The Safe-Debt Laffer Curve

Ricardo J. Caballero*

August 20, 2026

Link to the latest version

Abstract

Safe public debt can become a drag on aggregate demand even while the sovereign remains solvent. Its safety requires rollover support, market-making, and balance-sheet capacity. As these resources become scarce, the marginal cost of producing safe Treasury claims rises. The resulting fiscal adjustment reduces aggregate expenditure and can eventually outweigh the expansionary effect of increasing the supply of safe assets. Debt's stationary contribution to aggregate demand then peaks, and the equilibrium safe rate reaches its maximum at the same debt stock. Beyond this point, further issuance lowers aggregate demand and the safe rate even though the debt remains safe. Empirically, I combine cash-Treasury and real-duration premia with estimates of how additional private supply reprices the outstanding Treasury stock. The benchmark marginal cost more than doubled, from about 80 basis points in 2015Q1 to 187 basis points in 2026Q1. Using a 330-basis-point benchmark for the rate-equivalent wealth benefit, the safe-debt margin—the difference between this benefit and marginal cost—fell from about 250 to 143 basis points.

JEL Codes: E43, E44, E52, E62, H63, G12, G23.

Keywords: public debt, safe assets, aggregate demand, safe interest rates, Treasury intermediation, dealer balance sheets, rollover premia, term premia.

*MIT and NBER. I thank Matteo Maggiori, Annette Vissing-Jorgensen, and Chris Wolf for helpful comments and Ziwen Sun for outstanding research assistance. First draft: May 23, 2026.

1 Introduction

Public debt in the major developed economies is approaching levels rarely seen outside major wars, and it continues to rise. The familiar concern is fiscal stability: mounting debt may become unsustainable and end in implicit or explicit default. I ask a different question. Suppose the sovereign remains solvent and its debt remains safe. Can a larger stock of safe public debt nevertheless become a drag on aggregate demand? The canonical answer is no.

Government debt is the economy's reference safe asset, both a store of value and collateral for private contracting. Issuance expands the supply of safe financial wealth and normally raises the equilibrium safe rate required to place it. At the inherited rate, finite-horizon private owners spend part of the additional wealth, while debt service reduces government purchases. Over the admissible interest-rate range, the wealth-induced expenditure exceeds the service cost, so additional debt is initially expansionary.[1] A chronic demand for stores of value, as emphasized in the safe-asset-shortage literature, reinforces this conclusion.

That reasoning treats the production of safety as costless. In practice, government debt is liquid and safe because the financial system supplies market-making, warehousing, and rollover support. The fiscal authority supplies the promise; intermediaries supply the balance-sheet capacity that allows investors to treat it as a safe claim. I call their compensation the Treasury absorption premium. These payments enter the government budget even though households continue to receive a safe payoff. Safety can therefore be preserved while becoming more expensive to produce. As the debt stock grows, the same balance sheets must support more claims, and each active intermediary bears more exposure in states in which participation is limited.

The model builds this production mechanism around rollover episodes. Over any fixed interval, part of the debt stock comes due. A core group of balance sheets remains available in every state; other potential support providers participate only in some states. Broad participation allows the system to absorb refinancing needs with little stress. When participation is limited, fewer intermediaries share the shortfall, exposure per active balance sheet rises, and so does the compensation required to mobilize support. The marginal participant prices the claim by averaging the support costs it faces across participation states. Debt raises this rollover premium, financial capacity lowers it, and convex support costs make it steepen as rollover needs approach available capacity.

The rollover premium is an average cost. Marginal issuance is more expensive because it also changes the terms on which the stock will be refinanced. The stock effect is forward-looking: each unit of existing debt enters the refinancing flow when it matures, so a higher premium changes the cost of rolling over the stock. In steady state, the sequence of rollover payments is represented as an annualized cost on outstanding debt. One more unit generates its own rollover cost and increases the premium paid as existing debt comes due. The stock-wide marginal cost is the sum of these effects. This distinction between average and marginal cost produces the turning point. Even a moderate movement in the premium can have a large fiscal effect when refinancing needs are large.

[1] This is a stock experiment rather than a deficit-financed stimulus experiment: the object is the steady-state aggregate-demand contribution of the outstanding stock of safe public debt.

Producing safe public debt affects aggregate demand through the fiscal cost of sustaining its safety. Higher Treasury absorption payments require a fiscal adjustment. After accounting for the private expenditure generated by those payments, the adjustment produces a net reduction in aggregate expenditure. The model below provides one tractable mapping from Treasury absorption costs to this broader fiscal-demand effect.

An additional unit of debt increases private financial wealth but also raises the cost of supporting the debt stock as it rolls over. In the stationary allocation, the recipients of Treasury absorption income save part of that persistent flow and hold a larger share of the debt stock. At low debt, the stationary non-Ricardian wealth effect dominates: debt generates positive aggregate-demand pressure at the inherited safe rate, and the equilibrium safe rate rises to reduce human wealth and restore equilibrium. As financial capacity becomes scarce, the marginal cost of Treasury absorption increases and the expenditure effect weakens. Beyond the turning point, the fiscal drag exceeds the wealth effect, so additional debt generates negative aggregate-demand pressure and the equilibrium safe rate falls. The aggregate-demand Laffer curve describes this stationary contribution. The safe-rate schedule reaches its maximum at the same debt stock because both slopes depend on the difference between the rate-equivalent wealth benefit of an additional safe claim and its marginal production cost.

Financial capacity moves both the onset and the severity of rollover stress. More capacity allows the inelastic base of support providers to absorb a larger refinancing flow before a premium is required. Once the economy enters the stress region, additional capacity lowers exposure in every stressed participation state, reduces the premium, and flattens its response to debt. It raises the equilibrium safe rate at a given debt stock. Holding the rate-equivalent wealth benefit fixed, the flatter marginal-cost schedule also pushes the Laffer peak to a larger debt stock.

Maturity changes how the same rollover exposure is transmitted. Longer maturity reduces the share of debt that must be refinanced during a given interval, lowering the direct rollover shortfall. The claims that remain outstanding, however, have greater duration, so their prices respond more strongly to a change in the rollover-support premium. Duration is the multiplier on this price effect. The total premium is therefore determined jointly with this multiplier. Shorter maturity raises the quantity exposed to rollover while reducing the price amplification of each remaining claim; longer maturity does the reverse. Financial capacity lowers both the direct rollover premium and its transmission into bond prices.

This maturity structure supplies the empirical decomposition. I measure the stock-wide marginal cost using the levels of the cash-Treasury and real-duration premia and the repricing of the outstanding private Treasury stock through each component. The cash-Treasury premium captures direct financing and warehousing costs. The real-duration premium captures the price amplification associated with longer claims. In each case, the premium level values a marginal unit at prevailing prices, while the supply response measures how that unit changes the price of the stock already in private hands. This level-plus-repricing decomposition is the empirical counterpart of the model's average-versus-marginal-cost distinction.

Figure 1 summarizes the empirical results. Its upper panel reports the estimated stock-wide marginal cost and separates its rollover and real-duration components. The lower panel subtracts that cost from the benchmark rate-equivalent wealth benefit of additional safe debt, yielding the remaining safe-debt margin.

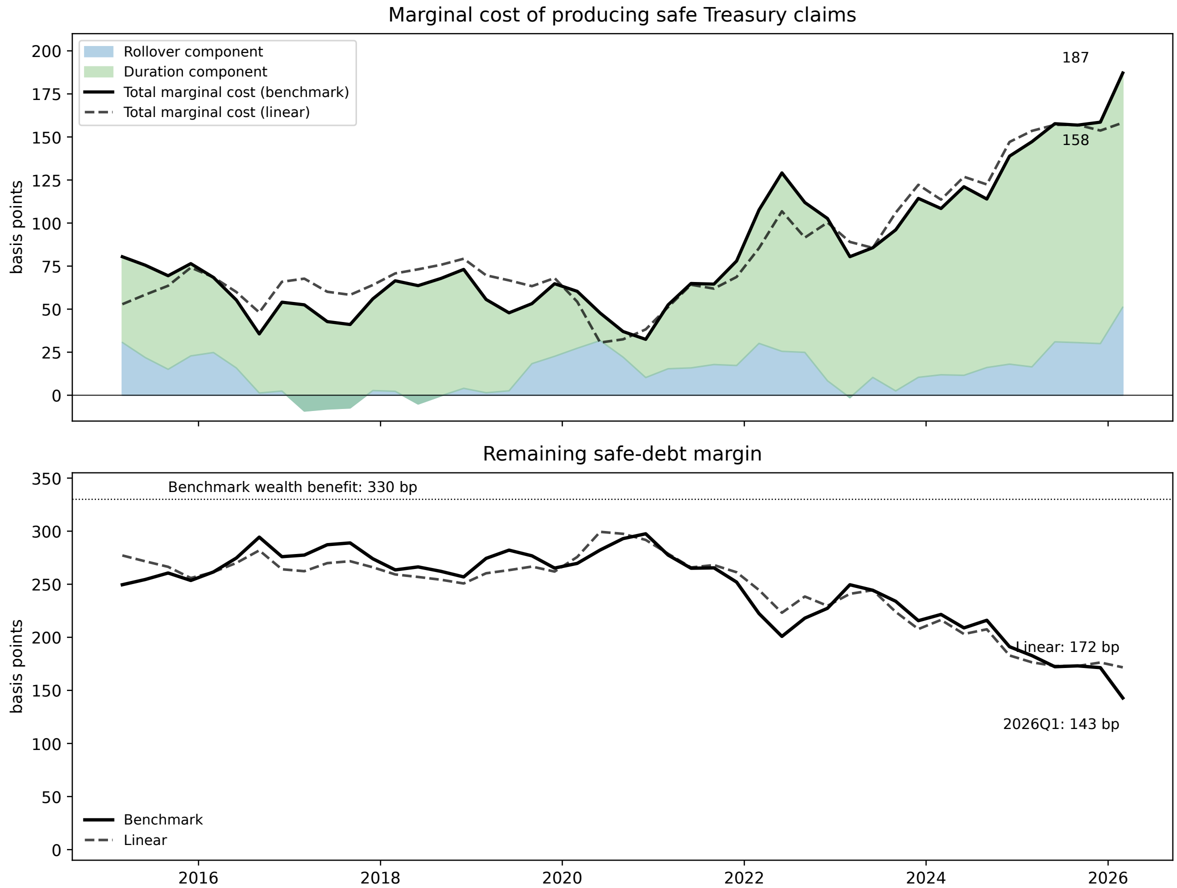

Figure 1: The marginal cost of producing safe Treasury claims and the remaining safe-debt margin. The solid series use the benchmark 2015+ OLS estimates, in which Treasury repricing varies with dealer-parent capacity under the 5 percent leverage mapping; dashed lines show the corresponding linear estimates. The upper panel stacks the rollover and real-duration components of marginal cost. The lower panel reports the rate-equivalent distance to the Laffer peak under the benchmark 330-basis-point wealth benefit. In 2026Q1, the benchmark marginal cost is approximately 187 basis points and the remaining margin is approximately 143 basis points.

For cash-Treasury repricing, I follow D'Amico and King (2013) and compare changes in the share held in the Federal Reserve's System Open Market Account (SOMA) with Treasury–swap spreads within maturity buckets. Start-month effects remove aggregate news and curve-wide repricing, while maturity effects absorb persistent differences across buckets. The within-bucket coefficient measures how a change in the private supply of a maturity segment changes its cash-Treasury premium. I then scale that response by the private float in each bucket and aggregate across maturities, producing a repricing term for the outstanding private Treasury stock.

For duration repricing, I combine Federal Reserve purchase and runoff announcements with the real term premium of D'Amico, Kim, and Wei (2018). The event-study coefficient measures how the real-duration premium responds when announced purchases change the quantity of Treasury duration held by private investors. Multiplying this response by the private coupon stock gives the corresponding stock-wide repricing term. The cash-Treasury and duration estimates are thereby placed on the same annualized basis-point scale and can be combined with their observed premium levels.

The benchmark stock-wide marginal cost more than doubles, rising from about 80 basis points in 2015Q1 to about 187 basis points in 2026Q1. The increase reflects both higher premium levels and the repricing of a much larger private Treasury stock. The supply response also steepens when dealer-parent capacity is scarce: the capacity interaction has the sign predicted by the model and is statistically significant under overlap-robust inference. Under the benchmark rate-equivalent wealth benefit, the safe-debt margin falls from about 250 to 143 basis points.

Literature Review. The natural starting point is the safe-asset-shortage literature. Caballero, Farhi, and Gourinchas (2008, 2016, 2017, 2021) and Caballero and Farhi (2018) show how scarcity depresses safe rates and contributes to global imbalances and recessions. The same view appears in work on Treasury safety and liquidity premia (Krishnamurthy and Vissing-Jorgensen, 2012), the supply and maturity of short safe claims (Greenwood, Hanson, and Stein, 2015; D'Avernas and Vandeweyer, 2024), and reserve-currency supply under limited intermediary capacity (Maggiori, 2017; Farhi and Maggiori, 2018). Recent formulations include Itskhoki and Mukhin (2026) and Brunnermeier, Merkel, and Sannikov (2024). These papers ask what happens when safe claims are scarce. I ask how far their supply can expand before the marginal cost of supporting them absorbs their contribution to demand.

Convenience-yield models supply part of the answer. When safety and liquidity services enter preferences or constraints, Treasury demand slopes downward and the convenience yield falls with supply (Krishnamurthy and Vissing-Jorgensen, 2012). In fiscal applications, dilution of the premium on outstanding debt can produce an interior revenue or debt-capacity peak (Brunnermeier, Merkel, and Sannikov, 2024; Mian, Straub, and Sufi, 2025; Li and Merkel, 2026). Here the corresponding stock effect operates through the cost of refinancing and supporting safe debt, and the outcome of interest is its marginal contribution to aggregate demand. This argument also belongs to the broader asset-shortage tradition linking a shortage of stores of value to low real rates, high valuations, bubbles, and macroeconomic fragility (Caballero, 2006, 2015). Auclert, Malmberg, Rognlie, and Straub (2025) provide a recent framework connecting asset supply, U.S. wealth, real rates, and fiscal sustainability.

On the demand side, Barro (1974) provides the Ricardian benchmark, while Blanchard (1985) and Weil (1989) show how finite horizons shape private wealth demand. Mian, Straub, and Sufi (2021, 2025) emphasize the aggregate consequences of differences in saving behavior across owners. Work on self-financing fiscal expansions (Angeletos, Lian, and Wolf, 2024), feedback among debt, distortions, and productivity (Fornaro and Wolf, 2025), and expansionary fiscal adjustment (Giavazzi and Pagano, 1990; Alesina and Perotti, 1997) studies other routes from fiscal policy to demand. Here the relevant fiscal composition is between government purchases and the compensation paid to the owners of Treasury-absorption capacity.

The supply side combines risk-bearing capacity with rollover support. Risk-centric macroeconomic models make asset prices, financial conditions, and intermediary capacity central to aggregate demand (Caballero and Simsek, 2020, 2021, 2023); Caballero, Caravello, and Simsek (2024) study the associated financial-conditions target. In Treasury markets, limited dealer balance sheets can turn a larger supply into market dysfunction (Kashyap, Stein, Wallen, and

Younger, 2025). Evidence from March 2020 and subsequent work on market resilience is provided by Duffie (2020); Duffie and Keane (2023), Schrimpf, Shin, and Sushko (2020), Kruttli, Monin, Petrasek, and Watugala (2021), He, Nagel, and Song (2022), and Du, Hébert, and Li (2023). The rollover structure draws on Calvo (1988) and Cole and Kehoe (2000), as well as the sovereign-spread and repayment literature (Eaton and Gersovitz, 1981; Arellano, 2008; Aguiar and Amador, 2014). He, Krishnamurthy, and Milbradt (2019) is especially close in making safety depend on the debt stock. I use these ingredients to put the cost of safety in the government budget and the goods-market equilibrium.

Empirically, the paper is related to estimates of the effect of debt and fiscal positions on government yields (Laubach, 2009; Gruber and Kamin, 2012; Gamber and Seliski, 2019; World Bank, 2026) and to high-frequency evidence on fiscal news, Treasury supply, term premia, inflation, and convenience yields (Gómez-Cram, Kung, and Lustig, 2024, 2025; Jiang, Richmond, and Zhang, 2025). The empirical object here is the stock-wide marginal cost of producing safe Treasury claims. The maturity extension connects that object to debt-management models (Angeletos, 2002; Buera and Nicolini, 2004; Faraglia, Marcet, Oikonomou, and Scott, 2019) and to the government's comparative advantage in supplying short safe debt (Greenwood, Hanson, and Stein, 2015). It separates the amount that must be rolled over from the price amplification created by duration.

Section 2 develops household demand and the fiscal closure. Section 3 derives the rollover premium from financial capacity and intermediary participation. Section 4 combines the two sides of the model and derives the Laffer curve. Section 5 introduces maturity, and Section 6 measures the U.S. marginal cost of producing safe Treasury claims. The appendices contain formal derivations, the maturity-choice extension, and the empirical construction details.

2 Household Demand and Fiscal Closure

This section maps the cost of producing safe debt into aggregate expenditure. Higher Treasury absorption payments require a fiscal adjustment. Government purchases fall by the full payment, while the recipients consume only part of the corresponding increase in persistent income and save the remainder. Treasury absorption therefore reduces aggregate expenditure at a fixed safe rate. The model below provides one tractable representation of this broader demand drag.

Output is an exogenous endowment normalized to one. The government has outstanding debt b, levies a fixed flow of taxes $\bar{\tau}$, and purchases the output left after servicing the debt and paying for the financial capacity that preserves its safety. For the moment, take the aggregate Treasury absorption payment S as given. Section 3 derives $S(b,N)$ from rollover needs and financial-system capacity N.

2.1 Private owners and their budget constraints

There are two unit-mass household groups, indexed by $j \in \{H, I\}$. Both receive endowment income and own financial assets. They may differ in their effective horizons and hence in their consumption and saving behavior. The financial firms that supply Treasury rollover support belong to the production side of the model developed in Section 3; their balance-sheet capacity and participation decisions are distinct from the consumption and saving decisions studied here. Group I owns these firms and receives their income.

Members of group j die with Poisson intensity $\lambda_j > 0$ and are replaced one-for-one by newborns with no initial holdings of safe government debt. They discount utility at rate $\rho > 0$ and have logarithmic flow utility. Actuarially fair annuity markets operate within each group: the safe-asset holdings of owners who die are transferred to the surviving members of the same group. Appendix A embeds these consumption rules in Epstein–Zin preferences and derives the limiting demand for supported safe claims when risk aversion to rollover losses is high.

Consider a surviving member of group j. Let $w_{j,t}$ denote safe financial wealth, $y_{j,t}$ nonfinancial income, and $c_{j,t}$ consumption. Because the annuity market pays the mortality credit $\lambda_j w_{j,t}$ to survivors, the individual budget constraint is

$$\dot{w}_{j,t} = (r_t + \lambda_j) w_{j,t} + y_{j,t} - c_{j,t}. \tag{1}$$

The safe return is r_t. The additional λ_j is not an aggregate return on wealth; it is the transfer from members of the group who die to those who survive.

Human wealth is the value of nonfinancial income received while the owner remains alive. At date t, it is

$$h_{j,t} = \int_t^{\infty} \exp\left[-\int_t^u (r_v + \lambda_j)\, dv\right] y_{j,u}\, du. \tag{2}$$

For a constant safe rate and a permanent income flow y_j, this becomes

$$h_j = \frac{y_j}{r + \lambda_j}. \tag{3}$$

The mortality hazard enters the discount rate because income received after death is of no value to the current owner.

Log utility and actuarially fair annuities yield the familiar consumption rule

$$c_{j,t} = (\rho + \lambda_j)(w_{j,t} + h_{j,t}). \tag{4}$$

Thus $\rho + \lambda_j$ is the annual propensity to consume out of total wealth. A longer effective horizon lowers this propensity and makes the group save more of a change in permanent income. This is the first object governing the incidence of Treasury absorption payments.

Aggregation removes the mortality credit from the group budget. Survivors receive the safe-asset holdings of deceased owners, but the same holdings leave the deceased owners' accounts. If a_j and c_j denote aggregate safe financial wealth and consumption of group j, respectively, then

$$\dot{a}_j = r a_j + y_j - c_j. \tag{5}$$

In a stationary allocation, aggregate wealth is constant, so

$$c_j = r a_j + y_j. \tag{6}$$

Combining this budget identity with the consumption rule gives stationary desired financial wealth:

$$a_j = \kappa_j(r) y_j, \qquad \kappa_j(r) = \frac{r - \rho}{(r + \lambda_j)(\rho + \lambda_j - r)}. \tag{7}$$

Positive wealth requires

$$\rho < r < \rho + \lambda_j. \tag{8}$$

The lower inequality makes saving attractive relative to immediate consumption. The upper inequality ensures that finite lives prevent wealth from growing without bound. Inside this interval,

$$\kappa_j'(r) = \frac{\rho\lambda_j + \lambda_j^2 + (r-\rho)^2}{(r+\lambda_j)^2(\rho+\lambda_j-r)^2} > 0. \tag{9}$$

A higher safe return therefore raises the stationary stock of safe wealth that the group wishes to hold per unit of permanent nonfinancial income.

Two distinct behavioral objects have now appeared. The propensity $\rho + \lambda_j$ determines how consumption responds to a change in total wealth. Holding financial wealth and the safe rate fixed, a permanent increase in nonfinancial income raises consumption according to

$$\left.\frac{\partial c_j}{\partial y_j}\right|_{a_j,r} = \frac{\rho+\lambda_j}{r+\lambda_j} \tag{10}$$

per dollar of additional income. The capitalization factor $\kappa_j(r)$ determines the stationary stock of safe wealth associated with the same income flow. These objects are linked by

$$\kappa_j(r)\big(\rho+\lambda_j-r\big) = 1 - \frac{\rho+\lambda_j}{r+\lambda_j} = \frac{r-\rho}{r+\lambda_j}. \tag{11}$$

The identity will connect the non-Ricardian wealth effect of additional debt to the fiscal drag generated by Treasury absorption payments.

Remark 1 (Why the model has two household groups)**.** The two-group structure separates the average wealth demand that places the existing debt stock from the stationary saving response to Treasury absorption income. Both groups' capitalization factors determine aggregate desired safe wealth and therefore the level and adjustment of the safe rate. The capitalization factor of group *I* additionally determines how much desired wealth is created by the persistent absorption income received by the owners of financial firms. This separation is useful for calibration: the observed debt stock and safe rate discipline average safe-wealth demand, while the saving response of group *I* determines how estimated marginal production costs map into the stationary distance from the Laffer peak.

2.2 Taxes, Treasury absorption income, and government purchases

Let S denote the government's aggregate annualized payment for Treasury absorption. The payment compensates financial firms for committing balance sheets to rollover support and is distributed to their ultimate household owners. In the steady state, this annual flow represents the sequence of payments made as outstanding claims mature and are refinanced. It does not imply that the government retroactively changes the return on bonds that have not yet rolled over.

The government levies the fixed tax flow $\bar{\tau}$. Its budget constraint is

$$g = \bar{\tau} - rb - S, \tag{12}$$

where g denotes government purchases. Taxes do not adjust when Treasury absorption becomes more expensive. At a given debt stock and safe rate, an additional dollar of S therefore reduces government purchases by one dollar.

Let $x_H > 0$ and $x_I > 0$ denote the groups' fixed after-tax endowment incomes. Both groups pay taxes, and their after-tax incomes satisfy

$$x_H + x_I = 1 - \bar{\tau}. \tag{13}$$

The households in group I own the financial firms that supply rollover support and receive the associated after-tax absorption income. The two groups' total nonfinancial incomes are therefore

$$y_H = x_H, \qquad y_I = x_I + S. \tag{14}$$

The analysis considers allocations in which both income flows remain positive. This condition also permits a signed empirical net premium: a convenience benefit can make the measured net absorption payment negative without making either group's total income negative.

Substituting the two income flows into (7) gives

$$a_H = \kappa_H(r)x_H, \qquad a_I = \kappa_I(r)(x_I + S). \tag{15}$$

Both groups demand safe wealth out of their baseline incomes. In addition, group I accumulates part of its recurring Treasury absorption income. One dollar of persistent absorption income therefore generates $\kappa_I(r)$ dollars of desired stationary safe wealth.

2.3 Asset-market and goods-market clearing

Government debt is the private sector's net safe financial asset. Asset-market clearing therefore requires

$$b = a_H + a_I = \kappa_H(r)x_H + \kappa_I(r)(x_I + S). \tag{16}$$

For given b and S, this equation determines the stationary safe rate. Desired safe wealth is strictly increasing in r over the admissible interval, so the equilibrium is locally unique whenever both groups remain in that interval.

Goods-market clearing follows from the same budgets. Writing the steps explicitly traces the absorption payment from government purchases into household income. In a stationary allocation, equations (5) and (14) imply the stationary budget constraints

$$0 = ra_H + x_H - c_H, \tag{17}$$

$$0 = ra_I + x_I + S - c_I. \tag{18}$$

Adding private consumption and government purchases gives

$$\begin{aligned} c_H + c_I + g &= r(a_H + a_I) + x_H + x_I + S + \bar{\tau} - rb - S \\ &= 1 + r(a_H + a_I - b). \end{aligned} \tag{19}$$

Consequently, asset-market clearing implies

$$c_H + c_I + g = 1. \tag{20}$$

The payment S cancels from the resource identity because it is a transfer from the government to private owners. Its incidence nevertheless matters for expenditure. Government purchases fall by the full payment, whereas the recipients consume only part of the associated increase in human wealth and save the remainder.

2.4 The stationary response to additional debt

The paper compares stationary equilibria indexed by the outstanding debt stock. Group H's permanent income is unchanged across these equilibria, while group I's income rises with the Treasury absorption payment. Differentiating the stationary wealth demands in (7) gives

$$da_H = \kappa_H'(r)x_H\,dr, \tag{21}$$

$$da_I = \kappa_I'(r)(x_I+S)\,dr+\kappa_I(r)\,dS. \tag{22}$$

Thus group H's stationary holdings change only through the safe rate. Group I's holdings also change because the absorption payment is a persistent source of income.

At a fixed safe rate, one additional unit of debt raises the supply of safe claims by one, while the associated absorption income raises stationary desired safe wealth by $\kappa_I(r)S_b(b,N)$. When the supply increase is larger, households attempt to exchange the excess claims for goods, generating positive aggregate-demand pressure. When the desired-wealth response is larger, they attempt to save more, generating negative aggregate-demand pressure. The stationary safe-debt margin is therefore

$$\frac{1}{\kappa_I(r)}-S_b(b,N). \tag{23}$$

The same comparison governs the equilibrium safe rate. Substituting $S(b,N)$ into (16) and differentiating yields

$$r_b=\frac{1-\kappa_I(r)S_b(b,N)}{\kappa_H'(r)x_H+\kappa_I'(r)(x_I+S(b,N))}. \tag{24}$$

The denominator is positive, and the numerator is $\kappa_I(r)$ times the margin in (23). When issuance creates positive stationary aggregate-demand pressure at the inherited rate, the equilibrium safe rate rises to increase desired wealth and reduce human wealth. When the fiscal drag dominates, the safe rate falls.

The portfolio adjustment is especially transparent at the turning point. There $r_b = 0$, and equations (21)–(22) imply

$$\frac{da_H}{db} = 0, \qquad \frac{da_I}{db} = \kappa_I(r) S_b(b,N) = 1. \tag{25}$$

At the peak, group H's desired holdings are unchanged and group I absorbs the entire marginal claim. Below the peak, $r_b > 0$ raises group H's desired holdings. Beyond the peak, $r_b < 0$ lowers them, so group I absorbs the new claim and part of group H's previous holdings.

Remark 2 (Initial incidence and the stationary allocation)**.** A newly issued bond may initially be distributed across both groups. Its immediate consumption effect then depends on the recipients' bond shares and marginal propensities to consume. The stationary comparison studied here is governed instead by the desired holdings $a_H = \kappa_H(r)x_H$ and $a_I = \kappa_I(r)(x_I + S)$. Portfolio adjustment removes the initial allocation from the stationary equilibrium condition. Both κ_H and κ_I determine aggregate desired wealth and the adjustment of the safe rate. Only κ_I enters the numerator of (24) because group I receives the persistent increase in absorption income. At the stationary turning point, that income induces group I to absorb the entire marginal claim.

Financial capacity works through the same accounting. Differentiating (16) with respect to N gives

$$r_N = -\frac{\kappa_I(r) S_N(b,N)}{\kappa_H'(r)x_H + \kappa_I'(r)(x_I + S(b,N))}. \tag{26}$$

Section 3 establishes that additional capacity lowers the absorption payment, $S_N < 0$. It therefore raises the equilibrium safe rate. At the fixed rate, the lower payment releases more government purchases than it subtracts from private consumption, so it also raises aggregate expenditure.

The macroeconomic turning point is

$$S_b(b^L, N) = \frac{1}{\kappa_I(r(b^L, N))}. \tag{27}$$

Section 3 now derives the absorption schedule and shows why its marginal cost rises with debt and falls with financial capacity. Section 4 combines that financial result with (27) to establish the existence and local shape of the safe-debt Laffer curve.

3 Producing Safe Public Debt

Section 2 took the Treasury absorption payment $S(b,N)$ as given. This section derives it. The economic event is a rollover. Treasury claims mature, the government places replacement claims, and financial institutions commit balance-sheet capacity to absorb the refinancing flow. The amount each active balance sheet must carry depends on how much debt comes due and how many potential providers participate. When participation is thin, a given refinancing need is concentrated on fewer balance sheets and becomes more costly to support.

The government compensates providers for this burden through the Treasury absorption premium. The premium is paid as claims roll over. In steady state, the sequence of future rollover payments can be represented as an annualized flow on the outstanding stock. This convention lets the macroeconomic block work with a stationary payment $S(b,N)$, while preserving the timing of the underlying mechanism: a change in the premium affects an outstanding claim when that claim is refinanced.

3.1 The rollover event and the Treasury absorption premium

The benchmark considers a unit-length rollover interval during which the full debt stock b must be refinanced. Section 5 replaces this benchmark with a finite rollover share. Let r be the safe return received by the ultimate holders of government debt, and let $r^G(b,N)$ be the government's financing rate. Their difference is

$$s(b,N) = r^G(b,N) - r. \tag{28}$$

The spread $s(b,N)$ is the Treasury absorption premium: the compensation per unit of debt paid for the balance-sheet support required at rollover. The government pays $r + s(b,N)$; the claim delivered to households earns the safe return r. The aggregate annualized payment is

$$S(b,N) = b\,s(b,N). \tag{29}$$

This payment is distributed to the households who own the financial firms supplying Treasury-intermediation capacity in Section 2.

The premium compensates financial firms for the consumption-equivalent opportunity cost of committing scarce balance-sheet capacity. This cost enters the firms' participation payoff below. The financial payment S is distributed to their owners as income, whose consumption and saving response was derived in Section 2. By carrying the refinancing exposure, providers deliver the safe payoff promised to households, and the associated opportunity cost is priced into the Treasury absorption premium.

3.2 Capacity, participation, and exposure per provider

Normalize the mass of potential support providers to one. A mass $\underline{\alpha} \in (0,1)$ is available in every rollover state. This inelastic base consists of support capacity whose participation does not respond within the rollover interval. The remaining providers participate when their private assessments make entry worthwhile. Let

$$\alpha \in [\underline{\alpha}, 1]$$

denote the total mass active in a particular state.

When all potential providers participate, the financial system can absorb N units of refinancing exposure without stress. If the active mass is α, available capacity is αN. Equivalently, each active provider can carry N units before incurring the additional balance-sheet burden modeled below. Because active providers share the refinancing flow, each carries b/α units of debt. Exposure beyond ordinary capacity per active provider is therefore

$$\max\left\{\frac{b}{\alpha} - N, 0\right\} = \frac{\max\{b - \alpha N, 0\}}{\alpha}. \tag{30}$$

The division by α is central. A refinancing gap of a given aggregate size is more demanding when it must be distributed across fewer balance sheets.

The function Γ maps this individual exposure into an annualized consumption-equivalent cost per unit of Treasury debt intermediated. It satisfies

$$\Gamma(0) = 0, \qquad \Gamma'(z) > 0 \text{ for } z > 0, \qquad \Gamma'_{+}(0) > 0, \qquad \Gamma''(z) \geq 0. \tag{31}$$

The cost is zero while ordinary capacity is sufficient. It rises once a provider must use its balance sheet more intensively, and convexity makes the burden increase more rapidly as individual exposure grows. In a state with active mass α, the per-unit support cost is consequently

$$\Gamma\left(\frac{\max\{b - \alpha N, 0\}}{\alpha}\right). \tag{32}$$

It is useful to express refinancing needs relative to full-system capacity:

$$q(b,N) = \frac{b}{N}. \tag{33}$$

A state is stressed when $b > \alpha N$, or equivalently when $\alpha < q(b,N)$. There are three regions. If $q \leq \underline{\alpha}$, the inelastic base can absorb the entire rollover and the support cost is zero in every state. If $\underline{\alpha} < q < 1$, the system is stressed only in states with sufficiently low participation. If $q \geq 1$, even full participation requires providers to operate beyond ordinary capacity. The debt stock at which stress first becomes possible is

$$\widehat{b}(N) = \underline{\alpha} N. \tag{34}$$

Additional capacity moves this threshold outward proportionally.

3.3 Participation and marginal-provider pricing

The elastic providers decide whether to participate after observing private signals about an aggregate condition θ. A larger θ indicates conditions under which other providers are less likely to enter. Provider i observes

$$x_i = \theta + \eta \varepsilon_i, \tag{35}$$

where $\eta > 0$ and the idiosyncratic noise has continuous distribution F. The monotone equilibrium has a cutoff x^*: providers with $x_i \leq x^*$ enter. Given the aggregate condition, the active mass is

$$\alpha(\theta; x^*) = \underline{\alpha} + (1 - \underline{\alpha}) F\left(\frac{x^* - \theta}{\eta}\right). \tag{36}$$

Worse aggregate conditions reduce participation and raise the exposure borne by each entrant.

An active provider receives the common premium s per unit of debt intermediated and bears the consumption-equivalent cost in (32). The provider at the cutoff is indifferent. Hence the premium equals the marginal provider's expected support cost across the participation states that it considers possible.

The flat-prior benchmark makes this expectation especially transparent. Conditional on observing the marginal signal $x_i = x^*$, define the fraction of elastic providers that enter as

$$F\left(\frac{x^*-\theta}{\eta}\right).$$

Under a flat prior over θ, the likelihood associated with the marginal signal is exactly offset by the Jacobian of this transformation. The implied participation fraction is uniform on $[0,1]$, and the total active mass is therefore uniform on $[\underline{\alpha},1]$. Appendix B gives the formal change of variables.

Lemma 1 (Marginal-provider pricing)**.** *Suppose elastic support providers use a monotone threshold strategy and have a flat prior over the aggregate condition. Conditional on the marginal signal, the active mass is uniformly distributed on* $[\underline{\alpha},1]$*. The Treasury absorption premium is independent of the cutoff and equals*

$$s(b,N) = \frac{1}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{1}\Gamma\left(\frac{\max\{b-\alpha N,0\}}{\alpha}\right)d\alpha. \tag{37}$$

The cutoff affects which providers enter in a realized rollover state. It does not affect the pricing distribution of the marginal provider under the flat-prior benchmark. The subsequent comparative statics therefore depend on debt and capacity through exposure and the set of stressed participation states, without requiring a separate selection among monotone cutoffs.

3.4 The premium schedule

The support cost is positive only when $\alpha < q(b,N)$. The pricing equation can therefore be written as

$$s(b,N) = \frac{1}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{\min\{q(b,N),1\}}\Gamma\left(\frac{b-\alpha N}{\alpha}\right)d\alpha, \tag{38}$$

with $s(b,N)=0$ when $q(b,N)\le\underline{\alpha}$. This expression averages the cost of balance-sheet support over precisely those participation states in which ordinary capacity is insufficient.

Consider first the interior stress region, $\underline{\alpha} < q(b,N) < 1$. Debt affects the premium by raising exposure in every state that is already stressed and by expanding the set of stressed states. Applying Leibniz's rule gives

$$s_b(b,N) = \frac{1}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N)}\frac{1}{\alpha}\Gamma'\left(\frac{b-\alpha N}{\alpha}\right)d\alpha > 0. \tag{39}$$

The moving upper limit makes no first-order contribution because exposure, and hence Γ, is exactly zero at $\alpha = q(b,N)$. At the margin, the newly stressed state begins with zero cost. The positive first derivative therefore comes from the intensive margin: more debt raises the exposure of every provider in every state that was already stressed.

Capacity has the opposite intensive effect:

$$s_N(b,N) = -\frac{1}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N)} \Gamma'\left(\frac{b-\alpha N}{\alpha}\right) d\alpha < 0. \tag{40}$$

Again the moving boundary contributes zero at first order. More capacity lowers exposure by one unit for each active provider in every stressed state.

The expansion of the stress region becomes visible in the second derivative. Differentiating (39) gives

$$s_{bb}(b,N) = \frac{1}{1-\underline{\alpha}}\left[\int_{\underline{\alpha}}^{q(b,N)} \frac{1}{\alpha^2}\Gamma''\left(\frac{b-\alpha N}{\alpha}\right) d\alpha + \frac{\Gamma'_+(0)}{b}\right] \geq 0. \tag{41}$$

The integral is the intensive-margin curvature: convex support costs make a given increase in debt more expensive when providers are already carrying large exposures. The boundary term is the extensive margin: more debt brings new low-participation states into the stress region, and $\Gamma'_+(0) > 0$ makes their costs rise immediately after entry. Either channel can steepen the premium schedule.

Capacity also flattens the response of the premium to debt. In the interior stress region,

$$s_{bN}(b,N) = -\frac{1}{1-\underline{\alpha}}\left[\int_{\underline{\alpha}}^{q(b,N)} \frac{1}{\alpha}\Gamma''\left(\frac{b-\alpha N}{\alpha}\right) d\alpha + \frac{\Gamma'_+(0)}{N}\right] < 0. \tag{42}$$

The first term lowers the sensitivity of cost within every stressed state. The second shifts the stress boundary inward. Thus greater capacity lowers both the level of the premium and its slope with respect to debt.

When $q(b,N) \geq 1$, every participation state is stressed and the upper integration limit is fixed at one. The boundary terms in (41) and (42) then disappear. Debt continues to raise the premium and capacity continues to lower it. The weak inequalities $s_{bb} \geq 0$ and $s_{bN} \leq 0$ remain; they are strict whenever Γ is strictly convex over a set of exposures with positive measure.

3.5 From the average premium to stock-wide marginal cost

The premium $s(b,N)$ is an average cost per unit of debt. The macroeconomic block requires the change in the government's total absorption payment when the debt stock rises. Since

$$S(b,N) = b\,s(b,N),$$

the stock-wide marginal cost is

$$S_b(b,N) = s(b,N) + b\,s_b(b,N). \tag{43}$$

The first term is the premium paid when the marginal unit rolls over. The second is the price effect on the refinancing stock: by increasing rollover pressure, the marginal unit raises the premium at which the previously outstanding debt will be refinanced as it matures. The term bs_b does not apply the new premium retroactively to unmatured bonds. It is the annualized steady-state cost of rolling the stock over at the higher premium in the future.

This distinction is quantitatively important. A modest change in the premium can have a large marginal fiscal effect when the stock that will eventually refinance at that premium is large. It is also the source of the paper's empirical stock multiplication: the estimated premium response is multiplied by the relevant privately held refinancing stock to obtain the repricing component of marginal cost.

The stock-wide marginal cost rises strictly with debt in the active stress region:

$$S_{bb}(b,N) = 2s_b(b,N) + b\,s_{bb}(b,N) > 0. \tag{44}$$

Even if Γ were locally linear within stressed states, the premium would rise with debt and the expanding refinancing base would make total marginal cost increase. Convexity of Γ and the entry of new stress states strengthen this result.

Capacity lowers both the total payment and its marginal response to debt:

$$S_N(b,N) = b\,s_N(b,N) < 0, \qquad S_{bN}(b,N) = s_N(b,N) + b\,s_{bN}(b,N) < 0. \tag{45}$$

The first inequality says that more capacity makes a given stock cheaper to support. The second says that it also makes additional issuance less costly at the margin.

Section 2 showed that the macroeconomic effect of debt is governed by the difference between the rate-equivalent contribution of a new safe claim and this stock-wide marginal cost:

$$\frac{1}{\kappa_I(r)} - S_b(b,N).$$

Before rollover stress begins, $S_b = 0$. Once debt exceeds $\widehat{b}(N)$, the premium becomes positive and the marginal absorption cost rises. The next section combines this increasing cost schedule with private wealth demand to derive the safe-debt Laffer curve and the corresponding turning point in the equilibrium safe rate.

4 Equilibrium and the Safe-Debt Laffer Curve

Section 2 derived the stationary non-Ricardian wealth effect of additional public debt and the fiscal drag created by Treasury absorption payments. Section 3 derived absorption payments from rollover pressure and financial-system capacity. This section combines the two blocks. At low debt, additional safe claims generate positive stationary aggregate-demand pressure. As marginal absorption cost rises, that pressure weakens and eventually reverses. The equilibrium safe rate reaches its maximum at the same stationary turning point.

The comparison turns on a marginal, rather than an average, cost. The government pays the premium $s(b,N)$ on the debt stock as it rolls over, so its total annualized absorption payment is $S(b,N) = bs(b,N)$. An additional unit of debt changes that payment by

$$S_b(b,N) = s(b,N) + bs_b(b,N).$$

As Section 3 emphasized, the second term is the future repricing of the refinancing stock. Unmatured claims retain their contracted terms until rollover; the higher premium enters as each claim is refinanced. The question in this section is whether the non-Ricardian wealth created by the new safe claim exceeds the fiscal drag generated by this additional payment.

4.1 The expenditure margin

Hold the safe rate and financial capacity fixed and compare stationary allocations. Group H's permanent income is unchanged, whereas group I receives the persistent increase in Treasury absorption income.

At the inherited rate, one additional unit of debt creates one unit of safe financial wealth. It also raises group I's desired stationary safe wealth by $\kappa_I(r)S_b(b,N)$. If the first quantity is larger, households attempt to exchange excess claims for goods and stationary aggregate-demand pressure rises. If the second is larger, desired saving rises by more than the supply of new claims and aggregate-demand pressure falls. The sign is therefore governed by

$$\frac{1}{\kappa_I(r)} - S_b(b,N). \tag{46}$$

The first term is the rate-equivalent non-Ricardian wealth benefit of the marginal safe claim after stationary portfolio adjustment. The second is the stock-wide marginal absorption cost.

Define the safe-debt margin

$$H(b,N) \equiv \frac{1}{\kappa_I(r(b,N))} - S_b(b,N). \tag{47}$$

Both terms are annual rates. Using (11), the first can be written as the net consumption effect of additional financial wealth divided by the share of persistent absorption income that group I saves. The empirical section estimates the second term and uses a benchmark for the first. A positive H means that additional debt generates positive stationary aggregate-demand pressure at the inherited rate; a negative H means that the fiscal drag is larger.

4.2 The equilibrium safe rate: an asset-market interpretation

The same comparison appears in the market for safe wealth. Recall the stationary asset-market condition:

$$b = \kappa_H(r)x_H + \kappa_I(r)(x_I + S(b,N)). \tag{48}$$

For given b and N, this equation determines the safe rate. Desired safe wealth is strictly increasing in r over the admissible interval because $\kappa_H'(r) > 0$ and $\kappa_I'(r) > 0$. Hence, whenever desired wealth passes through the outstanding debt stock inside that interval, the equilibrium is unique.

Differentiating (48) with respect to debt gives

$$r_b(b,N) = \frac{1 - \kappa_I(r)S_b(b,N)}{\kappa_H'(r)x_H + \kappa_I'(r)(x_I + S(b,N))}. \tag{49}$$

The denominator is positive. One more dollar of public debt adds one dollar to the supply of safe claims. It also raises persistent absorption income by S_b, generating $\kappa_I S_b$ dollars of desired stationary safe wealth. When the supply effect is larger, the safe rate rises to induce the private sector to hold the additional claim. When the income-induced demand for safe wealth is larger, the safe rate falls.

Multiplying the safe-debt margin in (47) by the positive term κ_I gives the numerator of (49). Stationary aggregate-demand pressure and the equilibrium safe-rate response therefore have the same sign. The expenditure formulation emphasizes the attempt to exchange excess financial wealth for goods; the asset-market formulation describes the same adjustment through the supply of and demand for safe claims.

4.3 The turning point

Before rollover stress begins, the inelastic base of financial capacity can absorb the refinancing need. The absorption premium and its marginal cost are then zero. Consequently, the safe-debt margin is positive and the safe rate rises with debt. Once b exceeds the stress threshold $\widehat{b}(N) = \underline{\alpha}N$, the premium becomes positive. Section 3 established that its stock-wide marginal cost rises strictly with debt:

$$S_{bb}(b,N) = 2s_b(b,N) + bs_{bb}(b,N) > 0.$$

The increasing marginal cost compresses the safe-debt margin. The Laffer peak is reached at the debt stock b^L satisfying

$$\frac{1}{\kappa_I(r(b^L,N))} = S_b(b^L,N) = s(b^L,N) + b^L s_b(b^L,N). \tag{50}$$

At this point, the stationary non-Ricardian wealth effect of the marginal claim is exactly offset by the fiscal drag from the additional absorption payment. Equivalently, the absorption income generated by marginal issuance induces exactly one additional unit of desired safe wealth. Stationary aggregate-demand pressure is zero and the new claim is absorbed without a change in the safe rate. Equation (25) shows how the initial broad incidence of the bond converges to this marginal stationary allocation.

Beyond the peak, public debt remains safe, but preserving that safety requires a larger fiscal and financial burden. Marginal issuance reduces government purchases and creates sufficiently persistent income for the owners of absorption capacity that their desired safe wealth rises by more than the supply of new claims. The equilibrium safe rate must then fall. The Laffer curve traces this reversal in the marginal contribution of safe debt to expenditure and the equilibrium safe rate.

Proposition 1 (Safe-debt Laffer curve)**.** *Suppose the stationary equilibrium in* (48) *exists and remains in the admissible interest-rate interval over the debt range under consideration. Let the economy enter the stress region at $\widehat{b}(N)$, and suppose $S_{bb}(b,N) > 0$ throughout that region. If there is a debt stock $\bar{b} > \widehat{b}(N)$ for which*

$$\kappa_I(r(\bar{b},N))S_b(\bar{b},N) > 1,$$

then there is a unique $b^L(N) \in (\widehat{b}(N), \bar{b})$ satisfying (50)*. Debt generates positive stationary aggregate-demand pressure and raises the equilibrium safe rate below $b^L(N)$, and reverses both effects above it. The equilibrium safe-rate schedule has a strict maximum at $b^L(N)$.*

The existence argument follows from continuity. At the stress threshold, $S_b = 0$, so the numerator of (49) equals one. By assumption, it is negative at $\bar{b}$, and hence it must cross zero. Strictly increasing marginal absorption cost makes the crossing unique. To see this directly, evaluate the derivative of the numerator at any candidate peak. Since $r_b = 0$ there,

$$\begin{aligned} &\frac{d}{db}\{1-\kappa_I(r(b,N))S_b(b,N)\}\Big|_{b=b^L} \\ &= -\kappa_I(r)S_{bb}(b^L,N) < 0. \end{aligned} \tag{51}$$

Every zero is therefore a crossing from the increasing branch to the decreasing branch. There cannot be a second crossing, because returning to the positive branch would require a zero with the opposite slope. Equation (51) also establishes that the peak is strict.

Figure 2 illustrates the mechanism using the linear-support-cost example. The left panel plots the increasing stock-wide marginal absorption cost against a fixed benchmark for the rate-equivalent wealth benefit. The right panel traces the corresponding equilibrium safe-rate schedule, holding its positive adjustment denominator fixed to isolate the sign change. The figure plots the safe-rate schedule rather than aggregate-demand pressure. Equations (46) and (49) show that the stationary turning points coincide.

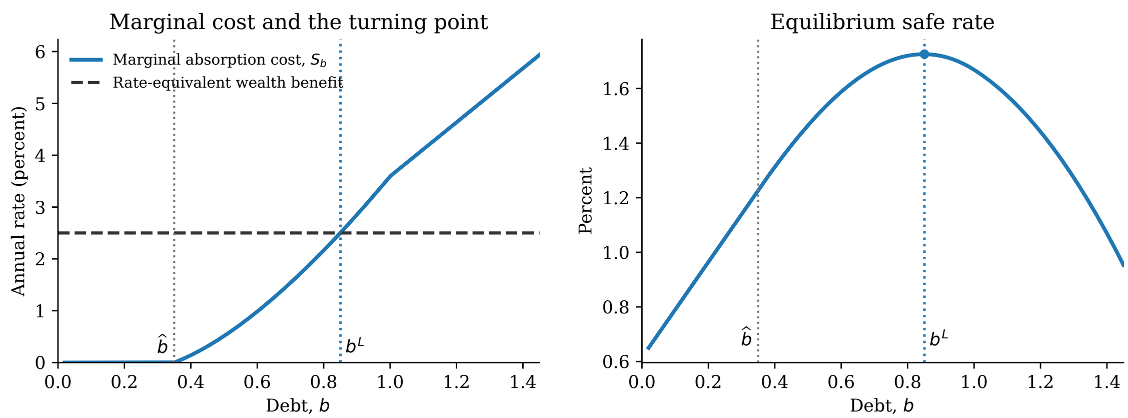

Figure 2: Marginal absorption cost and the equilibrium safe rate. The vertical dashed line marks the stress threshold $\widehat{b}$, and the dotted line marks the common stationary turning point b^L. Above $\widehat{b}$, stock-wide marginal absorption cost rises with debt. At b^L, it equals the rate-equivalent wealth benefit of additional safe debt, stationary aggregate-demand pressure is zero, and the safe-rate schedule in the right panel reaches its maximum. The positive denominator in (49) is held fixed to display the sign change.

The figure compares stationary equilibria indexed by b, with financial capacity held fixed. Rollover stress begins at $\widehat{b}$, but the safe rate continues to rise until marginal absorption cost reaches the rate-equivalent wealth benefit at b^L. At that same point, stationary aggregate-demand pressure changes sign.

4.4 Financial capacity

Financial capacity affects the equilibrium through the production schedule derived in Section 3. More capacity moves the stress threshold outward because $\widehat{b}(N) = \underline{\alpha}N$. Conditional on being in the stress region, it also lowers the absorption payment and flattens its marginal cost:

$$S_N(b,N) < 0, \qquad S_{bN}(b,N) < 0.$$

These are production-side results. They say that a given debt stock is cheaper to support when more balance sheets are available and that further issuance puts less pressure on the premium.

At a fixed debt stock, the equilibrium safe-rate effect is unambiguous. Differentiating (48) with respect to N gives

$$r_N(b,N) = -\frac{\kappa_I(r)S_N(b,N)}{\kappa_H'(r)x_H + \kappa_I'(r)(x_I + S(b,N))} > 0. \tag{52}$$

Greater capacity reduces absorption income. At the fixed rate, desired safe wealth falls while the debt stock is unchanged. The safe rate therefore rises to increase desired holdings. From the expenditure side, lower absorption payments release more government purchases than they subtract from private consumption.

Capacity also improves the safe-debt margin directly. Holding the safe rate fixed,

$$\left.\frac{\partial H(b,N)}{\partial N}\right|_r = -S_{bN}(b,N) > 0. \tag{53}$$

This is the nonlinear capacity interaction measured in the empirical section: when capacity is scarce, debt has a larger effect on the stock-wide marginal absorption cost.

The effect of capacity on the equilibrium location of the peak contains one additional term. Differentiating (50) yields

$$\frac{db^L}{dN} = \frac{-S_{bN}(b^L,N) - \dfrac{\kappa_I'(r)}{\kappa_I(r)^2} r_N(b^L,N)}{S_{bb}(b^L,N)}. \tag{54}$$

All terms are evaluated at the peak. The first term in the numerator is positive: more capacity flattens marginal absorption cost and pushes the peak outward. The second is negative: capacity raises the equilibrium safe rate, which increases group I's capitalization factor and lowers the rate-equivalent wealth benefit. Thus the peak moves outward when the direct flattening of marginal cost dominates this equilibrium effect. If the benefit-side benchmark is held fixed, as in the quantitative diagnostic and Figure 3, equation (54) reduces to

$$\frac{db^L}{dN} = -\frac{S_{bN}}{S_{bb}} > 0.$$

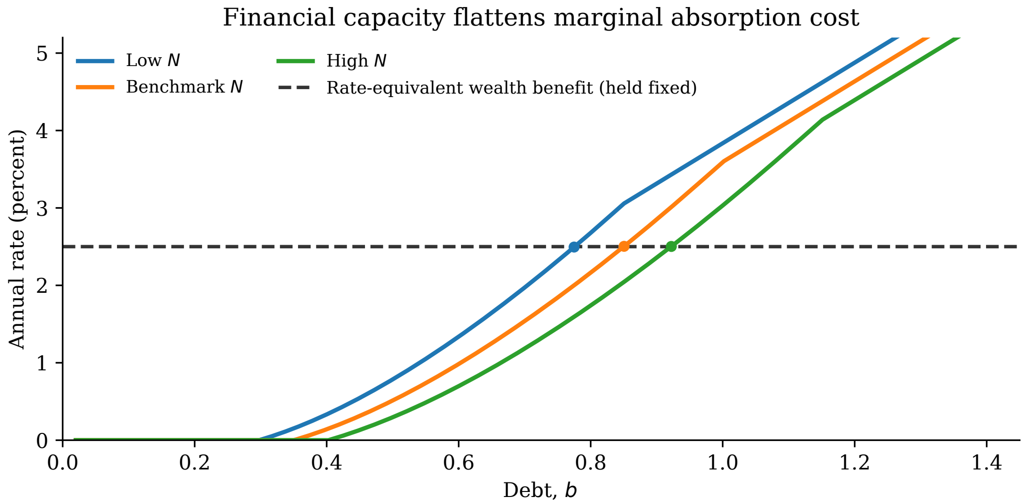

Figure 3: Financial capacity and marginal absorption cost. Higher N moves the stress threshold outward and flattens the marginal-cost schedule. The figure holds the benchmark rate-equivalent wealth benefit fixed, isolating the direct production-side effect through which greater capacity moves the peak to a larger debt stock.

The short-maturity benchmark used so far exposes the full debt stock to the rollover episode. Section 5 restores positive maturity. Only part of the stock then comes due during the episode, while the duration of the claims that remain outstanding amplifies the price response to the same rollover pressure. Maturity therefore changes the quantity and price dimensions of the mechanism without changing the equilibrium comparison between the rate-equivalent wealth benefit and stock-wide marginal absorption cost.

5 Quantity and Price Dimensions of Rollover Risk

The benchmark in Sections 3 and 4 exposes the full debt stock to a rollover episode. This normalization isolates the participation and capacity mechanism. Actual Treasury debt has positive maturity, so only part of the inherited stock reaches rollover during a given interval. At the same time, claims that remain outstanding have duration: their prices are more sensitive to changes in the premium when maturity is longer. Introducing maturity is therefore necessary for the empirical mapping: it separates the quantity reaching rollover from the price amplification measured by the real-duration premium.

These are two dimensions of the same rollover mechanism. Maturity determines the quantity of debt that requires refinancing support. It also determines how strongly a change in rollover compensation affects the price of outstanding claims. The second channel amplifies the first in prices. The total Treasury absorption premium must therefore be solved jointly with the duration of the claim.

The analysis first replaces full-stock rollover with the quantity of inherited debt that matures during the episode. It then derives the duration-price component and combines the two components in a fixed point. Throughout, derivatives of the premium hold the safe rate r fixed. Once the premium and the associated absorption payment have been determined, the asset-market condition in Section 4 gives the equilibrium adjustment in r.

5.1 The quantity reaching rollover

Let $\delta > 0$ be the Poisson maturity intensity of each unit of debt. Average maturity is $1/\delta$, so a larger δ means shorter maturity. Define a rollover episode as a unit-length interval and record whether each unit of the inherited stock matures during that interval. The probability that it does is

$$\mu(\delta) = 1 - e^{-\delta}, \qquad \mu'(\delta) = e^{-\delta} > 0. \tag{55}$$

The quantity of inherited debt exposed to rollover at least once during the episode is therefore $\mu(\delta)b$. Reissued claims enter the next episode. This timing counts each unit of the inherited stock once and makes the full-stock benchmark the limit $\mu(\delta) \to 1$.

The maturity intensity affects rollover pressure through this exposed quantity. The debt-capacity ratio relevant for the episode is

$$q(b, N, \delta) = \frac{\mu(\delta)b}{N}. \tag{56}$$

For a participation state with active mass α, active capacity is αN. When $\mu(\delta)b > \alpha N$, exposure per active provider is

$$\frac{\mu(\delta)b - \alpha N}{\alpha}.$$

Thus positive maturity changes the financial block of Section 3 only by replacing the full debt stock with the quantity that reaches rollover.

Let $s^R(b, N, \delta)$ denote the direct rollover component of the Treasury absorption premium. Marginal-provider pricing gives

$$s^R(b, N, \delta) = \frac{1}{1 - \underline{\alpha}} \int_{\underline{\alpha}}^{\min\{q(b,N,\delta),1\}} \Gamma\left(\frac{\mu(\delta)b - \alpha N}{\alpha}\right) d\alpha, \tag{57}$$

with $s^R = 0$ when $q \leq \underline{\alpha}$. The superscript R will also identify the cash-Treasury component in the empirical section.

Consider the interior stress region, $\underline{\alpha} < q(b, N, \delta) < 1$. The first derivatives follow directly from the derivatives in Section 3:

$$s_b^R(b, N, \delta) = \frac{\mu(\delta)}{1 - \underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha} \Gamma'\left(\frac{\mu(\delta)b - \alpha N}{\alpha}\right) d\alpha > 0, \tag{58}$$

$$s_N^R(b, N, \delta) = -\frac{1}{1 - \underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \Gamma'\left(\frac{\mu(\delta)b - \alpha N}{\alpha}\right) d\alpha < 0. \tag{59}$$

More debt raises exposure in every stressed state. More capacity reduces it. As before, the changing stress boundary contributes nothing to these first derivatives because the support cost is zero in the state that has just entered or left the stress region.

Shorter maturity also raises the direct premium:

$$s_{\delta}^{R}(b,N,\delta) = \frac{\mu'(\delta)b}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right) d\alpha > 0. \tag{60}$$

Debt and rollover intensity enter s^R through the same exposed quantity. In particular,

$$s_{\delta}^{R}(b,N,\delta) = \frac{\mu'(\delta)b}{\mu(\delta)} s_{b}^{R}(b,N,\delta). \tag{61}$$

This identity isolates the quantity channel. A larger δ raises the share of debt reaching maturity, exactly as an increase in b raises the exposed stock. Appendix E shows in addition that $s_{\delta N}^{R} < 0$: financial capacity weakens the incremental rollover pressure created by shorter maturity.

5.2 Duration and the price of rollover exposure

The units of debt that do not mature during the episode remain outstanding. The government pays the higher funding premium when these claims roll over, while their current market prices incorporate the expected future rollover terms. This price response is stronger for a longer-duration claim.

To make the duration calculation explicit, consider a unit flow paid while a claim remains outstanding. With Poisson maturity intensity δ and all-in yield $r+s(b,N,\delta)$, its price is

$$\frac{1}{r+s(b,N,\delta)+\delta}. \tag{62}$$

The price sensitivity to a change in the total premium is

$$-\frac{\partial}{\partial s}\left[\frac{1}{r+s+\delta}\right] = \frac{1}{(r+s+\delta)^2}. \tag{63}$$

The squared-duration term therefore has a direct pricing interpretation. Lower δ lengthens the claim and magnifies the price movement generated by a given change in rollover compensation. A higher total premium raises the discount rate and shortens the claim, attenuating its own price effect.

Let $\vartheta(N) > 0$ scale the compensation required for this price exposure, with $\vartheta_N(N) < 0$. A financial system with greater capacity can absorb a given bond-price movement at lower cost. The duration-price component is

$$s^{D}(b,N,\delta) = \frac{\vartheta(N)}{[r+s(b,N,\delta)+\delta]^2} s^{R}(b,N,\delta). \tag{64}$$

The direct rollover premium s^R measures the underlying pressure. The price sensitivity in (63) determines how strongly that pressure moves the value of a long claim, and $\vartheta(N)$ prices the resulting balance-sheet exposure. In particular, s^D vanishes when rollover pressure vanishes. Duration changes the price consequences of rollover risk; it leaves the underlying source of risk unchanged.

The total Treasury absorption premium is

$$s(b,N,\delta) = s^R(b,N,\delta) + s^D(b,N,\delta). \tag{65}$$

Substituting (64) yields the fixed point

$$s(b,N,\delta) = s^R(b,N,\delta)\left[1 + \frac{\vartheta(N)}{[r + s(b,N,\delta) + \delta]^2}\right]. \tag{66}$$

For a given s, the term in brackets is the price amplification of the direct rollover premium. Because s itself enters the denominator, the amplification is endogenous. An increase in rollover pressure raises the total premium, but the higher all-in yield shortens the claim and moderates the initial price response.

The fixed point has a unique nonnegative solution. At $s = 0$, its right-hand side is nonnegative. As s rises, the right-hand side decreases monotonically toward the finite value s^R, while the left-hand side increases without bound. The two sides therefore cross exactly once.

5.3 Comparative statics of the total premium

The fixed point preserves the debt and capacity signs of the direct rollover component. Differentiating with respect to debt gives

$$s_b = \frac{\left[1 + \dfrac{\vartheta(N)}{(r+s+\delta)^2}\right] s_b^R}{1 + \dfrac{2s^R\vartheta(N)}{(r+s+\delta)^3}} > 0. \tag{67}$$

The numerator is the direct response of the rollover premium, including its price amplification. The denominator captures the stabilizing feedback from endogenous duration: the higher total premium shortens the claim and reduces its price sensitivity.

Financial capacity lowers the total premium through both parts of the mechanism:

$$s_N = \frac{\left[1 + \dfrac{\vartheta(N)}{(r+s+\delta)^2}\right] s_N^R + \dfrac{s^R\vartheta_N(N)}{(r+s+\delta)^2}}{1 + \dfrac{2s^R\vartheta(N)}{(r+s+\delta)^3}} < 0. \tag{68}$$

The first term says that capacity lowers direct rollover pressure. The second says that it lowers the compensation required for the associated bond-price exposure. Both effects are divided by the same endogenous-duration feedback.

The effect of rollover intensity contains opposing forces:

$$s_\delta = \frac{\left[1 + \dfrac{\vartheta(N)}{(r+s+\delta)^2}\right] s_\delta^R - \dfrac{2s^R\vartheta(N)}{(r+s+\delta)^3}}{1 + \dfrac{2s^R\vartheta(N)}{(r+s+\delta)^3}}. \tag{69}$$

A higher δ exposes more debt to rollover, producing the positive first term. It also shortens the claim and reduces price amplification, producing the negative second term. The total premium can therefore rise or fall as maturity shortens even though the direct rollover component always rises.

Proposition 2 (Positive maturity and rollover-price amplification)**.** *Fix a finite maturity intensity δ and hold the safe rate fixed. In the active stress region, the direct rollover premium satisfies $s_b^R > 0$, $s_N^R < 0$, and $s_\delta^R > 0$. The total premium is the unique solution to (66) and satisfies $s_b > 0$ and $s_N < 0$. Its response to rollover intensity is given by (69). As $\delta \to \infty$, the exposed share approaches one, the duration-price component converges to zero, and the short-maturity premium of Section 3 is recovered.*

The limiting result connects the two versions of the financial block. As $\delta \to \infty$, every unit of the inherited debt stock reaches rollover during the episode, while the price sensitivity in (63) vanishes. The benchmark premium is therefore the full-stock, zero-duration limit of the positive-maturity model.

5.4 The positive-maturity safe-debt margin

For a fixed maturity structure, the government's annualized Treasury absorption payment is

$$S(b,N,\delta) = b\,s(b,N,\delta). \tag{70}$$

As in the short-maturity benchmark, the fiscal object relevant for marginal issuance is the change in the payment on the entire refinancing stock:

$$S_b(b,N,\delta) = s(b,N,\delta) + b\,s_b(b,N,\delta). \tag{71}$$

The first term is the total premium associated with the marginal unit. The second is the increase in the terms at which the outstanding stock will be refinanced as it rolls over, including the duration-price response generated by the same rollover pressure.

Using $s = s^R + s^D$, stock-wide marginal cost separates into

$$S_b = \left(s^R + b s_b^R\right) + \left(s^D + b s_b^D\right). \tag{72}$$

The first bracket is the direct rollover component: its premium level plus the repricing of the refinancing stock. The second is the duration-price component: its level plus the associated stock repricing. This decomposition is the bridge to the empirical section. The cash-Treasury spread is the counterpart of s^R, while the real-duration premium is the counterpart of s^D. Supply sensitivities of the two measured premia provide the corresponding repricing terms.

Positive maturity changes the components of marginal absorption cost, while the equilibrium comparison from Section 4 remains the same. The safe-debt margin is

$$H(b,N,\delta) = \frac{1}{\kappa_I(r(b,N,\delta))} - S_b(b,N,\delta). \tag{73}$$

At a fixed maturity intensity, the Laffer peak satisfies

$$\frac{1}{\kappa_I(r(b^L,N,\delta))} = s(b^L,N,\delta) + b^L s_b(b^L,N,\delta). \tag{74}$$

Thus maturity affects the location of the peak by changing both the amount of debt exposed to rollover and the price amplification embedded in the total marginal cost. The benefit side remains the rate-equivalent wealth effect derived in Sections 2 and 4.

5.5 Maturity choice as an extension

The empirical analysis evaluates (72) at the observed maturity structure. Appendix E studies the related debt-management problem at a given safe rate. For fixed b, minimizing total funding cost $b[r+s(b,N,\delta)]$ is equivalent to choosing δ to minimize $s(b,N,\delta)$. At an interior solution, the first-order condition is

$$\left[1+\frac{\vartheta(N)}{(r+s+\delta)^2}\right] s^R_\delta = \frac{2s^R\vartheta(N)}{(r+s+\delta)^3}. \tag{75}$$

The left side is the amplified increase in rollover cost caused by shorter maturity. The right side is the reduction in duration-price amplification. The optimum balances these two uses of maturity.

The capacity comparative static is conditional because capacity affects both sides of this balance. If the cross-derivative of the total premium satisfies $s_{\delta N} < 0$ at an interior optimum and the second-order condition holds, then

$$\frac{d\delta^*(b,N)}{dN} > 0.$$

Greater financial capacity then supports faster rollover and shorter maturity; scarce capacity tilts debt management toward a longer maturity structure. The appendix derives the full cross-derivative and the condition. The empirical section takes the observed maturity distribution as given and measures the two components of the resulting stock-wide marginal cost.

6 Measuring the U.S. Marginal Cost of Safe-Debt Production

This section brings the positive-maturity model to U.S. Treasury data. The object of interest is the cost of placing one more unit of debt when the terms required to absorb Treasury claims can change for the stock that reaches rollover. In the notation of Section 5, this stock-wide marginal cost is

$$S_b = s + b\frac{\partial s}{\partial b}.$$

The premium on the marginal claim contributes the first term. The repricing of the refinancing stock contributes the second. This distinction matters quantitatively: even a modest change in the premium per dollar can have a large marginal effect when it changes the terms on a large stock of debt as that stock rolls over.

I construct the empirical counterpart of both terms for privately held marketable Treasury notes and bonds. The exercise compares 2015Q1 with 2026Q1 and also produces a quarterly series over the intervening period. It has four inputs. Two measure the level of the Treasury absorption premium: a cash-Treasury component and a real-duration component. Two measure how those premia respond to private Treasury supply and therefore reprice the outstanding private stock. The cash-Treasury response is estimated from variation in Federal Reserve holdings across maturity buckets. Its benchmark specification allows that response to steepen when dealer-parent balance-sheet capacity is scarce. The real-duration response is estimated from changes in real term premia around Federal Reserve purchase and runoff announcements.

Equation (72) gives the theoretical decomposition. To implement it in the data, write the measured premium as

$$s_t = s_t^R + s_t^D, \tag{76}$$

where (R) denotes the direct rollover component and (D) the duration-price component. Let $\mathscr{M}_t^R$ and $\mathscr{M}_t^D$ denote the empirical counterparts of bs_b^R and bs_b^D. Substituting those estimates into (72) gives

$$S_{b,t} = s_t^R + s_t^D + \mathscr{M}_t^R + \mathscr{M}_t^D. \tag{77}$$

Every entry in this expression is measured in annualized basis points per dollar of additional privately held debt. The remainder of the section constructs the four entries in the order in which they appear, then adds them to obtain $S_{b,t}$.

6.1 The two premium levels

The cash-Treasury component, s_t^R, is the Treasury par yield minus the maturity-matched par swap rate. An interest-rate swap supplies duration exposure without requiring an investor to finance and warehouse a cash Treasury security. The Treasury–swap spread therefore isolates pricing that is specific to holding the cash claim. The measure is signed. A negative spread means that the cash-specific benefits of Treasuries, including their liquidity and convenience services, outweigh the balance-sheet cost of absorbing them; a rise in the spread means that this net advantage has diminished.

I construct a debt-weighted spread from the two-, three-, five-, seven-, ten-, and thirty-year maturity buckets. In each bucket, the swap-spread series follows a continuous chain of Treasury benchmarks and interest-rate swaps; the seven-year series is constructed from Treasury and swap par yields on a LIBOR-consistent basis. The weights are the shares of each bucket in privately held Treasury notes and bonds. The resulting cash-Treasury premium rises from −18.2 basis points in 2015Q1 to −0.5 basis points in 2026Q1.

The duration-price component, s_t^D, is the D'Amico–Kim–Wei real term premium (D'Amico, Kim, and Wei, 2018). DKW jointly prices nominal and inflation-indexed Treasuries and separates compensation for real-duration risk from inflation-risk compensation. I interpolate or extrapolate its five- and ten-year estimates to the six benchmark maturities and use the same private-debt weights as for the cash measure. This premium is the empirical counterpart of the price amplification in Section 5: when the terms expected at rollover change, a longer-duration claim experiences a larger current price movement. The debt-weighted real term premium rises from 8.4 basis points in 2015Q1 to 38.9 basis points in 2026Q1.

Table 1 reports the two levels and their sum. Together they rise by 48.2 basis points. This is the part of the increase in marginal cost that is visible directly in premium levels, before accounting for the repricing of the private stock.

Table 1: Measured Treasury absorption premia

Date	Cash-Treasury s_t^R	Real duration s_t^D	Total s_t
2015Q1	−18.2	8.4	−9.9
2026Q1	−0.5	38.9	38.3
Change	17.7	30.5	48.2

Notes: Entries are annualized basis points. Cash-Treasury premia are debt-weighted Treasury par yields minus maturity-matched par swap rates. Real-duration premia are debt-weighted DKW real term premia. Totals and changes use unrounded components, so displayed entries need not add exactly.

6.2 Cash-Treasury repricing

The next step estimates how the cash-Treasury premium changes when the private market must absorb more debt. The identifying variation comes from the Federal Reserve's holdings across the six maturity buckets. Following D'Amico and King (2013), at each horizon h I estimate the local projection

$$\Delta_h s^R_{m,t} = \phi_m + \psi_t + \beta_h \Delta_h share_{m,t} + \varepsilon_{m,t}, \tag{78}$$

where m indexes maturity buckets, t is the starting month, and

$$\Delta_h x_{m,t} = x_{m,t+h} - x_{m,t}.$$

The variable

$$share_{m,t} = 100 \times \frac{SOMA_{m,t}}{out_{m,t}}$$

is the percentage of Treasury debt in bucket m held in the Federal Reserve's System Open Market Account. Thus β_h is measured in basis points of spread per percentage-point increase in the SOMA share.

The two fixed effects give the regression a useful interpretation. Maturity-bucket effects, ϕ_m, absorb persistent differences in pricing and Federal Reserve holdings across tenors. Start-month effects, ψ_t, absorb aggregate news and curve-wide repricing between t and $t+h$. The coefficient is therefore identified from the relative change in Federal Reserve absorption and the relative change in the cash-Treasury spread across maturity buckets that share the same starting month. The sample excludes the first half of 2020, when emergency market interventions and market dysfunction make it difficult to interpret the persistent supply response measured here.

A negative β_h means that an increase in the SOMA share lowers the Treasury–swap spread in the affected maturity segment. Equivalently, moving debt from the Federal Reserve to private portfolios raises the spread. Converting this slope into the stock-repricing term in equation (77) requires three steps.

First, a one-percentage-point increase in $share_{m,t}$ removes one percent of the bucket's outstanding amount from private portfolios. Second, let

$$\varpi_{m,t} = \frac{SOMA_{m,t}}{out_{m,t}}.$$

The private share of the bucket is $1-\varpi_{m,t}$, so moving an amount equal to the bucket's entire current private float into private hands corresponds to a $100(1-\varpi_{m,t})$-percentage-point reduction in the SOMA share. Third, let $b_{m,t}$ be private float in bucket m, $b_t = \sum_m b_{m,t}$, and $w_{m,t} = b_{m,t}/b_t$. Weighting by $w_{m,t}$ converts the maturity-specific responses into the response for one dollar distributed in proportion to the current private maturity structure. These steps give

$$\mathscr{M}^R_{h,t} = -100\,\beta_h \sum_m w_{m,t}(1-\varpi_{m,t}). \tag{79}$$

The minus sign changes an increase in Federal Reserve holdings into an increase in private supply. The weighting sum is 0.781 in 2015Q1 and 0.837 in 2026Q1 in the Monthly Statement of the Public Debt notes-and-bonds universe.

I average the estimates over horizons $h = 12,\ldots,24$ months. This window asks whether a change in private supply produces persistent repricing rather than a short-lived announcement response. Adjacent observations have heavily overlapping outcome windows, so the reported inference uses Driscoll–Kraay/Newey–West standard errors with 24 monthly lags. In a linear specification beginning in 2015, the average slope is approximately -0.27. Substitution into equation (79) gives a 2026Q1 repricing term of

$$-100\times(-0.27)\times 0.837 \simeq 23 \quad \text{basis points.}$$

The corresponding estimates are 40 basis points for the sample beginning in 2010 and 31 basis points for the IV sensitivity in Table 2. These linear estimates will be useful for decomposing the change over time. The benchmark level allows the slope to depend on financial capacity.

Table 2: Linear cash-Treasury repricing estimates

Specification	Sample	Average β_h, $h = 12$–24	$\mathscr{M}^R_{2026Q1}$
OLS, no capacity interaction	2015+	−0.27	23
OLS, no capacity interaction	2010+	−0.48	40
IV sensitivity, no capacity interaction	2015+	−0.37	31

Notes: Coefficients are basis points per percentage point of the SOMA share within a maturity bucket; repricing terms are basis points. The labels 2015+ and 2010+ denote sample starting years. Inference allows for serial dependence across overlapping starting months. The IV row instruments the change in the SOMA share with lagged SOMA share and trailing issuance, an on-the-run proxy for recently issued securities. Appendix F gives the specification and inference details.

6.3 Financial capacity and the nonlinear benchmark

Equation (78) imposes the same supply response in every state of the financial system. The model instead predicts $s^R_{bN} < 0$: the premium rises more steeply with private supply when absorption capacity is scarce. I measure this capacity at the holding companies that own Treasury dealers, referred to below as dealer parents.

Let k denote the applicable leverage-ratio requirement. If dealer parent i has Tier 1 capital $T1_{i,t}$ and total leverage exposure $LE_{i,t}$, the regulatory requirement is

$$\frac{T1_{i,t}}{LE_{i,t}} \geq k. \tag{80}$$

Holding capital fixed, $T1_{i,t}/k$ is the largest total leverage exposure consistent with that requirement. The difference $T1_{i,t}/k - LE_{i,t}$ is therefore unused leverage-exposure room. A dealer parent can absorb Treasuries using this room, and it can continue to hold the Treasury inventory already on its balance sheet. I measure total potential Treasury absorption as

$$N_t(k) = \sum_i \left(\frac{T1_{i,t}}{k} - LE_{i,t} \right) + \sum_i UST_{i,t}, \tag{81}$$

where $UST_{i,t}$ is Treasury inventory already held by dealer parent i. This common-rule index summarizes the balance-sheet state that governs the cost of absorbing Treasury positions. It combines the inventory already supported by dealer parents with their unused leverage-exposure room.

The benchmark uses $k = 0.05$, the historical holding-company requirement for affected U.S. global systemically important bank parents. A tighter mapping uses $k = 0.06$, the historical well-capitalized requirement for their insured depository subsidiaries. For each mapping, I standardize $N_t(k)$ separately relative to its 2015–2024 mean and standard deviation. Denote the standardized series by $\widetilde{N}_t$, with higher values indicating more capacity. Because the regression uses the standardized series, a constant rescaling of the dollar measure cannot affect the estimates. The two regulatory mappings matter because they generate somewhat different capacity paths, especially in low-capacity states.

I estimate

$$\Delta_h s^R_{m,t} = \phi_m + \psi_t + \beta_h \Delta_h share_{m,t} + \gamma_h \left(\Delta_h share_{m,t} \times \widetilde{N}_t \right) + \varepsilon_{m,t}. \tag{82}$$

At capacity state $\widetilde{N}_t$, the response of the spread to the SOMA share is

$$\beta_h + \gamma_h \widetilde{N}_t. \tag{83}$$

Since an increase in the SOMA share reduces private supply, the model predicts $\beta_h < 0$ over the relevant states and $\gamma_h > 0$. Greater capacity then makes the negative slope less steep; scarce capacity makes it more steep. This is the empirical counterpart of the model's negative cross-partial between private debt and capacity.

Under the benchmark 5 percent mapping, the averages over $h = 12, \ldots, 24$ are $\beta_h = -0.24$ and $\gamma_h = 0.19$. The overlap-robust t-statistics on the interaction range from 2.4 to 2.9, with a mean of 2.6. Thus the data reject a state-invariant cash-Treasury supply response over this horizon window. The 2026Q1 capacity observation is approximately -2.0 standard deviations. Evaluating equation (83) at that state gives

$$-0.24 + 0.19 \times (-2.0) \simeq -0.62.$$

Using the 2026Q1 private-float weight, the stock-repricing term is therefore

$$\mathscr{M}^R_{h,t} = -100 \left(\beta_h + \gamma_h \widetilde{N}_t \right) \sum_m w_{m,t} (1 - \varpi_{m,t}) \simeq 52 \text{ basis points.} \tag{84}$$

The last number follows from $-100 \times (-0.62) \times 0.837$. Under the tight 6 percent mapping, the 2026Q1 state is -3.1 standard deviations and the corresponding repricing term is 72 basis points. The benchmark 2026Q1 capacity state lies below the 2015–2024 estimation range, so this endpoint uses the estimated nonlinear schedule at a state more adverse than those used to estimate it. The full coefficient paths and IV sensitivities are reported in Appendix F.

Table 3: Capacity-dependent cash-Treasury repricing

Mapping	Average β_h	Average γ_h	Overlap-robust t_{γ_h}	$\widetilde{N}_{2026Q1}$	$\mathscr{M}^R_{2026Q1}$
5 percent (benchmark)	−0.24	0.19	2.4–2.9	−2.0	52
6 percent (tight)	−0.34	0.17	1.0–1.6	−3.1	72

Notes: Coefficients are averages over $h = 12$–24. The t-statistics use Driscoll–Kraay/Newey–West standard errors with 24 monthly lags. Capacity is standardized separately for each mapping over 2015–2024. Repricing terms are basis points and apply equation (84) at the reported capacity state.

6.4 Real-duration repricing

The duration-price term requires a different source of variation. I use twelve Federal Reserve announcements that changed expected Treasury purchases or balance-sheet runoff. For each announcement e, I measure the two-trading-day change in the DKW real term premium and estimate

$$\Delta s^D_e = c_D + \beta_D Amount_e + u_e. \tag{85}$$

The variable $Amount_e$ is the announced change in expected cumulative Treasury purchases, in billions of dollars. A positive value represents greater Federal Reserve purchases and hence less debt left for private investors to absorb. Tapering and runoff announcements enter with negative values. Appendix F lists the events and explains the sizing of fixed and open-ended programs.

For the debt-weighted, average-maturity DKW real term premium, the estimate with an intercept is

$$\beta_D \simeq -0.0055$$

basis points per billion dollars. Greater Federal Reserve purchases lower the real term premium, so greater private supply raises it. Let b^{priv}_t denote the privately held stock of Treasury notes and bonds, measured in billions of dollars. Scaling the announcement response to that stock and reversing the sign gives

$$\mathscr{M}^D_t = -\beta_D b^{priv}_t. \tag{86}$$

Private coupon float rises from \$7.525 trillion in 2015Q1 to \$17.642 trillion in 2026Q1. The corresponding calculations are

$$0.0055 \times 7{,}525 \simeq 41.4 \quad \text{and} \quad 0.0055 \times 17{,}642 \simeq 97.1$$

basis points. This increase is driven by the growth of the private stock to which the estimated duration response applies.

The benchmark uses the debt-weighted, average-maturity real term premium and measures purchase announcements in dollars. Two alternatives change one choice at a time. Using the ten-year DKW premium produces a 2026Q1 duration repricing term of 116 basis points. Dividing announced purchase amounts by contemporaneous private float produces a 53-basis-point term.

The first alternative places more weight on long-duration price sensitivity; the second treats a purchase program of a given size as smaller when the market has more debt outstanding.

The event study contains only twelve observations, and several open-ended programs must be converted into expected cumulative purchases. The uncertainty is correspondingly substantial. With HC3 heteroskedasticity-robust inference and a small-sample t critical value, the 95 percent interval for β_D is $[-0.0109, -0.0001]$. At the 2026Q1 stock, this maps into a 1–193 basis point interval for $\mathscr{M}^D_{2026Q1}$. Leave-one-event-out estimates imply 74–116 basis points, while excluding the two March 2020 announcements gives 151 basis points. These ranges describe uncertainty in the duration component; they do not alter the cash-Treasury estimates.

6.5 Assembling the marginal cost

Equation (77) can now be evaluated directly. The benchmark combines the capacity-dependent 2015+ OLS estimate under the 5 percent mapping with the average-maturity duration estimate. In 2015Q1, the cash-Treasury premium is -18.2 basis points, the real-duration premium is 8.4, cash-Treasury repricing is approximately 48.9, and duration repricing is 41.4. Their sum is 80.4 basis points. In 2026Q1, the corresponding four entries are -0.5, 38.9, 51.7, and 97.1, for a total of 187.1 basis points. The stock-wide marginal cost therefore rises by approximately 107 basis points and more than doubles over the period.

The nonlinear cash-Treasury term is large at both endpoints because both are low-capacity states under the benchmark capacity measure. This fact explains why the linear specification gives almost the same increase in marginal cost even though it gives lower levels. In the 2015+ linear specification, $S_{b,t}$ rises from 52.9 to 158.3 basis points, an increase of 105.4 basis points. The 2010+ and IV linear specifications produce increases of 106.6 and 106.0 basis points.

The linear specification is especially useful for decomposing the change because its cash-Treasury slope does not vary with the capacity state. Measured premium levels account for 48.2 basis points of the 105.4-basis-point increase. Growth in the private coupon stock raises duration repricing by 55.7 basis points. The change in the cash-Treasury private-float weights raises cash-Treasury repricing by a further 1.5 basis points. The first two forces account for nearly all of the rise: Treasury absorption premia become less favorable, and the duration response is applied to a much larger private debt stock. The nonlinear estimate adds the economically important result that the level of cash-Treasury repricing is substantially greater when dealer-parent capacity is scarce.

Table 4: Marginal cost of producing safe Treasury claims

Date and specification	s_t^R	s_t^D	$\mathscr{M}_t^R$	$\mathscr{M}_t^D$	$S_{b,t}$	H_t
2015Q1, nonlinear OLS benchmark	−18.2	8.4	48.9	41.4	80.4	249.6
2026Q1, nonlinear OLS benchmark	−0.5	38.9	51.7	97.1	187.1	142.9
2015Q1, 2015+ linear OLS	−18.2	8.4	21.3	41.4	52.9	277.1
2026Q1, 2015+ linear OLS	−0.5	38.9	22.9	97.1	158.3	171.7
2015Q1, 2010+ linear OLS	−18.2	8.4	37.4	41.4	69.0	261.0
2026Q1, 2010+ linear OLS	−0.5	38.9	40.1	97.1	175.5	154.5
2015Q1, linear IV sensitivity	−18.2	8.4	29.3	41.4	60.8	269.2
2026Q1, linear IV sensitivity	−0.5	38.9	31.4	97.1	166.8	163.2

Notes: Entries are annualized basis points. The nonlinear benchmark uses the OLS capacity interaction under the 5 percent mapping. Linear rows use the indicated cash-Treasury specification. Repricing terms apply the date-specific mappings in equations (79), (84), and (86) to the MSPD notes-and-bonds universe. Sums use unrounded components. The last column uses the 330-basis-point wealth-benefit benchmark defined below.

Figure 1, introduced in the introduction, plots the quarterly benchmark and linear series. The benchmark series is the paper's preferred estimate of the level of marginal Treasury absorption cost. The linear series shows how much of the movement comes from the observed premium levels and debt stocks without the capacity interaction.

Table 5 varies the duration term while retaining the benchmark capacity-dependent cash-Treasury estimate. The share-normalized duration specification gives a 2026Q1 marginal cost of 143 basis points. The average-maturity benchmark gives 187 basis points, and the ten-year duration specification gives 206 basis points. Under the tight 6 percent capacity mapping, the average-maturity calculation instead gives approximately 207 basis points. All three duration constructions imply a high marginal Treasury absorption cost relative to its level a decade earlier.

Table 5: Duration sensitivity of the 2026Q1 marginal cost

Benchmark nonlinear specification	$S_{b,2026Q1}$	Safe-debt margin H_{2026Q1}
Share-normalized duration (53 bp)	143	187
Average-maturity duration (97 bp)	187	143
Ten-year duration (116 bp)	206	124

Notes: Entries are basis points and use the cash-Treasury premium and 52-basis-point repricing term from the benchmark 5 percent capacity mapping. The safe-debt margin uses the 330-basis-point wealth-benefit benchmark. Differences of one basis point can arise from unrounded regression coefficients and premium levels.

6.6 From marginal cost to the safe-debt margin

The preceding subsections estimate the marginal cost of producing safe Treasury claims, $S_{b,t}$. Translating that cost into distance from the stationary Laffer peak requires a benchmark for the rate-equivalent wealth benefit of an additional safe claim. Let

$$m_I(r) = \frac{\rho + \lambda_I}{r + \lambda_I}$$

denote group I's MPC out of a persistent income flow at a fixed safe rate. Since group I receives Treasury absorption income, equations (7) and (11) imply that the stationary benefit is

$$B(r) = \frac{1}{\kappa_I(r)} = \frac{\rho + \lambda_I - r}{1 - m_I(r)}. \tag{87}$$

The denominator is the share of persistent absorption income that group I saves. The numerator is the net consumption effect of additional stationary financial wealth. Group H remains important for the level of the safe rate and the size of its adjustment through aggregate desired wealth. It does not enter the stationary turning-point condition because its permanent income is unchanged; at the peak, its desired holdings are therefore unchanged as well. Group I, which receives the absorption income, absorbs the entire marginal claim.

I use $B = 330$ basis points as the benchmark. Its magnitude is guided by a suggestive back-of-the-envelope calculation for households likely to own the financial firms that supply intermediation capacity. The top 0.1 percent receives close to 10 percent of national income, with slightly less than one-third coming from non-capital income (Piketty, Saez, and Zucman, 2018; Saez and Zucman, 2020). Applied to the \$26.36 trillion of national income reported for 2026Q1 (Bureau of Economic Analysis, 2026), these shares imply \$0.83–\$0.88 trillion of non-capital income. The Federal Reserve's Distributional Financial Accounts report \$25.1 trillion of net worth for the top 0.1 percent in 2026Q1 (Federal Reserve Board, 2026). The corresponding non-capital-income-to-net-worth ratio is approximately 330–350 basis points. I use the lower end of this range. This benchmark converts the estimated marginal cost into the safe-debt margin

$$H_t = 330 - S_{b,t} \tag{88}$$

in basis points. The calculation disciplines the scale of the benchmark; the estimates of $S_{b,t}$ come from the preceding subsections.

Under the benchmark nonlinear estimates, the safe-debt margin falls from approximately 250 basis points in 2015Q1 to 143 basis points in 2026Q1. The linear specification gives 277 and 172 basis points at the same dates. The tight 6 percent capacity mapping, combined with the average-maturity duration estimate, places the 2026Q1 margin at approximately 123 basis points.

For a fixed wealth-benefit benchmark,

$$\Delta H_t = -\Delta S_{b,t}. \tag{89}$$

Thus the estimated increase in marginal absorption cost reduces the remaining margin one-for-one. The benchmark calculation indicates that the United States still has a positive safe-debt margin in 2026Q1, but that producing an additional safe Treasury claim has become much more costly: the stock-wide marginal cost has risen from about 80 to 187 basis points over eleven years.

7 Conclusion

Safe public debt is jointly produced. The fiscal authority supplies the promise, but the financial system supplies the market-making, warehousing, and rollover capacity that makes the promise liquid and safe. Solvency therefore leaves open a second limit to issuance: the cost of mobilizing that capacity.

The safe-debt Laffer curve follows from the interaction between this cost and the fiscal adjustment required to finance it. Higher Treasury absorption payments reduce aggregate expenditure after accounting for the private spending generated by the recipients' additional income. The benchmark model captures this net demand drag through lower government purchases and partial saving of absorption income. At low debt, the additional safe claim expands supply by more than the resulting absorption income expands desired safe wealth. As rollover pressure builds, the premium rises and so does the cost of refinancing the stock. The turning point is reached when marginal absorption income generates exactly one additional unit of desired safe wealth. Beyond it, further issuance generates negative stationary aggregate-demand pressure and lowers the equilibrium safe rate, even though the claims remain safe.

The Laffer curve describes debt's stationary contribution to aggregate-demand pressure at a fixed safe rate. The equilibrium safe-rate schedule describes the rate that clears the market across stationary allocations. Their peaks coincide because the same safe-debt margin determines the sign of both effects: the rate-equivalent non-Ricardian wealth benefit of the marginal claim minus its stock-wide marginal absorption cost.

Maturity changes the form of the exposure. Short debt puts more of the stock through rollover; long debt reduces that flow but magnifies the price response of the claims left outstanding. Financial capacity dampens both the direct rollover premium and this duration-price amplification. These are the two pieces measured empirically by the cash-Treasury and real-duration components.

For the United States, the benchmark stock-wide marginal cost more than doubles, rising from about 80 basis points in 2015Q1 to about 187 basis points in 2026Q1. Under the benchmark 330-basis-point wealth benefit, the safe-debt margin falls from about 250 to 143 basis points. The estimated interaction between Treasury supply and dealer-parent capacity has the sign predicted by the model and is statistically significant under overlap-robust inference. In the linear decomposition, premium levels account for 48.2 basis points of the increase, expansion of the private coupon stock raises duration repricing by 55.7 basis points, and cash-Treasury repricing contributes 1.5 basis points.

Financial capacity is shared across issuers. Simultaneous debt expansion by major sovereigns and large corporations can reduce the capacity available to support any one government's safe claims. The relevant pressure is therefore debt relative to shared financial capacity, rather than the debt stock of one sovereign viewed in isolation.

Greater Treasury-intermediation capacity lowers marginal absorption costs and raises the equilibrium safe rate. Holding the rate-equivalent wealth benefit fixed, it also pushes the peak outward. Past the peak, additional debt lowers the equilibrium safe rate; if that rate cannot adjust, the burden appears directly as weaker demand (Caballero and Farhi, 2018). The supply of safe public debt and the financial capacity that sustains its safety are therefore two sides of the same object.

Appendix

A Safe-Asset Demand and Perpetual-Youth Wealth

This appendix derives the household block used in Section 2. It first shows how high aversion to rollover losses leads both household groups to hold supported safe claims. It then derives their consumption and stationary wealth demands. Financial firms supply rollover support at the opportunity cost studied in Section 3; households choose between claims with and without that support.

A.1 Rollover losses and supported claims

Members of group $j \in \{H, I\}$ have Epstein–Zin preferences with unit intertemporal elasticity and risk aversion σ_j across aggregate payoff states (Epstein and Zin, 1989). If ω indexes rollover conditions and $c_{j,i,t}(\omega)$ is state-contingent consumption, the within-date certainty equivalent is

$$C_{j,t}^{\sigma_j} = \left(\mathbb{E}_t[c_{j,i,t}(\omega)^{1-\sigma_j}]\right)^{1/(1-\sigma_j)}. \tag{90}$$

The logarithmic time aggregator delivers the consumption rules used below, while σ_j governs aversion to rollover-state risk.

Consider a household with total government-bond holdings $b_{j,i}$. It chooses the supported position $\iota_{j,i} \in [0, b_{j,i}]$ and hence the unsupported exposure

$$u_{j,i} = b_{j,i} - \iota_{j,i} \geq 0. \tag{91}$$

A supported claim carries the government's gross financing rate r^G and pays the support premium π to the financial firms, leaving the household with the safe return r. An unsupported claim avoids the support payment but incurs the rollover-state loss $\ell(\omega) \geq 0$. Other financial wealth is invested in the safe annuitized asset. The state-contingent wealth law of a surviving household is therefore

$$\dot{w}_{j,i}(\omega) = (r+\lambda_j)w_{j,i} + y_j - c_{j,i}(\omega) + (r^G - r - \pi)\iota_{j,i} + (r^G - r - \ell(\omega))u_{j,i}. \tag{92}$$

The normalized state-price weight induced by (90) is

$$\Lambda_j^{\sigma_j}(\omega) = \frac{c_{j,i}(\omega)^{-\sigma_j}}{\mathbb{E}[c_{j,i}(\omega)^{-\sigma_j}]}. \tag{93}$$

An interior unsupported position satisfies

$$\mathbb{E}\left[\Lambda_j^{\sigma_j}(\omega)(r^G - r - \ell(\omega))\right] = 0, \tag{94}$$

with a weak inequality in the direction of lower unsupported exposure at the corner $u_{j,i} = 0$. A positive supported position satisfies

$$r^G - r - \pi = 0. \tag{95}$$

Holding total government-bond holdings fixed, one more unit of support pays $\ell(\omega)-\pi$. At an interior support choice, the corresponding first-order condition is

$$\mathbb{E}\left[\Lambda_j^{\sigma_j}(\omega)(\ell(\omega)-\pi)\right]=0, \tag{96}$$

with the usual one-sided inequalities at the two portfolio corners.

Let $\ell^{\max}=\operatorname{ess\,sup}_\omega \ell(\omega)$. The benchmark assumes that the support premium is below the worst rollover loss, $\pi<\ell^{\max}$, and studies the limit as σ_j becomes large. For any fixed positive unsupported position, the certainty equivalent converges to consumption in the worst rollover state. Replacing that position with support raises worst-state resources by $\ell^{\max}-\pi>0$ per unit. The optimal unsupported position therefore converges to zero as $\sigma_j\to\infty$. In the limiting allocation, $u_{j,i}=0$ and $\iota_{j,i}=b_{j,i}$. Equation (95) implies $r^G-\pi=r$: households hold government debt as a safe claim, while financial firms receive the Treasury absorption premium for supplying the support that makes it safe. In equilibrium, $\pi=s(b,N)$, and the aggregate payment to the firms is $b\pi=S(b,N)$.

A.2 Individual and aggregate wealth

A member of group $j\in\{H,I\}$ dies with Poisson hazard λ_j, discounts at rate ρ, and has logarithmic utility. An annuity market redistributes the safe-asset holdings of those who die to surviving members of the same group. If $w_{j,t}$ is individual safe financial wealth and $y_{j,t}$ is nonfinancial income, the wealth law of a surviving individual is

$$\dot{w}_{j,t}=(r_t+\lambda_j)w_{j,t}+y_{j,t}-c_{j,t}. \tag{97}$$

Human wealth is

$$h_{j,t}=\int_t^\infty \exp\left[-\int_t^v (r_u+\lambda_j)\,du\right] y_{j,v}\,dv. \tag{98}$$

The consumption rule is therefore

$$c_{j,t}=(\rho+\lambda_j)(w_{j,t}+h_{j,t}). \tag{99}$$

The mortality credit in (97) accrues only to survivors. Aggregating over the whole group, including the wealth removed when members die, gives

$$\dot{a}_{j,t}=r_t a_{j,t}+y_{j,t}-c_{j,t}, \tag{100}$$

where $a_{j,t}$ denotes group j's aggregate safe financial wealth.

A.3 Stationary wealth demand

In a stationary allocation with constant r and y_j,

$$h_j=\frac{y_j}{r+\lambda_j}.$$

Combining this expression with (99) and setting $\dot{a}_j = 0$ in (100) yields

$$a_j = \kappa_j(r) y_j, \qquad \kappa_j(r) \equiv \frac{r-\rho}{(r+\lambda_j)(\rho+\lambda_j-r)}. \tag{101}$$

The relevant stationary equilibrium has $\rho < r < \rho + \lambda_j$. Over this interval,

$$\kappa_j'(r) = \frac{\rho\lambda_j + \lambda_j^2 + (r-\rho)^2}{(r+\lambda_j)^2(\rho+\lambda_j-r)^2} > 0. \tag{102}$$

A higher safe rate raises the desired stock of financial wealth relative to nonfinancial income.

A.4 Fiscal accounting and market clearing

Let $\bar{\tau}$ be aggregate tax revenue. After paying taxes, the two groups receive positive endowment incomes x_H and x_I, with

$$x_H + x_I = 1 - \bar{\tau}. \tag{103}$$

Group I receives aggregate Treasury absorption income $S(b,N) = b\,s(b,N)$. The groups' stationary incomes are consequently

$$y_H = x_H, \qquad y_I = x_I + S(b,N).$$

Because public debt is the private sector's net financial asset, asset-market clearing is

$$b = \kappa_H(r) x_H + \kappa_I(r)\left[x_I + S(b,N)\right]. \tag{104}$$

The government's stationary budget is

$$g = \bar{\tau} - rb - S(b,N). \tag{105}$$

Using (100) in stationarity, $c_j = ra_j + y_j$. Summing across groups, applying asset-market clearing, and using (103) and (105) gives

$$c_H + c_I + g = rb + x_H + x_I + S + \bar{\tau} - rb - S = 1.$$

Thus the asset-market condition and the two group budgets imply goods-market clearing. The accounting traces absorption payments from lower government purchases into the income of financial-firm owners.

Implicit differentiation of (104) gives

$$r_b = \frac{1 - \kappa_I(r) S_b(b,N)}{\kappa_H'(r) x_H + \kappa_I'(r)[x_I + S(b,N)]}, \tag{106}$$

and

$$r_N = \frac{-\kappa_I(r) S_N(b,N)}{\kappa_H'(r) x_H + \kappa_I'(r)[x_I + S(b,N)]} > 0. \tag{107}$$

The denominator is positive. The safe-rate response to debt changes sign when

$$S_b(b,N) = \frac{1}{\kappa_I(r)}.$$

B Threshold Participation and the Rollover Premium

Lemma 1 rests on the signal and participation game developed here. The main-text lemma sets the rollover share to one and calls the resulting premium $s(b,N)$. To cover the maturity extension at the same time, this appendix retains $\mu(\delta) = 1 - e^{-\delta}$ and calls the direct rollover premium $s^R(b,N,\delta)$. On the active stress region $q(b,N,\delta) > \underline{\alpha}$, it is

$$s^R(b,N,\delta) = \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{\min\{q(b,N,\delta),1\}} \Gamma\left(\frac{\mu(\delta)b - \alpha N}{\alpha}\right) d\alpha, \tag{108}$$

where

$$q(b,N,\delta) = \frac{\mu(\delta)b}{N}. \tag{109}$$

When $q(b,N,\delta) \leq \underline{\alpha}$, the inelastic base absorbs rollover exposure and $s^R(b,N,\delta) = 0$. Setting $\mu(\delta) = 1$ in (108) gives (3.4). The derivation first defines active providers' state-contingent costs, then prices them using the marginal participant's belief over participation. Under a flat prior—a constant prior density over the aggregate fundamental—this belief is uniform over active participation states. The private-signal environment follows Morris and Shin (1998).

B.1 Opportunity cost in a rollover-stress state

The government has debt stock b. Over the unit-length rollover interval used in the main text, the rollover share is $\mu(\delta) = 1 - e^{-\delta}$, so the amount that must be refinanced is $\mu(\delta)b$. Normalize the mass of potential rollover support providers to one. A fraction $\underline{\alpha} \in (0,1)$ is inelastic and always participates. Among the remaining providers, the participating share is $\tilde{\alpha} \in [0,1]$, so the active mass is

$$\alpha = \underline{\alpha} + (1-\underline{\alpha})\tilde{\alpha} \in [\underline{\alpha}, 1]. \tag{110}$$

With aggregate financial-system capacity N, the active providers can absorb αN. A rollover-stress shortfall arises when $\mu(\delta)b > \alpha N$, or equivalently when $\alpha < q(b,N,\delta)$. Each active provider is assigned $\max\{\mu(\delta)b - \alpha N, 0\}/\alpha$. The support schedule maps this exposure into a consumption-equivalent flow cost per unit of supported debt:

$$\Gamma(0) = 0, \quad \Gamma'(z) > 0 \text{ for } z > 0, \quad \Gamma'_+(0) > 0, \quad \Gamma''(z) \geq 0. \tag{111}$$

The schedule Γ gives each active provider's opportunity cost from its assigned exposure.

The state-contingent cost is

$$a(b,N,\alpha) = \Gamma\left(\frac{\max\{\mu(\delta)b - \alpha N, 0\}}{\alpha}\right). \tag{112}$$

For $\alpha < q(b,N,\delta)$, the cost is $\Gamma((\mu(\delta)b - \alpha N)/\alpha)$; otherwise it is zero. Low participation raises exposure per active provider. Both the cost and the premium are flow spreads per unit of supported debt; multiplying the premium by the debt stock gives aggregate premium income.

B.2 The participation game

An aggregate fundamental $\theta \in \mathbb{R}$ shifts the remaining, price-responsive providers' willingness or ability to participate; larger θ denotes worse conditions. Provider i observes

$$x_i = \theta + \eta \varepsilon_i, \tag{113}$$

where $\eta > 0$, and ε_i has continuous distribution F with density $f > 0$ on $\mathbb{R}$. Conditional on θ, idiosyncratic noises are independent across providers.

Each elastic provider chooses

$$e_i \in \{0,1\}, \tag{114}$$

where $e_i = 1$ denotes entry. Payoff per unit of supported debt is

$$\Pi_i(e_i, \alpha; b, N) = e_i \left[s^R - a(b,N,\alpha) \right]. \tag{115}$$

A symmetric threshold strategy is a cutoff x^* such that

$$e_i(x_i) = 1 \quad \Longleftrightarrow \quad x_i \leq x^*. \tag{116}$$

Because higher signals indicate worse conditions, entry occurs below the cutoff. Elastic participation is

$$\tilde{\alpha}(\theta; x^*) = F\left(\frac{x^* - \theta}{\eta} \right), \tag{117}$$

and total participation is

$$\alpha(\theta; x^*) = \underline{\alpha} + (1 - \underline{\alpha}) F\left(\frac{x^* - \theta}{\eta} \right). \tag{118}$$

B.3 Threshold pricing and price invariance

Given a candidate threshold x^*, the expected payoff from participation after signal x_i is

$$V(x_i; x^*, s^R, b, N) = s^R - \mathbb{E}\left[a\big(b, N, \alpha(\theta; x^*)\big) \mid x_i \right]. \tag{119}$$

The marginal provider is indifferent when

$$V(x^*; x^*, s^R, b, N) = 0, \tag{120}$$

or equivalently

$$s^R = \mathbb{E}\left[a\big(b, N, \alpha(\theta; x^*)\big) \mid x_i = x^* \right]. \tag{121}$$

Because $a(b,N,\alpha)$ falls with α and $\alpha(\theta; x^*)$ falls with θ, the payoff advantage of entry falls monotonically with the signal. This is the single-crossing property used below.

The following argument proves Lemma 1 and establishes that the common price supports every monotone threshold equilibrium.

Proof of Lemma 1. Under the flat prior, the posterior density of θ conditional on observing the marginal signal $x_i = x^*$ is proportional to

$$f\left(\frac{x^*-\theta}{\eta}\right). \tag{122}$$

This posterior is proper because

$$\int_{-\infty}^{\infty} f\left(\frac{x^*-\theta}{\eta}\right) d\theta = \eta \int_{-\infty}^{\infty} f(u)\, du = \eta.$$

Define

$$z = F\left(\frac{x^*-\theta}{\eta}\right). \tag{123}$$

Then $z \in [0,1]$ and

$$\frac{dz}{d\theta} = -\frac{1}{\eta} f\left(\frac{x^*-\theta}{\eta}\right). \tag{124}$$

The Jacobian from θ to z offsets the likelihood, leaving transformed density proportional to

$$f\left(\frac{x^*-\theta}{\eta}\right)\left|\frac{d\theta}{dz}\right| = \eta,$$

which is constant. Thus $z \mid x_i = x^*$ is uniform on $[0,1]$. Since

$$\alpha = \underline{\alpha} + (1-\underline{\alpha})z, \tag{125}$$

α is uniform on $[\underline{\alpha},1]$, independently of x^*. The zero-profit premium is therefore the uniform average of the state-contingent opportunity cost, as in (108); setting $\mu(\delta)=1$ gives (1). This average does not depend on the cutoff. Every cutoff consequently satisfies marginal indifference, and single crossing makes agents below it enter and those above it stay out. Each cutoff is therefore a monotone threshold equilibrium with the same premium. □

Substituting (112) into the uniform-average pricing equation gives

$$s^R(b,N,\delta) = \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{1} \Gamma\left(\frac{\max\{\mu(\delta)b-\alpha N, 0\}}{\alpha}\right) d\alpha \tag{126}$$

$$= \frac{1}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{\min\{q(b,N,\delta),1\}} \Gamma\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right) d\alpha, \tag{127}$$

This is the finite-rollover-intensity counterpart of (3.4).

B.4 Interpretation

The resulting $s^R(b,N,\delta)$ is the marginal provider's zero-profit compensation for $a(b,N,\alpha)$. Low participation concentrates a given rollover shortfall on fewer balance sheets, and the marginal provider averages their opportunity costs uniformly over $[\underline{\alpha},1]$. In the short-maturity benchmark, $s^R = s$; because higher b changes the premium on the rolled-over stock, marginal fiscal cost contains both s and the stock-wide term bs_b.

C Comparative Statics of the Rollover Premium

To obtain the comparative statics, differentiate the direct rollover premium $s^R(b,N,\delta)$ derived in Appendix B. The short-maturity benchmark sets $\mu(\delta)=1$, in which case $s^R=s$ and the formulas below establish the results used in Sections 3 and 4. On the active region $\underline{\alpha}<q(b,N,\delta)<1$, exposure per active provider in state α is $(\mu(\delta)b-\alpha N)/\alpha$. This exposure is zero at the boundary $\alpha=q$, and $\Gamma(0)=0$, so the changing upper integration limit contributes nothing to first derivatives:

$$s_b^R(b,N,\delta)=\frac{\mu(\delta)}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)}\frac{1}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha>0, \tag{128}$$

$$s_N^R(b,N,\delta)=-\frac{1}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha<0. \tag{129}$$

Debt raises stress in every active state; capacity lowers it.

The fiscal cost of the direct rollover component is

$$S^R(b,N,\delta)=bs^R(b,N,\delta), \tag{130}$$

so

$$S_b^R(b,N,\delta)=s^R(b,N,\delta)+bs_b^R(b,N,\delta). \tag{131}$$

In steady-state rollover accounting, the marginal unit adds s^R and raises the cost of refinancing the stock by bs_b^R. In the short-maturity benchmark these are simply s and bs_b.

The curvature result follows from differentiating (128). The interior derivative is

$$\frac{\partial}{\partial b}\left[\frac{\mu(\delta)}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)\right]=\frac{\mu(\delta)^2}{\alpha^2}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right), \tag{132}$$

and the term from the changing upper integration limit is

$$\frac{\mu(\delta)}{1-\underline{\alpha}}q_b(b,N,\delta)\frac{\Gamma'_+(0)}{q(b,N,\delta)}. \tag{133}$$

Hence

$$\begin{aligned}s_{bb}^R(b,N,\delta)=\frac{1}{1-\underline{\alpha}}\Bigg[&\mu(\delta)^2\int_{\underline{\alpha}}^{q(b,N,\delta)}\frac{1}{\alpha^2}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha\\&+\mu(\delta)q_b(b,N,\delta)\frac{\Gamma'_+(0)}{q(b,N,\delta)}\Bigg]\geq 0.\end{aligned} \tag{134}$$

Within-state curvature accounts for the integral term, while the boundary term reflects expansion of the set of stress states. Since

$$S_{bb}^R(b,N,\delta)=2s_b^R(b,N,\delta)+bs_{bb}^R(b,N,\delta), \tag{135}$$

it follows that

$$S_{bb}^R(b,N,\delta)>0. \tag{136}$$

The cross-derivative with respect to debt and financial capacity uses

$$S_{bN}^R(b,N,\delta) = s_N^R(b,N,\delta) + b s_{bN}^R(b,N,\delta). \tag{137}$$

Differentiating (128) with respect to N, the interior derivative is

$$\frac{\partial}{\partial N}\left[\frac{\mu(\delta)}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)\right] = -\frac{\mu(\delta)}{\alpha}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right), \tag{138}$$

and the term from the changing upper integration limit is

$$\frac{\mu(\delta)}{1-\underline{\alpha}} q_N(b,N,\delta)\frac{\Gamma'_+(0)}{q(b,N,\delta)}. \tag{139}$$

Because

$$q_N(b,N,\delta) = -\frac{1}{N}q(b,N,\delta), \tag{140}$$

this gives

$$\begin{aligned} s_{bN}^R(b,N,\delta) = &-\frac{\mu(\delta)}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)}\frac{1}{\alpha}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha \\ &-\frac{\mu(\delta)}{(1-\underline{\alpha})N}\Gamma'_+(0) < 0. \end{aligned} \tag{141}$$

Together with $s_N^R < 0$, this gives

$$S_{bN}^R(b,N,\delta) < 0. \tag{142}$$

Hence capacity lowers the marginal fiscal cost of safe-debt production.

Finally, $\mu'(\delta) = e^{-\delta}$. In the active stress region,

$$s_\delta^R(b,N,\delta) = \frac{e^{-\delta}b}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)}\frac{1}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha > 0, \tag{143}$$

and

$$\begin{aligned} s_{b\delta}^R(b,N,\delta) = &\frac{e^{-\delta}}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)}\frac{1}{\alpha}\Gamma'\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha \\ &+\frac{\mu(\delta)e^{-\delta}b}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)}\frac{1}{\alpha^2}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha + \frac{e^{-\delta}\Gamma'_+(0)}{1-\underline{\alpha}}. \end{aligned} \tag{144}$$

The first two terms hold the set of stress states fixed; the last comes from its changing boundary. Since $S_b^R = s^R + b s_b^R$,

$$S_{b\delta}^R(b,N,\delta) = s_\delta^R(b,N,\delta) + b s_{b\delta}^R(b,N,\delta) > 0. \tag{145}$$

This is the rollover-intensity comparative static used in Section 5. When $\mu(\delta) = 1$, the superscripts can be dropped and the same derivations give $s_b > 0$, $s_N < 0$, $S_{bb} > 0$, and $S_{bN} < 0$ in the short-maturity benchmark.

D Closed-Form Illustration

The linear stress-cost schedule admits a closed form and generates the figures in the main text. It also illustrates the comparative statics in Appendix C.

With active capacity αN, assume the linear support-cost schedule

$$\Gamma(z) = \varphi z, \qquad \varphi > 0. \tag{146}$$

The parameter φ is the support cost per unit of stress. The stress threshold is the debt level at which the inelastic base just absorbs rollover exposure. The benchmark figures set $\mu(\delta) = 1$, whereas the formulas retain finite δ:

$$\widehat{b}(N) = \frac{\underline{\alpha} N}{\mu(\delta)}. \tag{147}$$

For $b \leq \widehat{b}(N)$, no stress state is reached and $s^R(b,N,\delta) = 0$. For

$$\widehat{b}(N) < b < \frac{N}{\mu(\delta)}, \tag{148}$$

the interior region satisfies $\underline{\alpha} < q(b,N,\delta) < 1$. Substituting the linear stress-cost schedule into (127) gives

$$\begin{aligned}
s^R(b,N,\delta) &= \frac{\varphi}{1-\underline{\alpha}} \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{\mu(\delta)b - \alpha N}{\alpha} \, d\alpha \\
&= \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta) b \int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha} \, d\alpha - N \int_{\underline{\alpha}}^{q(b,N,\delta)} 1 \, d\alpha \right] \\
&= \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta) b \log\left(\frac{\mu(\delta) b}{\underline{\alpha} N} \right) - \mu(\delta) b + \underline{\alpha} N \right].
\end{aligned} \tag{149}$$

The derivative with respect to debt is obtained term by term:

$$\begin{aligned}
s^R_b(b,N,\delta) &= \frac{\varphi}{1-\underline{\alpha}} \left[\mu(\delta) \log\left(\frac{\mu(\delta) b}{\underline{\alpha} N} \right) + \mu(\delta) - \mu(\delta) \right] \\
&= \frac{\varphi \mu(\delta)}{1-\underline{\alpha}} \log\left(\frac{\mu(\delta) b}{\underline{\alpha} N} \right) > 0.
\end{aligned} \tag{150}$$

The middle $+\mu(\delta)$ cancels the derivative of $-\mu(\delta)b$. The second derivative is

$$s^R_{bb}(b,N,\delta) = \frac{\varphi \mu(\delta)}{(1-\underline{\alpha}) b} > 0. \tag{151}$$

The derivative with respect to financial capacity is

$$s^R_N(b,N,\delta) = -\frac{\varphi}{1-\underline{\alpha}} \left[q(b,N,\delta) - \underline{\alpha} \right] < 0. \tag{152}$$

The marginal safety-production cost is

$$
\begin{aligned}
S^R_b(b,N,\delta) &= s^R(b,N,\delta) + bs^R_b(b,N,\delta) \\
&= \frac{\varphi}{1-\underline{\alpha}}\left[\mu(\delta)b\log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right) - \mu(\delta)b + \underline{\alpha}N\right] + \frac{\varphi}{1-\underline{\alpha}}\left[\mu(\delta)b\log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right)\right] \\
&= \frac{\varphi}{1-\underline{\alpha}}\left[2\mu(\delta)b\log\left(\frac{\mu(\delta)b}{\underline{\alpha}N}\right) - \mu(\delta)b + \underline{\alpha}N\right]. \qquad (153)
\end{aligned}
$$

The log term appears once through the average premium and once through stock-wide repricing. When rollover exposure exceeds full inelastic-support capacity, $b \geq N/\mu(\delta)$, the upper integration limit is one and

$$
S^R_b(b,N,\delta) = \frac{\varphi}{1-\underline{\alpha}}\left[2\mu(\delta)b\log\left(\frac{1}{\underline{\alpha}}\right) - N(1-\underline{\alpha})\right], \qquad (154)
$$

which grows without bound as $b \to \infty$.

In the interior stress region, the debt-capacity cross-derivative is

$$
S^R_{bN}(b,N,\delta) = \frac{\varphi}{1-\underline{\alpha}}\left[-\frac{2\mu(\delta)b}{N} + \underline{\alpha}\right] < 0. \qquad (155)
$$

Because $\mu(\delta)b > \underline{\alpha}N$, the sign is negative. When even full participation leaves a shortfall—the saturated region—$S^R_{bN} = -\varphi < 0$. Capacity therefore lowers marginal cost throughout the stress region.

In the short-maturity benchmark, the direct rollover premium is the total premium. If its Laffer peak lies in the interior stress region, it solves

$$
\frac{1}{\kappa_I(r(b^L,N))} = \frac{\varphi}{1-\underline{\alpha}}\left[2b^L\log\left(\frac{b^L}{\underline{\alpha}N}\right) - b^L + \underline{\alpha}N\right]. \qquad (156)
$$

If the peak is saturated, its right-hand side is (154) at $b = b^L$ and $\mu(\delta) = 1$. Greater N or $\underline{\alpha}$ lowers marginal cost at a given b; holding the rate-equivalent wealth benefit fixed, it pushes the peak outward. Greater unit support cost φ pulls the peak inward. At finite maturity, rollover intensity raises the direct marginal cost in (153); the total Laffer peak must also include the duration-price component from Section 5.

E Duration-Price Amplification and Maturity Choice

The capacity and maturity comparative statics below complete the analysis of the total-premium fixed point in (66). The government then chooses rollover intensity δ; higher δ means shorter maturity, faster rollover, and lower duration.

E.1 Rollover component

The direct rollover component is $s^R(b,N,\delta)$ from (57). Because exposure is $\mu(\delta)b$, $s^R_\delta > 0$ in the stress region. Differentiating with respect to N gives

$$s^R_{\delta N}(b,N,\delta) = -\frac{e^{-\delta}b}{1-\underline{\alpha}}\int_{\underline{\alpha}}^{q(b,N,\delta)} \frac{1}{\alpha}\Gamma''\left(\frac{\mu(\delta)b-\alpha N}{\alpha}\right)d\alpha$$
$$-\frac{e^{-\delta}b}{(1-\underline{\alpha})N}\Gamma'_+(0) < 0. \tag{157}$$

Capacity reduces the marginal rollover stress created by shorter maturity both within each state and by shrinking the set of states under stress.

E.2 Duration-price multiplier and fixed-point derivatives

For a Poisson-maturity bond, a unit flow paid while the bond remains outstanding has price and duration $1/[r+s(b,N,\delta)+\delta]$. The same rollover-support risk that generates s^R has a larger price effect when the bond has longer duration. The duration-price component is

$$s^D(b,N,\delta) = \frac{\vartheta(N)}{[r+s(b,N,\delta)+\delta]^2}s^R(b,N,\delta). \tag{158}$$

For the capacity dependence, parameterize

$$\vartheta(N) = \bar{\vartheta}N^{-\eta_\vartheta}, \qquad \bar{\vartheta} > 0, \qquad \eta_\vartheta > 0. \tag{159}$$

where $\bar{\vartheta}$ scales duration-price transmission and η_ϑ is its capacity elasticity. Higher capacity dampens the price effect of a given rollover opportunity cost. Combining (158) with the direct rollover premium gives

$$s = s^R(b,N,\delta)\left[1+\frac{\vartheta(N)}{(r+s+\delta)^2}\right]. \tag{160}$$

On the nonnegative-premium domain, the right-hand side is nonnegative at zero and converges to finite s^R, while the left-hand side is increasing and unbounded. Continuity and monotonicity imply a unique solution.

Implicit differentiation gives the capacity and maturity derivatives in (68) and (69). Because $s^R_N < 0$ and $\vartheta_N < 0$, $s_N < 0$. The maturity derivative balances the amplified increase in rollover exposure from shorter maturity against the reduction in its duration-price multiplier.

Differentiating (69) with respect to N yields the full cross-derivative of the premium with respect to rollover intensity and capacity:

$$s_{\delta N} = \left\{\left[1+\frac{\vartheta}{(r+s+\delta)^2}\right]s^R_{\delta N} + \left[\frac{\vartheta_N}{(r+s+\delta)^2} - \frac{2\vartheta s_N}{(r+s+\delta)^3}\right]s^R_\delta \right.$$
$$\left. -(1+s_\delta)\frac{2s^R\vartheta}{(r+s+\delta)^3}\left[\frac{\vartheta_N}{\vartheta}+\frac{s^R_N}{s^R}-\frac{3s_N}{r+s+\delta}\right]\right\}\left[1+\frac{2s^R\vartheta}{(r+s+\delta)^3}\right]^{-1}. \tag{161}$$

Capacity reduces the incremental rollover exposure created by shorter maturity, making the first term negative. The other terms reflect its direct and endogenous effects on the price multiplier. At an interior maturity optimum, where $s_\delta = 0$, the expression simplifies to

$$\begin{aligned} s_{\delta N} = \Bigg\{ & \left[1+\frac{\vartheta}{(r+s+\delta)^2}\right] s^R_{\delta N} + \left[\frac{\vartheta_N}{(r+s+\delta)^2} - \frac{2\vartheta s_N}{(r+s+\delta)^3}\right] s^R_\delta \\ & - \frac{2s^R\vartheta}{(r+s+\delta)^3}\left[\frac{\vartheta_N}{\vartheta} + \frac{s^R_N}{s^R} - \frac{3s_N}{r+s+\delta}\right]\Bigg\}\left[1+\frac{2s^R\vartheta}{(r+s+\delta)^3}\right]^{-1}. \end{aligned} \tag{162}$$

E.3 Maturity choice

For fixed b, minimizing total service cost bs is equivalent to minimizing s. The government chooses

$$\delta^*(b,N) \in \arg\min_{\delta>0} s(b,N,\delta). \tag{163}$$

At an interior optimum,

$$s_\delta(b,N,\delta^*) = 0, \qquad s_{\delta\delta}(b,N,\delta^*) > 0. \tag{164}$$

Using (69), the first-order condition is

$$\left[1+\frac{\vartheta(N)}{(r+s+\delta)^2}\right] s^R_\delta = \frac{2s^R\vartheta(N)}{(r+s+\delta)^3}. \tag{165}$$

It equates the amplified increase in rollover exposure from shortening maturity with the resulting reduction in duration-price amplification.

Proposition 3 (Maturity choice and financial capacity)**.** *Let $\delta^*(b,N)$ be an interior solution of the maturity problem and suppose the second-order condition in* (164) *holds. If*

$$s_{\delta N}(b,N,\delta^*) < 0, \tag{166}$$

then

$$\frac{d\delta^*(b,N)}{dN} > 0. \tag{167}$$

Thus higher financial capacity permits a higher optimal rollover intensity; equivalently, lower financial capacity lengthens maturity.

The comparative static follows by differentiating the first-order condition:

$$s_{\delta\delta}(b,N,\delta^*)\frac{d\delta^*}{dN} + s_{\delta N}(b,N,\delta^*) = 0. \tag{168}$$

Consequently,

$$\frac{d\delta^*}{dN} = -\frac{s_{\delta N}}{s_{\delta\delta}}. \tag{169}$$

Under the second-order condition, $s_{\delta N} < 0$ implies $d\delta^*/dN > 0$. Higher capacity then permits shorter maturity, while scarcity induces the government to term out.

F Empirical Construction

The data, inference, and sensitivity exercises below support the cash-Treasury and real-duration components in Section 6.

F.1 Data construction

Dealer-parent Treasury capacity comes from FR Y-9C filings and the leverage-ratio mapping in (81). The measure assigns all unused leverage-exposure room to potential Treasury intermediation and adds Treasury inventory already held. It is a broad system-capacity diagnostic. Constant rescalings leave the standardized interaction unchanged. Each leverage-ratio mapping produces its own capacity series, so I standardize the 5 and 6 percent measures separately over 2015–2024.

The 5 percent series applies the historical U.S. GSIB holding-company threshold, and the 6 percent series applies the tighter historical well-capitalized threshold for insured depository subsidiaries. Effective April 1, 2026, a final rule tied U.S. GSIB standards to firm-specific systemic-risk surcharges and capped the depository-subsidiary enhancement at one percentage point (Federal Reserve Board, FDIC, and OCC, 2025). The 2026Q1 diagnostic extends each historical mapping through the final observation, preserving a fixed-rule comparison with the estimation sample.

The monthly interaction state interpolates the quarterly regulatory measure using dealer-parent market equity. Equity supplies the monthly movement and is less mechanically tied to Treasury positions within individual maturity buckets than dealer holdings or sector size. The interpolation preserves each quarterly regulatory mean, so leverage-room data determine the level of N_t.

The cash-Treasury component s^R is the Treasury par yield minus the maturity-matched par swap rate. The two-, three-, five-, ten-, and thirty-year buckets use continuous benchmark-versus-IRS spread chains available from 2003. The seven-year bucket uses a LIBOR-consistent spread constructed from Treasury and swap par yields. This produces a continuous six-bucket panel through the transition in the underlying swap benchmark. Higher s^R denotes a higher cash-Treasury premium.

The duration component s^D is the DKW real term premium (D'Amico, Kim, and Wei, 2018). Jointly pricing nominal and inflation-indexed Treasuries removes inflation-risk compensation within one term-structure model. This distinction matters: a debt-related inflation premium could reflect dilution or fiscal risk, whereas the empirical object is the real price component associated with duration. The DKW series is available at five and ten years; the debt-weighted level interpolates between these anchors and extrapolates beyond them. An alternative subtracts an inflation-risk-premium estimate from the Adrian–Crump–Moench (ACM) nominal term premium.

Table 6 lists the data sources used to construct the debt, capacity, premium, Federal Reserve holdings, and dealer-parent equity series.

Table 6: Data sources for the empirical section

Series	Source	Frequency
Marketable debt, maturing share, and maturity buckets	Monthly Statement of the Public Debt and Treasury Fiscal Data	Monthly
Federal Reserve Treasury holdings by Committee on Uniform Securities Identification Procedures (CUSIP) identifier	Federal Reserve SOMA holdings	Weekly/monthly
Tier 1 capital, leverage exposure, and Treasuries held by dealer parents	FR Y-9C regulatory filings	Quarterly
Treasury and swap par yields, matched by maturity	Treasury and swap-curve data	Daily
Real term premia	DKW estimates (D'Amico, Kim, and Wei, 2018)	Daily
Nominal term premia and robustness checks	Adrian, Crump, and Moench (2013) estimates and inflation-risk-premium estimates	Daily/monthly
Dealer-parent equity values from the Center for Research in Security Prices (CRSP)	CRSP dealer-parent equity data	Daily

F.2 Cash-Treasury repricing details

The cash-Treasury design is a monthly maturity panel. Equations (78) and (82) include maturity and start-month fixed effects and exclude the first half of 2020, focusing on persistent repricing outside the emergency-market episode. Because adjacent $h = 12$–24 changes overlap, inference uses Driscoll–Kraay/Newey–West standard errors with 24 monthly lags. I refer to the resulting confidence bands as overlap-robust bands; the figures also show month-clustered bands to display the inference correction.

The conversion in (79) accounts for percentage-point units and weights each bucket by its private debt. A one-point increase in the SOMA share removes one percent of the bucket's outstanding stock from private balance sheets. The factor $(1-\varpi_{m,t})$ converts this amount into a share of private float, and $w_{m,t} = b_{m,t}/b_t$ aggregates across the maturity distribution of marginal issuance. The sum $\sum_m w_{m,t}(1-\varpi_{m,t})$ is 0.781 in 2015Q1 and 0.837 in 2026Q1.

Start-month effects absorb curve-wide repricing, so $\mathscr{M}^{R}_{h,t}$ measures maturity-specific movements around the Treasury curve. The interaction sample begins in 2015Q1, when regulatory capacity becomes available. The linear estimates use the average supply slope, while the interacted estimates evaluate how that slope changes at the scarce 2026Q1 capacity state.

Figure 4 reports the linear estimates for the 2015+ and 2010+ samples. The average slopes over the persistent $h = 12$–24 window are -0.27 and -0.48, respectively. The overlap-robust bands are wider than the month-clustered bands, especially in the shorter sample, while the negative point-estimate profile is preserved in both samples.

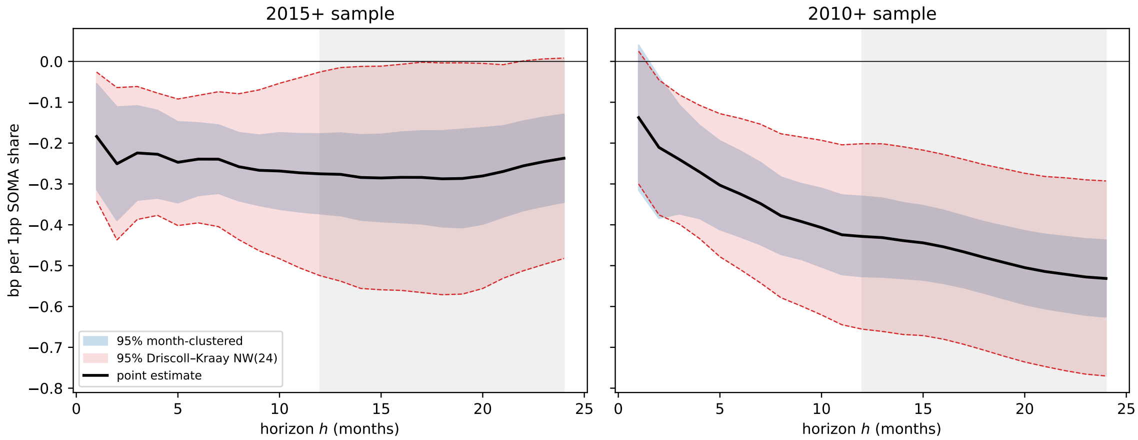

Figure 4: Linear cash-Treasury repricing estimates. The panels report β_h for the 2015+ and 2010+ samples. Black lines are point estimates; blue bands are month-clustered 95 percent confidence intervals; red dashed bands use Driscoll–Kraay/Newey–West inference with 24 monthly lags. The shaded region is the $h = 12$–24 averaging window. The average slopes in that window are -0.27 and -0.48, respectively.

Figure 5 reports the benchmark nonlinear estimates and both sets of confidence bands. The interaction remains positive throughout the persistent window and is statistically significant under overlap-robust inference.

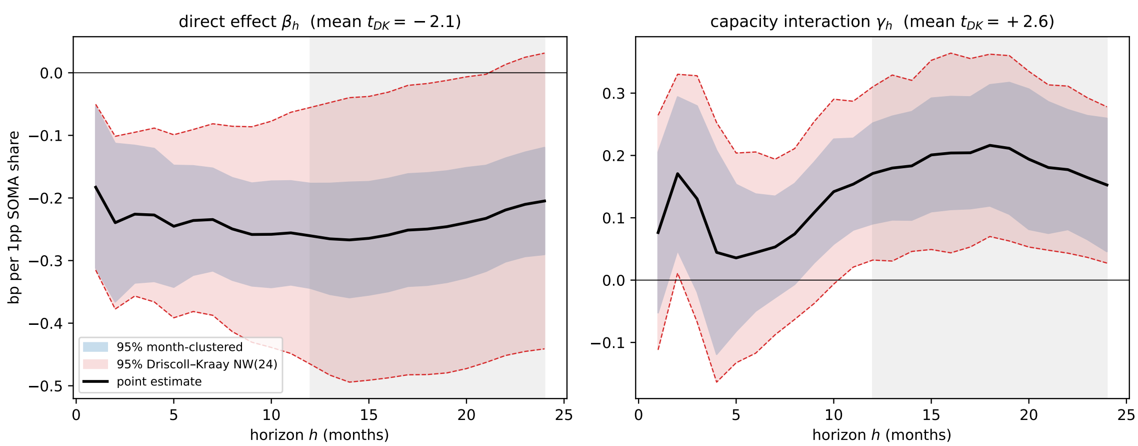

Figure 5: Benchmark nonlinear cash-Treasury repricing estimates. The left panel reports the direct slope β_h, and the right panel reports the capacity interaction γ_h, using the 5 percent leverage mapping. Black lines are point estimates; blue bands are month-clustered 95 percent confidence intervals; red dashed bands use Driscoll–Kraay/Newey–West inference with 24 monthly lags. Over $h = 12$–24, the interaction averages 0.19 and its overlap-robust t-statistics average 2.6.

Figure 6 compares the 5 and 6 percent leverage-ratio mappings. The interaction path is similar under both; the larger 2026Q1 reading under 6 percent chiefly reflects a more extreme standardized capacity state.

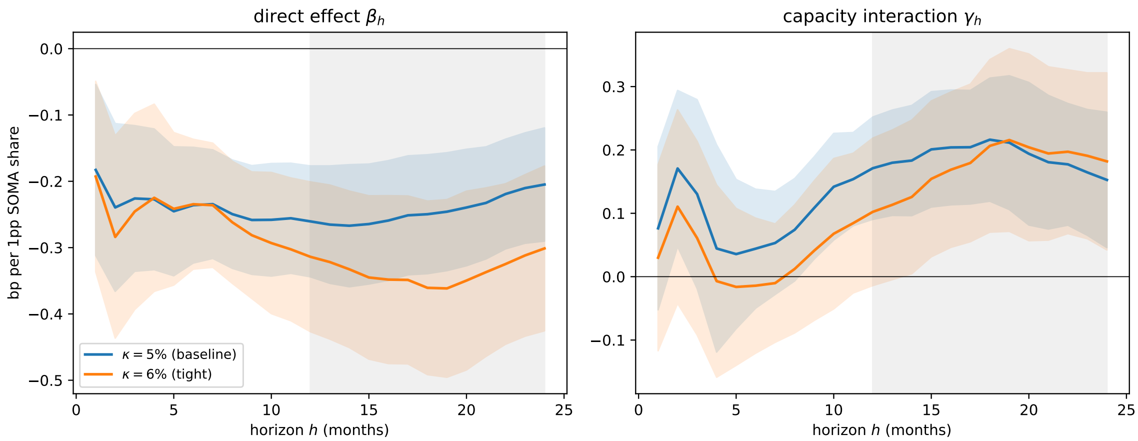

Figure 6: Nonlinear estimates under 5 and 6 percent leverage-ratio mappings. Point estimates and bands are shown for the direct slope β_h and capacity interaction γ_h; the shaded region is the $h = 12$–24 averaging window. The interaction paths are similar, while the standardized 2026Q1 capacity states differ. Evaluated at those states, the implied cash-Treasury repricing terms are approximately 52 and 72 basis points, respectively.

The IV sensitivity specification instruments the SOMA share within each maturity bucket with its $t-6$ level and trailing six-month issuance into the bucket relative to bucket outstanding, a proxy for the on-the-run status of recently issued securities. These variables adapt the pre-program security characteristics in D'Amico and King (2013); maturity effects absorb time-invariant maturity attributes. Lagged holdings and recent issuance can also reflect persistent security-specific pricing conditions, so OLS remains the empirical anchor. A pricing residual from a fitted Nelson–Siegel yield curve is excluded because overidentification tests reject when it is included. Figure 7 reports the linear two-instrument sensitivity specification alongside the linear OLS comparison.

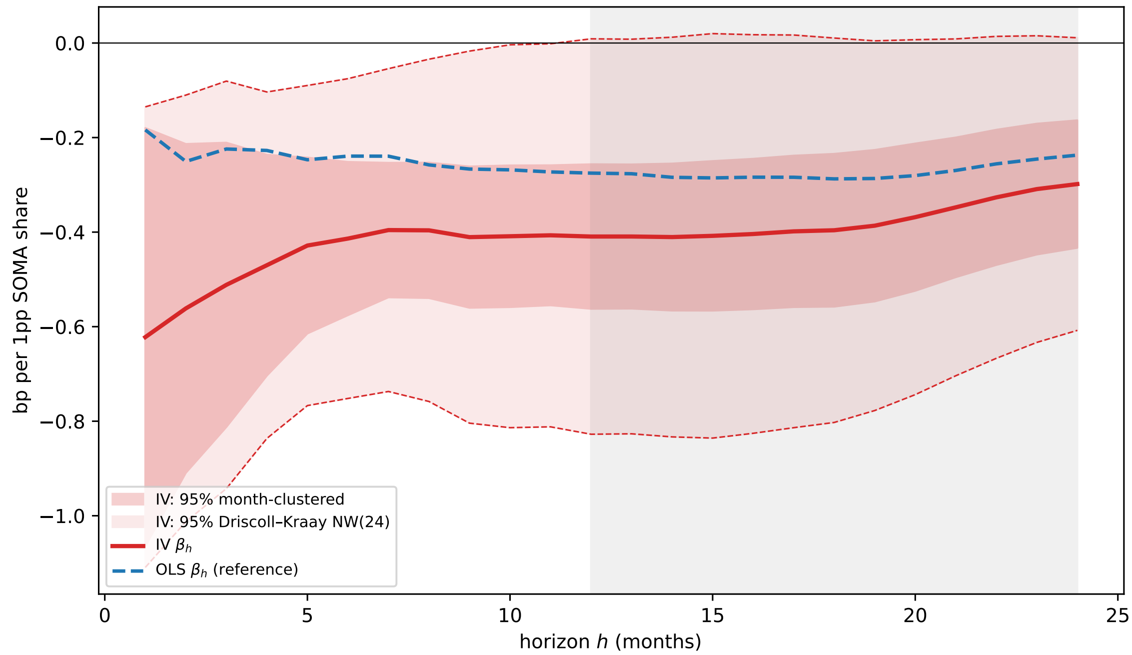

Figure 7: Linear IV sensitivity with overlap-robust inference. The instruments, measured at $t-6$, are lagged maturity-bucket SOMA share and a trailing-issuance on-the-run proxy. The solid red line is the IV estimate and the dashed blue line is the OLS benchmark; the bands show month-clustered and Driscoll–Kraay/Newey–West 95 percent confidence intervals. The shaded region is the $h = 12$–24 averaging window, over which the IV and OLS slopes average -0.37 and -0.27.

F.3 Real-duration repricing details

The duration event study regresses two-trading-day changes in the DKW real term premium on announced changes in expected cumulative Treasury purchases or runoff. The Maturity Extension Program changed SOMA's maturity composition while leaving face amount roughly unchanged, so it belongs to a maturity-composition design rather than the aggregate private-supply experiment in (85).

The baseline with-intercept slope is $\beta_D = -0.0055$ basis points per billion dollars. Applied to the 2026Q1 MSPD private coupon float of \$17.642 trillion, it gives $\mathscr{M}^D_{2026Q1} = 97$ basis points. A ten-year specification gives 116 basis points, while scaling announcements by contemporaneous private float gives 53 basis points. These are alternative specifications of the same twelve-event design; Table 8 reports statistical uncertainty. Table 7 lists the announcements and sizing assumptions.

Table 7: Treasury-relevant asset-purchase and runoff announcements

Date	Event	Amount	Date	Event	Amount
2009-03-18	QE1 Treasury purchases	+300	2013-12-18	Taper start	-200
2010-08-10	QE2 pre-announcement	+300	2020-03-15	COVID QE	+500
2010-11-03	QE2	+600	2020-03-23	COVID open-ended QE	+1000
2012-12-12	QE3	+540	2021-11-03	Taper	-120
2013-05-22	Taper hint	-300	2022-05-04	QT plan	-540
2013-06-19	Taper timeline	-200	2024-05-01	QT slowdown	+210

Notes: Amounts are announced changes in expected cumulative Treasury purchases, in billions of dollars, with negative values denoting tapering or balance-sheet runoff. QE denotes quantitative easing; QT denotes quantitative tightening. The Maturity Extension Program left the face amount of Treasury supply approximately unchanged and therefore belongs to a maturity-composition design. Open-ended and pace-change announcements require judgment in translating statements into cumulative amounts.

Table 8 reports HC3 heteroskedasticity-robust standard errors with small-sample t critical values. The baseline slope barely excludes zero at the 5 percent level, producing a wide stock-level confidence interval. Removing the two March 2020 announcements increases the point estimate, while the leave-one-event-out estimates bracket the baseline. The point estimate is stable to removing any single announcement.

Table 8: Duration event-study inference

Specification	β_D	HC3 standard error	95 percent interval	$\mathscr{M}^{D}_{2026Q1}$
Baseline, with intercept	-0.0055	0.0024	$[-0.0109, -0.0001]$	97 $[1, 193]$
Excluding March 2020	-0.0086	0.0051	$[-0.0204, 0.0032]$	151
Leave-one-event-out range	$[-0.0066, -0.0042]$	—	—	$[74, 116]$

Notes: Slopes are basis points per billion dollars; repricing terms are basis points. The baseline and March-2020 intervals use HC3 standard errors and t critical values with 10 and 8 degrees of freedom. The baseline intercept is 0.10 basis points with an HC3 standard error of 1.57. The leave-one-event-out row reports ranges of point estimates, not confidence intervals.

References

Aguiar, Mark, and Manuel Amador (2014). "Sovereign Debt." In Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Volume 4. Elsevier, 647–687.

Alesina, Alberto, and Roberto Perotti (1997). "Fiscal Adjustments in OECD Countries: Composition and Macroeconomic Effects." *IMF Staff Papers*, 44(2), 210–248.

Angeletos, George-Marios (2002). "Fiscal Policy with Noncontingent Debt and the Optimal Maturity Structure." *Quarterly Journal of Economics*, 117(3), 1105–1131.

Angeletos, George-Marios, Chen Lian, and Christian K. Wolf (2024). "Can Deficits Finance Themselves?" *Econometrica*, 92(5), 1351–1390.

Arellano, Cristina (2008). "Default Risk and Income Fluctuations in Emerging Economies." *American Economic Review*, 98(3), 690–712.

Auclert, Adrien, Hannes Malmberg, Matthew Rognlie, and Ludwig Straub (2025). "The Race Between Asset Supply and Asset Demand." NBER Working Paper No. 34470. Prepared for the Federal Reserve Bank of Kansas City's Jackson Hole Economic Policy Symposium.

Barro, Robert J. (1974). "Are Government Bonds Net Wealth?" *Journal of Political Economy*, 82(6), 1095–1117.

Blanchard, Olivier J. (1985). "Debt, Deficits, and Finite Horizons." *Journal of Political Economy*, 93(2), 223–247.

Brunnermeier, Markus K., Sebastian Merkel, and Yuliy Sannikov (2024). "Safe Assets." *Journal of Political Economy*, 132(11), 3603–3657.

Bureau of Economic Analysis (2026). "National Income by Type of Income." National Income and Product Accounts, Table 1.12, 2026Q1 release.

Buera, Francisco, and Juan Pablo Nicolini (2004). "Optimal Maturity of Government Debt without State Contingent Bonds." *Journal of Monetary Economics*, 51(3), 531–554.

Caballero, Ricardo J. (2006). "On the Macroeconomics of Asset Shortages." In Andreas Beyer and Lucrezia Reichlin, eds., *The Role of Money: Money and Monetary Policy in the Twenty-First Century*, 272–283. Frankfurt: European Central Bank.

Caballero, Ricardo J. (2015). "A Caricature (Model) of the World Economy." In Ricardo J. Caballero and Klaus Schmidt-Hebbel, eds., *Economic Policies in Emerging-Market Economies: Festschrift in Honor of Vittorio Corbo*, Central Banking, Analysis, and Economic Policies Book Series, vol. 21, 61–77. Santiago: Central Bank of Chile.

Caballero, Ricardo J., and Emmanuel Farhi (2018). "The Safety Trap." *Review of Economic Studies*, 85(1), 223–274.

Caballero, Ricardo J., and Alp Simsek (2020). “A Risk-Centric Model of Demand Recessions and Speculation.” *Quarterly Journal of Economics*, 135(3), 1493–1566.

Caballero, Ricardo J., and Alp Simsek (2021). “A Model of Endogenous Risk Intolerance and LSAPs: Asset Prices and Aggregate Demand in a ‘COVID-19’ Shock.” *Review of Financial Studies*, 34(11), 5522–5580.

Caballero, Ricardo J., and Alp Simsek (2023). “A Monetary Policy Asset Pricing Model.” NBER Working Paper No. 30132, revised March 2023.

Caballero, Ricardo J., Tomás E. Caravello, and Alp Simsek (2024). “Financial Conditions Targeting.” NBER Working Paper No. 33206.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas (2008). “An Equilibrium Model of ‘Global Imbalances’ and Low Interest Rates.” *American Economic Review*, 98(1), 358–393.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas (2016). “Safe Asset Scarcity and Aggregate Demand.” *American Economic Review: Papers and Proceedings*, 106(5), 513–518.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas (2017). “The Safe Assets Shortage Conundrum.” *Journal of Economic Perspectives*, 31(3), 29–46.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas (2021). “Global Imbalances and Policy Wars at the Zero Lower Bound.” *Review of Economic Studies*, 88(6), 2570–2621.

Calvo, Guillermo A. (1988). “Servicing the Public Debt: The Role of Expectations.” *American Economic Review*, 78(4), 647–661.

Cole, Harold L., and Timothy J. Kehoe (2000). “Self-Fulfilling Debt Crises.” *Review of Economic Studies*, 67(1), 91–116.

Adrian, Tobias, Richard K. Crump, and Emanuel Moench (2013). “Pricing the Term Structure with Linear Regressions.” *Journal of Financial Economics*, 110(1), 110–138.

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

D’Amico, Stefania, Don H. Kim, and Min Wei (2018). “Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices.” *Journal of Financial and Quantitative Analysis*, 53(1), 395–436.

D’Avernas, Adrien, and Quentin Vandeweyer (2024). “Treasury Bill Shortages and the Pricing of Short-Term Assets.” *Journal of Finance*, 79(6), 4083–4141.

Du, Wenxin, Benjamin Hébert, and Wenhao Li (2023). "Intermediary Balance Sheets and the Treasury Yield Curve." *Journal of Financial Economics*, 150(3), 103722.

Duffie, Darrell (2020). "Still the World's Safe Haven? Redesigning the U.S. Treasury Market After the COVID-19 Crisis." Brookings Institution Working Paper.

Duffie, Darrell, and Frank M. Keane (2023). "Market-Function Asset Purchases." Federal Reserve Bank of New York Staff Report No. 1054.

Eaton, Jonathan, and Mark Gersovitz (1981). "Debt with Potential Repudiation: Theoretical and Empirical Analysis." *Review of Economic Studies*, 48(2), 289–309.

Epstein, Larry G., and Stanley E. Zin (1989). "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 57(4), 937–969.

Federal Reserve Board (2026). "Distribution of Household Wealth in the U.S. since 1989." Distributional Financial Accounts, 2026Q1 release. https://www.federalreserve.gov/releases/z1/dataviz/dfa/distribute/table/index.html.

Faraglia, Elisa, Albert Marcet, Rigas Oikonomou, and Andrew Scott (2019). "Government Debt Management: The Long and the Short of It." *Review of Economic Studies*, 86(6), 2554–2604.

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

Federal Reserve Board, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency (2025). "Agencies Issue Final Rule to Modify Certain Regulatory Capital Standards." Joint press release, November 25. https://www.federalreserve.gov/newsevents/pressreleases/bcreg20251125b.htm.

Fornaro, Luca, and Martin Wolf (2025). "Fiscal Stagnation." CEPR Discussion Paper No. 20149.

Gabaix, Xavier, and Matteo Maggiori (2015). "International Liquidity and Exchange Rate Dynamics." *Quarterly Journal of Economics*, 130(3), 1369–1420.

Gamber, Edward, and John Seliski (2019). "The Effect of Government Debt on Interest Rates." Congressional Budget Office Working Paper 2019-01.

Giavazzi, Francesco, and Marco Pagano (1990). "Can Severe Fiscal Contractions Be Expansionary? Tales of Two Small European Countries." *NBER Macroeconomics Annual*, 5, 75–122.

Gómez-Cram, Roberto, Howard Kung, and Hanno Lustig (2024). "Government Debt in Mature Economies: Safe or Risky?" Federal Reserve Bank of Kansas City Economic Policy Symposium, Jackson Hole.

Gómez-Cram, Roberto, Howard Kung, and Hanno Lustig (2025). "Can U.S. Treasury Markets Add and Subtract?" NBER Working Paper No. 33604.

Greenwood, Robin, Samuel G. Hanson, and Jeremy C. Stein (2015). "A Comparative-Advantage Approach to Government Debt Maturity." *Journal of Finance*, 70(4), 1683–1722.

Gruber, Joseph W., and Steven B. Kamin (2012). "Fiscal Positions and Government Bond Yields in OECD Countries." *Journal of Money, Credit and Banking*, 44(8), 1563–1587.

He, Zhiguo, Arvind Krishnamurthy, and Konstantin Milbradt (2019). "A Model of Safe Asset Determination." *American Economic Review*, 109(4), 1230–1262.

He, Zhiguo, Stefan Nagel, and Zhaogang Song (2022). "Treasury Inconvenience Yields During the COVID-19 Crisis." *Journal of Financial Economics*, 143(1), 57–79.

Itskhoki, Oleg, and Dmitry Mukhin (2026). "Global Imbalances: A Progress Report." In Hélène Rey, Beatrice Weder di Mauro, and Jeromin Zettelmeyer (eds.), *P398 Paris Report 4: The New Global Imbalances*, CEPR Press, Paris and London.

Jiang, Zhengyang, Robert J. Richmond, and Tony Zhang (2025). "Convenience Lost." NBER Working Paper No. 33940.

Kashyap, Anil K., Jeremy C. Stein, Jonathan L. Wallen, and Joshua Younger (2025). "Treasury Market Dysfunction and the Role of the Central Bank." *Brookings Papers on Economic Activity*, Spring, 221–265.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen (2012). "The Aggregate Demand for Treasury Debt." *Journal of Political Economy*, 120(2), 233–267.

Kruttli, Mathias S., Phillip J. Monin, Lubomir Petrasek, and Sumudu W. Watugala (2021). "Hedge Fund Treasury Trading and Funding Fragility: Evidence from the COVID-19 Crisis." Finance and Economics Discussion Series 2021-038, Board of Governors of the Federal Reserve System.

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

Li, Wenhao, and Sebastian A. Merkel (2026). "Quantitative Easing and Government Debt Sustainability." NBER Working Paper No. 35421.

Maggiori, Matteo (2017). "Financial Intermediation, International Risk Sharing, and Reserve Currencies." *American Economic Review*, 107(10), 3038–3071.

Mian, Atif, Ludwig Straub, and Amir Sufi (2021). "Indebted Demand." *Quarterly Journal of Economics*, 136(4), 2243–2307.

Mian, Atif, Ludwig Straub, and Amir Sufi (2025). "A Goldilocks Theory of Fiscal Deficits." *American Economic Review*, 115(12), 4253–4291.

Morris, Stephen, and Hyun Song Shin (1998). "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review*, 88(3), 587–597.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). "Distributional National Accounts: Methods and Estimates for the United States." *Quarterly Journal of Economics*, 133(2), 553–609.

Saez, Emmanuel, and Gabriel Zucman (2020). "The Rise of Income and Wealth Inequality in America: Evidence from Distributional Macroeconomic Accounts." *Journal of Economic Perspectives*, 34(4), 3–26.

Schrimpf, Andreas, Hyun Song Shin, and Vladyslav Sushko (2020). "Leverage and Margin Spirals in Fixed Income Markets During the Covid-19 Crisis." BIS Bulletin No. 2.

Weil, Philippe (1989). "Overlapping Families of Infinitely-Lived Agents." *Journal of Public Economics*, 38(2), 183–198.

World Bank (2026). *Global Economic Prospects, June 2026*. Washington, DC: World Bank.