction sizes that Treasury collects before each refunding would give the expectation against which the news is measured; and a split of the foreign-and-international allotment class into its official and private parts, which Treasury holds but does not publish, would let the marginal share of the official sector be followed past 2008, when the class stopped being the reserve managers' and the paper's flow test at the auction stops being sharp.

References

- Ahmed, Rashad and Alessandro Rebucci**, “Dollar Reserves and U.S. Yields: Identifying the Price Impact of Official Flows,” NBER Working Paper 30476, National Bureau of Economic Research 2022.
- Beltran, Daniel O., Maxwell Kretchmer, Jaime Marquez, and Charles P. Thomas**, “Foreign Holdings of U.S. Treasuries and U.S. Treasury Yields,” *Journal of International Money and Finance*, 2013, 32, 1120–1143.
- Bernanke, Ben S., Vincent R. Reinhart, and Brian P. Sack**, “Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment,” *Brookings Papers on Economic Activity*, 2004, 2004 (2), 1–100. With Comments and Discussion.
- Bertaut, Carol and Ruth Judson**, “Estimating U.S. Cross-Border Securities Positions: New Data and New Methods,” International Finance Discussion Paper 1113, Board of Governors of the Federal Reserve System August 2014.
- Bi, Huixin, Maxime Phillot, and Sarah Zubairy**, “Treasury Supply Shocks: Propagation Through Debt Expansion and Maturity Adjustment,” 2026. Working paper draft dated February 23, 2026; first version May 2025. No series number printed.
- Board of Governors of the Federal Reserve System**, “Factors Affecting Reserve Balances (H.4.1): Marketable U.S. Treasury Securities Held in Custody for Foreign Official and International Accounts,” n.d. Weekly Wednesday levels at face value; FRED series WMTSECL1.
- , “Financial Accounts of the United States (Z.1),” n.d.
- Bordo, Michael D. and Robert N. McCauley**, “Triffin: Dilemma or Myth?,” BIS Working Papers 684, Bank for International Settlements December 2017.
- Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas**, “Safe Asset Scarcity and Aggregate Demand,” NBER Working Paper 22044, National Bureau of Economic Research February 2016. *American Economic Review Papers and Proceedings*, May 2016.

- Cahill, Michael E., Stefania D’Amico, Canlin Li, and John S. Sears**, “Duration Risk versus Local Supply Channel in Treasury Yields: Evidence from the Federal Reserve’s Asset Purchase Announcements,” Finance and Economics Discussion Series 2013-35, Board of Governors of the Federal Reserve System April 2013.
- Chaudhary, Manav, Julie Zhiyu Fu, and Haonan Zhou**, “Anatomy of the Treasury Market: Who Moves Yields?,” 2024.
- Chinn, Menzie D., Jeffrey A. Frankel, and Hiro Ito**, “The Dollar versus the Euro as International Reserve Currencies,” NBER Working Paper 32387, National Bureau of Economic Research April 2024.
- Choi, Jason, Rishabh Kirpalani, and Diego J. Perez**, “The Macroeconomic Implications of US Market Power in Safe Assets,” NBER Working Paper 30720, National Bureau of Economic Research 2022.
- D’Amico, Stefania and Thomas B. King**, “Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply,” *Journal of Financial Economics*, 2013, 108 (2), 425–448.
- Du, Wenxin, Benjamin M. Hébert, and Wenhao Li**, “Intermediary Balance Sheets and the Treasury Yield Curve,” NBER Working Paper 30222, National Bureau of Economic Research May 2023.
- , **Joanne Im, and Jesse Schreger**, “The U.S. Treasury Premium,” *Journal of International Economics*, 2018, 112, 167–181.
- Duffie, Darrell**, “Resilience Redux in the US Treasury Market,” Symposium Paper, Federal Reserve Bank of Kansas City, Jackson Hole Economic Policy Symposium August 2023.
- Engen, Eric M. and R. Glenn Hubbard**, “Federal Government Debt and Interest Rates,” in “NBER Macroeconomics Annual 2004,” Vol. 19, MIT Press, 2004, pp. 83–138. With comments and discussion.
- Eren, Egemen, Andreas Schrimpf, and Fan Dora Xia**, “The Demand for Government Debt,” 2023. Working paper, February 2025 draft.

Fang, Xiang, Bryan Hardy, and Karen K. Lewis, “Who Holds Sovereign Debt and Why It Matters,” NBER Working Paper 30087, National Bureau of Economic Research May 2022. Revised September 2024.

Farhi, Emmanuel and Matteo Maggiori, “A Model of the International Monetary System,” *Quarterly Journal of Economics*, 2018, 133 (1), 295–355.

Federal Reserve Bank of New York, “Administration of Relationships with Primary Dealers,” 2016. Operating policy; primary dealers are expected to bid in every Treasury auction for at least their pro rata share at reasonably competitive prices.

—, “System Open Market Account Holdings,” n.d.

Fleming, Michael J., “Who Buys Treasury Securities at Auction?,” Current Issues in Economics and Finance 1, Federal Reserve Bank of New York January 2007.

— **and Weiling Liu**, “Intraday Pricing and Liquidity Effects of U.S. Treasury Auctions,” 2017. Working paper, September 2017.

Gabaix, Xavier and Ralph S. J. Koijen, “In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis,” NBER Working Paper 28967, National Bureau of Economic Research June 2021.

Gorodnichenko, Yuriy and Walker Ray, “The Effects of Quantitative Easing: Taking a Cue from Treasury Auctions,” NBER Working Paper 24122, National Bureau of Economic Research December 2017. Published as Droste, Gorodnichenko, and Ray, “Unbundling Quantitative Easing: Taking a Cue from Treasury Auctions,” *Journal of Political Economy* 132(9): 3115–3172, 2024.

Gourinchas, Pierre-Olivier, Hélène Rey, and Maxime Sauzet, “The International Monetary and Financial System,” NBER Working Paper 25782, National Bureau of Economic Research April 2019. Prepared for the Annual Review of Economics.

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

- , **Samuel G. Hanson, Joshua S. Rudolph, and Lawrence H. Summers**, “Government Debt Management at the Zero Lower Bound,” Working Paper 5, Hutchins Center on Fiscal and Monetary Policy, Brookings Institution September 2014.
- He, Zhiguo, Arvind Krishnamurthy, and Konstantin Milbradt**, “A Model of Safe Asset Determination,” *American Economic Review*, 2019, 109 (4), 1230–1262.
- Jansen, Kristy A. E., Wenhao Li, and Lukas Schmid**, “Granular Treasury Demand with Arbitrageurs,” NBER Working Paper 33243, National Bureau of Economic Research December 2024. Revised June 2026.
- Jiang, Zhengyang, Arvind Krishnamurthy, and Hanno Lustig**, “Foreign Safe Asset Demand and the Dollar Exchange Rate,” *Journal of Finance*, 2021, 76 (3), 1049–1089.
- , **Robert J. Richmond, and Tony Zhang**, “Convenience Lost,” NBER Working Paper 33940, National Bureau of Economic Research June 2025.
- Kaminska, Iryna and Gabriele Zinna**, “Official Demand for US Debt: Implications for US Real Rates,” *Journal of Money, Credit and Banking*, 2020, 52 (2–3), 323–364.
- Koijen, Ralph S. J. and Motohiro Yogo**, “A Demand System Approach to Asset Pricing,” NBER Working Paper 21749, National Bureau of Economic Research July 2019. Published in the *Journal of Political Economy*, 2019.
- **and** — , “Exchange Rates and Asset Prices in a Global Demand System,” NBER Working Paper 27342, National Bureau of Economic Research June 2020. Revised June 2026.
- Krishnamurthy, Arvind and Annette Vissing-Jorgensen**, “The Aggregate Demand for Treasury Debt,” 2008. Working draft, February 12, 2008.
- **and** — , “The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy,” *Brookings Papers on Economic Activity*, 2011, pp. 215–265. Fall 2011, with Comments and Discussion.
- **and** — , “The Aggregate Demand for Treasury Debt,” *Journal of Political Economy*, 2012, 120 (2), 233–267.

Laubach, Thomas, “New Evidence on the Interest Rate Effects of Budget Deficits and Debt,” *Journal of the European Economic Association*, June 2009, 7 (4), 858–885.

Lou, Dong, Hongjun Yan, and Jinfan Zhang, “Anticipated and Repeated Shocks in Liquid Markets,” *Review of Financial Studies*, 2013, 26 (8), 1891–1912.

Mian, Atif R., Ludwig Straub, and Amir Sufi, “A Goldilocks Theory of Fiscal Deficits,” NBER Working Paper 29707, National Bureau of Economic Research January 2022.

Phillot, Maxime, “U.S. Treasury Auctions: A High Frequency Identification of Supply Shocks,” *American Economic Journal: Macroeconomics*, 2025, 17 (1), 245–273.

Somogyi, Fabricius, Jonathan Wallen, and Lingdi Xu, “What Treasury Auctions Reveal About Investor Demand,” Working Paper 26-033, Harvard Business School 2025. This version December 2, 2025.

Stedman, Karlye Dilts, “The Changing Investor Composition of U.S. Treasuries, Part 1: Foreign Treasury Sales Could Raise U.S. Yields,” Economic Bulletin, Federal Reserve Bank of Kansas City July 2025.

Tabova, Alexandra M. and Francis E. Warnock, “Foreign Investors and US Treasuries,” NBER Working Paper 29313, National Bureau of Economic Research September 2021. Revised April 2022.

U.S. Department of the Treasury, “Increased Noncompetitive Bid and Award Limits for Foreign and International Monetary Authority Accounts,” March 2023. Press release, March 9, 2023: the per-account noncompetitive limit for FIMA accounts rises from \$100 million to \$500 million and the aggregate award limit from \$1 billion to \$2 billion, for auctions announced from March 14, 2023.

—, “Foreign Portfolio Holdings of U.S. Securities,” 2025. Treasury International Capital (TIC) System annual survey report.

—, “Daily Treasury Par Yield Curve Rates,” n.d.

—, “Investor Class Auction Allotments,” n.d.

— , “Quarterly Refunding Statements,” n.d.

— , “Treasury International Capital System: Major Foreign Holders of Treasury Securities,” n.d.
Monthly holdings of Treasury securities by foreign official institutions, notes and bonds separately; benchmarked to the annual survey.

U.S. Department of the Treasury, Fiscal Data, “Monthly Statement of the Public Debt,” n.d.

— , “Treasury Securities Auctions Data,” n.d.

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

Vissing-Jorgensen, Annette, “The Treasury Market in Spring 2020 and the Response of the Federal Reserve,” NBER Working Paper 29128, National Bureau of Economic Research August 2021.

Warnock, Francis E. and Veronica Cacad Warnock, “International Capital Flows and U.S. Interest Rates,” *Journal of International Money and Finance*, 2009, 28 (6), 903–919.

A Proofs and the two-bucket closed forms

Every expression in this appendix is a fragment emitted by the replication code’s derivation check, which verifies each claim in Section 3 symbolically or, where symbolic simplification fails, by sampling the assumption domain. The report of that check accompanies the replication package.

Proposition 1. From the clearing condition with the surprise, (5), dividing by $(1 - \theta_b)Q_b$ and taking logarithms gives (6). Differentiating with respect to s_b at $s_b = 0$ gives (7); multiplying by Q_b gives (8); differentiating with respect to θ_b gives (9), which is positive on the domain $\theta_b < 1$, $\varepsilon > 0$. The derivation check also confirms (7) against implicit differentiation of (5).

Proposition 2. The ratio of (9) to (8) is (11). The elasticity of (7) with respect to Q_b is -1 by inspection. Under the level alternative, $Q_b + s_b = \theta_b Q_b + A_b + \lambda y_b$ gives $dy_b/ds_b = 1/\lambda$, (12), whose derivative with respect to θ_b is zero.

Proposition 3. The two-bucket system is (13), (14), (3), its twin for bucket 2, and $dS = dX_1 + dX_2$, five linear equations in $dy_1, dy_2, dX_1, dX_2, dS$. Solving, with the common

denominator

$$D = F_1 F_2 a^2 \varepsilon^2 v^2 + 2F_1 F_2 a^2 \varepsilon^2 v \sigma^2 + F_1 a \varepsilon v + F_1 a \varepsilon \sigma^2 + F_2 a \varepsilon v + F_2 a \varepsilon \sigma^2 + 1 \quad (25)$$

gives the own response

$$\frac{dy_1}{ds_1} = \frac{a(F_2 a \varepsilon v^2 + 2F_2 a \varepsilon v \sigma^2 + v + \sigma^2)}{D} \quad (26)$$

the spillover

$$\frac{dy_2}{ds_1} = \frac{a \sigma^2}{D} \quad (27)$$

and the relative response

$$\frac{d(y_1 - y_2)}{ds_1} = \frac{a v (F_2 a \varepsilon v + 2F_2 a \varepsilon \sigma^2 + 1)}{D} \quad (28)$$

Every term of D and of each numerator is positive on the domain, which gives the three signs.

The derivative of the relative response with respect to F_1 is negative, its elasticity with respect to F_1 lies in $(-1, 0)$, and its derivative with respect to v is positive; all three are verified by the derivation check, the first two symbolically and the third by sampling the domain, and their expressions are in the replication package. The limits as $v \rightarrow \infty$ and $v \rightarrow 0$ are (15) and

$$\lim_{v \rightarrow 0} \frac{d(y_1 - y_2)}{ds_1} = 0 \quad (29)$$

The ratio of the relative response to the segmented response,

$$\varepsilon F_1 \frac{d(y_1 - y_2)}{ds_1} = \frac{F_1 a \varepsilon v (F_2 a \varepsilon v + 2F_2 a \varepsilon \sigma^2 + 1)}{D} \quad (30)$$

lies in $(0, 1)$ because its numerator is D less a sum of positive terms. Because the elasticity of the relative response with respect to F_1 is above -1 , the within-day ratio β_2/β_1 is below the segmented $1/(1 - \bar{\theta})$.

Proposition 4. With only the fast sector responding at the auction and the official sector taking the quantity G_b , clearing on the auction day is $Q_b + s_b = \theta_b Q_b + G_b + (1 - \theta_b) Q_b \exp(\varepsilon_0 \Delta y_b^{(0)})$, whose first-order solution is (16); its derivative with respect to G_b is negative, its derivative with respect to θ_b positive, and the average share G_b/s_b falls with s_b while the marginal share dG_b/ds_b is zero by construction. Writing $G_b = \eta_b s_b$ gives (17); with the whole private sector responding to the surprise component, ε_1 replaces ε_0 and gives (18); the difference is

$$\Delta y_b^{(0)} - \Delta y_b^{(\infty)} = \frac{s_b (1 - \eta_b) \left(-\frac{1}{\varepsilon_1} + \frac{1}{\varepsilon_0} \right)}{Q_b (1 - \theta_b)} \quad (31)$$

The derivative of (17) with respect to θ_b is

$$\frac{\partial \Delta y_b^{(0)}}{\partial \theta_b} = -\frac{s(\eta - 1)}{Q\epsilon_0(\theta - 1)^2} \quad (32)$$

and is positive; the ratio of (17) to (18) is ϵ_1/ϵ_0 ; dividing (17) by $q_b = s_b/Q_b$ gives

$$\frac{\Delta y_b^{(0)}}{q_b} = \frac{1 - \eta_b}{\epsilon_0(1 - \theta_b)} \quad (33)$$

An official sector with semi-elasticity ϵ_{off} would change its holdings by $\theta_b Q_b \epsilon_{\text{off}} \Delta y_b^{(0)}$ to first order, and with $d\Delta y_b^{(0)}/ds_b = 1/(\epsilon_0(1 - \theta_b)Q_b)$ its marginal share of the issue is (19), which is positive and rises with ϵ_{off} .

The cost decomposition. Taking logarithms of (7) gives (20); the check confirms the identity on the domain.

B Data construction

The composition series. The annual survey reports carry the maturity distribution of foreign official and foreign private holdings of long-term Treasuries in percentage tables for the survey years 2004 to 2020 and, in the reports for 2023 to 2025, in dollar tables by holder covering 2017 to 2025; the earliest reports do not split holders. The replication code parses each table from the report's text layer, checks that each distribution sums to one hundred, and converts percentages to dollars with the Financial Accounts' end-June foreign official and foreign private long-term Treasury positions (series FL263061130 less FL263061160, and FL263061105 less FL263061110 less the official position). Where the percentage-based and the dollar-table series overlap, the two agree within 1.2 percent at the market level. The stock by bucket is the sum of par values of marketable notes, bonds, inflation-indexed, and floating-rate securities with remaining maturity in the bucket at the June month-end from the Monthly Statement of the Public Debt's security-level table; the Federal Reserve's holdings by bucket are the par value plus inflation compensation of the SOMA's Treasury holdings by CUSIP at the as-of date nearest the month-end, and are computed for every month-end from 2003. The bucket of a tenor is the bucket whose upper bound equals the tenor. Months between surveys carry linear interpolations of the official share and are flagged; the surveys date from June and the robustness section restricts to April through September.

The auction and allotment data. The auction universe is Treasury's auction query, filtered to notes and bonds with original terms of 2, 3, 5, 7, 10, 20, and 30 years, excluding inflation-indexed and floating-rate securities and a handful of records of \$25 million that are not regular auctions. The investor-class allotments come in two spreadsheets, the earlier one, which ends with the auctions of the third quarter of 2009, in millions of dollars and the later one in billions, matched to auctions one record per auction. The spreadsheets carry the issue date and the term of each issue but no CUSIP, and the pair is not unique: Treasury has reopened several off-the-run securities of one original term to settle on one day (4 off-the-run 10-year notes of \$10 billion each in October 2008), and it has sold a note as a legal reopening of an older CUSIP with a longer original term (a 2-year note reopening a 5-year CUSIP). The tenor of such an auction is the term of the instrument sold, which is how Treasury announces it and how the allotment release records it, and it is treated as a new issue of that tenor; within an issue date and tenor with as many allotment rows as auctions, rows are assigned to auctions by rank of accepted amount, one row to several auctions is prorated by accepted amount, and several rows to one auction are summed. The private issue is the total issue less the SOMA add-on. A slot is a tenor and a new-issue or reopening status; consecutive pairs within a slot more than two hundred days apart are dropped. The allotment totals are checked against the auction records' accepted amounts, and the units are consistent across the spreadsheets' change from millions to billions.

The holdings records. The monthly series is the Treasury International Capital system's major-foreign-holders table, foreign official institutions' holdings of notes and bonds, in billions of dollars at each month-end; the monthly coupon stock is the sum of outstanding notes, bonds, inflation-indexed, and floating-rate securities in the Monthly Statement of the Public Debt at each month-end, taking one line per security (the statement's per-payment lines, which carry no maturity date, are dropped, as in the composition series); the Federal Reserve's monthly coupon holdings are the SOMA file's par value plus inflation compensation summed across buckets. The weekly series is the H.4.1 memorandum item for marketable Treasury securities, bills included, held in custody for foreign official and international accounts, Wednesday levels at face value, from July 2007. Coupon settlements are assigned to the first Wednesday on or after the auction's issue date, summing the total accepted amounts; coupon maturities are each security's last recorded outstanding amount assigned to the first Wednesday on or after its maturity date; net settlements are the difference; bill settlements are built from the bill auctions the same way. The

horizon regressions are local projections of the cumulative change in the holdings series on the current change in the stock (monthly) or the current net settlements (weekly), by period, with Newey-West standard errors using $h + 1$ lags; the monthly regressions add the cumulative change in the ten-year constant-maturity yield over the same horizon; the weekly regressions add the week's bill settlements, indicators for the week of the month, and a linear trend within the window, and report the gross-settlement coefficient, the no-control coefficient, and a placebo in which the dependent variable is the custody change over the eight weeks before the settlement.

The surprises. The naive surprise is the offering amount less the offering amount of the previous auction in the same slot. The guidance-netted surprise uses the most recent quarterly refunding statement before the announcement: where the statement carries a monthly size table, the table's entry for the tenor and month; where it states increases of a given number of billions per month for named tenors, the previous auction's size plus one step; where it states that sizes will be unchanged, the previous auction's size; otherwise the naive surprise. On the statement's own day, its sizes are the announcement and the surprise is measured against the previous statement's table where one exists and against the previous auction otherwise. The path news for a refunding statement is, for each tenor, the sum of the sizes of the tenor's auctions over the three months beginning with the statement's month less the sum the previous statement's table implied or, where it implied nothing, the previous same-slot sizes held constant. The parse of every statement is listed in the replication package with its classification.

The registration memo. The replication package reproduces the identification memo committed before the first results, with its five dated amendments: the first recording the failure of the announcement design and the move to the auction-absorption design; the second recording the trend sensitivity of the concession result and pre-specifying the marginal-absorption design at auction; the third recording the changes made after a review of the paragraph outline; the fourth recording the absorption designs at the weekly and monthly horizons and the diagnostics added after the first round of review; and the fifth recording the bill-settlement and trend controls, the net-settlement regressor, the Federal Reserve's share in the accounting, and the corrected monthly stock added after the second round, each written before it was run.

Table C1. Marginal absorption at auction by tenor

Tenor	Pairs	Foreign avg.	Marginal	(s.e.)	Dealers avg.	Marginal	(s.e.)	Funds avg.	Marginal	(s.e.)
2-year	300	0.19	0.09	(0.14)	0.49	0.59	(0.21)	0.29	0.32	(0.18)
3-year	219	0.18	0.30	(0.11)	0.43	0.39	(0.19)	0.38	0.29	(0.16)
5-year	276	0.18	0.20	(0.13)	0.43	0.51	(0.17)	0.37	0.23	(0.14)
7-year	207	0.18	0.31	(0.20)	0.31	0.11	(0.18)	0.50	0.61	(0.18)
10-year	265	0.18	0.06	(0.04)	0.42	0.64	(0.04)	0.38	0.30	(0.02)
20-year	66	0.13	0.11	(0.08)	0.17	0.30	(0.05)	0.69	0.60	(0.09)
30-year	200	0.12	0.18	(0.11)	0.34	0.37	(0.15)	0.53	0.52	(0.11)

Notes: First-difference regressions of the change in each class's allotment on the change in the privately placed issue between consecutive auctions of the same slot, by tenor; standard errors clustered by auction month. Average shares are means of the class's share of the private issue. Source: Treasury investor-class auction allotments and auction records; authors' calculations.

Table C2. Official custody holdings and coupon settlements at the weekly horizon

Window	Net coupons, $h = 0$	(s.e.)	$h = 4$	(s.e.)	$h = 8$	(s.e.)	Gross coupons, $h = 8$	Bills, $h = 8$	Placebo, prior 8 weeks	Gross, no controls, $h = 8$
July 2007 to August 2008	0.25	(0.07)	0.52	(0.19)	0.59	(0.17)	-0.10 (0.13)	0.26 (0.13)	-0.03 (0.18)	0.21 (0.11)
September to December 2008	0.15	(0.13)	0.17	(0.05)	-0.03	(0.04)	-0.10 (0.11)	0.06 (0.03)	-0.44 (0.10)	-0.08 (0.06)
2007–2008	0.28	(0.05)	0.47	(0.13)	0.32	(0.21)	-0.14 (0.11)	0.19 (0.04)	0.12 (0.17)	0.47 (0.14)
2009–2013	0.09	(0.02)	0.00	(0.07)	-0.01	(0.09)	-0.02 (0.06)	0.48 (0.19)	-0.11 (0.06)	0.03 (0.02)
2014–2019	0.07	(0.06)	-0.17	(0.13)	-0.28	(0.21)	-0.03 (0.05)	-0.18 (0.12)	-0.05 (0.05)	-0.00 (0.02)
2020–2025	-0.00	(0.02)	-0.06	(0.06)	-0.02	(0.10)	-0.06 (0.05)	0.14 (0.06)	0.05 (0.05)	-0.02 (0.02)
All	0.04	(0.02)	0.04	(0.05)	0.09	(0.08)	0.00 (0.03)	0.16 (0.04)	0.07 (0.03)	-0.08 (0.02)

Notes: Local projections of the cumulative change in marketable Treasury securities held in custody for foreign official and international accounts (H.4.1, face value, bills included) over h weeks on the week's net coupon settlements (new issues less maturities), with the week's bill settlements, indicators for the week of the month, and a linear trend within the window; Newey-West standard errors with $h + 1$ lags. The last columns give, at eight weeks, the coefficient on gross coupon settlements in the same regression, the coefficient on bill settlements, a placebo in which the dependent variable is the custody change over the eight weeks before the settlement, and the gross-settlement coefficient with no controls. Source: Federal Reserve Board (H.4.1), Treasury auction records, Fiscal Data; authors' calculations.

C Additional results

This appendix carries the exhibits the body refers to and does not discuss at length: the marginal shares by tenor, the weekly custody regressions, the concession and the SOMA coefficient by period, the heteroskedasticity check on announcement days, the largest refunding-statement path-news days, the event path by tercile in tabular form, the windows around the auction, the official-share coefficient by event day, the allotment-share regressions, the full battery for the concession, the specification curve and leave-one-out plots, the placebo distributions, and the nested-exponent profile.




Figure C1. The official-share coefficient at each event day

Notes: The coefficient on the foreign official share in a regression of the relative one-day yield change at each event day on the share, with tenor and month effects and the volatility control, with 95 percent intervals clustered by month. Source: as Table 6; authors' calculations.




Figure C2. The specification curve for the concession

Notes: The coefficient on the foreign official share across every specification in the battery, sorted, with 95 percent intervals; the baseline and the tenor-specific-trend specifications are marked. Source: as Table 6; authors' calculations.




Figure C3. Leave-one-year-out estimates

Notes: Left: the official-share coefficient in the concession regression with each year dropped in turn, with 95 percent intervals. Right: the foreign class's first-difference marginal share at auction with each year dropped, against the class's full-sample average share. Source: authors' calculations.




Figure C4. Placebo distributions

Notes: Left: the coefficient on the official share when the share is permuted across tenors within each month, 500 permutations, with the actual estimate marked; the permutation keeps each tenor's time path, so it tests the cross-tenor relation and not the within-tenor trend. Right: the news interaction when each announcement day is re-dated to a random non-announcement day of the same year, 300 re-datings, with the actual estimate marked. Source: authors' calculations.

Table C3. The concession by period


Period	θ_b	(s.e.)	θ_b^F	(s.e.)	Mean θ	Mean concession (bp)	Mean amplitude (bp)	N
2000–2008	-10.30	(8.17)	1.37	(5.70)	0.43	-0.25	2.16	252
2009–2013	-9.97*	(5.27)	0.88	(3.59)	0.40	0.90	2.25	357
2014–2019	-0.37	(2.19)	0.58	(1.29)	0.31	0.43	1.35	432
2020–2025	-11.57***	(4.28)	9.43***	(2.55)	0.19	-0.91	-1.06	500
2000–2013	-2.92	(3.55)			0.41	0.40	2.21	609
2014–2025	6.21***	(2.05)			0.25	-0.29	0.06	932

Notes: Each row estimates equation (22) on the period's auctions with tenor and month effects and the volatility control, once with the official share and once with the Federal Reserve's exact monthly share; standard errors clustered by month; the SOMA sample begins in 2003. Source: Treasury auction records, Treasury par yields, the composition series, SOMA holdings; authors' calculations.

Table C4. Heteroskedasticity on announcement days

Tenor	Announcement days	Variance, announced	Variance, other days	Ratio	p	Statement days	Ratio, statement days
2-year	302	8.05	7.72	1.04	0.300	104	0.98
3-year	221	5.07	4.65	1.09	0.173	104	1.31
5-year	278	1.97	1.96	1.01	0.461	104	1.27
7-year	203	1.20	1.51	0.80	0.016	104	1.17
10-year	264	1.77	1.77	1.00	0.480	104	0.85
20-year	68	8.04	4.70	1.71	0.000	104	1.24
30-year	211	6.65	6.48	1.03	0.385	88	2.03

Notes: Variance of the tenor's daily yield change relative to the mean change across the tenors, on the days the tenor's auction size was announced and on all other days (refunding-statement days excluded from the latter), with the F -test p -value of the ratio; the last columns use refunding-statement days. Source: Treasury par yields and auction records; authors' calculations.

**Figure C5.** The nested exponent on the elastic float

Notes: Residual sum of squares, relative to its minimum, of the concession regression with the auction size scaled by $(1 - \theta_b)^{-\psi}$, over a grid of ψ ; the model is $\psi = 1$ and the level alternative $\psi = 0$. Source: authors' calculations.

Table C5. The largest refunding-statement path news

Statement	Tenor	Path news (\$ bn)	Percent of stock	Yield change (bp)
November 5, 2025	5-year	43	2.20	7
November 4, 2020	7-year	18	1.89	-10
August 2, 2023	2-year	18	0.58	-4
August 5, 2020	10-year	18	3.83	3
November 1, 2023	5-year	18	0.98	-15
May 6, 2020	7-year	18	2.15	4
February 2, 2022	7-year	-18	-1.49	-2
November 1, 2017	7-year	16	2.17	-1

Notes: The eight statements with the largest absolute path news for any tenor; path news is the quarter's announced issuance path for the tenor less the path implied by the previous statement, in billions of dollars and in percent of the bucket's stock; the yield change is the tenor's par-yield change from the close before the statement to the close of the statement day. Source: Treasury refunding statements, auction records, and par yields; authors' calculations.

Table C6. The event path by official-share tercile

Event day	Low θ	(s.e.)	Middle	(s.e.)	High θ	(s.e.)
-5	0.07	(0.12)	0.17	(0.07)	0.27	(0.12)
-4	0.15	(0.13)	0.01	(0.07)	0.30	(0.11)
-3	-0.01	(0.11)	0.10	(0.07)	0.25	(0.12)
-2	0.00	(0.11)	0.06	(0.07)	0.05	(0.11)
-1	0.30	(0.11)	0.05	(0.07)	0.04	(0.09)
0	-0.72	(0.15)	0.16	(0.09)	0.53	(0.14)
+1	-0.30	(0.12)	-0.43	(0.08)	-0.52	(0.12)
+2	0.22	(0.12)	-0.13	(0.07)	-0.33	(0.11)
+3	-0.01	(0.10)	-0.01	(0.07)	-0.09	(0.11)
+4	0.15	(0.13)	-0.11	(0.06)	-0.04	(0.11)
+5	0.10	(0.12)	0.07	(0.07)	-0.10	(0.11)

Notes: Mean one-day change in the auctioned tenor's yield relative to the other tenors at each business day from five before to five after the auction, by tercile of the bucket's official share, with standard errors of the mean; all auctions. Source: Treasury par yields, auction records, and the composition series; authors' calculations.

Table C7. The official share and the windows around the auction

Outcome (bp, relative to the curve)	θ_b	(s.e.)	θ_b^F	(s.e.)	N
Run-up, $t - 2$ to t	3.22	(2.42)	-0.27	(1.50)	1541
Auction day, $t - 1$ to t	4.33**	(1.84)	-0.24	(1.13)	1541
Reversal, t to $t + 2$	-4.87**	(2.30)	2.04	(1.50)	1541
Amplitude (run-up minus reversal)	8.09***	(3.02)	-2.31	(2.02)	1541
Run-up, $t - 5$ to t	1.44	(3.84)			1541
Reversal, t to $t + 5$	-4.71	(3.87)			1541
Five-day amplitude	6.15	(5.30)	-2.95	(3.31)	1541

Notes: Each row is equation (22) with the named window as the dependent variable (relative to the curve, in basis points), with tenor and month effects and the volatility control; the SOMA columns replace the official share with the Federal Reserve's exact monthly share. Standard errors clustered by month. Source: as Table 6; authors' calculations.

Table C8. The concession and who took the issue (descriptive)

Regressor (outcome: auction-day concession, bp)	Coefficient	(s.e.)	N
Dealers' allotment share	6.38***	(1.19)	1540
Foreign and international allotment share	-4.09***	(1.47)	1540
Investment funds' allotment share	-5.09***	(1.38)	1540
Indirect-bidder share (auction results)	-6.70***	(1.51)	1319

Notes: The dependent variable is the auction-day concession; each regressor is a class's share of the privately placed issue (allotments) or of accepted amounts (indirect bidders), with tenor and month effects and the volatility control. An instrumental-variables version that uses the bucket's official share as the instrument is not reported, because the official share enters the concession directly in Proposition 4 and the exclusion restriction does not hold. Standard errors clustered by month. Source: Treasury investor-class auction allotments, auction records, and par yields; authors' calculations.

Table C9. The full battery for the concession

Specification	θ_b	(s.e.)	N
No volatility control	4.14	(1.75)	1541
Add auction size (percent of stock)	4.35	(1.84)	1541
Add tenor-specific linear trends	0.01	(1.90)	1541
Tenor and year effects	5.55	(1.78)	1567
Tenor effects only	8.38	(1.08)	1567
Month effects only	-1.02	(0.80)	1541
θ includes SOMA (inelastic share)	3.02	(1.40)	1541
Survey months only (April to September)	2.75	(2.42)	769
Percent-table years only (2004-2020)	1.28	(2.20)	1054
Drop extrapolated months before June 2004	4.45	(1.85)	1505
Drop 2008-2009	4.14	(1.84)	1432
Drop 2020	4.37	(1.88)	1461
Drop 20- and 30-year	4.74	(1.86)	1262
Drop 2-year	3.60	(2.00)	1243
Drop FOMC and first-Friday days	4.84	(1.87)	1505
New issues only	2.07	(2.33)	1162
Reopenings only	-5.95	(6.03)	330
2000-2013	-2.92	(3.55)	609
2014-2025	6.21	(2.05)	932
High-volatility half	8.25	(3.11)	736
Low-volatility half	-0.52	(2.22)	750
Outcome winsorized at 1/99	3.93	(1.67)	1541
Outcome winsorized at 2.5/97.5	3.64	(1.54)	1541
Trim loutcomel > 10 bp	3.97	(1.67)	1529
Weighted by auction size	5.26	(1.82)	1541
Clustered by year	4.33	(2.81)	1541
Outcome: two-day run-up (t-2 to t)	3.22	(2.42)	1541
Outcome: five-day run-up (t-5 to t)	1.44	(3.84)	1541
Outcome: five-day reversal (t to t+5), sign flipped	4.71	(3.87)	1541
Outcome: cycle amplitude (run-up minus reversal)	6.15	(5.30)	1541
Outcome: own yield change (not relative)	⁶⁹ -0.32	(3.48)	1541
Placebo: pseudo-auction day (t-10)	-0.11	(1.50)	1541

Notes: The coefficient on the official share in equation (22) under each variant; the baseline has tenor and month effects, the

D Calibration details

The calibration in Table 11 uses the market-level composition at the June 2008 peak of the official share and at June 2025, the corresponding coupon stocks and elastic floats, and four values of the private sector's semi-elasticity: the holdings-weighted yield semi-elasticity of Eren et al. (2023), which is already in the model's units; the price elasticities of Jansen et al. (2024) and Chaudhary et al. (2024), converted to yield units by multiplying by a duration of 6.5 years, the duration Chaudhary et al. (2024) use for their own translation; and the value at which the model's cost per percent of stock at the 2025 composition equals the low-frequency supply effect of Laubach (2009) and Engen and Hubbard (2004), taken as 3.5 basis points per percent of the stock. The ratios in the table are independent of the semi-elasticity; the levels are proportional to its inverse. The decomposition of the change in the log cost per dollar between the two dates uses the identity (20). The ratio under the nested exponent ψ is $[(1 - \theta_{2008})/(1 - \theta_{2025})]^\psi$. The model-implied sensitivity of the auction concession to the inelastic share, used for the power calculation in Section 6, is the pre-2014 mean concession divided by one minus the sample mean of the inelastic share, which is the derivative of (17) with respect to θ_b evaluated at the mean concession. The implied fast-sector semi-elasticity ϵ_0 is the auction size in percent of the stock times one minus the foreign class's marginal share, divided by the mean concession in percentage points and by one minus the mean inelastic share; the marginal share an elastic official sector would show at the auction is (19) evaluated at that ϵ_0 , the 2008 official share, and each candidate value of ϵ_{off} .